CHAPTER:
Changing Utility Functions and Two-System Economic Models

Hersh Shefrin

Santa Clara University

January 2025

Draft chapter for second edition of *Handbook of Utility Theory*, to be published by Springer Nature.[1]

[1]I wish to thank the editors and two anonymous referees for their helpful remarks in the preparation of the version of this chapter which appeared in the first edition of this volume; and I also thank Orazio Attanasio, Nick Barberis, Shlomo Benartzi, Johannes Boehm, Etienne Fize, Drew Fudenberg, Jaime Gimeno-Ribes, Tobin Hanspal, Takashi Hayashi, Martin Holm, Xavier Jaravel, Greg Kaplan, Peter Maxted, Gisle Natvik, Abigail Sussman, Norio Takeoka, Richard Thaler, Barbara Tversky, and Alexander Zentefis. I especially thank the editors of the second edition of this volume, Ani Guerdjikova and Marcus Pivato, for valuable comments.

1 Introduction and Outline

When economists speak of changing utility functions, they typically mean that over time people appear to change their minds. People might do so by putting off until tomorrow what they could do today, but when tomorrow arrives continuing to procrastinate. They might also form contingency plans about risks they would accept tomorrow, but change their minds tomorrow as the contingencies play out. Both of these behaviors exhibit "dynamic inconsistency." This chapter focuses on economic frameworks for analyzing dynamic inconsistency.

Intertemporal choice concerns the tradeoffs faced by an individual between consumption at different points of time. The standard way to evaluate such tradeoffs is by means of a utility function. This chapter addresses issues which arise when the utility function changes over time. Changes to utility functions can occur in conjunction with the instantaneous utility functions, as well as the function which evaluates the entire intertemporal profile.

This chapter deals with three ways in which utility functions can change over time:

1. changes in respect to *intertemporal* substitution that involve overconsumption associated with present bias, the excessive weighting of the present relative to the future;

2. changes in respect to substitution between *different commodities*; and

3. changes in respect to substitution between *alternative risky prospects*.

A core notion underlying all three is the fact that changes take place in a moving frame of reference. In this respect, there are two central issues. The first issue is to identify the characteristics which define the frame of reference. The second issue concerns a *relativity* principle. Relativity requires that relative relationships among events appear to be preserved as the frame of reference moves. The change to the utility function can be said to be germane when the relativity principle is violated.[2]

The following questions underlie the discussion throughout the chapter. As an empirical matter, do utility functions change over time? If they do change, then does the relativity principle hold? If the relativity principle is violated, then what are the positive implications associated with the way that individuals make intertemporal choices? And what are the welfare implications attached to their choices, given that there may well be conflicts between the preferences held by a single individual at different points of his or her life?

The previous paragraph articulates the defining issues and questions addressed in the first edition of this handbook chapter. The issues associated with present bias were analyzed using two approaches, a game theoretic approach and a behavioral approach. The issues associated with substitution between commodities were analyzed using an approach focusing on habituation to past consumption levels. The issues associated with risk were analyzed using the approach known as prospect theory.

This revised version discusses issues and questions which have arisen in the three decades separating the first and second editions of this handbook. In respect to present bias, what is new in the second version is a comparative discussion of the game theoretic approach and the behavioral approach. This discussion focuses on the degree to which the empirical evidence which has emerged during the three decades separating the two editions supports various aspects of the two approaches. In respect to substitution between commodities and habituation, what is new in the second edition is a discussion of how the theoretical behavioral approach extends to a multicommodity setting, and how the presence of many commodities in theoretical models bears on the issues raised by economists developing the game theoretic approach. In respect to issues of risk, the second edition introduces an alternative approach to capturing the essential features of prospect theory.

The game theoretic approach to present bias is neoclassical, meaning decision mak-

[2]The terms frame of reference and relativity are borrowed from physics. Although they can be interpreted in analagous fashion, the parallel should not be taken too far. For instance, there need be no counterpart to the Lorentz transformations, which in physics, reconcile variations across all frames of reference.

ers are assumed to act as rational agents. The approach is elegant and intellectually interesting. That said, it is important to be mindful, especially in a volume on utility theory, that the purpose of theory is to shed insight about how the world works, not to do theory for theory's sake. This means that theorists need to pay careful attention to empirical evidence, rather than rely on a mix of stylized facts and intuition. In this regard, there are good reasons for concluding that the neoclassical perspective on changing utility functions, which treat individuals as utility maximizers, is at odds with the manner in which humans make decisions, not to mention the decisions themselves. The new material in this second edition includes an extensive section about empirical evidence. Notably, the evidence supports the contention that the behavioral approach more closely captures the manner in which humans make decisions and the actual behavior generated by these decisions, than does the game theoretic approach.

Since the publication of the first edition of this handbook, two Nobel prizes have been awarded for behavioral economics, one to Daniel Kahneman and the other to Richard Thaler. The Nobel committee emphasized prospect theory in the 2002 award to Kahneman, and emphasized the two-system behavioral approach in the 2017 award to Thaler.

The first formal economic model of a two-system approach to individual decision making appears in Thaler and Shefrin (1981). Subsequently, Kahneman used the two-system framework to organize the impact of psychology on general decision making. Thaler and Shefrin use the terms "doer" and"planner" for the two systems. Kahneman's terminology features "fast" and "slow" thinking. Fast thinking System 1 pertains to subconscious, automatic mental processes, which mostly feel effortless. Slow thinking System 2 pertains to conscious, deliberative processes which feel as if they require effort.[3]

The first ten sections of the chapter closely follow the first edition. The remaining sections make a case for developing *TF&S*-economics models, where *TF&S* stands for "thinking, fast and slow;" and this issue constitutes the overarching theme of the chapter. As mentioned above, this chapter deals with three ways in which utility functions can change: intertemporal substitution associated with present bias, the impact of self-control challenges on coincident substitution between different commodities, and the effect of changing mental states in respect to substitution between alternative risky prospects. The chapter presents a general two-system TF&S approach for modeling all three.[4]

Arguably, the game theoretic approach is effectively a one-system framework where individuals engage in deliberative System 2 thinking, but which they do quickly and effortlessly. The chapter discussion makes the case that the game theoretic approach lacks the choice architectural features associated with fast thinking System 1. This is important for at least two reasons. The first reason pertains to issues associated with welfare and nudging behavior. Here the term nudging refers to specific behavioral interventionist measures, mostly aimed at fast thinking System 1 [Thaler and Sunstein (2021)].[5] The second reason pertains to the fact that the implications of the two-system behavioral approach are consistent with the broad empirical evidence, but not the one-system neoclassical approach.

The remainder of the chapter is organized as follows. Section 2 contains a discussion

[3]Thaler and Shefrin (1981) discuss various two-system frameworks in the academic literature, for example the work of psychologist William James and economic philosopher Adam Smith. Kahneman was clear to say that the two-system perspective did not originate with him. On p. 20 of his book, Kahneman (2011) attributes the terms "System 1" and" System 2" to psychologists Keith Stanovich and Richard West. In 1977 two psychologists, Richard Shiffrin and Walter Schneider, reported a series of experiments whose results they interpreted using a two-system perspective. Their article, "Controlled and Automatic Human Information Processing: I. Detection, Search and Attention," appeared in *Psychological Review.* In private correspondence, Barbara Tversky mentions others such as: "Zacks and Hasher's work, Neisser; in cognitive science, the procedural-declarative distinction (Norman and Bobrow), processes that initially take attention like reading music, tennis serves, driving, proof-reading newspapers, multiplication, the list goes on, become automatic with practice." As for a two-system approach to self-control in the psychology literature, there is Metcalfe J. and W. Mischel, "A Hot/Cool System Analysis of Delay of Gratification: Dynamics of Willpower," which in 1999 appeared in *Psychological Review.* Gigerenzer (2007) provides a psychological two-system framework in his book *Gut Feelings.*

[4]Present bias is also a major factor in humans' inadequate response to global warming [Shefrin (2023)].

[5]For a critique of the nudge approach, see Viale (2022).

of the historical development of intertemporal choice theory, which led to the discounted utility model (section 3). This sets the stage for the first formal treatment of changing utility, namely Strotz's pioneering paper on dynamic consistency and exponential discounting (section 4). Experimental evidence suggests that the discount function is hyperbolic rather than exponential. Section 5 is devoted to the properties of the hyperbolic discount function. Section 6 describes the properties of a consistent plan for the hyperbolic discount function case. Section 7 sets out conditions under which a consistent plan exists. The existence issue involves some subtle aspects concerning indifference. Having discussed existence, I turn next to the welfare implications associated with consistent plans. Section 8 establishes that the consistent plan is Pareto–inefficient in the hyperbolic discount function case. That is, there are plans which are judged to dominate the consistent plan, from the perspective of *all* utility functions held by the individual over his lifetime. Section 9 indicates how the character of the consistent plan is affected when the utility function depends on past consumption. Most of the analysis in this chapter assumes that there is a single consumption good. Section 10 also discusses the influence of past consumption on utility, but in a multi–commodity setting. Section 11 describes a behavioral approach to self-control and willpower, based on two-system framework involving simultaneous internal conflict. Section 12 discusses a series of issues relating to the applicability of the intrapersonal game theoretic approach. Section 13 surveys the empirical evidence associated with present bias. Section 14 extends the discussion of changing tastes from the certainty framework to the uncertainty framework. Section 15 describes the structure of prospect theory, a descriptive theory of decision making under uncertainty developed by cognitive psychologists. A key feature of the analysis concerns the connection between uncertainty framing effects and dynamic inconsistency. This section also describes how to embed prospect theory within a two-system TF&S framework. Section 16, the final section is a conclusion. The conclusion is intended to serve three purpose: 1) summarize the main discussion points; 2) emphasize how the various strands of the literature relate to one another; and 3) highlight the differences between the intrapersonal game theoretic approach and the behavioral two-system approach.

2 Historical Development

This section contains a brief overview of the way that ideas about intertemporal choice have evolved in economics. This evolution has featured a pattern of ebbs and flows in its emphasis on the relative importance of particular issues. The writings of early economists stressed the importance of will in delaying gratification, as well as cognitive imperfections in weighing future benefits against present ones. Both of these concepts were later downplayed, as formal utility framework was applied to intertemporal choice. However, as we shall see, issues of will and cognition returned as central issues in the consideration of changing utility functions.

Early economists, such as Adam Smith (1776), focused on capital accumulation as the critical ingredient for explaining differences in the standard of living across demographic groups and over time. Rae (1834) was especially concerned with understanding the determinants of differences in the rate at which capital has accumulated. He argued that the rate at which capital accumulates depends on the willingness of the public to defer gratification from the present to the future. In particular, Rae identified four factors which he felt underlie this willingness to defer. The four are: *1) uncertain lifetime; 2) abstinence, meaning the psychological discomfort of deferral; 3) the bequest motive; and 4) prudence*. The fourth category dealt with "the extent of intellectual powers, and the consequent power of habits of reflection" (1834, p. 58). Notice that the first two of Rae's factors favor present consumption over future consumption, with the third and fourth providing some offset to this tendency.

The economist N.W. Senior (1836) first suggested that abstinence was priced in the market, noting that this seemed to be the reason why the rate of interest was positive rather than zero. In Senior's treatment, a positive interest rate provided compensation for abstaining from current consumption, which he described as "among the most painful exertions of human will" (1836, p. 60). It is interesting to note the central role played by the psychological factor abstinence in the writings of Rae and Senior, as these were to be downplayed as the theory of intertemporal choice was developed by those who followed.

Using a Benthamite perspective, in which individuals are regarded as being both self–

centered and centered in the present relative to other points in time, Jevons (1871) asked why individuals do not act completely myopically? Why, he wondered, do they attach any weight at all to future consumption? Jevons suggested that the weight which individuals attach to future consumption stems from the *present anticipation* of future pleasure. In other words, at any point in time, the individual decides how to allocate his wealth between present pleasure (currently experienced) and the anticipation of future pleasure (*also* currently experienced). However, Jevons did not go so far as to argue that the ability to anticipate future pleasure was perfectly calibrated, stating instead that future pleasures do not all "act upon us with the same force as if they were present" (1871, p. 76). Unlike Senior and Rae, who viewed abstinence as resulting from the psychological pain of delaying gratification, Jevons viewed the asymmetry between present and future consumption in terms of a cognitive imperfection: the anticipated pleasure from future consumption is a downward biased indicator of the future pleasure that it prefigures.

Böhm–Bawerk (1889) modified Jevon's approach in a way which set the stage for modern intertemporal choice theory, in which there is a single intertemporal utility function. Böhm–Bawerk rejected Jevon's view that individuals are capable of feeling emotions in advance, and suggested instead that a person acts as a referee in respect to the intertemporal allocation of pleasure over time. Böhm–Bawerk's individuals trade off satisfactions at different points in time, and these are regarded as being placed upon a single cognitive plane on which they can be compared. This view provides the conceptual framework underlying the single utility function theory of intertemporal choice. Nevertheless, Böhm–Bawerk also recognized the presence of a cognitive imperfection concerning the treatment of present consumption relative to future consumption, noting a "systematic tendency to underestimate future wants" (1889, p. 268-269).

Unlike Rae and Senior who emphasized the contribution of abstinence as a factor in intemporal choice, Jevons and Böhm–Bawerk rested their arguments on the cognitive weighting of present and future consumption. Yet Böhm–Bawerk was well aware that abstinence was an important issue, and invoked it in his welfare analysis. Loewenstein (1992) poiints out that Böhm–Bawerk discusses an injudicious decision by an individual to favor the present over the future, as a result of giving in to weakness, with the immediate knowledge that when tomorrow arrives, he will regret having done so. Ironically, although it is Böhm–Bawerk who first puts forward the conceptual framework underlying intertemporal choice theory based upon a single utility function, his recognition of anticipated "regret" clearly reflects his appreciation of changing utility.

The quantification of Böhm–Bawerk's conceptual framework is due to Irving Fisher (1930). Fisher presented the tradeoff between present and future consumption in a two–date model involving indifference curves over consumption at the two dates. These indifference curves (or "willingness lines" as Fisher called them), were not symmetric, with the asymmetry capturing the notion of "impatience." The reasons put forward by Fisher to explain why an individual would be impatient for present consumption are similar to the ones discussed by his predecessors, although he did emphasize imperfect foresight and a failure to display self–control.

3 The Discounted Utility Model

Fisher's indifference curves can be regarded as the level lines of an intertemporal utility function $U(c_1, c_2)$, where c_1 denotes present consumption and c_2 denotes future consumption. Samuelson (1937) developed the notion of intertemporal utility in the form of a discounted utility model over intertemporal consumption profiles of the form $c = (c_1, c_2, \cdots, c_T)$. In Samuelson's formulation, lifetime discounted utility has the form:

$$U(c_1, c_2, \cdots, c_T) = \sum_{t=1}^{T} u(c_t)\delta^{t-1}, \tag{1}$$

where $u(c_t)$ is the instantaneous utility derived at date t from consuming c_t and δ^t is the discount factor, $0 < \delta < 1$. In this formulation, it is assumed that $u' > 0$ and $u'' < 0$. Equation (1) has become the standard utility functional for comparing alternative consumption profiles, especially for continuous time where c_t is replaced by $c(t)$, and (1) takes the integral form:

$$U(c) = \int_0^T e^{-pt} u(c(t))dt. \quad (2)$$

where p is the interest rate per annum.[6]

Equation (1) indicates that intertemporal utility is additively separable in its arguments. Samuelson recognized that this assumption is quite restictive, in that instantaneous utility depends only on consumption at that moment, not on the memory of past consumption or the anticipation of future consumption. Moreover, future utility is discounted back to the present as if it were the same as future cash flow, and δ is treated as if it were the interest rate.[7] Nevertheless, the additively separable form has been used extensively in the literature on changing tastes. Although this has largely been to simplify the analysis, it does raise questions about the extent to which the various conclusions from the model can be generalized. As we shall see, the central results in this literature apply outside the additively separable framework.

4 A Changing Tastes Model

The seminal article on changing utility functions is by Robert Strotz (1956). In the single utility function intertemporal model, an individual's utility function $U()$ remains invariant throughout his lifetime. The tradeoffs evaluated at age twenty are evaluated in exactly the same way at age seventy. This feature struck Strotz as being inconsistent with casual observation. He noted that some individuals appeared to have difficulty demonstrating abstinence. Moreover, he noted that other individuals engage in *precommitment* behavior, whereby they constrained their future choice set of consumption profiles.

Strotz provided a wide ranging set of examples to illustrate what he meant by pre-

[6]The expression e^{-pt} arises from the equation $lim_{n\to\infty}$ $(1 + p/n)^n = e^p$. Here, the interest rate applied to a period of length $1/n$ is p/n.

[7]Koopmans (1960) provided a set of axioms which establish conditions for the intertemporal utility function to be additive in instantaneous utilities $u_t(c_t)$. His axioms can be organzized into three groups. The first, *independence* states that if two profiles share a common consumption value at any date t, then preference between the two profiles is determined by consumption levels at the other dates. The second, *stationarity* assumes that if date 1 consumption is the same for two profiles, then preference between them will be preserved by dropping the first period, and shifting the profiles forward in time. The third group, *completeness* are a set of technical axioms.

commitment. He began his article with the following well known quotation from Homer's classic *The Odyssey* in which Ulysses precommits himself, lest he be tempted by the wail of the Sirens into wrecking his ship upon the rocks:

> [B]ut you must bind me hard and fast, so that I cannot stir from the spot where you will stand me ... and if I beg you to release me, you must tighten and add to my bonds.

In the body of his article, he mentions a variety of precommitment strategies, such as: the use of "insurance policies and Christmas Clubs which may often be hard to justify in view of the low rates of return"; having one's academic salary dispensed on a twelve month basis instead of nine; and joining the army. These are all viewed as devices to cope with the *mañana* effect, continually putting off until tomorrow what is unpleasant to do today.

In what might be viewed as a "back to Jevons" argument, Strotz reintroduced the notion that individuals are centered in the present, and he relaxed the requirement that the discount function be exponential, as it is in (1). Imagine that the present means date τ, with date t occuring after τ, i.e. $t > \tau$. In Strotz's framework, instantaneous utility stemming from date t consumption continues to be $u(c_t)$. However, when the present is date τ, the discount factor[8] applied to $u(c_t)$ is $\lambda_\tau(t)$ rather than δ^{t-1}. Strotz assumes that the discount function $\lambda_1()$ is maintained, but shifts forward in time. In other words, $\lambda_\tau(t)$ is an invariant function of the difference $|t - \tau|$. Since τ parameterizes the moving frame of reference, it can be termed a *reference date.*

With a forward shifting discount function, past consumption is also discounted relative to present consumption. In this sense, the utility function changes whenever the discount function $\lambda_\tau()$ is nonconstant in τ. However, because time is unidirectional, the individual at date τ must accept his past consumption choices.

What interested Strotz was whether the forward shifting discount function resulted in a meaningful change in utility function, in respect to the consumption choices from

[8]Which is indexed by τ.

date τ on? Would the shifting discount function induce a change in the rank ordering of consumption profiles from τ through T?

Formally, consider an individual whose problem is to allocate lifetime wealth q over T consumption dates. The individual's utility function $U_\tau(c)$ has the form:

$$U_\tau(c) = \sum_{t=1}^{T} u(c_t)\lambda_\tau(t) \tag{3}$$

where the shifting discount function feature is captured by the condition:

$$\lambda_\tau(t) = \lambda_1(|t - \tau + 1|) \tag{4}$$

and $\lambda_1(1)$ is normalized to unity. Budget feasibility is specified by the constraint:

$$\sum_{t=1}^{T} c_t \leq q. \tag{5}$$

For simplicity, a zero interest rate is assumed in (5). In addition, equality is assumed rather than weak inequality because of the nonsatiation condition $u' > 0$.

Consider the individual at date 1. Denote by c^1 the profile which maximizes lifetime utility (3) (for $\tau = 1$), subject to the budget constraint (5). The standard first order optimizing condition for c^1 is that for all t and s:

$$\lambda_1(t)u'(c^1_t) \geq \lambda_1(s)u'(c^1_s) \tag{6}$$

with equality in (6) if and only if c^1_t and c^1_s are both strictly positive.

Assume that at date 1, the individual is free to select date 1 consumption c_1, but must defer the implementation of later consumption c_τ until date τ. Suppose that at date 1, the individual selects $c_1 = c^1_1$, believing that at date τ, he will select $c_\tau = c^1_\tau$. Of course, this belief may be naïve, because the individual's utility function at date τ will have changed from U_1 to U_τ. For instance, at date 2, the individual will have remaining wealth $q_2 = q - c^1_1$. He is free to compute the profile c^2 which maximizes $U_2(c)$ subject to the constraints:

$$c_1^2 = c_1^1$$

and

$$\sum_{t=2}^{T} c_t^2 = q_2.$$

With the passage of time, this procedure results in the naïve consumption profile $c^N = (c_1^1, c_2^2, \cdots, c_T^T)$.

Strotz was interested in characterizing the conditions under which the shifting discount function did not cause the individual to change his mind about the profile selected at date 1. In other words, under what conditions will $c_t^\tau = c_t^1$? He established:[9]

Theorem 1 *The equality $c_t^\tau = c_t^1$ holds for all τ if and only if $\lambda_\tau()$ has the exponential form:*

$$\lambda_\tau(t) = \delta^{t-\tau}. \tag{7}$$

Proofs of all theorems appear in the appendix. In this regard, the discussion in the body of the chapter occasionally references equations and figures which appear in the appendix.

Notice that exponential discounting gives rise to a relativity principle. That is, people with exponential discount functions whose past plans were formulated by maximizing their initial utility functions U_1 would have no desire to revise those plans with the passage of time, even though their utility functions have changed. Put somewhat differently, such people would have no motive to enter into commitments which would constrain their own behavior in the future. Their future behavior and present utility functions are *dynamically consistent*. The passage below indicates Strotz's thinking on the subject:

> Special attention should be given, I feel, to a discount function ... which differs from a logarithmically linear one in that it "overvalues" the more proximate satisfactions relative to the more distant ones. p. 177.

[9]Strotz made the argument for a continuous time model.

The people who have the potential to display dynamically inconsistent behavior are the ones whose discount functions are not exponential. For instance, suppose an individual discounts the future in such a way that:

$$\lambda_2(3) < \frac{\lambda_1(3)}{\lambda_1(2)}.$$

Such an individual weighs the relative merits of date 3 consumption and date 2 consumption quite differently at date 2 than he does at date 1. For him, date 2 consumption is relatively more important at date 2 than it is at date 1. The rate at which he discounts the future decreases as a function ot the time delay. Consequently, when date 2 arrives, he will wish to consume more c_2 than he had planned under c^1. This feature has come to be known as the *common difference effect* [Loewenstein and Prelec (1992)]. As the following passage demonstrates, for Strotz, this is the cause, not only of overspending on present consumption relative to the future, but a wide range of social ills.

> [T]here are no doubt some who, either through lack of training or insight, have never learned to behave consistently and for whom the intertemporal tussle remains unsolved. These people we call "spendthrifts." By contrast, those who have taken on log–linear discount functions have learned to be "thrifty."
>
> Spendthriftiness, in the general sense of inconsistent or imprudent planning, is by no means insignificant. It is especially among the lower–income classes, where education and training are commonly blighted, that one would expect to find imprudent behaviour of this sort. In America, lower income people tend to gorge themselves with food after pay–day; overheat their homes when they have money for a bucket of coal; are extravagant, going on sprees on pay–day, not budgeting their money, and engaging in heavy instalment buying; do not keep their children in school; and are freer in the expression of their sexual and agressive impulses. Their high birth rate is well–known. All these behaviour characteristics can be explained as a failure to cope intelligently with the problem of the intertemporal tussle. p. 178

5 Hyperbolic Discounting

Strotz' work stimulated both economists and psychologists to examine the way that people discount the future. The most compelling experimental evidence that utility functions change over time involves the common difference effect. Some representantive articles which address this issue are Ainslie (1974, 1975, 1985), Thaler (1981), Horowitz (1988), Benzion, Rapoport, and Yagil (1989), and Loewenstein and Prelec (1992). Interestingly, Ainslie (1974) discusses experimental evidence involving animals.

Consider a mutually exclusive choice between receiving *A) $1 today; and B) $2 tomorrow*. Suppose an individual prefers A over B. Now let us ask that same individual to choose between two other mutually exclusive choices involving the receipt of *C) $1 25 days from now; and D) $2 26 days from now.* One of the Koopmans (1960) axioms underlying discounted expected utility is *stationarity*, and it stipulates that A is preferred to B if and only if C is preferred to D. The violation of stationarity leads to the common difference effect. The common difference effect gets its name from the fact that the time difference in the two choices, in this example one day, is common.

More formally, imagine that an individual is consuming an amount c at two dates t and s, $t < s$. Suppose that he is indifferent to receiving incremental amounts x at date t and $y > x$ at date s. Then with exponential discounting, we have:

$$u(c+x)\delta^t + u(c)\delta^s = u(c)\delta^t + u(c+y)\delta^s,$$

which implies that:

$$u(c+x) - u(c) = (u(c+y) - u(c))\delta^{s-t},$$

so that the only aspect of time which is relevant here is the delay $s - t$, but not the dates t and s themselves. In other words, exponential discounting is inconsistent with the common difference effect.

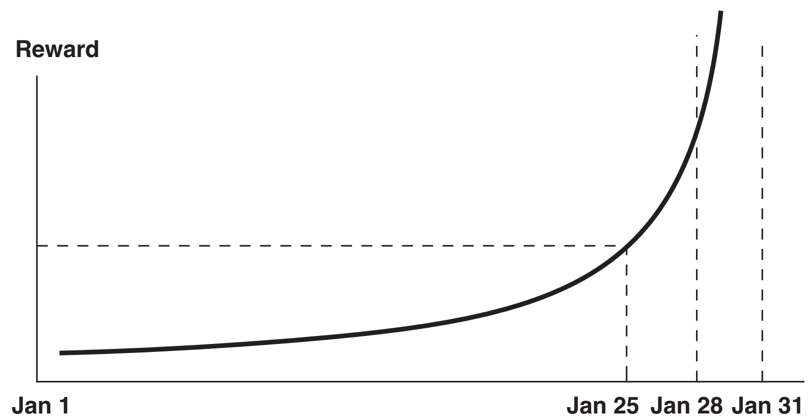

Figure 1: Reward value at t as a function of time delay.

A useful way of depicting the difference between exponential and nonexponential discounting is to graph reward value against time delay: see Figure 1. Consider a reward of Z to be received at the current date $\tau = 0$. The X–axis represents time. As we move to the right in the figure, there is a delay until a reward is received. The curve in the figure reprents the size of reward which is required to leave the individual indifferent between Z at $\tau = 0$ and a larger reward at each future time t. Imagine that the current date τ is January 1, and that an individual is indifferent between a \$1.00 reward on January 25 and \$1.20 reward on January 28.[10] Then the points (25, 1) and (28, 1.2) will lie on the same curve in Figure 1. Let time advance and consider how the individual feels on January 25 when the \$1 reward is due. An individual who discounts exponentially will continue to be indifferent between (25, 1) and (28, 1.2). Such an individual would not shift his indifference curve in Figure 1. However, a nonexponential discounter would exhibit a different indifference curve on January 25, and for example, would regard \$2 on January 28 as exact compensation for waiting another 3 days. This individual's indifference curves would cross, with the January 25 indifference curve lying above the January 1 curve to the right of January 25. See Figure 2.

[10]This example is taken from Frank (1992).

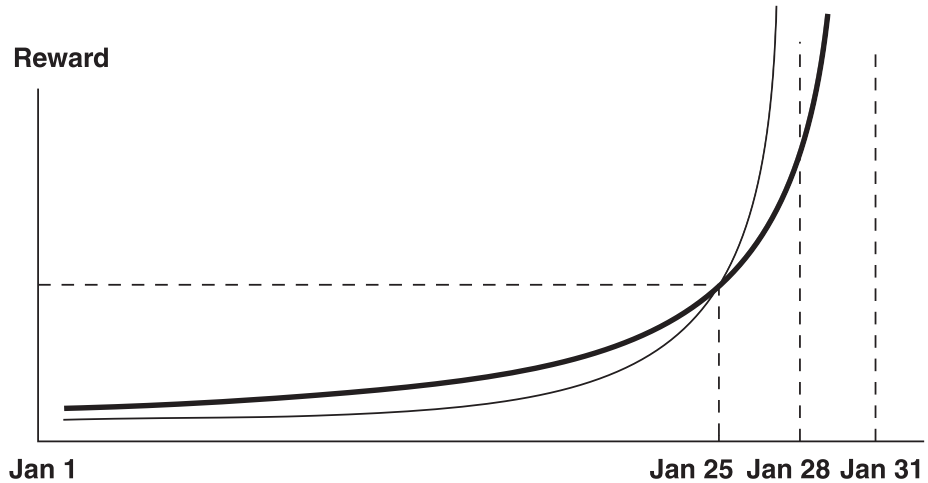

Figure 2: Indifference curves which cross: dynamic inconsistency. The thicker line is the indifference curve on January 1, and the thinner curve is the indifference curve on January 25. Notice that the reward required on January 28 is higher on the January 25 indifference curve than it is on the January 1 indifference curve.

One discount function which has received considerable attention in the psychology literature is known as *Herrnstein's matching law*, [Herrnstein (1961)]. This discount function has performed well in experiments, with both human and with animal subjects. Economist Robert Frank (1992) describes the matching law as being hard wired into the nervous systems of most animals. The function in question has, as its arguments the present value V, size of reward A, and delay $(t-\tau)$ until the reward is received. The function itself is:

$$V = \frac{A}{K + \alpha(t-\tau)}, \tag{8}$$

where K and α are parameters. The role of K is to prevent V from approaching infinity as the delay goes to zero. V denotes the present value at date τ of reward A to be received after a delay of $(t-\tau)$; see Ainslie and Haslam (1992). The discount function is, of course,

$$\lambda_\tau(t) = \frac{1}{K + \alpha(t-\tau)}. \tag{9}$$

Equation (8) is known as hyperbolic discounting because it traces out hyperbolic curves. Under the matching law,[11] indifference curves cross in the manner of Figure 2:

[11]The term matching law arises from Herrnstein's prediction that effort will be allocated across rewards in order that effort match reward.

Future rewards are heavily discounted, and present ones loom large. Notice from (8) that, except for delays close to zero, V is roughly inversely proportional to the length of the delay.

Loewenstein and Prelec (1992) provide an enlightening discussion of the generalized hyperbolic discount function, based upon Prelec (1989). The generalized hyperbola has the form:

$$\lambda(t) = (1 + \alpha t)^{-\beta/\alpha}. \tag{10}$$

Notably, the exponential function $e^{-\beta t}$ occurs as the limit of (10) when $\alpha \to 0$.

Consider an individual who is indifferent between reward x at $\tau = 0$ and reward $y > x$ at some later date s. Formally,

$$u(x) = u(y)\lambda(s)$$

Suppose that both rewards are delayed by the same amount t. The inequality

$$u(x)\lambda(t) \leq u(y)\lambda(t + s)$$

indicates that preference has shifted to the later reward, with the inequality being strict for the nonexponential case. That is, the indifference before the delay results in strict preference for the later reward y in the presence of the delay. Therefore restoration of indifference requires that the time of arrival of y be delayed even further from $(t + s)$ to $(kt + s)$, where k depends only on x and y, but not on t or s.

Theorem 2 *If for all x and y,*

$$u(x) = u(y)\lambda(s)$$

implies that for all t,

$$u(x)\lambda(t) = u(y)\lambda(kt + s) \tag{11}$$

then

$$\lambda(t) = (1 + \alpha t)^{-\beta/\alpha}.$$

6 Dynamic Consistency as Nash Equilibrium

If hyperbolic discounting is the norm, then the individual will not be able to implement the plan c^1 which is optimal under his preferences at date 1. Pollak (1968) develops a particularly insightful example of a naïve consumption plan. His example features the case of logarithmic utility. Consider the date τ plan which is optimal for the situation when the individual's discount function is given by $\lambda_\tau()$. As is well known, maximizing logarithmic utility leads to the constant budget share condition. Define

$$\phi_{\tau,t} = \frac{\lambda_\tau(t)}{\sum_{s\geq\tau}^{T} \lambda_\tau(s)}. \tag{12}$$

Let $t \geq \tau$. At date τ, the individual's naïve plan calls for date t consumption to be allocated the share $\phi_{\tau,t}$ of remaining wealth q_τ. As Theorem 1 established, these budget shares will be the same at each date if and only if the discount function is exponential. In particular, if the discount function is hyperbolic, then the individual will readjust the budget shares every period. It is straightforward to use recursion in order to compute the naïve plan c^N. Date 1 consumption is just $\phi_{1,1}q$. Date 2 consumption is $\phi_{2,2}q_2$, which is $\phi_{2,2}(1-\phi_{1,1})q$. Date 3 consumption is $\phi_{3,3}q_3$, which is

$$\phi_{3,3}(1-\phi_{2,2})q_2 = \phi_{3,3}(1-\phi_{2,2})(1-\phi_{1,1})q, \tag{13}$$

and so on.

An individual who has a hyperbolic discount function would prefer to precommit the choices of his future selves, and select c^1 as their consumption profile. However, precommitment may not be a feasible option, in which case a far sighted individual will consider the behavior of his future self when selecting current consumption.

A convenient way to think about this situation is as a game with T players, one for each date τ. In this case, player τ's preferences are represented by U_τ. Given lifetime wealth q, player 1's strategy is to choose c_1. Given remaining wealth $q_1 = q - c_1$, player 2's strategy is to choose c_2. In general, a strategy for player τ is a function $\rho_\tau()$ which specifies the rate of consumption $c_\tau = \rho_\tau(q_{\tau-1})$ from remaining wealth $q_{\tau-1}$. A plan c^* which corresponds to a *subgame perfect Nash equilibrium* for this game is known as a *Strotz–Pollak equilibrium*. Such a plan is dynamically consistent. As I discuss in the next section, the conditions under which a Strotz–Pollak equilbrium exist are quite general.

For purposes of exposition, I confine attention to the case $T = 3$ in the remainder of this section.[12] At date 1, player 1 considers the reaction of the other players to a particular selection c_1. If he knows $\rho_2()$ and $\rho_3()$, then he can compute both c_2 and c_3 as functions of c_1. That is, $c_2(c_1) = \rho_2(q - c_1)$ and $c_3(c_1) = \rho_3(q_1 - c_2)$. Therefore player 1's perfect Nash equilibrium behavior is to select c_1 to maximize:

$$u(c_1) + \sum_{t=2}^{3} \lambda_1(t) u(c_t(c_1)).$$

Put somewhat differently, player 1 trades off two utilities, $u(c_1)$ from current consumption c_1 and $V_C(q_1)$ from his bequest $q_1 = q - c_1$. Here $V_C()$ is given by:

$$V_C(q_1) = \sum_{t=2}^{3} \lambda_1(t) u(c_t(c_1))$$

where the subscript 'C' in $V_C()$ refers to a *consistent* plan. A similar maximizing statement holds for player 2, and player 3 simply sets $c_3 = q_2$. Perfection follows because the reaction functions are computed in reverse order, from $T = 3$ back to $\tau = 1$.

Consider the dynamically consistent plan in Pollak's logarithmic utility example. Given q_2, we already know that consumption levels at dates 2 and 3 are determined by the budget shares $\phi_{2,2}, \phi_{2,3}$ respectively. At date 1, the individual's dynamically consistent choice of c_1 maximizes:[13]

$$\lambda_{1,1} \ln(c_1) + \lambda_{1,2} \ln(\phi_{2,2}(q - c_1)) + \lambda_{1,3} \ln(\phi_{2,3}(q - c_1)). \tag{14}$$

[12] The arguments are easily generalized.

[13] The constant elasticity of intertemporal substitution function has the form $u(c) = c^{1-\theta}/(1-\theta)$. Logarithmic utility corresponds to the degenerate case $\theta = 1$. Applying the argument associated with (14) to $u(c)$ establishes that logarithmic utility is a knife edge case. When $\theta > 1$ the ratio $c_1/(g_1 - c_1)$ is higher for the naïve case than the sophisticated case. But the inequality reverses when $\theta < 1$

By (12), notice that maximizing (14) is the same as maximizing:

$$\phi_{1,1} \ln(c_1) + (\phi_{1,2} + \phi_{1,3}) \ln(q - c_1). \tag{15}$$

Hence, the maximizing choice of c_1 is actually invariant to the budget shares selected at future dates. By repeating the argument recursively, we can conclude that when the utility function is logarithmic, the dynamically consistent plan actually coincides with the naïve plan.[14] Of course, this does not mean that at date 1, the individual is indifferent to the budget shares chosen by his future self. The level of date 1 total utility will vary with the budget shares used at future dates, and will lead to maximum date 1 utility when they coincide with $[\phi_{1,t}]$. However, while the total utility function will shift, the marginal utility function will not, and this is why the maximizing choice of c_1 does not vary with the budget shares used at future dates.

The case of logarithmic utility can also be used to provide some insight into a hypothesis proposed by Strotz (1956). Strotz put forward a condition which he thought characterized a dynamically consistent plan in a continuous time model. This condition involves, at each moment t, the local approximation of the true discount function $\lambda(z-t)$ by an exponential function e^{-pt}. Recall that the continuous time version of Theorem 1 indicates that naïve choice will be dynamically consistent when the discount function takes this form. Strotz hypothesized that at each moment, the consistent rate of consumption would be given by the rate along the naïve path associated with the local exponential approximation. Pollak (1968) suggested that this hypothesis is counterintuitive. He reasoned that at any given moment, the rate of consumption at a given moment along a sophisticated plan would take into account the entire future time path. In contrast, Strotz's hypothesis indicates that the rate of consumption is determined only by the local features of the discount function. Pollak used a counterexample with logarithmic utility to demonstrate that Strotz's hypothesis is false. In the counterexample, $\dot{c}(t)/c(t) = -p$ at

[14] From this example, one could draw the conclusion that a naïve plan need not be dynamically inconsistent. However, the naïve plan should be understood as a process. Certainly, the process of choosing a plan naïvely leads to sequential revision of the plan selected at date 1.

moment t along the path obtained by using the local exponential approximation. However, $\dot{c}(t)/c(t) \neq -p$ along the sophisticated plan. Rather:[15]

$$\frac{\dot{c}(t)}{c(t)} = \frac{\lambda(T-t) - 1}{\int_t^T \lambda(z-t)dz} \tag{16}$$

along the sophisticated plan, which is also the naïve plan in the case of logarithmic utility. The last equation reflects the fact that the rate of consumption is determined by taking the entire future path into consideration, not just the local discount rate.[16]

A consistent plan is self–enforcing. Even an individual with a hyperbolic discount function will not find himself tempted to deviate from a consistent plan over time. However, a consistent plan is not the same as the precommited plan c^1. An individual typically consumes at entirely different rates along the two profiles. For example, in the logarithmic utility example with hyperbolic discounting, date 2 consumption is higher along the consistent plan than along the precommited plan. Therefore we need to ask whether it is generally the case that the act of following a consistent plan, although it may resolve Strotz' "intertemporal tussle", will still lead to spendthrift behavior? When the discount function is hyperbolic, a key question is whether the consistent value c_1^* is always larger than its precommited counterpart c_1^1? The next theorem provides a negative answer to this question. In addition, the proof describes structural features which play a part in later discussion.

Theorem 3 *There are utility functions and parameter specifications for which $c_1^* > c_1^1$ and parameter specifications for which $c_1^* < c_1^1$.*

When $T = 3$, (70) is the Euler equation for a standard intertemporal optimization involving an unchanging utility function. The Euler equation for the sophisticated plan is given by $u'(c_1) = V_C'(q_1)$. Laibson (1995, 1997) has developed the Euler condition associated with a dynamically consistent plan when the discount function is a quasi-

[15]This is one of several places in the chapter where counterexamples to claims are presented. In such cases I have decided to omit proofs. See also Goldman's discussion of the Peleg–Yaari example in section 7, and Pollak's discussion of von Weizsäcker's model in section 10.

[16]Notice that (16) is the continuous time counterpart to the relation implied by (12) and (13).

hyperbolic function. Specifically, he analyzes the following discount function which was introduced by Phelps and Pollak (1968):

$$\lambda_\tau(t) = \beta\delta^{t-\tau}, \tag{17}$$

where $0 \leq \beta \leq 1$ and $t > \tau$. As before, $\lambda_\tau(t) = 1$. Call this structure the $\beta\delta$ model.

Observe that relative to current date τ consumption, the discount factor applied to date t consumption is $\beta\delta^{t-\tau}$. However for two future dates t and s, where $s > t$, the relative weights attached to consumption at these two dates is δ^{t-s}. Notice that exponential discounting is associated with the condition $\beta = 1$. However, if $\beta < 1$, then the individual always discounts the future more heavily relative to the current date than he does relative to any future date. This is the essential feature of the hyperbolic discount function.

Laibson argues that the Euler equation associated with the preceding discount function has the following form:

$$u'(c_t) = \delta u'(c_{t+1})[1 + ((\beta - 1)\partial c_{t+1}/\partial q_{t-1})] \tag{18}$$

It follows that when $\beta < 1$, $u'(c_\tau) > u'(c_t)$ for all $\tau > t$.

Laibson has consistently emphasized that impatience applies to very short periods of time, such as being caught in traffic or waiting in a checkout line at a store. Harris and Laibson (2013) develop a continuous time quasi-hyperbolic time preference model, which they call the *instantaneous-gratification model* or the IG model. In the IG model, people are only impatient about consumption within an instant *dt* of the current moment, but they discount exponentially after *dt*.

An important implication of the IG model is that its solution mimics the features of an exponential discounter. In this regard, the value function of the IG model is identical to the value function of an optimization problem with the same long-run discount rate as the IG model, but with a different flow utility function that depends on both the level of consumption and the level of financial assets. Harris and Laibson explain that the IG model provides what they call "a tractable niche between dynamically inconsistent models and dynamically consistent models."

For the IG model, consider the case of certainty and a complete market without liquidity constraints. In this case, the valuation functional at date s is given by:

$$U_s(c) = \int_s^{\infty} \beta e^{-p(t-s)} u(c(t))dt \tag{19}$$

where $u()$ is a power utility function. Think about (19) as the limit of a discrete time model based on (19). In this case, the period $t = 1$ corresponds to a short interval of time dt, for which $\beta = 1$. Subsequent to dt, preferences are quasi-hyperbolic, meaning that $\beta < 1$. It is with this interpretation that for $t > s$ in (19), the value of the discount function is $\beta e^{-p(t-s)}$. Specifically, for the instantaneous moment dt at $t = s$, the value of the discount function being unity rather than $\beta < 1$, exerts a negligible effect on the value of the integral, which goes to zero in the limit. However, within the short interval dt, for the case when $\beta < 1$ for $t > s$, the unity value at $t = s$ exerts a non-negligible impact on the value of $c(t)$; and this induces overconsumption relative to the case when $\beta = 1$; and because the utility function changes from one dt to the next, the net effect on $c()$ is not negligible.

Keep in mind that exponential discounting corresponds to the case $\beta = 1$, and quasi-hyperbolic discounting corresponds to the case $\beta < 1$. For the IG model described above, let the counterpart variables associated with the exponential discounting case feature the hat symbol. For example, the counterpart to u is $\hat{u}$ and the counterpart to $c(t)$ is $\hat{c}(t)$. Using a reverse engineering argument, Harris and Laibson establish that for the version of the problem with certainty and no liquidity constraints, $\hat{u}$ is a positive affine transformation of the power utility function $u(c)$, with $\hat{c}(t) = c(t)$. Specifically, they state that the IG model, which is dynamically inconsistent, has the same valuation functional as a standard dynamically consistent optimization problem. Recall that the $\hat{u}-$agent maximizes $\hat{U}$, in contrast to the time-inconsistent IG-agent who chooses a supoptimal solution by overconsuming. All else being the same this situation would imply $\hat{U}(c) > U(c)$. However, Harris and Laibson's construction seeks to equate $\hat{U}$ and U, an equality which requires that $\hat{u}(c) < u(c)$.

Notably, the IG model features a single welfare criterion, even though the model involves dynamic inconsistency. This is because at each instant, overconsumption applies only to the current instant. As a result, the IG model features both a single welfare function and overconsumption.

Harris and Laibson introduce interest rate risk into their framework, with the valuation functional being expressed as:

$$U_s(c) = E\{\int_s^{\infty} \beta e^{-p(t-s)} u(c(t))dt\} \tag{20}$$

In the IG model, as long as the individual is not constrained in the amount to be borrowed, the quasi-hyperbolic solution will continue to mimic an exponential $\hat{u}-$solution. However, Harris and Laibson show that a binding borrowing constraint will interrupt the equivalence.

Harris and Laibson provide a detailed characterization of the IG-consumption function. They establish that when the expected rate of return is below a key threshold, the equilibrium consumption function exhibits a discontinuity at the liquidity constraint. As a result, when the individual hits the liquidity constraint, consumption falls discontinuously. They note that this feature is not present in a dynamically-consistent consumption model, which of course includes the $\hat{u}-$case discussed above.

7 Existence and Indifference

A discussion appearing in the appendix, involving Figure 10, analyzes the possibility of there being multiple local maxima (q_1^L and q_1^U).[17] Indifference between these two maxima occurs along the consistent plan when, in Figure 10, the area of triangle B equals that of triangle A. The presence of indifference can be important. From the perspective of early preferences, an individual will typically care about how his later self would resolve future points of indifference. Specifically, his early choices may be premised on the tie–breaking mechanism his future self will use. Therefore it is important to be clear about the extent

[17] Notably, Harris and Laibson establish a uniqueness theorem for the IG model.

to which the current self is aware of the way its later self will resolve future indifference.

In the absence of indifference, the existence of a Strotz–Pollak equilibrium follows from backward induction, under the assumption that the utility function is continuous. However, the presence of indifference interferes with the use of this argument as a means of establishing existence. Nevertheless, as Goldman (1980) demonstrated, indifference does not prevent the existence of a Strotz–Pollak equilibrium. His result is discussed below.

Previous to Goldman's work, there was some confusion about whether an equlibrium can be guaranteed to exist in this type of game. Recall that in the model under discussion, the date t self only games forward. That is, it takes the past as given, and bases its choice upon the anticipated reaction of the selves which follow. Peleg and Yaari (1973) present an example in which the date t-self games in both directions. That is the date t-self chooses its reaction function with a view to influencing the choice of its earlier self, as well as its future self. Peleg and Yaari argue that an equilibrium does not exist for their example. In their general model, Peleg and Yaari use the Nash concept of equilbrium. That is, each date t self chooses a reaction function, and in a Nash equilibrium, the chosen functions are best responses to the reaction functions of all the other selves. Goldman argues that the equilibrium concept which Peleg and Yaari discuss in the context of their example is non–Nash, and hence at odds with their general model. Moreover, he asserts that the example does actually possess a Nash equilibrium.[18]

Goldman's existence theorem is quite general. In addition to admitting indifference, the utility function at each date need not be additively separable, and can have as its domain the entire consumption path. Consequently, the proof will also cover the case of endogenously changing tastes, which is discussed in section 9. Below is a sketch of Goldman's existence proof.

Theorem 4 *Let the utility function at each date be continuous in c. Then a Strotz–Pollak equilibrium exists.*

Notice that in Goldman's proof, indifference is resolved by punishing past preferences

[18] In the Peleg–Yaari model, there is no requirement that the Nash equilibrium be perfect. Goldman raises the concern that as a result, when considering the implications of choosing a non–equilbrium strategy, the date t self may make erroneous assumptions about the behavior of its future self.

for deviations from a particular reaction function. In this respect, von Auer (1999, 2004) explores some of the structural issues arising from indifference. His work deals with the characteristics of four specific sophisticated choice mechanisms.

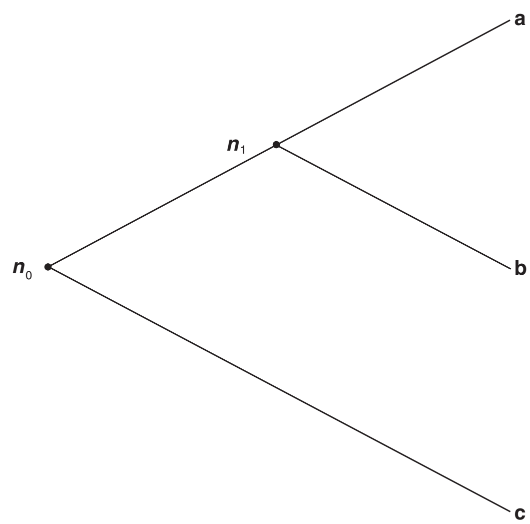

Figure 3: Hammond Decision Tree.

In order to define these mechanisms, consider the following example due to Hammond (1976a). In Figure 3, a decision maker faces a two period decision tree with two decision nodes n_0 and n_1, and three possible intertemporal outcomes a, b, and c. At the initial date (0), the decision maker's preferences are $bPcPa$. However, date 1 preferences feature $aPbPc$. At each decision node, the decision maker must select exactly one branch. Observe that a naïve decision maker would select the top branch at node n_0 with the intent of achieving outcome b, only to find themselves selecting a at node n_1. Hence a represents naïve choice (c^N), and b represents precommitment(c^1). A sophisticated decision maker will understand the full implications of selecting the top branch at n_0, and will choose the lower branch instead, thereby leading to outcome c. That is, c represents consistent choice (c^*).

Consider the use of backward induction to arrive at the choice sets attached to each node. For instance, when n_1–preferences feature aPb, then $C(n_1) = \{a\}$. In such a case, von Auer (1999) states that b can be termed *irrelevant* for the choice at n_0: only a and c are *relevant*. If n_0–preferences feature cPa, (written $cP(n_0)a$), then by induction, $C(n_0) = \{c\}$.

The generation of choice sets through backward induction enables us to specify which outcomes are eliminated through strict preference. In the absence of indifference, all choice sets are singletons. This is not generally true when indifference is involved.

Consider an example involving indifference. For instance, let $aI(n_1)b$, $bP(n_0)c$, and $cP(n_0)a$. Here $C(n_1) = \{a, b\}$. What will an individual do if his earlier self chose the top branch, thereby confronting him with indifference at n_1? If he resolves the indifference by assigning his earlier (or earliest) preferences seniority, then von Auer labels his choice mechanism as being *dogmatic*. Hence dogmatic choice in this example stipulates that the tie at n_1 be broken in favor of b, since b is strictly preferred at n_0.

Dogmatic choice is the first of four sophisticated mechanisms described by von Auer. It uses seniority to achieve resolution among relevant options. The second is *sagacious choice*. This rule uses the seniority of $bP(n_0)a$ to eliminate a from $C(n_0)$ *only if* a and b split at n_0. In this example, a and b emanate from the same branch at n_0, and so a would not be eliminated by appeal to b. Of course, $cP(n_0)a$, and c does split from a at n_0.

A third mechanism, even less restrictive is *lenient choice*. Here a will not be eliminated as long as it is relevant, and it emanates from the same branch as some other option which is undominated in the set of relevant options. In this case $C(n_0)$ would contain both a and b, meaning that the decision maker understands that he could end up at either terminal node when he chooses the upper branch.

The last mechanism proposed by von Auer is *cautious choice*. When this mechanism is used, the decision maker at n_0 proceeds as follows. Branch by branch, he ascertains what would be the worst possible way, from the perspective of his n_0 preferences, that his future self might break an indifference tie. He then selects the branch offering the best of the worst cases. In other words, he invokes the maximin principle.

Of the four sophisticated choice mechanisms, von Auer establishes that three produce choice sets consistent with subgame perfect equilibria. The exception is leniency, and it is easy to see how subgame perfection would be violated in this case. Let $aP(n_0)cP(n_0)b$ and $aI(n_1)b$. Then lenient choice permits b to belong to $C(n_0)$. But the decision maker would definitely select c at n_0, if he anticipated that the upper branch would lead to b.

The preceding analysis can be extended to cover the case of uncertainty, for which von Auer establishes a series of general results. In Section 11, I discuss some of the issues which arise when tastes change in an uncertainty framework.

8 Welfare Implications

In the proof of Theorem 3 we encountered the choice problem of an individual who takes full account of how his future self will behave. The proof involved an explicit comparison of the dynamically consistent plan c^* with the plan c^1 which is optimal from the perspective of the date 1 utility function U_1. Whereas the individual judges c^1 to be superior to c^* under utility function U_1, his future self will typically disagree, since his future self uses different utility functions. Consequently, it is no longer possible to speak unambiguously about *the individual's* choice as being optimal.

One may wonder whether the consistent choice can be regarded as optimal relative to some meaningful aggregate of the utility functions experienced over the individual's lifetime? The general answer is no. Hammond (1976a) makes the argument that when an individual has an interest in precommitting his future self, then his choice function does not conform with an underlying preference ordering.

Consider the sophisticated decision maker's choice function $C()$ in Hammond's example (Figure 3). Recall that at the initial date (0), the decision maker's preferences feature $bPcPa$, but at date 1 preferences feature $aPbPc$. Confronted with the full set $\{a, b, c\}$ of alternatives, the decision maker selects $C(\{a, b, c\}) = \{c\}$. However, consider a subset of alternatives, say $\{b, c\}$. In this case, the sophisticated decision maker is precluded from selecting a at n_1. Hence the sophisticated decision maker can initially select the top branch, knowing that the choice will conclude with b. That is, $C(\{b, c\}) = \{b\}$.

If the decision maker's choice is to conform with an underlying long run preference ordering, then it is necessary that his choice function conform with Sen's condition of α–rationality [Sen (1971)]. In the context of this example, α–rationality states the following: if c is in the choice set $C(\{a, b, c\})$, and c belongs to a subset X of $\{a, b, c\}$, then c must lie in the choice set $C(X)$. In the preceding example α–rationality is violated, with $X = \{b, c\}$.

Hammond's example illustrates an additional point about the decision maker's choice function. It may feature Pareto–inefficiency. Notice that from the perspective of *both* preference orderings, the precommited choice b is preferred to sophisticated choice $C\{a, b, c\} = \{c\}$. As we shall now see, Pareto–inefficiency is a general feature of intertemporal choice when the discount function is hyperbolic. Specifically, we have:

Theorem 5 *Under hyperbolic discounting, there exists a consumption path c' which satisfies $\sum_{t=1}^{T} c'_t = q$, involves lower consumption at all dates before T, and Pareto–dominates the consistent plan c^*. That is, for all $\tau \leq T$, c' is judged superior to c^* under utility function U_τ.*

Goldman (1979) proves a result similar to Theorem 5, but under different hypotheses. In particular, his framework relaxes the assumption of additive separability, and replaces it with a set of restrictions on the reaction functions. One of these restrictions states that "an increase in the first generation's consumption would result in a decrease in the consumptions of all later generations at a rate bounded away from zero" (p. 622). As we saw earlier, this restriction does not hold as a general condition in the present framework.

9 Endogenously Changing Utility Functions

The discounted utility model assumes that instantaneous utility $u(c_t)$ is not directly affected by past consumption levels. An intriguing model with endogenously changing utility functions was proposed by Yaari (1977). In his model, the utility function at each date τ has the form

$$U_\tau = \int_\tau^\infty e^{-p(\tau)t} u(c(t))dt, \tag{21}$$

where $p(\tau)$ is endogenously determined. Specifically, $p(0)$ is given at $\tau = 0$, and $p(\tau)$ is assumed to evolve according to:

$$p(\tau) = p(0) + \int_0^\tau c(t)dt. \tag{22}$$

In consequence the individual begins with impatience parameter $p(0)$ and concludes with an impatience parameter $p(0) + q$, after having consumed the entire wealth q. Since $\dot{p} = c(\tau)$, the level of impatience is endogenous: it not only rises continuously as wealth is depleted, but it is rises with the size of each "bite." For this reason Yaari describes q as an "appetite–arousing cake."

Notably, the appetite–arousing cake model features the common difference effect, and in this respect is similar to the hyperbolic discounting model. Yaari discusses this property in terms of naïve planning. Let $\hat{T}(p, q)$ be the length of time an individual with impatience parameter p would plan to consume a cake of size q, under the naïve belief that his impatience parameter will remain invariant at p over time. Yaari assumes that $0 < u' < \infty$, which implies that $\hat{T}(p, q) < \infty$. Let $T(p, q)$ be the actual time it takes for a naïve individual to consume a cake of size q, when his initial impatience parameter is p. Yaari demonstrates that $T(p(\tau), q(\tau)) < \hat{T}(p(\tau), q(\tau))$ for every τ featuring a positive consumption rate. In other words, the cake will be consumed faster than was called for by any of the naïve plans formulated by the individual.[19]

In a discrete time version of the appetite arousing cake problem, let $\lambda_\tau(t) = (1/y_\tau)^{t-\tau}$, where y_τ is a positive, monotone increasing function of past consumption $\sum_{t<\tau} c_t$. As established below, a version of Yaari's result holds in the discrete version of the model. In the argument below, keep in mind that what makes a plan naïve is the individual's assumption that he or she will be able to follow the precommitment plan, but without precommiting. Therefore, the issue at hand is that a naïve individual completely consumes

[19]In Yaari's continuous time model, $c(t)$ is given by $(\partial T/\partial q)^{-1}$. This relationship implies that the precommitment plan features slower consumption rates at each moment than the consistent plan.

the cake at an earlier date than an individual who is able to precommit at $t = 1$.[20]

A related issue pertains to the difference at $t = 1$ between the precommitment plan and a consistent plan. In this respect, notice that the discount factor $1/\tilde{\lambda}_1$ applied to the utility of consumption at $t = 2$, as well as the size of cake q, are the same for both plans. Theorem 6 below establishes that when appetite-arousal is sufficiently weak, the individual following a consistent plan will not consume the cake more slowly than the same individual following a precommitment plan.

The proof of Theorem 6 below treats the case in which discount factors are fixed, and hence exogenous, before analyzing the impact of discount rates being functions of past consumption. Yaari's equation (22) is quite specific, and for large q implies that impatience becomes extreme as the cake is being consumed. In contrast, the discussion below focuses on a more general specification for the sensitivity of the discount rate to past consumption levels. The more general approach allows for a discussion of how the sensitivity of discount rates to past consumption impacts the respective times for the cake to be consumed completely under both plans.

Theorem 6 *1. Let T_P and T_C be the minimum length of time needed to consume wealth q completely along the precommitment and consistent plans respectively. Then for a sufficiently small degree of appetite-arousal, $T_C \leq T_P < \infty$. There are parameter specifications for which the inequality is strict, and there are parameter specifications for which equality holds.*

2. Let T_N be the minimum length of time need to consume wealth q completely along a naïve plan. Then $T_N \leq T_P$.

The naïve investor who at $t = 1$ sets out to follow the precommitment plan, but without precommitment constraints, will be surprised by what next transpires. Relative to anticipations at $t = 1$, the surprises will include the higher degree of impatience at

[20]Yaari also established the Pareto–inefficiency theorem for his model. An interesting feature of his result is that the comparison of the dynamically consistent path and Pareto dominating comparison path is accomplished from the utility functional corresponding to each impatience parameter p, rather than the utility function at each time t. Under the endogeneity property, these are not the same.

$t = 2$, the corresponding higher ratio c_2/q_1, the lower bequest q_2 at the outset of $t = 3$, the higher ratio c_3/q_2, and so on. If the future discount rate is high enough, then $T_N \leq T_P$, given that $T_P \geq 3$. Yaari's point is that when the individual is naïve, and recomputes the precommit after experiencing the surprises just described, going forward, he or she will continue to be surprised in the same way.

As was noted in the proof of Theorem 6, for the naïve investor, the endogeneity of appetite-arousal exerts no impact on the time it takes to consume the cake completely, relative to the case of exogenous hyperbolic discounting. In contrast, the proof identified conditions under which endogeneity reduced T_C relative to the exogenous case. Indeed, for the consistent plan, it is possible that when endogeneity a sufficiently strong, the inequality $T_C > T_P$ holds. To see why, consider the following example with the assumption that q is sufficiently large that $T_P \geq 4$. Focus attention on two plans, both of which involve precommitment with constraints.

- The constraint in the first plan is that $c_t = 0$ for all $t > 2$. Denote this constrained precommitment plan by c^2.
- To define the constraint in the second plan, focus attention on date $T_P - 1$, and consider the bequest q_{T_P-1} that is available at the beginning of $t = T_P - 1$. This bequest is not large enough to induce positive consumption at date $T_P + 1$. For $\epsilon > 0$, consider the smallest value of bequest at the beginning of $t = T_P - 1$ that would result in utility maximizing consumption at $t = T_P + 1$ in the amount ϵ. Here utility maximization is undertaken under date 1 preferences, and the requirement that consumption at $T_P + 1$ be equal to ϵ is a constraint. Define $q^{\epsilon}_{T_P-1}$ to be the bequest at $t = T_P - 1$ associated with this constraint, and the corresponding constrained maximizing consumption plan to be the relatively small shift from c_P to c^{ϵ}. Implementing c^{ϵ} requires a redistribution from consumption at all dates before $T_P - 1$ to $T_P - 1$.

Suppose that preferences are exponential at every date, but possibly featuring different rates of time preference. In addition, suppose that the rate of time discount at date t is a function of cumulative consumption through date $t-1$. Consider an appetite-arousal process for which the impact of past cumulative consumption at t is zero, when cumulative consumption is less than or equal to $\sum_{\tau=1}^{t-1} c_{\tau}^{\epsilon}$. Assume that when this is the case for the entire consumption plan, the rate of time discount will be the same at all dates. Now suppose that past cumulative consumption at t exceeds past cumulative consumption along c^{ϵ} at t by ϵ/T_P. In this case, assume that the rate of time preference at t will be will be so high that any positive bequest will be consumed at t. In particular, if consumption at $t=1$ exceeds date 1 consumption along c_{ϵ} by more than $\epsilon/2T_P$, then the entire bequest q_1 will be consumed at $t=2$. In particular, these assumptions imply that if an individual attempts to implement a plan that is maximizing under date 1 preferences, the resulting appetite-arousal will cause the cake to be completely consumed at $t=2$.

As precommitment plans, both c^2 and c^{ϵ} are suboptimal under date 1 preferences. Choose total consumption q so that c_{T_P} is within ϵ of a tipping point at T_P+1. This choice makes $q_{T_P-1}^{\epsilon}$ and q_{T_P-1} within ϵ of each other, and therefore small. That is, the magnitude of the required reallocation from c_P to c^{ϵ} to implement the increase in bequest at T_P-1 is small. As a result, in absolute value, the decline in total utility that stems from the relatively small shift from c_P to c^{ϵ} can be chosen to be less than than the corresponding decline in utility from the large consumption shift from c_P to c^2.

It follows that under date 1 preferences, the plan c^{ϵ} is preferred to the plan c^2. To avoid choosing c^2 as the consistent plan, consumption at $t=1$ will have to be no more than $\epsilon/2T_P$ greater than c_1^{ϵ}. Indeed, a similar statement applies at every $t \leq T_P-1$: At every t, the maximizing choice is to keep c_t close to c_t^{ϵ} in order to avoid the highly sensitive appetite-arousal process from causing the remaining bequest to be consumed in short order. Hence, the consistent plan will lie in an ϵ/T_P-neighborhood of c^{ϵ}. By construction, $T_C = T_P + 1$. Therefore, in this example, $T_C > T_P$.

10 Changing Tastes and Consumer Theory

Until now, I have focused on the case in which there is a but a single commodity. However, changing tastes have also played a role in the development of consumer choice theory. In Section 6, I discussed an example due to Pollak in which utility is logarithmic in intertemporal consumption. In this section, I describe a multi–commodity model in which utility is logarithmic across commodities.

Let p_i be the price of commodity i, x_i the quantity consumed of the $i-$th commodity, and M the consumer's total expenditure (income). Consider a utility function of the form:

$$U(x_1, x_2, \ldots, x_n) = \sum_{k=1}^{n} a_k ln(x_k - b_k), \tag{23}$$

where for all k, $a_k > 0$, $x_k > b_k$, and $\sum a_k = 1$. Notice that with this utility function, the consumer's utility becomes infinitely negative as x_k approaches b_k. For this reason, b_k can be regarded as a subsistence level. With this interpretation of b, we can view the preceding utility function as being logarithmic (*i.e.*, Cobb–Douglas) in increments over and above subsistence. That is, the consumer first allocates the portion $\sum_k b_k p_k$ of income M to subsistence. He then allocates the residual $M - \sum_k b_k p_k$ according to the constant budget share rule which characterizes maximum logarithmic utility. Hence, the demand function corresponding to the preceding utility function is:

$$x_i = b_i - \frac{a_i}{p_i} \sum_k b_k p_k + \frac{a_i}{p_i} M, \tag{24}$$

which has come to be known as the Klein–Rubin–Stone–Geary linear expenditure system. Observe that this demand system exhibits linear Engel curves, a feature which plays a prominent role in one of the key results described below.

It appears that Stone (1954) may have been the first to suggest that the linear expenditure system could accomodate habit formation by letting some of the parameters depend on past consumption. In Pollak's (1971) model of habit formation, the subsistence level $b_{i,t}$ in the linear expenditure system takes the following form:

$$b_{i,t} = b_i^* + \beta_i x_{i,t-1}. \tag{25}$$

In Pollak's formulation, the individual becomes habituated to an amount which varies as a linear function of the previous level at which he consumed commodity i. Pollak suggests that the reference consumption level b_i^* can be thought of as the "physiologically necessary " component of $b_{i,t}$, and $\beta_i x_{i,t-1}$ as the "psychologically necessary" component. With this modification, the linear expenditure system gives rise to demand functions which are contingent on the consumption levels at the previous date. These demand functions are termed *short term*. Notice that in the multi–commodity setting, $b_{i,t}$ plays the role of a *reference bundle* in respect to the individual's moving frame of reference.[21]

The literature on habit formation has been concerned with the question of whether demands $x_{i,t}$ converge to steady state levels x_i when prices and expenditure levels remain invariant over time. Such steady state values are termed *long term demands*. Recall that for the short term demand functions, the reference consumption level for commodity i depends upon the past consumption of no other commodity save i. However, this is not so for long term demands, where reference consumption for *every* commodity i depends on the previous consumption levels of all commodities. This is because in the steady state, the budget constraint forces the process of habituation for each commodity to depend on the consumption levels of the other commodities.

The existence of long run demand functions gave rise to the question of whether there might be a long run utility function which would rationalize these long run demands. Peston (1967) considered a two–commodity model with logarithmic (Cobb–Douglas) utility, and derived a long run utility function which rationalizes his long run demand function. Gorman (1967) considered a more general model with n commodities in which short run

[21]Stigler and Becker (1977) propose a model of habit formation (addiction) in which past consumption leads to the creation of capital. For example, listening to music in the past leads to the production of music capital. This capital increases the current utility produced from a given quantity of current musical input, just as more physical capital enables labor to be more productive. Their model is similar in structure to the habit formation models described here, although they offer a different interpretation of the factors underlying addiction.

utility depends on both a quantity vector and a vector of taste parameters. He presents a set of conditions under which a long run demand function can be represented by a long run utility function.

The existence of long run utility functions immediately raises questions about long term welfare. Von Weizsäcker (1971) addressed this issue in a two commodity model. His short run demand system can be expressed as:

$$x_{i,t} = f_i(p_1, p_2, M, x_{1,t-1}, x_{2,t-1}) \tag{26}$$

for $i = 1, 2$. Let

$$g_{i,j} = \partial x_{i,t} / \partial x_{j,t-1}. \tag{27}$$

Von Weizsäcker proves if $|g_{1,1} + g_{2,2}| < 1 - \epsilon$ for some $\epsilon > 0$, then there is a unique steady state point (x_1, x_2). In this case, there is a long run utility function over the set of steady state solutions.

Changing tastes can give rise to myopic behavior in which a sequence of short term improvements leads to a worsening long run situation. Figure 4 (taken from von Weizsäcker (1971)) illustrates a sequence of consumption points, and the indifference curves on which they appear. Notice that the consumer appears to move to a higher indifference curve from one date to the next. However, the sequence of points is moving to the origin. It is because of pathological examples such as these that von Weizsäcker suggests using long run preferences as the *real* indicators of welfare, rather than the short term preferences. In particular, he shows the following. Let $0 < g_{1,1} + g_{2,2} < 1$, and consider two steady state points x^1 and x^2. If x^2 has higher long run utility than x^1, then it will be possible to move from x^1 to x^2 through a sequence of short run choices, in which short run utility improves at every transition along the sequence.

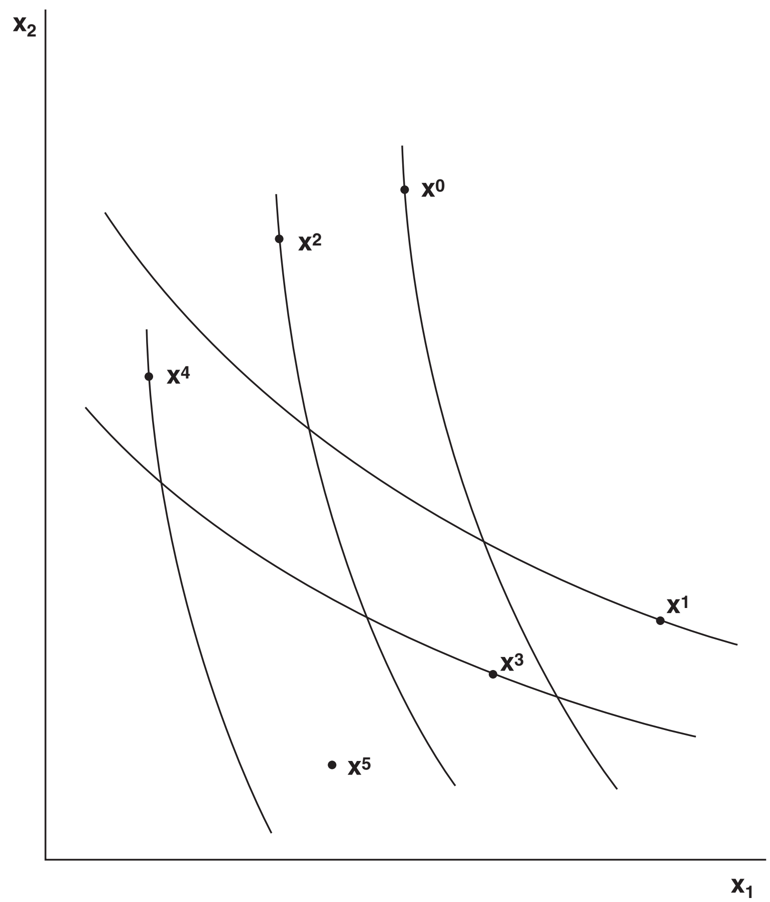

Figure 4: This figure depicts short run improvement, but long run deterioration.

Pollak (1976) argues that there are at least two serious difficulties with the welfare portion of von Weizsäcker's (1971) argument. The first difficulty is that x^2 may not appear to be a short run improvement over x^1 when x^1 is the previous consumption point. To illustrate his concern, Pollak provides an example using the linear expenditure habit formation system in which:

$$b_{i,t} = 1/2x_{i,t-1} \tag{28}$$

and $a_k = 1/2$ for all k. The long run utility function turns out to be the product of the two demands. Let $x^1 = (6, 4)$ and $x^2 = (4, 7)$. It is easily verified that x^2 is the preferred steady state bundle, but it appears inferior in the short run if x^1 is consumed. Pollak contends basing welfare comparisons on the existence of a particular history leading from

x^1 to x^2 is simply too fragile.

The second of Pollak's concerns involves the special nature of the assumption that there are only two commodities. Pollak argues that once there are three or more commodities, a long run utility function which rationalizes the long run demand function may not exist. He makes the argument for the class of linear Engel curve systems with linear habit formation, which includes the Klein–Rubin–Stone–Geary linear expenditure system as a special case. If the short run utility functions are additively separable, then the long run demand functions can, in fact, be rationalized by a long run utility function which also is additively separable. However, the case of additively separable short run utility curves is the *only* case for which there is a corresponding long run utility function. For linear Engel curve systems generated by short run utility functions which are not additively separable, such as the CES, no long run utility function will exist to rationalize the long run demand function. Welfare analysis based upon a long run utility function will just not be possible.

Hammond (1976b) uses choice theory to analyze the general relationship among short term preferences exhibiting path dependence, the acyclicity of long run preferences, and the global stability of long run choice. One of the key concepts in Hammond's analysis is that of a *conservative sequence*. To understand what is meant by such a sequence, suppose that there is some date at which the individual cannot do better than make the same choice that he did at the preceding date. If he does so, then the sequence is conservative. Hammond shows that the long run choices from any compact set A are the cluster points of all conservative sequences. However, this result is consistent with multiple long run choices because "any cluster point of a conservative sequence is a long run choice. Oscillations between the neighborhoods of multiple long–run choices are still possible." (p. 338). This can be so despite the long run preference relation being acyclic. Indeed Hammond provides an example in which this property holds despite there being a continuous long run utility function. In the example, the objects of choice are real numbers. The utility of an object is its distance from the interval $[2, 4]$. Imagine a sequence of points which alternately approach $[2, 4]$ from the left of 2, and the right of 4. Let each point be closer to the interval $[2, 4]$ than its predecessor, and let 2 and 4 be cluster points of the sequence. Then both cluster points are in the choice set, and the preference relation is acyclic.

11 TF&S-Economics: Simultaneous Internal Conflict and Two-System Thinking

11.1 Planner-Doer Framework

In the two-system psychological framework, System 1 operates quickly (fast thinking), is automatic, intuitive, and requires little effort. In contrast, System 2 operates slowly, is deliberate, and requires intentional effort. The neoclassical framework is effectively a one-system framework, and that system is System 2. The intrapersonal game theoretic approach to present bias is neoclassical

The planner-doer approach to present bias [Thaler and Shefrin (1981), Shefrin and Thaler (1988)] is a two-system approach. System 1 is identified with a set of "doers" and automatic "rules." System 2 is identified with a "planner" optimization, involving the exercise of willpower, along with rule selection.

There are several differences between the two-system planner-doer approach to present bias and the intrapersonal game theoretic approach. The main technical difference is that the planner-doer approach embodies present bias within the marginal utility functions, while, as discussed above, the intrapersonal approach embodies present bias within the time preference functions.

The planner-doer approach associates a neoclassical utility function to the planner, to capture the idea of the planner being a rational actor. In this regard, the analysis below uses log-utility; however the approach works just as well with other utility functions, such as power utility. Log-utility is useful because of its tractable features, and also because it is the fittest function (equivalent to being entropy minimizing).

The behavioral life cycle hypothesis (BLC) is a planner-doer approach for analyzing saving and consumption over the life cycle. The drive to consume at a given date t is captured by the doer. The date t does has a utility function Z_t, to be interpreted as *experienced utility*. Z_t has several arguments, one of which is the consumption level c_t. In the planner-doer framework the impulse to consume originates in the region of a person's brain known as the limbic system. In this regard, the impulse to increase consumption beyond level c_t is represented by marginal utility $\partial Z_t/\partial c_t$. This impulse registers in the prefrontal cortex of the brain, which is the seat of System 2 thinking. A planner-doer conflict, stemming from present bias, involves $\partial Z_t/\partial c_t > dln(c_t)/dc_t$. In this model, self-control is about *impulse control*.

Consumption at date t might be bounded by above, by either an external boundary such as a liquidity constraint, or an internal rule. In the BLC, internal rules are enforced automatically through System 1, using habit. In this regard, Adriaanse, Kroese, Gillebaart, and De Ridder (2014) discuss the psychological elements associated with the difference between habitual behavior and conscious effort. If $\partial Z_t/\partial c_t > 0$ at c_t, and there is no boundary in place at that level of consumption, then the planner will need to ascertain, consciously, whether to inhibit the doer's impulse to increase consumption beyond c_t at t. In the BLC, such inhibition takes the form of increased willpower, an activity assumed to require effort, and to be costly. The willpower variable at date t is denoted θ_t.

In this approach, doer impulses and habituation are treated as System 1 elements, which is consistent with the psychology literature on two-system frameworks described in the introduction. Habitual behavior in economic choice is a key feature in Thaler and Shefrin (1981), who analyze decisions about whether to rely on rules or discretion. This decision is also analyzed by Hayashi and Takeoka (2022) who use a recursive utility framework.[22]

If the planner uses willpower to induce the doer's drive for additional consumption beyond c_t, then the planner will need to induce Z_t to achieve a local maximum at c_t. At a local maximum, $\partial Z_t/\partial c_t = 0$, so that the drive for additional consumption is curbed:

[22] In the planner-doer framework, the planner changes the doer's utility at a cost. By way of contrast, in the model developed by Hayashi and Takeoka, the decision-maker changes the future utility at a self-control cost.

that is, so that doer t feels satiated. In the BLC, increasing θ_t is assumed to reduce both consumption c_t and, because willpower is costly, to reduce doer utility Z_t as well. Notably, present bias occurs when costly willpower and the inability to impose consumption boundaries lead to excessive consumption at date t from the wealth remaining at the outset of date t. The discussion below will clarify what is meant by "excessive consumption."

For sake of illustration, assume that when the planner does not exercise willpower, the Z_t function is logarithmic in c_t. Figure 5 below illustrates how the exercise of willpower folds back the logarithmic utility function in order to induce a desired level of consumption, c_t^*, and reduces the Z_t function as a result of the exercise of willpower (θ_t). In Figure 5, the curve labeled $Z-$reduced form is the locus of local maxima in Z_t for varying levels of θ_t. A critical feature of Figure 5 is the presence of a variable m which measures the amount of temptation being confronted from a choice set known a *mental account*.[23] The concepts of temptation and mental accounts are discussed in detail below. At this stage, the point to note is that greater temptation leads to higher willpower costs.

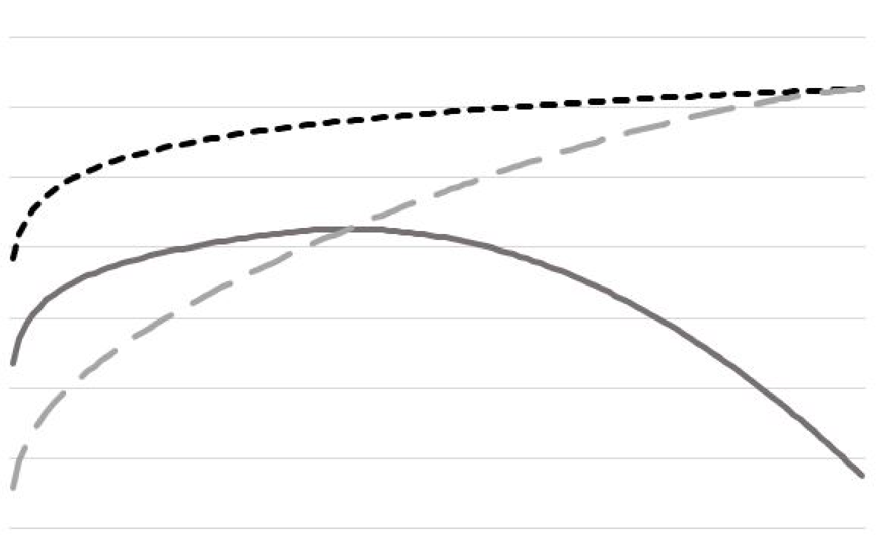

Figure 5: This figure displays the logarithmic utility function (dotted black curve), the doer utility function $Z_t(c_t, \theta_t, m_t)$ (solid gray curve), and the reduced form function $Z_t^R(c_t, m_t)$ (dotted gray curve), all associated with a particular mental account, in this case *the income account*. The difference between Z_t^R and the logarithmic function reflects decision fatigue associated with the effort to reduce consumption below the temptation

[23]The theory features several mental accounts indexed by i, so that m is indexed by both i and t.

level m_t associated with the income account. In respect to fatigue, System 2 mental activity requires much more effort than System 1 mental activity.

Upper bounds on consumption provide an alternative to the exercise of costly willpower. These might take the form of external liquidity constraints, or they might take the form of internal constraints which are imposed through habitual rules of thumb. Habit is a System 1 activity, and in the BLC, not all habits are enforceable. For instance, consider the rule of thumb "don't dip into capital." This rule of thumb is a prohibition that pertains to the portion of wealth an individual classifies as capital, rather than income. The Modigliani-Miller principle holds that in the absence of taxes and transaction costs, the form in which people hold wealth is irrelevant. The behavioral approach instead holds that even in the absence of taxes and transaction costs, the form of wealth can be relevant.

The BLC assumes that enforceable habits are connected to natural *mental accounts* associated with classifications such as income and capital. In this respect, the BLC focuses on three distinct mental accounts, respectively relating to income, liquid assets, and future income. Funding consumption from the income account, but not from the asset account or the future income account is a potential way to reduce willpower costs.

The discussion below uses credit card borrowing to illustrate the concept of invading a future income account. Paying the full monthly credit card balance is a way to avoid associated borrowing. On the other hand, paying the monthly minimum leads to borrowing. Individuals for whom willpower costs are low will not be tempted to overspend significantly using their credit cards. On the other hand, individuals who face high willpower costs and use their credit cards will be tempted to overspend significantly. If the degree of overspending is sufficiently great, such individuals might be better off cutting up their credit cards, to prevent access to their future income accounts.

Willpower is about *restraint relative to a reference point.* A central tenet of the BLC is that the cost of exercising willpower reflects the difficulty in restraining consumption from mental accounts. Large mental account balances can be tempting, and the greater the temptation, the greater the restraint required to achieve a particular level of consumption. Saving is the act of restraining consumption from the income account, with the reference point being income. This means that willpower costs are associated with the effort required to save, meaning reducing consumption below income. As a result, consumption is determined as a residual variable, based on how much saving takes place from the income account. The larger the amount saved, the greater the amount of willpower required. The same principle holds for asset accounts such as credit cards, where the reference point is the credit limit.

The residual variable issue might seem minor; however it is a major feature of the *psychology of saving and spending*, and of profound importance for the nature of predictions about life cycle behavior patterns.[24]

The next subsection describes a formal BLC framework. At the heart of the framework is a willpower penalty function. The penalty function is displayed in Figure 5 as the difference between the logarithmic function at the top of the graph and the reduced form doer utility function below it. Notice that the penalty is zero at the far right, where the individual consumes the full balance of the associated mental account. Moving from right to left along the graph involves increasing degrees of restraint, increasing degrees of willpower cost, and lower consumption. As displayed, the marginal cost of willpower has a convexity property: it increases with the amount of willpower exercised.[25] The gap between the two curves at the far left is an entry fee associated with the mental account. Entry fees play the role of guardrails. Guardrails inhibit, but might not prevent mental accounts from being invaded. Because of the entry fee, the individual, once an account is invaded, will consume a discrete chunk of the account, and will not nibble. In the credit card analogy, individuals will be inclined to avoid paying the full monthly balance, save for a dollar. That is, they will not intentionally borrow a dollar, but instead borrow zero

[24]In this respect, Gul and Pesendorfer (2001) develop an axiomatic theory around this concept, emphasizing the role of temptation in self-control issues. Shefrin (2020) describes why the introduction into Shefrin and Thaler (1988) of two concepts, the temptation variable m and a mental accounting structure described below, constitute the most important modifications to the planner-doer model presented in Thaler and Shefrin (1981). To the best of the author's knowledge, Shefrin and Thaler (1988) provides the first formal economic model describing the role of temptation in self-control issues. Although Gul and Pesendorfer cite neither Thaler and Shefrin (1981) nor Shefrin and Thaler (1988), the Gul-Pesendorfer temptation framework shares important similarities with the planner-doer approach. Gul and Pesendorfer do cite Rabin (1998), whose discussion about self-control and intertemporal choice, on pages 39-41, describes the two-system approach developed by Thaler and Shefrin (1981).

[25]See Muraven, Tice, and Baumeister (1998).

by paying the full balance.

11.2 Specific Functional Forms

Shefrin and Thaler developed the models in Thaler and Shefrin (1981) and Shefrin and Thaler (1988) in terms of particular assumptions, but not specific functional forms. Fudenberg and Levine (2006) describe their self-control model as being in the spirit of Thaler and Shefrin (1981). Notably, the Fudenberg-Levine model does feature specific functional forms, such as logarithmic utility and CRRA (constant relative risk aversion). Fudenberg and Levine[26] also note that their model is closely connected to the temptation feature in Gul and Pesendorfer's framework. In this respect, they model the cost of self-control using a penalty function which shares important similarities with the framework introduced in Shefrin and Thaler (1988).[27]

There are important ways in which the Fudenberg-Levine model is in the spirit of the planner-doer model. Fudenberg and Levine point out that their model provides an explicit micro-foundation for the treatment of mental accounting developed in Shefrin and Thaler (1988). They also relate their model to the neuroeconomics literature, citing McClure, Laibson, Loewenstein, and Cohen (2004): As was mentioned above, the planner-doer framework is neurologically-based. Finally, they note that their model can explain why, for some individuals, the marginal propensity to consume from income might be unity:

[26]Fudenberg and Levine (2012) extend their self-control framework to discuss self-control issues related to delay, arguing that "making decisions simultaneously is different from making them with an arbitrarily small delay." (p. 2).

[27]See Assumption 2 in Fudenberg and Levine (2006). Notably, Theorem 7 in this chapter can be used to provide a planner-doer interpretation of the Fudenberg-Levine model. On a related issue, the papers by Gul and Pesendorfer (2007) and Fudenberg and Levine (2006) provide the starting point for Hayashi and Takeoka (2022) who discuss habit formation in a valuation model having a penalty function structure. Habit formation and addiction are key elements in Thaler and Shefrin (1981). Thaler (1985) discusses addiction in the context of seductive goods, a concept described later in the chapter: see pages 207-208 of this article. Gul and Pesendorfer (2007), which cites neither Thaler and Shefrin (1981) nor Shefrin and Thaler (1988), is a dynamic extension of Gul and Pesendorfer (2001) and incorporates habit formation within the temptation utility function underlying their penalty function. In this framework, the temptation to consume an addictive product induces consumption followed by stronger addiction. Whereas in Gul and Pesendorfer (2007) temptation utility only pertains to current consumption, in Hayashi and Takeoda (2022) temptation utility is also forward looking. On account of this forward looking feature, the decision maker might be tempted by the entire subsequent path of the consumption stream.

This last issue is a central point in Shefrin and Thaler (1988).[28]

To develop a formal structure for the BLC with specific functional forms, consider the following formal structure for mental accounts. Index the mental accounts I (income), A (Asset), and F (future income) by $i = 1, 2, 3$ respectively. Consider a mental account i, and let $c_{t,i}$ be the amount of consumption at date t funded by account i. Total consumption at date t, c_t is given by $\sum_i c_{t,i}$. Given a mental account *pecking order* for funding, mental account i is active when consumption is not funded by a higher ranking account. Let account i be active, and denote by $m_{t,i}$ the maximum total consumption c_t. For sake of notation, set $m_{0,t} = 0$. Then $m_{i-1,t} \leq c_t \leq m_{1,t}$. Assume that consumption is bounded away from zero by ϵ, a minimum survival condition. For sake of discussion, focus on the case $\epsilon = 1$.

Consider a doer utility function $Z_t(c_t, \theta_t, m_{i,t})$, which is defined with respect to mental accounts, when active. For account i at date t, Z_t is given by the function:

$$Z_t(c_t, \theta_t, m_{i,t}) = ln(c_t) - \varphi_{0,i,t}$$

for $c_t < 1/\theta_t$, and

$$Z_t(c_t, \theta_t, m_{i,t}) = \zeta_{0,i,t} + \zeta_{1,i,t}c_t - \zeta_{2,i,t}c_t^2 \tag{29}$$

for $c_t \geq 1/\theta_t$.[29]

Here $\zeta_{0,i,t}$ is a function of θ_t and $m_{i,t}$, while $\zeta_{1,i,t}$ and $\zeta_{2,i,t}$ are functions of θ_t. Their specific forms, along with $\varphi_{0,i,t}$, are defined below.

$$\zeta_{1,i,t} = \frac{\theta_t^2}{1 + \theta_t} \tag{30}$$

$$\zeta_{2,i,t} = \frac{\theta_t}{2} \tag{31}$$

[28]Fudenberg and Levine (2006) develop an interesting banking model to explain why people might exhibit less impatience in intertemporal tradeoffs involving large sums of money. This issues surfaces in the discussion below about windfalls from lottery winnings.

[29]The log-utility function is used for illustrative purposes. It is a simple matter to substitute a CRRA function into (29) for the $ln-$function, with a similar statement applying to the marginal utility function $1/c_t$.

$$\zeta_{0,i,t} = ln(\frac{1+\theta_t)}{\theta_t}) - [(\varphi_{1,i,t}(\varphi_{2,i,t}(m_{i,t} - (\frac{1+\theta_t}{\theta_t})))^{\varphi_{3,i,t}}) + \varphi_{4,i,t}] - \zeta_{1,i,t}(\frac{1+\theta_t}{\theta_t}) + \zeta_{2,i,t}(\frac{1+\theta_t}{\theta_t})^2 \quad (32)$$

$$\varphi_{0,i,t} = ln(\frac{1}{\theta_t}) - (\zeta_{0,i,t} + \zeta_{1,i,t}(\frac{1}{\theta_t}) - \zeta_{2,i,t}(\frac{1}{\theta_t})^2) \quad (33)$$

The variables $\varphi_{1,i,t}$, $\varphi_{2,i,t}$ and $\varphi_{3,i,t}$ are nonnegative. In the discussion below, they are treated as constants, although they can be functions of $m_{t,i}$, to allow willpower costs to vary with account size.[30] The variable $\varphi_{4,i,t}$ is an impulse function of c_t, which is zero for $c_t \neq m_{i,t}$ and nonzero when $c_t = m_{i,t}$ for given i. In most of the discussion below, $\varphi_{4,i,t}$ is equal to zero.[31]

The following characterization theorem pertains to the reduced form for Z_t which is depicted in Figure 5.

Theorem 7 *1. For each account i, and fixed θ_t and $m_{i,t}$, the function Z_t defined by (29) is continuous and differentiable.*
2. For each account i, and fixed $m_{i,t}$, the value of θ_t which leads Z_t to be maximized at c_t is given by:

$$\theta_t = \frac{1}{c_t - 1} \quad (34)$$

3. The reduced form function $Z^R_t(c_t, m_{i,t}) = Z_t(c_t, \theta_t(c_t, m_{i,t}), m_{i,t})$ is given by:

$$Z^R_t(c_t, m_{i,t}) = ln(c_t) - [(\varphi_{1,i,t}(\varphi_{2,i,t}(m_{i,t} - c_t))^{\varphi_{3,i,t}}) + \varphi_{4,i,t}] \quad (35)$$

Z^R_t is the difference between a log-utility function and a hedonic penalty function which expresses the utility cost of exercising willpower.

[30]For example, $\varphi_{1,i,t}$ can be the product of a constant and the inverse of account size. If $\varphi_{2,i,t}$ is a constant, then it will effectively be redundant, and can be set equal to uniity.

[31]There is a short discussion below about the role played by the variable $\varphi_{4,i,t}$ when it is nonzero. That discussion pertains to miserly behavior during retirement. A nonzero value also plays a role in the case when the entry fee to an account is small, but the cost of willpower is high. This combination leads to account invasion and subsequent high consumption from the account. The text below describes another way to model this combination, involving the merger of mental accounts.

4. The first and second partial derivatives of $Z_t^R(c_t, m_{i,t})$ with respect to c_t are given by:

$$\frac{\partial Z_t^R(c_t, m_t)}{\partial c_t} = \frac{1}{c_t} + [\varphi_{1,i,t}\varphi_{1,2,t}\varphi_3(\varphi_{2,i,t}(m_{i,t} - c_t))^{\varphi_{3,i,t}-1}] \geq \frac{1}{c_t} > 0 \tag{36}$$

$$\frac{\partial^2 Z_t^R(c_t, m_{i,t})}{\partial c_t^2} = -(\frac{1}{c_t^2}) - [(\varphi_3 - 1)\varphi_{1,i,t}\varphi_{1,2,t}^2\varphi_3(\varphi_{2,i,t}(m_{i,t} - c_t))^{\varphi_{3,i,t}-2}] < 0 \tag{37}$$

for $\varphi_{1,i,t} > 0$, $\varphi_{2,i,t} > 0$ and $\varphi_{3,i,t} > 1$. The inequality in (36) is strict for $c_t < m_{i,t}$.

5. For $c_t < m_{i,t}$ where i is the active account, the partial derivative of Z_t^R with respect to c_t is:

$$\frac{\partial Z_t^R}{\partial c_t} = \frac{\partial Z_t^R}{\partial \theta_t}\frac{\partial \theta_t}{\partial c_t} \tag{38}$$

6. When willpower is employed, the sign of $\frac{\partial Z_t^R}{\partial \theta_t}$ is negative:

$$\frac{\partial Z_t^R}{\partial \theta_t} < 0 \tag{39}$$

7. The utility entry fee $E_{i,t}$ for account i at date t is given by:

$$E_{i,t} = ln(m_{i-1,t}) - [(\varphi_{1,i,t}(\varphi_{2,i,t}(m_{i,t} - m_{i-1,t}))^{\varphi_{3,i,t}}) + \varphi_{4,i,t}] \geq 0 \tag{40}$$

8. For $\varphi_{1,i,t}$, $\varphi_{2,i,t}$, $\varphi_{3,i,t}$ all positive, and $0 < m_{i-1,t} < m_{i,t}$, the left limit of $\frac{\partial Z_t^R}{\partial c_t}$ at $c_t = m_{i,t}$ is strictly less than the right limit.

9. For $\varphi_{1,i,t}$, $\varphi_{2,i,t}$, $\varphi_{3,i,t}$ all positive, and $0 < m_{i-1,t} < m_{i,t}$, as $\varphi_{1,i,t} \to 0$, the right limit and left limit of $\frac{\partial Z_t^R}{\partial c_t}$ at $c_t = m_{i-1,t}$, converge to the same value, as do the right and left limits of $Z_t^R(c_t, m_{i-1,t})$. In addition, if $\varphi_{4,i,t} = 0$, then $E_{i,t} \to 0$. If $\varphi_{4,i,t} \neq 0$, then $E_{i,t} \to -\varphi_{4,i,t} \neq 0$.

10. For general $\epsilon > 0$, $\theta_t = 1/(c_t^* - \epsilon)$. Furthermore, in equation (32), the general form of the terms $1/(1 + \theta_t)$ appearing in this expression are $1/(1 + \epsilon\theta_t)$.

Figure 6 below depicts the logarithmic utility function, and associated reduced form Z_t^R function, across consumption for all three mental accounts. The penalty function in Figure 6 is the difference between the (top) logarithmic function and the reduced form

Z_t^R. The penalty function is:

$$\Pi_{i,t} = [(\varphi_{1,i,t}(\varphi_{2,i,t}(m_{i,t} - (\frac{1+\theta_t}{\theta_t})))^{\varphi_{3,i,t}}) + \varphi_{4,i,t}] \quad (41)$$

for $m_{i-1,t} < c_t \leq m_{i,t}$.[32]

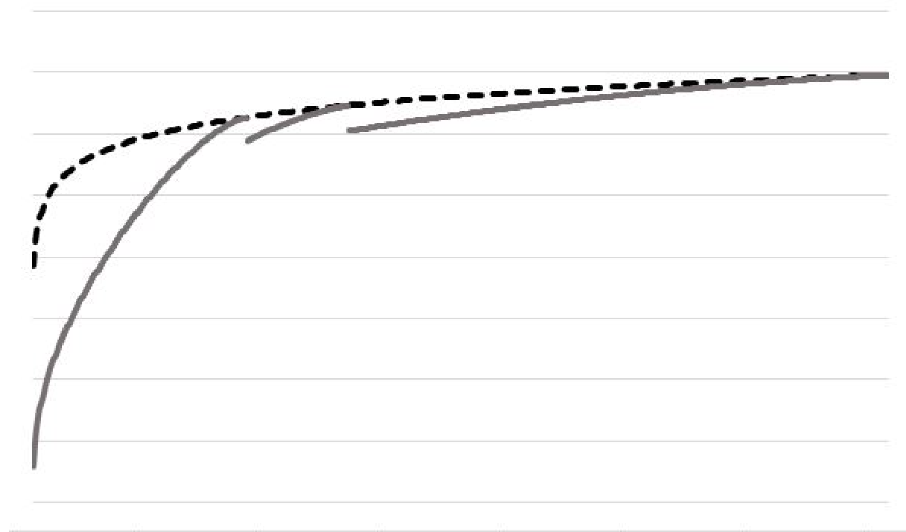

Figure 6: This figure displays, for all three mental accounts, the logarithmic utility function $ln(c_t)$ (dotted black curve) and the reduced form doer utility function $Z_t^R(c_t, m_t)$ (solid gray curve).

Focus attention on the form (35) of $Z_t^R()$ which is linear in the logarithmic function and the *nonlinear* penalty function $\Pi()$ (given by (41)). Notably, $\Pi()$ (given by (41) is quadratic, not linear. The quadratic form implies that the cost of willpower associated with an additional unit of consumption-abstinence strictly increases with the amount of abstinence. By way of contrast, with a linear penalty function, the cost associated with the marginal willpower does not vary with the amount of abstinence. A linear penalty function generates a linear reduced doer utility function, with constant marginal utility. Linear utility produces highly unstable solutions, as small changes in parameters give rise to large changes in maximizing solution.

[32]Technically, each account can be associated with its own penalty function, which is defined by (41) within the boundaries of the account, and defined to be zero outside those boundaries. Doing so allows the marginal utility function to be well defined at the upper bound of the account, and avoids the need to use a Dirac delta function in Figure 7 below.

The model developed in Gul and Pesendorfer (2001) pertains to lotteries in a risk framework.[33] The Gul-Pesendorfer axioms pertain to the case in which restricting the choice set can lead to an increase in overall utility. Their Theorem 1 establishes conditions in which overall utility $U(A)$, when facing choice set A, has the form

$$max_{x \in A}\{u(x) - [max_{y \in A}v(y) - v(x)]\} \tag{42}$$

where U, u, and v are continuous *linear* functions. Think of the second term in the above equation, $max_{y \in A}v(y) - v(x)$, as a penalty function for which $max_{y \in A}v(y)$ serves as a reference point, akin to $m_{i,t}$.[34] By way of contrast, the objective function in the BLC is the sum of discounted differences between logarithmic functions and nonlinear penalty functions. Similarly, Fudenberg and Levine (2006) effectively assume that the penalty function is convex and not linear.[35]

Figure 7 displays the marginal utility functions for both logarithmic planner utility and the doer's reduced form Z_t^R. Notice that marginal doer utility lies above its logarithmic utility counterpart, except at the boundaries of mental accounts. Statement 5 of Theorem 7 is extremely important for interpreting the meaning of marginal utility in the BLC. In the neoclassical framework, the marginal utility of consumption measures the additional enjoyment derived from consuming a marginal unit of a commodity. However, in the BLC, marginal utility measures the reduction in willpower costs associated with a marginal decrease in "saving," meaning the exercise of less restraint.

Figure 7 also displays the utility entry fee associated with invading an account. It is fixed for given $[m_{i,t}]$ but is endogenous and a function of the penalty function, including the values of $[m_{i,t}]$. The entry fee is akin to an investment, for which there is a payback period counterpart associated with the integration of the marginal utility function. If the entry fee is excessive, no amount of consumption utility will pay back the entry fee.

[33]The discussion below describes the generalization of the planner-doer model to include risk.
[34]The discussion in Gul and Pesendorfer (2004), which analyzes a consumption-saving problem, makes this point clearly and explicitly.
[35]Nonlinearity implies that individual preferences do not satisfy the Gul-Pesendorfer axioms. From a behavioral perspective, people can agree in principle to axioms being reasonable, but because of bounded rationality, fail to behave in ways that lead to the axioms being satisfied.

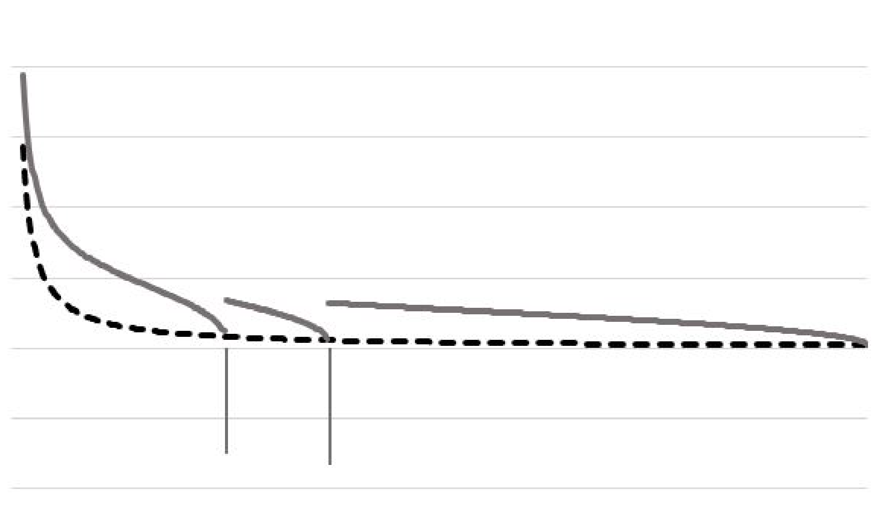

Figure 7: This figure displays marginal utility for all three mental accounts, for both the logarithmic planner utility function (dotted black curve) and reduced form doer utility function $Z_t^R(c_t, m_t)$ (solid gray curve), along with the entry fees for the accounts (vertical lines at points of discontinuity). The fundamental theorem of calculus implies that doer utility is the given by the integral of marginal utility minus the entry fee(s).

11.3 The Planner's Optimization Problem

The planner's utility function U is given by

$$U = \sum_{\tau=1}^{T} \lambda(t) Z_\tau \tag{43}$$

which is the present value of the doer-utility stream. The planner's optimization problem involves the selection of a rule R, income profile y, and willpower variables $[\theta_t]$ to maximize U subject to a set of constraints.

A rule can feature both internal and external components. The internal components relate to a mental accounting structure and choice architecture. The BLC focuses on three accounts, namely income, asset, and future income. The variables and equations associated with a mental accounting structure are as follows.

Define $c_{t,I}$, $c_{t,A}$, and $c_{t,F}$ as the amount of consumption c_t respectively funded by the income, asset, and future income accounts. Total consumption c_t is given by the sum of

$c_{t,I}$, $c_{t,A}$, and $c_{t,F}$. The account balance equations are as follows. The balance for the income account at date t is:

$$I_t = y_t \tag{44}$$

For the asset account, the balance is:

$$A_t = (1 + i_t)A_{t-1} + (y_{t-1} - c_{t-1,I}) - c_{t-,1,A} - c_{t-,1,F} = 0 \tag{45}$$

and for the future income account, the balance is:

$$F_t = \sum_{\tau > t} \frac{y_\tau}{(1 + i_\tau^{\tau - t})} \tag{46}$$

In a world without liquidity constraints, the budget constraint at date t is

$$c_t \leq I_t + A_t + F_t \tag{47}$$

In this formulation, A_t can be negative, to account for borrowing against future income. To ensure repayment of such borrowing, A_{T+1} is constrained to be nonnegative.

The internal component of one mental accounting rule is:

$$c_{t,I} \leq I_t$$

$$c_{t,A} \leq A_t$$

$$c_{t,F} \leq F_t \tag{48}$$

where (48) refers to the three inequalities above.

The mental accounting choice architecture refers to the ordering with which accounts are invaded to fund consumption. In this section, the focus is on the order I, A, F where account i is exhausted before account $i+1$ is invaded. However, other architectures are possible, which allow for more than one account to be active simultaneously.

A rule featuring the $[I, A, F]$-architecture requires discipline to enforce, a feature associated with System 1. A rule featuring less discipline might involve two accounts, such as $[I + A, F]$, where the income and asset accounts are merged. Here the individual has no psychological boundary between income and assets. The first two inequalities in (48) collapse to

$$c_{t,I} + c_{t,A} \leq I_t + A_t \tag{49}$$

A rule describing a complete lack of discipline might involve just one account, $[I+A+F]$ with (48) replaced by the budget constraint (47).

Rules can feature external components in the form of additional constraints, imposed externally, such as the following:

$$c_{t,I} \leq (1 - d_{t,I})y_t \tag{50}$$

$$c_{t,A} \leq max\{0, d_{t,A}A_t\} \tag{51}$$

$$c_{t,F} \leq d_{t,F}F_t \tag{52}$$

where $d_{t,I}$, $d_{t,A}$, and $d_{F,t}$ are nonnegative parameters: $d_{t,I}$ is a deduction rate for income, and $d_{t,A}$ and $d_{t,F}$ are liquidity constraint parameters for the asset and future income accounts respectively. Some d variables might be at the discretion of the individual, such as deduction rates in retirement plans. Others might be liquidity constraints imposed by the market. In practice, most retirement saving is accomplished through the use of institutional programs, which are external components. These programs feature accumulation and protection, and in many cases are effective substitutes for having to use willpower to save for retirement.

The combination of internal and external constraints define the upper bound temptation variables $m_{i,t}$. External constraints can be hard, and feature high values of $[\varphi_{4,i,t}]$, which effectively preclude the invasion of account i. For the $[I, A, F]$-case discussed above, the $[m_{i,t}]$-variables are given by (50), (51), and (52). The case in which $d_{t,I} = 0$ and $d_{t,A} = d_{t,F} = 1$ correspond to the internal rule described by (48). Unless specifically indicated, assume that in the discussion below, this is the case.

Rules are feasible if they are both available and psychologically implementable by System 1. Denote by D_R the set $\{d_{t,I}, d_{t,A}, d_{t,F}\}$ which are both available to the individual and psychologically implementable. In a neoclassical setting, all available rules might be implementable. However, as discussed in Thaler and Benartzi (2004), status quo bias often prevents people from implementing "good" choices that are available.

An income profile is denoted $y = [y_1, y_2, ..., y_T]$. Individuals might have choices about the shapes of their income profile. For example, some individuals might choose to work for firms that have low base pay but offer generous bonuses. Some individuals might choose equity portfolios that are concentrated in high dividend paying stocks. These features alter the flows through the income account. Denote by Y_R the set of income profiles from which the planner can choose.

A rule has three components:

1. a mental accounting structure m_R with a choice architecture;
2. a set of parameters d_R from D_R for (50), (51), and (52); and
3. a choice of income profile y_R from Y_R.

Denote by S_R the set of feasible rules $[m_R, d_R, y_R]$, those that are available and implementable by the planner. Each element of S_R has an associated set of upper bounds for consumption, denoted by $f = [f_{i,t}]$. The planner's problem can be described as choosing a rule $[m_R, d_R, y_R] \in S_R$, consumption variables $[c_{t,I},\ c_{t,A}, c_{t,F}]$, and willpower variables $[\theta_t]$ to maximize $U(Z)$ subject to the survival constraint $[c_t > \epsilon]$ for all t, and

$$[\frac{\partial Z_{i,t}}{\partial c_{i,t}} \cdot (c_{i,t} - f_{i,t}(c, y)) = 0] \tag{53}$$

(53) implies that for any active account (featuring positive consumption $c_{i,t}$), either the planner induces the date t doer to be satiated at an internal optimal solution, or else the constraint $(c_{t,i} \leq f_{i,t})$ is binding.

The optimization described above does not impose a pecking order on using mental accounts, such as "use I before A before F." To induce this pecking order into the optimization problem, add the constraint

$$c_t = c_{t,I} + (\chi\{c_{t,I} = y_t\}c_{t,A}) + (\chi\{c_{t,A} = A_t\}\chi\{c_{t,I} = y_t\}c_{t,F}) \tag{54}$$

where $\chi()$ is the characteristic function. This constraint stipulates that breaking the pecking order aspect of the rule brings no additional consumption benefits.

For the mental accounting structure $[I, A, F]$, with pecking order, the optimization problem can be written as a maximization with respect to consumption c (instead of θ) of $U(Z^R)$. Doing so casts the problem in terms of the marginal utility of consumption. Doing so requires a slightly different definition of the temptation variable m: In the above formulation, m is defined by the value of the asset and future income accounts respectively. In the Z^R-formulation, m is measured in terms of total consumption rather than the balance in a specific account.

11.4 Optimization Nuances

Consider the following nuances associated with the planner's optimization problem.

1. The planner's optimization problem features the selection of consumption variables $[c_{t,I}, c_{t,A}, c_{t,F}]$ and willpower variables $[\theta_t]$ to maximize $U(Z)$. In the context of the model, the planner chooses $[\theta_t]$ but the doers choose $[c_{t,I},\ c_{t,A}, c_{t,F}]$ to maximize Z_t. This means that if the individual chooses to breach a rule at date t, then it is in the interest of the date t doer to do so.

2. The Bellman functional equation associated with the planner's optimization problem is as follows. Define

   $$U_T = u(c_T)$$

   where $c_T = y_T + A_T + F_T$, and the arguments of U_T are y_T, A_T, F_T. The value function at date t is given by:

   $$U_t = u(c_t) + \lambda(1)U_{t+1} - \Pi(c_t, m_{i,t}) \tag{55}$$

where Π is the penalty function associated with the exercise of willpower and $\lambda(1)$ is the one-period time preference discount factor. Recall that $\Pi()$ is given by (41), a power function in $m_{i,t} - c_t$.

3. The imposition of an external component to a rule alters the temptation variables $[m_{i,t}]$ in the doer utility function Z_t^R. An external liquidity constraint typically reduces the maximum amount the individual can consume from a given mental account. Such a reduction can decrease the amount of willpower cost required to implement a given level of consumption, by reducing the amount of discretionary saving (restraint). At the same time, excessively restrictive rules will negatively impact planner utility, by overly constraining the planner's choice set.

4. A one unit marginal reduction in date t consumption, increases the value of the asset account at all future dates, ceteris paribus. At date $\tau > t$, the marginal increase in the asset account balance will be $(1+i)^{\tau-t}$. If the asset account is active at date τ, the additional account balance will generate a reduction in doer utility at that date by $\frac{\partial Z_\tau^R}{\partial m_{A,\tau}}(1+i)^{\tau-t}$. The value of $\frac{\partial Z_\tau^R}{\partial m_{A,\tau}}$ can be deduced from (35), which by (41) is the negative of the marginal penalty function with respect to consumption, at τ. Therefore, saving a marginal unit at date t creates a future temptation cost at all subsequent dates where the asset account is active. The overall magnitude of the effect is obtained by summing the discounted values of $\frac{\partial Z_\tau^R}{\partial m_{A,\tau}}(1+i)^{\tau-t}$ over all dates ($\tau > t$) when the asset account is active. This feature essentially generates an additional endogenous source of discounting, as it provides an extra disincentive at the margin for the planner to save.

5. Positive willpower costs can induce hyperbolic discounting. Consider the example associated with Figure 2. With exponential discounting, the compensation required to induce indifference for delaying the payment of a reward by three days is the same at $t = 1$ as at $t = 25$. In contrast, under hyperbolic discounting the compensation is higher at $t = 1$ than at $t = 25$. In the planner-doer model, this feature is partly driven by the fact that at $t = 1$, the date 1 doer is involved in the assessment.

Rejecting a \$1 reward at $t = 1$ requires the exercise of costly willpower at $t = 1$. Future doers have no voice at $t = 1$. When, at $t = 1$, the planner computes the amount of compensation required to delay the reward until $t = 4$, it uses the strength of doer utility at $t = 1$ and its estimate of discounted doer utility at $t = 4$. When at $t = 1$ it repeats the compensation computation for delaying the reward from $t = 25$ to $t = 28$, it uses its estimate of doer utilities for these two respective dates. Because the doer at $t = 1$ is present and salient, a naïve planner might underestimate the intensity of future doer utilities, and therefore provide a lower compensation amount for the delay from $t = 25$ to $t = 28$ than from $t = 1$ to $t = 4$. In this respect, remember that the example associated with Figure 2 involves a change on January 25 in delay-compensation, relative to the plan developed on January 1. The January 25 doer has a voice on January 25 but not on January 1.

There are circumstances when it is reasonable for the two compensation amounts computed at $t = 1$ to differ. For example, if the reward flows through the income account at $t = 1$, but through the asset account at all other dates, then the doer at $t = 25$ might have a weaker voice at $t = 25$ than its counterpart at $t = 1$. Remember that the marginal propensity to consume from income is typically higher than than the marginal propensity to consume from wealth, and why this is the case.

6. Some individuals' System 1 might lack the discipline to follow rules such as "don't dip into capital." One way to model this feature is through the mental accounting structure $[I + A, F]$. In this situation, at each date, the individual would face the temptation to spend from the sum of current income and assets. Gone would be the opportunity to constrain consumption below income without having to exercise willpower, and experience the associated willpower costs.[36]

7. In some circumstances, the planner will choose to use the pecking order rule "don't dip into capital" associated with $[I, A, F]$, but breach that rule by invading the asset account. Doing so is different from using the pecking order rule $[I + A, F]$,

[36] The discussion about marginal propensity to consume below discusses why individuals using rules featuring less discipline typically involve more present bias, but lower marginal propensities to consume.

as the latter involves a lack of sufficient discipline to follow the rule "don't dip into capital," where capital includes liquid assets.

8. As discussed above, the pecking order rule can be imposed as an explicit constraint, equation (54). However, it can also emerge as a feature of an optimizing solution. If the entry fee associated with the asset account is lower than for the future income account, then it might be optimal to invade the asset account before invading the future income account. That said, if temptation and willpower costs vary across accounts, then there are cases where the optimal solution might feature the future income account being invaded while the asset account balance is still positive. This can occur when it is relatively costly to abstain a little when invading an account, but very costly to abstain a lot. This means that several mental accounts can be active at the same time. This coincidence can occur when the temptation to spend from the asset account is very strong, relative to the future income account.[37]

9. The formal optimization framework discussed above does not incorporate status quo bias. Incorporating status quo bias formally can be accomplished by incorporating fixed costs attached to rule selection. A *default rule* would feature a zero cost. If fixed costs for other rules are sufficiently high, the planner will stick with the default rule. For sake of exposition, the notation associated with status quo bias is omitted here.

10. Most of the discussion involves the case $d = 0$. However, it is important to understand the importance of nonnegative d_1, which relates to active saving from the income account. This can be accomplished through a formal savings program with automatic deduction. It can also be accomplished through a routinized savings habit such as contributing to a Christmas club, that is, by making regular payment into

[37]The discussion below extends the framework to incorporate additional mental accounts, where similar issues arise in respect to having several active mental accounts. In respect to policy, Shefrin (2009b) suggests that instead of using conventional tax reduction to inject fiscal stimulus into the economy, governments use vouchers, essentially generalized version of food vouchers,but with short expiration dates. In a mental accounting framework, entry fees serve as obstacles to prevent reductions in spending in other accounts, to offset increased spending in the voucher account; and the expiration date provides a strong incentive to spend the voucher before it expires worthless.

a bank account dedicated to funding holiday gift purchases. Routine behavior is typically exercised through System 1, thereby avoiding the need to incur willpower costs. Rules can have both components to accumulate assets and components to protect assets. "Don't dip into capital" is a component designed to protect assets. The joint goal to *accumulate and protect* is a critical feature of an effective rule.

11. In the BLC, the planner is interpreted as being the rational component of the individual. However, Kahneman (2011) makes the point that System 2 is not perfectly rational. The discussion below discusses evidence suggesting that most people are unable to engage in complex intertemporal optimization. In this respect, Thaler and Shefrin (1981) and Shefrin and Thaler (1988) discuss rules such as not dipping into assets, as heuristics. Rules place constraints on spending, saving, and borrowing. The optimization framework is largely intended as a guide for characterizing sensible heuristics, which allow for limited optimization. Mental accounting-based rules provide guardrails to assist individuals in choosing behavior patterns, that while non-optimal, constitute reasonable approximations to optimizing solutions. A subtle issue in the optimization is that current savings can increase temptation in the future. This would factor into the optimization solution. There is little reason to believe that individuals are far sighted enough to take this feature into account.

11.5 Numerical Example

To analyze the nature of the nature of BLC solutions, consider a numerical example with the following parameter values:

- $\varphi_{1,I,t} = \varphi_{1,A,t} = 0.15$
- $\varphi_{2,I,t} = 1, \varphi_{1,A,t} = 0.05$
- $\varphi_{3,I,t} = \varphi_{3,A,t} = 1.5$
- $\varphi_{4,i,t} = 0$ for all i and t
- $F_1 = 50$

- $i_t = 0$ for all t
- $\lambda(t) = 1$ for all t
- $\epsilon = 0.001$
- $T = 3$: there are three dates, referred to as young, middle age, and old.
- Two potentially active accounts, income and future income, with the asset account only being used to record saving and borrowing. This means that the cost of dipping into assets is a function of the sum total of the future income account and asset account.

The assumption of a zero interest rate and zero time preference parameter simplifies the example. Consider a series of cases, which differ primarily in the income profile $[y_1, y_2, y_3]$. The first of these cases features a flat income profile, with the other four featuring hump shaped profiles. In moving from one case to the next, date 1 income declines, even as wealth remains unchanged.

In the first case, the income stream is flat, with $y_t = 50/3 = 16.67$ for all t. Because the planner utility function is symmetric, the optimal solution is for the individual to consume its income exactly, in each period. In doing so, the individual need not employ any willpower, and therefore need not incur any willpower costs. This solution has an associated planner utility of 8.44, with the marginal utility of consumption being equal to 0.06 at each date.

In the second case, the income profile $[y_1, y_2, y_3]$ is $[12.5, 35.5, 2.0]$. The optimal planner solution for this income profile is $[c_1, c_2, c_3]$ equal to $[11.65, 33.88, 4.47]$. Compared to the first case, the individual chooses to underconsume at dates 1 and 3 and to overconsume at date 2. Notably, the individual consumes less than income at dates 1 and 2 by 0.85 and 1.62 respectively, and consumes 2.47 more than income at date 3. The marginal utility of consumption at all three dates is 0.22. This equality reflects the fact that consumption is below temptation levels at dates 1 and 2 and equal to the temptation level at date 3. The associated planner utility is 7.05.

The character of this solution is that the income stream is concentrated during the middle years ($t = 2$). This is when the individual has the opportunity to save for retirement. However, spending out of assets is tempting and saving is costly. As a result, consumption while old ($t = 3$) is low. The unwillingness to save more during the middle years induces the individual, while young, to save a bit out of income (for retirement), even though income and consumption while young are one third of their respective values during the middle years.

In the second case, the individual has the option to consume at the first best level. However, doing so will require the individual to invade the future income account at date 1, and save 53% of income at date 2. Both behaviors require large amounts of willpower. As a result, planner utility from this solution, for the second case, is -4.14, much worse than the optimal solution. The associated marginal utilities for the three dates, for the equi-consumption solution are $[0.25, 0.71, 0.06]$. For the first two periods, marginal consumption reflects, not incremental utility of consumption per se, but incremental utility associated with willpower.

The remaining cases feature retirement income at $t = 3$ being equal to 9.2. Consider the character of the planner's optimum as income shifts from $t = 1$ to $t = 2$. If $y_1 > 9$, the planner's solution will be similar to the second case, with saving taking place at both $t = 1$ and $t = 2$. When $y_1 = 9$, the young individual will consume the entire income y_1, not invade the asset account, and the equi-marginal utility condition will hold.

When y_1 lies between 1.1 and 9.0, the planner's optimum will feature the young individual consuming all income. This is hand-to-mouth behavior. Notably, when $y_1 = 1.1$, the planner is indifferent between invading and not invading the asset account. This is a tipping point knife edge case. For $y_1 < 1.1$, the planner's optimum will feature invasion of the asset account. Notably, invasion involves a discrete jump in consumption, from 1.1 to 9.6. The difference $8.5 = 9.6 - 1.1$ is the amount borrowed from the asset account. In middle age, the individual will continue to save, but not enough to repay the debt in full. Therefore, the individual when old will be forced to repay the remaining debt, and experience lower consumption as a result, 3.6, well below retirement income of 9.2.

This individual has a modest self-control problem, which only becomes apparent when income is low, while young. In this situation, the individual splurges, and pays the price during retirement. When income at $t = 1$ is not quite as low, the individual engages in prudent behavior and does not borrow. For the rest of the y_1- range discussed above, the individual acts as if following the simple rule “do not dip into the asset account.” Moreover, for much of that range, the individual engages in hand-to-mouth behavior, totally consuming income y_1. Only when y_1 is sufficiently high, will the individual begin to save while young.

One additional case involves peak income occurring while young, which is often the case for professional athletes. This case is similar to the situation with a single mental account. In the absence of an effective rule, costly willpower will lead to overconsumption while young and underconsumption during middle age and retirement. A key driver of this feature is that relative to middle age and retirement, the degree of temptation is larger at younger ages, leading to higher willpower costs while young. In the general model, discounting also contributes to this phenomenon. However, it is important to understand the contribution of relative willpower costs. In general, both willpower costs and discounting, including hyperbolic discounting, are key variables underlying impatience and present bias.[38]

The variables $[\varphi_{j,i,t}]$ for $j = 1, 2, 3$ control the degree to which willpower is costly. In particular, reducing the value of $\varphi_{1,i,t}$ for each i decreases the magnitude of the penalty function, moving Z_t^R closer to log-utility. In doing so, marginal doer utility declines, reflecting the decrease in marginal cost of exercising willpower. In addition, the mental accounting entry fees decrease, facilitating entry if warranted, and the size of splurge upon invasion.

[38] Consider a comparison of two cases. The first case features full income while young and the second case features very low income while young. These two cases share an important similarity. In the latter case, the young individual who invades the asset account has access to the full wealth, just as in the former case. However, the individuals in the two cases might not face the same willpower costs, if consuming from the asset account is different, say less tempting, than consuming from the income account. Indeed, in the fourth case in the numerical example discussed above, the individual at $t = 1$ does invade the asset account, but chooses to consume less than the same individual at $t = 2$. In this respect, in the latter case, marginal consumption at $t = 1$ is from the asset account, while in the former case, marginal consumption at $t = 1$ is from the income account.

11.6 Differential Marginal Propensities to Consume

The marginal propensity to consume (MPC) is defined as the fraction of a small, one-time windfall which a household spends during a given time period. The windfall can arrive as either income or wealth. The following theorem focuses on how willpower costs influence the marginal propensity to consume from income, compared to the marginal propensity to consume from wealth.

Theorem 8 *Assume* $\varphi_{1,i,t} > 0$*,* $\varphi_{2,i,t} > 0$ *and* $\varphi_{3,i,t} > 1$*.*
1. Consider two alternative marginal changes at $t = 1$*. The first is an increase in income* y_t*. The second is an increase in income* y_τ *at a future date* τ *which leads to a marginal increase in wealth equal in magnitude to the increase in income. Consider an internal solution to the planner's maximization problem. If the marginal increases in income and wealth preserve the choice of active accounts at* t*, then the marginal propensity to consume out of the income account is greater or equal to the marginal propensity to consume out of wealth. At an internal solution for the planner's problem, featuring equal marginal utilities across* t*, and* $\varphi_{1,i,t} > 0$ *for all* i *and* t*, then the inequality is strict. If the marginal increase in income induces the planner to shift from* $c_{t,A} > 0$ *to* $c_{t,A} = 0$*, then the marginal propensity to consume from income will be negative, while the marginal propensity to consume from wealth will be positive.*

2. For given $[\varphi_{j,A,t}]$*,* $j = 2, 3, 4$*, a sufficiently high value of* $\varphi_{1,A,t}$ *will induce the planner to refrain from invading the asset account. A sufficient condition for noninvasion is* $E_{i,t} > ln(m_{i,t}) - ln(m_{i-1,t})$*. In particular, for* $t = T$*, the planner will be miserly, and choose to refrain from consuming from the asset account, for a sufficiently high value of* $\varphi_{1,A,T}$*.*

Consider some insights from the numerical example described above about the marginal propensity to consume. For this purpose, focus on how the planner reacts to a marginal increase in income, say at date 1. If the planner's optimal solution features equi-marginal utilities, then the marginal income will be consumed over time, in amounts determined by the impact on marginal planner utility from changes to consumption at each date. This is because the planner seeks to distribute marginal income to get the "biggest bang for the buck." A doer for whom marginal planner utility drops off slowly with consumption will be favored over a doer for whom marginal utility drops off quickly.[39]

Recall that at an internal solution, the marginal utility of consumption completely reflects willpower cost. As a result, the planner will choose to allocate a smaller share of marginal income to dates where the extra income induces a relatively sharp decline in marginal willpower costs. The planner prefers to allocate larger share of marginal income to dates where, at an internal solution, marginal willpower cost falls off slowly with increased consumption.

The discussion above pointed out how the model accommodates different degrees of discipline in respect to choice of rules. In this regard, consider an individual for whom the planner finds it optimal to consume 100% of current income without invading the asset account. Doing so requires no willpower (cost). Now imagine that the income account and asset account are merged, and that this merger forces the individual to exercise willpower in order to achieve the same consumption level. This change will increase marginal utility at the original consumption level, inducing the individual to consume more as a result, in contrast to the original situation.

The new solution will typically feature an internal solution requiring the exercise of willpower. The marginal propensity to consume from income in the first situation will be unity, as in the example above. However, in the second situation, the marginal propensity to consume from income will be less than unity, as is the case for internal solutions. In the original situation, the individual has the discipline to implement a mental accounting rule without dipping into assets. In the revised situation, the individual lacks the discipline to implement a mental accounting rule that prohibits dipping into assets. Notice that the marginal propensity to consume from income is greater for the more disciplined case. Hence, the less disciplined individual saves less than the more disciplined individual, but has a lower marginal propensity to consume.

[39]This is a second derivative issue, similar to (71) and (72).

11.7 Heterogeneity

Shefrin and Thaler (1988) present a representative agent-based discussion of the differential MPC hypothesis. Greater insight into the differential MPC issue can be had by focusing on household heterogeneity. Consider the series of cases in the numerical example above. Imagine that the economy consists of a heterogeneous mix of these individual types, at different stages of their lives. Some will be severely liquidity constrained, especially while young. Others might be like the middle aged individual in the example, who is unconstrained. Others will be retired. Another might be an individual at the tipping point featuring indifference between invading the asset account and not invading. In the numerical example, all have the same initial wealth, but different income profiles.

The difference in income profile, and age, is a major factor underlying why marginal propensities to consume differ across the population, for both income and wealth. For the unconstrained young person in case 2 of the numerical example, the marginal propensity to consume from income is 17% while for the constrained young person with low income, it is 6.5%. For the borderline tipping point case with asset account invasion, the marginal propensity to consume from income is negative. Technically, values of respective second derivatives at the planners' optimal solutions are the drivers of the differences. For the population as a whole, the marginal propensity to consume from income (and wealth) can be computed as an income (wealth) weighted average. Unless the borderline tipping point case is dominant, the marginal propensity to consume from income will exceed the marginal propensity to consume from wealth, for the population as a whole.

Just to clarify: an example of a borderline tipping point case involves an individual with low income who is wavering between whether or not to assume credit card debt. Borrowing simply entails paying less than the full monthly balance. Being indifferent between the two, an extra dollar of income would induce the individual to pay in full. However, in the absence of that dollar, the individual takes on the debt. Now suppose that a month later, the individual finds a dollar bill on the ground, and therefore receives the additional income. As a result, the individual decides to abstain from borrowing, and therefore reduces consumption. This behavior produces a negative marginal propensity to consume from income. As discussed below, empirically, the borderline tipping point case occurs with relatively low relative frequency.

Heterogeneity in the ability to exercise willpower, in terms of utility cost, is another factor. In the above discussion, all individuals are treated as having the same willpower cost parameters. As discussed above, the BLC model is built to accommodate such heterogeneity through variation in the values of the parameters $[\varphi_{j,i,t}]$.

With the penalty function (41), individuals who invade their asset accounts take a discrete bite: they do not just nibble into assets. For some, the bite is relatively small, and for others it amounts to a splurge. The values of $[\varphi_{j,i,t}]$ are critical for determining the relative size of the bite.

The population is a mix of some who are financially constrained and some who are unconstrained. The mix includes individuals live hand-to-mouth and engage in no discretionary saving. Some of these individuals might dip into their asset accounts by borrowing against future income, for example by using credit card debt. Others will use credit cards for transactions only, and pay the full balance on their accounts every month so as not to incur interest. As was mentioned above, some individuals also lack the discipline to structure efficient mental accounting-based rules. Others are disciplined, and when borrowing, invade their asset accounts without splurging; others also invade their asset accounts but cannot prevent themselves from splurging.

Although the formal model treats retirement as a single date, T, the general framework allows for retirement to occur over a series of dates, say from T_r through T. In this regard, Thaler and Shefrin (1981) include a discussion about "miserliness" among the elderly, with miserly behavior leading to underconsumption during retirement.[40] In the BLC, this issue pertains to the internal enforcement mechanisms associated with rules, what Chambers et al. (2009) describe as "neural mechanisms of response inhibition." There is an adage that "habits become nature." Some people establish rules with high account entry fees

[40] See footnote 7 of Thaler and Shefrin (1981), which explains miserliness as a habit, with Z_t achieving an internal maxiimum when $\theta_t = 0$. The model in Thaler and Shefrin (1981) lacked a mental accounting structure and associated temptation variable. Shefrin and Thaler (1988) extended the 1981 model to include both. Doing so provides a theory of miserliness associated with status quo bias, with the mental entry fees associated with invading the asset account becoming excessively high because of past discipline.

as mental guardrails. However, they then find it too painful to break these rules, even when circumstances change, especially in respect to the flows in their income accounts. This is especially the case for retirees with sufficient wealth to substitute asset account consumption-funding for income account consumption-funding. This issue arises in the discussion below, in respect to two specific issues, namely consumption funded by dividend income and bequests.

11.8 Model Extension: Multiple Commodities

The working paper version of Thaler and Shefrin (1981) included a discussion about a multicommodity planner-doer framework. The journal editor handling the submission suggested removing this discussion. Subsequently, Thaler (1985) described the basic approach, and Shefrin (2020) discussed the modeling issues.[41]

In respect to the modeling multicommodity issues, assume there are J types of commodities indexed by j. Extending the formal single commodity framework is relatively straightforward. Every (j, t)-combination will have its own doer $Z_{j,t}$. The J commodities are partitioned into groups, for the purpose of budgeting. Each group will be associated with its own consumption-mental account. This means that there will be two types of mental accounts int the model, consumption-mental accounts and funding-mental accounts such as I, A, and F. The two types give rise to a *matrix mental accounting* structure for which funding mental accounts finance which consumption mental accounts.

The multicommodity framework features both types of accounts, as well as rules for how funding-accounts are allocated across commodity-accounts. Formally, denote by $C_{j,t}$ the balance in consumption-mental account j at date t. Denote by $C_{j,i,t}$ the portion of $C_{j,t}$ that derives from funding-mental account i, at t.

Let $\{k\}_j$ be the commodities associated with consumption-mental account j. Denote by $p_{k,t}$ the price of commodity k at time t. A typical budgeting constraint in the mul-

[41] Galberti (2019) develops a theory of personal budgeting based on this approach. Galberti states: "This paper uses the term 'personal budgeting' rather than 'mental accounting' because the latter has the much broader meaning of a general process whereby people frame events, outcomes, and decisions. This also includes choice bracketing, narrow framing, and gain–loss utility, which differ from budgeting."

ticommodity framework posits that the value of purchases associated with account j, $\sum_k p_{k,t}c_{k,t}$, cannot exceed $C_{j,t}$.

Similar to the single commodity framework, the planner's control variables will include $[C_{j,i,t}]$, that specify how consumption will be funded. The planner will need to decide how much purchasing power to assign to consumption-mental account j at date t; and then how much of account j's purchasing power to assign to commodity k, again at date t. Achieving within account allocations will involve the use of discretion, meaning θ. As with the pure intertemporal model, the planner can use $\theta_{k,t}$ as an instrument to induce any desired choice of consumption level $\bar{c}_{k,t}$ (on the part of the (k,t) doer) by causing this doer to be satiated: $\partial Z_{k,t}/\partial c_{k,t} = 0$ at $\bar{c}_{k,t}$. Assume that higher levels of $\theta_{k,t}$ lead to both greater modification (e.g., $\partial c_{k,t}/\partial\theta_{k,t} < 0$) and higher costs ($\partial Z_{k,t}/\partial\theta_{k,t} < 0$).[42] The upper bound rule-related function f embodies commodity prices and budget constraints in the multicommodity setting, just as it embodies present value prices in the single commodity setting.[43]

The situation for a single commodity-account in the multicommodity setting is structurally similar to the single commodity setting that involves just one funding-mental account $[I + A + F]$. In the one-account single commodity setting, the planner uses discretion to choose consumption at all dates prior to T, with date T consumption determined as a residual. A pecking order can be assigned to the commodities in each account, leading commodities within an account to be prioritized.[44] Doing so produces temptation variables $m_{i,j,t}$, defined as residual purchasing power in the account, based on consumption levels chosen for higher priority commodities.

A discussion of household budgeting appears later in the chapter. Households which use mental accounts to engage in budgeting vary in the degree of detail in their budgets.

[42] The first order conditions associated with this problem appear in section A.4 of Shefrin (2020).

[43] In the single commodity version of the model discussed above, prices are given by the discount factors associated with the interest rate. When the interest rate is positive, the most expensive commodity is date 1 consumption. Present bias encourages overconsumption of expensive commodities. This is why the young are encouraged to begin saving for retirement early. The same issue, meaning overconsumption of high priced items, can surface in the multicommodity framework. The next subsection contains a discussion about budget constraints and associated rules in an uncertainty setting.

[44] In the single commodity one-account setting, time provides the priority ordering of different doers.

In this case, mental accounts group different commodities together. Typically, the more granular (and sophisticated) the budget, the greater the number of mental accounts. This point is reflected in the discussion above about $[I, A, F]$ and $[I + A, F]$, as the latter features fewer mental accounts and a less disciplined budgeting framework.

The multicommodity planner-doer model recognizes that some commodities are more tempting than others. Commodities that are particularly tempting, but whose consumption does not impact future self-control costs, are called *seductive*. Commodities whose consumption generate direct impacts in the future are termed *investment goods*, and if current consumption makes future willpower more difficult are termed *addictive goods*.

General saving from income features an increase in the asset account, and so saving can be viewed as the purchase of assets. Durable commodities are by nature investment goods. Durable goods are modeled by assigning a mental account to each such good, and allowing future doers to invade the account over time, when consuming the durable good. The magnitude of an account entry fee, if any, will reflect the character of any willpower issues associated with consumption of the durable good in question.[45]

One issue that arises in the multicommodity setting, but not the single commodity BLC, is that willpower is required to increase consumption of some goods, rather than decrease consumption. For example, some sedentary people find it difficult to engage in physical exercise [DellaVigna and Malmendier (2004, 2006)]. Modeling this issue involves some minor modifications to the structure. The modifications basically involve the following:

- The relationship between the amount consumed $c_{k,t}$ and willpower variable $\theta_{k,t}$ will be $c_{k,t} = (1 + \theta_{k,t})$ instead of $c_{k,t} = (1 + \theta_{k,t})/\theta_{k,t}$.
- When $\theta_{k,t} = 0$, $Z_{k,t}$ will be a quadratic function whose maximizing value of $c_{k,t}$ occurs at $\epsilon = 1$. This means that willpower will be required to increase the value of $c_{k,t}$ above its minimum.

[45]Spending too much time driving a new automobile would be an example of a durable good related self-control problem.

- The planner's utility function is set to zero at $\epsilon = 1$. For example, this is automatically the case for log-utility, meaning $ln(1) = 0$.

- To accommodate the use of willpower to increase consumption of a commodity, the structural equations (29) through (33) involve two changes. The first is to substitute $(1+\theta_{k,t})$ for $(1+\theta_{k,t})/\theta_{k,t}$. The second is to modify the penalty function underlying $Z^R_{k,t} = ln(c_{k,t}) - \Pi_{k,t}(c_{k,t}, m_{k,t})$. The new doer utility function would have the form

$$Z^R_{k,t} = \frac{1}{(\zeta_{5,k,t} - 1)} + \frac{c_t^{1-\zeta_{5,k,t}}}{(1 - \zeta_{5,k,t})} \tag{56}$$

  where

$$\Pi_{k,t}(c_{k,t}, m_{k,t}) = ln(c_{k,t}) - [\frac{1}{(\zeta_{5,k,t} - 1)} + \frac{c_t^{1-\zeta_{5,k,t}}}{(1 - \zeta_{5,k,t})}] \tag{57}$$

  Notice that when $c_{k,t} = 1$, $Z^R_{k,t} = 0$. This function is the subject of discussion later in the chapter, and displayed as Panel C in Figure 9 below.

- The derivative of $Z^R_{k,t}$ with respect to c_{kt} is $c_{k,t}^{-\zeta_{5,k,t}}$. Given any price for commodity k at date t, a high enough value of $\zeta_{5,k,t}$ will lead the marginal utility at zero to be arbitrarily close to zero. This is the case in which willpower alone will be too costly to induce the individual to engage in exercise, no matter how low the price.

In his book, Kahneman (2011) begins with a discussion about two *systems* and concludes with a discussion about two *selves*. Systems and selves are different. The concept of two-systems is about thinking, fast and slow. The concept of two-selves is about the *experiencing self* and the *remembering self*. Specifically, two-self issues are important, and connected to the idea of investment goods, meaning goods that generate experienced utility both currently and in the future.

11.9 Model Extension: Risk

Shefrin (2009a) describes how to extend the model in Shefrin and Thaler (1988) to accommodate risk, in order to provide an integrated formal framework for the mental accounting approach to dividends presented in Shefrin and Statman (1984). To describe the nature of the extension, consider an uncertainty tree with a typical node at date t being denoted by n_t and having associated probability $\Pi(n_t)$. Let $c(n_t)$ and $y(n_t)$ be the respective consumption and income amounts associated with the occurrence of event n_t.

Shefrin and Thaler (1988) informally analyze the impact of risk by focusing on (anticipated) variations in income over time. Smoothing over states is similar to smoothing over time. However, a proper analysis of risk involves an uncertainty tree. Although more thorough treatment of risk appears later in the chapter, at this stage it suffices to discuss a few points that bear on specific issues below.

In the planner-doer framework, choices involving risk involve stochastic income, stochastic willpower costs, stochastic temptation variables, and stochastic consumption. In this respect, consider two possible events at $t = 2$, one featuring low $y(n_2)$ and the other featuring high $y(n_2)$. If this is the only source of exogenous risk in the model, then the planner can compute expected utility as the probability-weighted utility associated with each of the two events.[46] This is the expected utility for a single risk. If the planner faces a choice between two risks at $t = 1$, then the planner will choose the risk having the higher expected utility. In this respect, keep in mind that the curvature of the planner's utility function reflects the cost of exercising willpower. Therefore, a high marginal cost of willpower gives rise to high risk aversion.

Fudenberg and Levine (2011) note that some decisions might feature the choice between a certain reward at $t = 1$ and a risky reward at $t = 2$. In the planner-doer model, a certain reward at $t = 1$ will increase temptation in the $t = 1$ income account. This introduces present bias into the planner's computation. Because present bias is manifest in the curvature of the doers' utility functions, increased temptation for the income account will lead to higher risk aversion (at $t = 1$). Because the risky reward in this example impacts future income, it also impacts the value of the future income account at $t = 1$. Therefore, the choice between the certain immediate reward and delayed risky reward manifests the marginal propensities to consume, both from income and from wealth. In

[46] Keep in mind that utility in the expression for expected utility refers to doer utility, which is a function of willpower effort and temptation, as well as consumption.

some circumstances, the differential MPC hypothesis will be germane for these kinds of choices.

Fudenberg and Levine (2011) also discuss the impact of making the reward at $t = 1$ risky instead of certain. The type of risk they analyze involves reducing the probability of a fixed reward. Because willpower costs are convex, the lower reward-probability leads to a reduction in the reward's certainty equivalent, as well as the expected value of the temptation variable for the income account. Therefore, for any internal choice of consumption at $t = 1$, the lower probability reduces requisite willpower costs, and therefore lowers the degree of risk aversion.

The discussion about dividends below pertains to a securities market. The additional notation required for this discussion is as follows. Denote by W the individual's lifetime portfolio wealth. Consider a set of securities $\{S_j\}$ that is available for trade over time. $D_j(n_t)$ denotes the number of units of physical consumption, which a single unit of security j pays at node n_t. At date t security S_j pays $D_j(n_t)$, and its ex-dividend price on the n_t-market is given by $Q_j(n_t)$. $D_j(n_t)$ corresponds to interest income in the case of bonds and cash dividends in the case of stocks.

Suppose that a household holds $x_j(n_{t-1})$ units of security j at the end of period t–1 through the beginning of period t. The value of the associated portfolio $\{x_j\}$ and its associated dividends will fund both the household's consumption in event n_t and portfolio choices $\{x_j(n_t)\}$. That is, the household faces the following budget constraint:

$$c(n_t) + \sum_j Q_j(n_t)x_j(n_t) \leq \sum_j (Q_j(n_t) + D_j(n_t))x_j(n_{t-1}) \tag{58}$$

The model in Shefrin (2009a) assumes a single consumption good. In line with the modeling approach in the working paper, it is a simple matter to extend the (2009) framework to include many commodities. For example, $c_j(n_t)$ would denote the amount consumed of commodity j should event n_t occur.

12 Game Theoretic Models, Two-System Models, and Human Nature

The two-system BLC framework has a different perspective on how to model human nature from the game theoretic approach. The BLC framework treats present bias as the result of imperfect mental technology in which the planner associated with the prefrontal cortex activity has difficulty restraining the doer-impulses generated by the limbic system. In this respect, the BLC treats the planner as having dynamically consistent preferences, and the time sequence of doers as being extremely myopic. By way of contrast, the game theoretic approach treats present bias as the outcome of a single system in which the same self, but at different times, disagree about the relative discount factors associated with the future relative to the time-shifting present.

The two approaches share similarities, but also major differences. This section and the next focus on the differences. To set the stage for this discussion, consider Maxted (2024), which builds on the IG model developed by Harris and Laibson (2013).[47] Recall that in the IG model, when the borrowing constraint does not bind, the dynamically inconsistent IG model corresponds exactly to a particular dynamically consistent model. However, when the borrowing constraint binds, the exact correspondence breaks down.

Maxted introduces an interest penalty structure which induces the individual to choose the "same" solution as in the IG model, but without the need for a borrowing constraint. This innovation preserves the exact correspondence, and provides closed form expressions which greatly facilitates the characteristics of the decisions made by present-biased individuals, along with the welfare consequences of those decisions. Using this approach, Maxted examines overconsumption-based issues such as the use of short-term borrowing on unsecured accounts such as credit cards and payday loans. He also investigates the interaction of present bias with asset illiquidity.

Maxted discusses an IG model having two assets. The first asset is liquid, and it pays a certain yield. The second yield is illiquid but features a stochastic rate of return with a

[47]As discussed above, IG stands for instantaneous-gratification.

Brownian motion component. An individual who wishes to engage in borrowing does so by holding a negative balance in the liquid asset, and paying a premium rate to do so.

An important feature of the model is the tractability it offers in analyzing welfare implications. The model compares two outcomes. The first outcome features a perfect commitment device corresponding to the case $\beta = 1$, with the commitment device having an associated perpetual consumption tax. The second outcome pertains to the actual value of β. Consider setting the consumption tax at a level which equates lifetime utility for the individual. Then the welfare cost of present bias can be expressed as the associated consumption tax.

Consider the following four points from Maxted (2024).

1. "[P]resent bias does not necessarily engender a demand for illiquid assets. Provided that the borrowing limit on liquid wealth does not bind in equilibrium, present-biased consumers do not seek out illiquidity because illiquid assets do not actually limit overconsumption. Intuitively, the illiquid asset is never needed to fund current consumption, because the agent can always increase their consumption by adjusting their holdings of the liquid asset instead. Retirement systems around the world rely on illiquidity to incentivize retirement savings (Beshears et al., 2015). However, the results in this paper cast doubt on the benefits of such policies."

2. "[A]nother boundary of the results presented thus far is that they may fail to hold under less restrictive equilibrium assumptions and/or when present bias is combined with boundedly rational optimization. For example, non-Markov or boundedly rational equilibria may feature "pseudo hard constraints" such as mental accounts (Thaler, 1985) and personal rules (Ainslie, 1992; Bernheim et al., 2015) that act similarly to binding hard constraints in placing bright-line restrictions on agents' overconsumption."

3. The IG model "speaks to the puzzle of why present-biased agents do not use commitment devices (Laibson, 2015; Bernheim and Taubinsky, 2018).[48] This literature

[48]By this, Maxted means *external* commitment devices, in contrast to *internal* commitment devices which are central to the BLC approach, a point he made in private correspondence.

has concluded that commitment is often hampered by an important tradeoff between commitment and flexibility (Amador et al., 2006; Laibson, 2015)."[49] There is also a "complementary explanation: it is hard to design devices that can even generate commitment in the first place... There are many margins that can potentially be adjusted to bring utility into the current period, ranging from the consumption of unhealthy food to decreasing exercise to staying up too late. Unless a commitment device can block all of these sources of temptation, there is no reason for the agent to choose an ineffective 'commitment device' that actually serves only to limit flexibility."

4. "Despite the evidence that consumers exhibit present bias, the modeling of consumption-saving behavior has been slow to incorporate these insights. This is because the literature is stuck at an impasse: while present bias can introduce a variety of novel and economically relevant behaviors, models with present bias are often difficult to characterize in practice."

These four points serve to highlight key commonalities and differences between the game theoretic approach and two-system BLC approach. The first point identifies an "irrelevance result" in line with the Modigliani-Miller principle. In the IG model, imposing self-control constraints on the illiquid asset brings no benefits, and might actually bring harm. This contrasts with the BLC which focuses on the relevance of self-imposed constraints associated with mental accounts for future income. This difference is the essence of the second point, which mentions "pseudo hard constraints" related to mental accounts and rules. Commitment devices are a major feature of BLC-based self-control rules, and the third point notes that such rules do not receive a great deal of use. This is an important issue, and an empirical one. The fourth point is about the difficulty of developing realistic models that incorporate present bias. This might be the case with the game theoretic approach, which involves the characterization of complex equilibrium

[49]During the time this literature developed, the $\beta\delta$ approach dominated the way in which self-control issues were modeled. Fudenberg and Levine (2006) discuss a series of contributions that lie outside the $\beta\delta$ approach, which for lack of space are not covered in this chapter.

behavior. It is less of an issue in the BLC, which focuses on the nature of rules people use to reduce self-control difficulties.

The differences between the game theoretic approach and the BLC approach involve issues which are partly theoretical and partly empirical. The empirical issues are the subject of the next section. The theoretical issues pertain to modeling human nature. The game theoretic approach treats individuals as being able to solve, or acting as if they were able to solve, very sophisticated optimization problems, given their beliefs, possibly incorrect, about how their future selves will react (value of β). In contrast, the BLC approach treats individuals as choosing fairly simple rules with limited optimization which they can implement psychologically, and which describe behavioral limits and boundaries. In the BLC, people only optimize *at the margin* when they judge that the limits imposed by their rules are inadequate. In this case, they use willpower to face temptation. The two approaches feature a very different view of human nature and people's capabilities.

13 Empirical Evidence

This section surveys empirical evidence about people's choices or biological factors which influence those choices. Biological factors pertain to interpersonal differences stemming from variations in genetic makeup or neurological activity. Actual behaviors pertain to rates of saving, character of household budgeting, reliance on dividends to finance consumption, and effectiveness of managing debt. The main theme of this section is to focus on whether the BLC approach to overconsumption is more in line with the evidence than the intrapersonal game theoretic approach.[50] A summary of the main points of this section appears in the conclusion.

13.1 Neuroeconomics

The BLC is fundamentally a neuroeconomics model involving the interaction between different components of the human brain. When it comes to self-control, the central focus

[50]The first three subsections below draw on Shefrin (2021).

of the planner-doer model, some people exhibit more willpower others.

The neuroscientific literature offers insights into why this is the case. Hare, Camerer, and Rangel (2009) report on an experiment in which a group of volunteers, all self-reported dieters, were connected to neuroscientific monitoring equipment and shown photos of fifty different foods. The monitoring equipment recorded brain activity as experimental volunteers examined and made decisions about these foods. Examples of foods were tempting items such as candy bars and healthy, staple items such as vegetables. The experimenters asked the volunteers to rate each food first on taste and then on health benefits. The experimenters then selected for each participant an index food that the participant rated in the middle for both taste and health benefits.

Subsequently, the experimenters presented participants with a series of choices, asking them to choose between accepting either the index food or one of the other foods. The question is how did the self-reported dieters choose, especially when faced with foods that they rated high on taste and low on health benefits?

Did the self-described dieters always choose the index food over foods that were low in health benefit by high in taste? Some did, and others not. What differentiated the natures of people who always made the healthy choices from the natures of those whose decisions were guided by taste? The experimenters discovered that the differentiating factor turned out to be brain region activity. If we think about our brain as a circuit board with lights, we can ask what parts of our brain light up in key circumstances.

The prefrontal cortex is the region at the front of the human brain associated with conscious thinking. The food choice experiment identified two regions in this part of the brain. For ease of discussion, let us call one region V (for Ventromedial) and the other region D (for Dorsolateral). Everyone uses region V to make value-laden decisions. When the individual faces a choice, region V lights up.[51] However, region D only lights up for individuals who balance health benefits against taste. This is because region D appears to modulate the activity of region V.[52] Region D is akin to an internal voice, forcefully

[51]In the context of the BLC, facing a choice induces impulses associated with doer activity, as reflected in the function Z.

[52]In the context of the BLC, region D is where willpower θ and a rule-based constraint f operate on doer utility Z. Rules involve inhibitory behavior [Chambers et al. (2009)].

asking, as a person considers reaching for a candy bar, whether he needs the extra sugar, or would instead benefit by choosing a more nutritious alternative.

Viewed through the lens of the BLC-approach, weak D-activity corresponds to an individual having to incur high utility costs when using willpower to reduce consumption, which discouraging the individual from exercising willpower. Viewed through the lens of the $\beta\delta$ approach, weak D-activity corresponds to a low value of β; however, in the $\beta\delta$ approach, there appears to be no natural notion of willpower as an intrapersonal struggle.

13.2 Nature and Nurture

The BLC is structured to accommodate heterogeneity in people's ability to exhibit self-control. Nature and nurture both play a factor in the way people save. As for nature, meaning generic differences, Kuhnen (2015) discusses experimental evidence about the importance of specific genes, especially the gene known as COMT.[53] Notably, different people can have different variants of this gene. The COMT gene plays an important role both in learning and financial behavior. Indeed, one particular variant of this gene, known as Met/Met, significantly increases the availability of dopamine in the prefrontal cortex of the brain.[54] Dopamine is a neurotransmitter, which a large body of neuroscience research has shown to be critical for learning from new outcomes. The favorable variant of the COMT gene impacts cognition mediated by the prefrontal cortex, specifically executive control and working memory, with the favorable variant being generally associated with better performance.

The experimental evidence indicates that people having the COMT-gene variant which significantly increases the availability of dopamine in the prefrontal cortex of the brain, appear to make better financial decisions. Because this same gene variant is also associated with better learning, it is plausible that people with this genetic variant have learned to be more financially literate than others, and to make better financial decisions.

In respect to nurture, Cronqvist and Siegel (2015) analyze whether savings propensity

[53]COMT is short for catechol-O-methyltransferase.

[54]The Met/Met genotype is associated with greater financial acumen and learning. The genotype breakdown was 22% for Met/Met, and 78% for the two others.

is governed by parents instilling preferences into their children, or instead is the result of individual-specific life experiences. They report that there is evidence that suggests that some parents make efforts to teach their children good savings habits. These efforts entail activities such as providing a piggy bank, opening a savings account, and emphasizing the benefits of a frugal lifestyle. However, the evidence is effectively anecdotal rather than systematic and comprehensive; therefore, we lack good evidence about the effectiveness of parental efforts. Regardless, we know that parents tend to pass their savings propensities onto their children. In this regard, their research helps us to understand the extent to which this similarity is genetic as opposed to the result of social transmission of behavior from parents to children.

Cronqvist and Siegel analyze data from Sweden about the savings behavior of identical and fraternal twins. Their goal was to decompose individuals' propensity to save into two components, one genetic and the other environmental. In this regard, their research methodology relied on the fact that identical twins share 100% of their genes, whereas fraternal twins typically share only 50%. Cronqvist and Siegel reasoned that if identical twins save in a way that is more similar to their parents than is the case for fraternal twins, then they could conclude that savings behavior is partly genetic in nature.

Cronqvist and Siegel report that genes explained about 33% of savings behavior for the individuals in their sample. Moreover, the effect is long lasting, and does not disappear later in life. An important finding is that parenting explains 40 to 50% of savings behavior for those aged 20 to 25 years. However, the parenting effect decays significantly and reaches zero for people in their 40s. Therefore, parenting appears not to have a lifelong impact on the savings propensities of children. Not surprisingly, the effect of parenting on saving behavior is smaller when time for parenting is scarce. The research findings suggest that the wealth of parents and current socioeconomic status moderate genetic predispositions to particular savings behaviors. Cronqvist and Siegel summarize their main findings by saying that both the genetic component and lack of self-control are the primary drivers of savings behavior.

In respect to financial acumen, age is also a factor. Specifically, the frequency of financial mistakes varies with age and follows a U-shaped pattern. That is, financial mistakes decrease with age until people reach their early 50s, and then begin to increase. The declining pattern up to the early fifties is consistent with people acquiring more assets during their early and middle years. However, the increase in mistakes at older ages highlights the natural limits on individuals' financial decision-making that come with aging.

Cronqvist and Siegel's research focuses on good savings habits, and habits play a central role in the BLC. There is nothing in their discussion to suggest that the subjects in their database engage in solving complex dynamic optimization problems, or more complex problems involving intrapersonal game theoretic equilibria over their respective life cycles.

13.3 Financial Literacy

The optimization problem in the BLC serves as a guide for structuring sensible savings heuristics requiring limited optimization. Financial literacy relates to basic financial knowledge, knowledge that underpins the ability to make optimizing decisions. This section documents empirical findings about the relationship between financial literacy and the quality of households' decisions. A key issue is the extent to which these findings support the notion that households' behaviors are consistent with the intrapersonal game theoretic approach or with the BLC.

Findings described in Lusardi and Mitchell (2014) identify three major issues associated with financial literacy: "(i) numeracy and capacity to do calculations related to interest rates, such as compound interest; (ii) understanding of inflation; and (iii) understanding of risk diversification." (p. 10). Notably, the literature on financial literacy is largely concerned with identifying positive relationships between the quality of household decisions and the degree of financial literacy. In this respect, the evidence suggests that in relative terms, households with greater financial literacy are less likely to make investment mistakes, more likely to engage in retirement planning, more likely to set aside emergency funds, less likely to have costly mortgages, less likely to use high-cost sources of borrowing, and less likely to hold excessive debt. With respect to credit cards, Lusardi and Mitchell note that "while less knowledgeable individuals constitute only 29% of the cardholder population, they accounted for 42% of these charges ... Indeed, the average fees paid by those with low knowledge were 50% higher than those paid by the average cardholder." (2014, p. 8). In addition, Lusardi et al. (2017) document that financial knowledge accounts for between 30 and 40% of inequality in retirement wealth, an issue directly related to undersaving/overconsumption.

Lusardi (2010) measures degree of financial literacy using a quiz consisting of five specific questions about the three issues mentioned in the previous paragraph. In her survey of 28,146 individuals, the mean and mode for number of correct answers were 3 and 4 respectively.[55]

In Lusardi's sample, the proportion of individuals having all correct answers is 15%, which is close to the 22% associated with the Met/Met figure for the COMT gene. These values are consistent with most individuals lacking the capability to carry out complex optimization. Additional support for this statement comes from Ambuehl, Bernheim, and Lusardi (2014) who analyze people's knowledge about present value and future value. In this regard, recall that the objective function in the intrapersonal game theoretic approach, at each moment in time, is a sum of expected discounted utilities.

Ambuehl, Bernheim, and Lusardi (2014) use a set of five questions to assess people's comprehension of present value and future value. They focus on the length of time it takes for the value of an investment to double, as a function of the compound rate of return. Ambuehl et al. find that the average subject correctly answers just under two of five questions.[56]

[55] The histogram for number of correct answers is as follows:
0. 7%; 1. 11%; 2. 17%; 3. 24%; 4. 27%, 5. 15% .

[56] The associated bias is known as "extrapolation growth bias." Ambuehl et al. study whether it is possible to educate people to improve their ability to compute the length of time it takes for an investment to double in value. To do so, they introduce an intervention whereby they teach subjects the "rule of 72." This rule is a heuristic which stipulates that dividing the number 72 by the interest rate produces a good estimate of the doubling time. For example, at an annual interest rate of 10%, the rule stipulates that an investment will double in 7.2 years (72/10 = 7.2 years). Ambuehl et al. report that this intervention dramatically increases the average score, by about 1.4 additional correct answers.

13.4 Save More Tomorrow

The above discussion made the point that the joint goal to *accumulate and protect* is a critical feature of an effective rule.[57] The BLC provides the theoretical basis for one of the most successful savings programs in the U.S. In 2019, the value of the U.S. National Retirement Risk Index (NRRI), which measures the state of Americans' retirement preparedness stood at 47%. This suggests that almost half of Americans undersave/overconsume. Between 2019 and 2022, this value decreased to 37%, reflecting soaring home values, government pandemic stimulus programs, strong employment, and rising asset prices. Most of these factors appear to be temporary.[58]

With these figures as context, consider the impact of a behavioral savings program called "Save More Tomorrow" (SMT). SMT is a sophisticated extension of mental accounting heuristics. In 2006, the principles underlying SMT were incorporated into U.S. law as part of the Pension Protection Act of 2006. The Act encouraged firms to adopt the core principles of SMT. As of 2022, more than 25 million Americans are estimated to participate in SMT-based savings programs. This is 30% of U.S. workers who have access to 401(k) plans, and 15% of the U.S. civilian workforce.[59]

SMT has four main features, as described by Thaler and Benartzi (2004), who developed the program . The four features are as follows:

1. The program is voluntary, and any pain associated with increased saving is deferred. That is, employees who participate agree to increase their contributions months before they are called on to actually do so. This feature is a circumvention to remove temptation. That is, by asking for precommitment, plan sponsors avoid stimulating fast thinking processes associated with the immediate gratification generated by

[57] See the discussion about the planner's optimization problem featuring constraints (50), (51), and (52).

[58] See the discussion of the Federal Reserve's 2022 Survey of Consumer Finances in "The National Retirement Risk Index: An Update from the 2022 SCF," February 27, 2024; Center for Retirement Research at Boston College, https://crr.bc.edu/the-national-retirement-risk-index-an-update-from-the-2022-scf/. The CRR discussion states the following: "Though these results are encouraging, most people do not tap their home equity in retirement and prices may not stay at such high levels... Importantly, even with what may be a temporary improvement, the NRRI shows that 2 in 5 of today's working households could fall short."

[59] See http://www.shlomobenartzi.com/save-more-tomorrow. In addition, the source of some of this information is Shlomo Benartzi.

spending in the moment. Therefore, the need for the slow thinking D-region of the brain to promote self-control is diminished.

2. Employees begin with the first paycheck after a raise. People hate losses. The work of Tversky and Kahneman (1992) indicates that for the average person, a small loss feels about twice as intense as a small gain of the same magnitude. When people become accustomed to particular levels of spending, a sudden decrease in spending tends to register, mentally, as a loss. Losses are registered by activation of particular regions in the brain. By waiting for a raise, and taking the increased saving only from the raise, SMT avoids people feeling a loss associated with the increased saving, because their spending does not go down.

3. The contribution rate continues to increase on each scheduled pay raise until it reaches a preset maximum.

4. Employees can opt out at any time. Therefore people know they are not trapped by the plan. If their circumstances change, and they need to reduce their saving as a result, they are free to do so. However, psychological inertia, sometimes known as "status quo bias," causes most participants to continue to participate in the plan.

SMT has been extremely effective. In its early stages, for the 80% of households who agreed to participate, saving rates increased from 3% to 11% over the course of 28 months. In 2024, Vanguard, one of the largest financial firms in the U.S., reported that employees in automatic enrollment plans they offer saved an average of 12.3%, this when taking into account both employee and employer contributions. In contrast, the report notes that employees in voluntary enrollment plans saved an average of 7.4%, this because of significantly lower overall participation.[60]

[60]The report states the following: "Automatic enrollment designs: Historically, employees in a 401(k) or 403(b) plan have had to make an active choice to join the plan, but this trend is shifting. In voluntary enrollment plans, decisions were framed as a positive election: "Decide if you'd like to join the plan." Why have some employees failed to take advantage of their employer's plan? Research in the field of behavioral finance provides explanations:
1) Lack of planning skills...
2) Default decisions. ..

With this as context, consider how the evidence about the characteristics and outcomes associated with SMT relate to the following three questions. First, are commitment devices so inflexible that most individuals will not select them? Second, is it rare for individuals to choose commitment devices? Third, for those individuals who use commitment devices, are those devices ineffective, because individuals find ways to circumvent them?

In respect to these three questions, consider the following. First, SMT is a commitment device which is flexible in the sense of providing an opt out feature.[61] Second, over fifteen million Americans are enrolled in SMT programs. That said, it is not so much that they choose to participate in SMT as they are initially opted in by default. Third, as mentioned above, participation in SMT has been very effective at increasing saving.

There is no evidence supporting the idea of major demand for strong externally enforced commitment devices; and such devices by their nature are inflexible. Yet, this state of affairs is why people choose rule-based internally enforced commitment devices, possibly with external components such as SMT and Christmas Clubs.

In the BLC, penalties for invading mental accounts induce status quo bias. Opting into a plan might also require effort, which might also induce status quo bias. In the BLC, the System 2 planner considers mental strategies to manipulate the System 1 doers. By way of contrast, in the intrapersonal game theoretic approach, intrapersonal rules play no role. There is no System 1 for System 2 to manipulate. Soft constraints embedded in internally enforced self-control rules and external commitment devices with opt out features do not "bite" in the neoclassical intrapersonal game theoretic approach. In this approach, self-control soft constraints do not work because people will always find ways around them. Rather, providing an individual with the choice to opt out will lead to opting out, unless he or she would have voluntarily chosen the same outcome as opting in provides.

3) Inertia and procrastination."

[61] The is also true for Christmas Clubs.

13.5 Household Budgeting

The multicommodity extension of the BLC focuses on mental accounting-based budgeting. Recall the discussion above about the four issues raised in Maxted (2024). Maxted makes an important point about self-control, when he notes the presence of "many margins" which provide challenges for exercising self-control. This issue pertains specifically to a multicommodity setting in which the central issue is *budgeting.* Zhang, Sussman, Wang-Ly, and Lyu (2022) present evidence about the way consumers budget. The discussion below highlights some of their findings which provide insight into the "many margins" issue raised by Maxted.

Zhang et al. find that roughly 66% of their sample of households report that they currently budget. Of those who do not currently budget, 42% report that they budgeted at some point in the past. When Zhang et al. initially asked respondents whether or not they budget, they intentionally refrained from defining what budgeting entails. They did so in order that respondents avoid considering whether their budgeting activities were formal or informal. When Zhang et al. limit the definition to formal budgeting, they find that about 38% of their sample report engaging in formal budgeting, which is consistent with findings in other surveys they reviewed.[62]

In respect to overconsumption, the study finds that nearly 75% of those who budget state that the predominant reason for doing so is the desire to avoid overspending their income. This finding relates directly to Maxted's point about "many margins." Moreover, this motivation for budgeting does not appear to vary substantially with income or wealth. Zhang et al. suggest that these findings stem from individuals' concerns about problems with self-control, citing the mental accounting issues stressed by the BLC. In addition, individuals appear to link their motivation for budgeting to their establishment of long-term goals. That said, Liao (2021b) finds a positive relationship between budgeting and saving activity to establish emergency ("rainy day") funds.

[62]Similar to the results for financial literacy, the analysis of budgeting indicates that the propensity to budget varies with with age, exhibiting a hump-shaped profile over the life-cycle. The propensity to budget is also greater for those with higher levels of education, who are currently employed, and who live with a spouse or partner. Notably, the propensity to budget does not differ by gender.

The early literature on overconsumption discussed the case of Christmas Clubs as a device for holiday shopping [Strotz (1956), Thaler and Shefrin (1981)]. Shefrin (2020) reports that approximately 55% of households set holiday budgets. Most do so during November and December, setting vague "plans" in early November and more formal budgets thereafter. Approximately 36% of those who set holiday spending budgets overspend those budgets. Households who keep to their budgets spend 63% less than those who spend more than they had planned (in November). The least accurate holiday budgeters plan to spend either less than $100 or more than $749; and on average they overspend by about 30%.

In line with the above discussion about the multicommodity version of the planner-doer framework, Zhang et al. state the following: "Though the particular budgeting process may vary from household to household, a growing body of empirical evidence suggests that people allocate and track their spending within distinct budgetary categories... For instance, a household might maintain a separate 'entertainment budget' or 'clothing budget.' Furthermore, households appear to treat these categories as rigid, with funds not fully substitutable across categories." Category-based budgeting is intended to address the challenge of "many margins."

Consider customers who use a particular credit card almost exclusively to purchase items in a particular category.[63] When Zhang et al. investigate customers who use credit cards in this way, they find that the two most common top spending categories are "Shopping and Entertainment," which may reflect the use of a separate card for discretionary purchases. The next three most common categories are "Transport," "Utilities," and "Groceries."

The findings in Zhang et al. lend support to the fact that many households use different cards for different categories of expenses, potentially as a way of organizing their spending. This is consistent with the analysis in Shefrin and Nicols (2014) who note that this practice is a variant of the traditional "envelope heuristic" whereby households place cash to be spent over the subsequent month in separate envelopes, with each envelope

[63] By almost exclusively is meant 95% of purchases pertain to a particular category.

earmarked for a specific consumption category. Shefrin and Nicols find that 30% of their sample use multiple credit cards. Part of their study used focus groups, whose participants report using multiple credit cards for budgeting purposes.

In the multicommodity BLC model, having specific goals tends to increase the granularity of a household's budget, and the number of commodity-mental accounts. Shefrin and Nicols apply cluster analysis to segment their sample into four distinct clusters. For three of the clusters, about 30% of cardholders report that their budgets reflect specific goals, while about 50% report that their budgets reflect only general goals. The fourth cluster is distinctly different, with fewer than 10% reporting that their budgets reflect specific goals. The cardholders in this cluster, which comprise 20% of the sample, report being less confident than the others, in less control of their finances, and are more prone to paying only the minimum monthly balance on their credit card statements.

Shefirn and Nicols find that the use of multiple cards is especially pronounced in the cluster with a relatively large component of cardholders aged 55 and above. This cluster comprises 20% of the sample, and 85% of its member cardholders use multiple credit cards. Notably, 56% of this cluster's cardholders classify themselves as deliberate spenders, and 28% classify themselves as spontaneous spenders. These proportions feature less budgeting and more impulse spending than 53% of the sample (for whom about 65% self report as deliberate spenders). As a point of contrast, one of the other clusters, which comprises 33% of the sample, 99% use a single credit card.

Gathergood, Mahoney, Stewart, and Weber (2019) document that people carrying positive balances on multiple credit cards with varying interest rates engage in nonoptimal repayment behavior. Specifically, instead of first paying down the balance on the card featuring the highest interest rate, many instead pay down the balances proportionally across cards, in what Gathergood et al. call the "balance-matching" heuristic.[64]

Some cardholders who subscribe to the balance matching heuristic might be acting suboptimally. At the same time, the balance-matching heuristic might constitute "reasonable" budgeting behavior in a BLC setting. By reasonable, I mean an individual doing

[64]The balance-matching heuristic matches the ratio of card payments to the ratio of card balances.

the best that he or she can, being adaptive, when facing a given task, external environment including degree of nudging, and own capabilities. Within the context of the current discussion, it means being able to maintain desired expenditure proportions across commodities; and interest minimizing behavior will typically interfere with that when it limits relative credit card spending on cards bearing low interest.[65] People who use multiple credit cards to structure mental accounts for multicommodity budgeting purposes will be willing to pay a bit more interest to achieve control of their expenditures. This is similar to the way people used Christmas Clubs in the past, in the sense of being willing to forgo interest and perceived flexibility by depositing money into Christmas Clubs instead of savings accounts that paid higher interest rates [(Strotz, 1956), (Thaler and Shefrin, 1981)].[66]

In the multicommodity BLC framework, a rule provides the choice architecture for the budgeting process. A completely specified budget that is implementable corresponds to something like the planner's first best solution in the single commodity case. Such a solution represents a full response to the challenge of "many margins."However, with self-control issues involved, let alone transaction costs, feasible budgetary rules will typically reflect partial budgets. As in the single commodity case, implementation is accomplished either through reliance on temporary satiation of needs, (meaning $\partial Z/\partial c = 0$), or by coming up against a limit associated with a rule (f).

In most situations, people can only do one thing at a time, so that only one type of impulse (Z) is operative at any moment. To take an illustrative example, consider the consumption of food items, say appetizer, main course, dessert, and beverage. These food items will typically be consumed consecutively in time, not simultaneously. A person might have budgeted for a single meal, but not for the individual items making up that meal. Therefore, the formal choice problem for the meal is akin to the single commodity choice problem with pure discretion. In this case, the expenditures on the components

[65] Although a subset of consumers do use credit cards in this way, many do not. For this reason, it is not clear to identify the degree to which the balance matching heuristic reflects self-control behavior relative to suboptimal balance payment behavior.

[66] Recall that Christmas Clubs have an opt out feature, and so the inflexibility is perceived rather than real. The issue is more about inducing a mental state than imposing a constraint.

of the meal all reside within the same mental account. Notably, the appetizer doer goes first, followed by the main course doer, and so on, just as in the single commodity case the period 1 doer goes before the period 2 doer. If the appetizer is especially tempting, and the main course is delayed, then the agent will face Thaler's "cashew nut appetizer" self-control challenge [Thaler (2015)]. Either natural satiation (meaning $\partial Z/\partial c = 0$, as the agent has his fill of cashews), or willpower-induced satiation provide the only solutions to stop eating cashews before exhausting the amount budgeted for the whole meal. If cashews are that tempting, then they need to be placed into a mental account of their own.[67]

13.6 Heterogeneity Statistics: Credit Cards and Retirement Saving

There is considerable heterogeneity in savings behavior, and mental account invasion. Credit card debt is a key way for individuals invade their future income accounts. Shefrin and Nicols (2014) report that 53% of American credit cardholders carry a credit card balance from month to month. This means that at least half of cardholders routinely invade their future income accounts.

As was discussed above, Shefrin and Nicols (2014) identify four distinct clusters of credit card users. Members of clusters 1 and 2, together, predominantly pay only the minimum monthly balance, thereby borrowing on their cards and paying (high) interest rates. Cluster 1 comprises 20% of the sample, and 42% of this cluster pays the minimum balance. The members of this cluster also feature the lowest response rate for setting specific budget goals (9%), and express the lowest confidence about being in control of their finances. In respect to impulsiveness, this cluster exhibits the greatest tendency to engage in spontaneous spending (53%). Relative to cluster 1, the members of cluster 2, which comprise 27% of the sample, are less prone to be spontaneous spenders (16%), more prone to set specific budget goals (31%), and are much more confident about managing

[67] In addition, filling up on cashews will make exercising willpower easier when it comes to main courses and dessert.

their finances.

The members of the other two clusters avoid paying the minimum balance, and more that 90% pay the full balance each month. The members of cluster 3 tend to be older, with 45% being aged 55 and above, compared to 35% for cluster 4. Cluster 4 are more prone to set specific budget goals than cluster 3 (32% compared to 23%), and to engage in less spontaneous spending (22% compared to 28%).

Many of the features of credit card usage in the U.S. appear to persist over time. Consider the following:[68]

- Data from 2024 indicate that roughly half of American credit cardholders continue "to carry a credit card balance from month to month."[69]

- The Fed's research shows younger borrowers and those who live in lower-income areas are much more likely to be maxed out than those who tend to keep a lower utilization rate: 15.3% of Gen Z borrowers and 12.1% of millennials have maxed out their cards, compared to 9.6% of Gen X and 4.8% of baby boomers.[70]

- 36% of U.S. adults have more credit card debt than emergency savings, according to Bankrate's 2024 Emergency Savings Report. 47% of U.S. adults who say money negatively affects their mental health, at least occasionally, cite being in debt as a reason why, according to Bankrate's latest Money and Mental Health Survey.[71]

- Prolonged debt also plays a factor... as more Americans fall behind on their credit card bills. About 7.18% of cardholders fell into delinquency in the second quarter [of 2024], up from 5% in the previous quarter, Fed statistics show.[72]

[68] These are direct quotes from sources listed in the associated footnotes below.

[69] https://www.bankrate.com/credit-cards/news/credit-card-debt-survey/. The survey also points out: "That's compared to 44% in January 2024 and 60% in March 2020." In addition, "60% of U.S. adults with credit card balances have had that balance for at least a year."

[70] https://fortune.com/2024/05/14/americans-debt-credit-cards-inflation-interest-rates/. https://libertystreeteconomics.newyorkfed.org/2024/05/delinquency-is-increasingly-in-the-cards-for-maxed-out-borrowers/.

[71] https://www.bankrate.com/credit-cards/news/credit-card-debt-report/.

[72] https://www.cbsnews.com/news/credit-card-debt-total-us-2024/

In the BLC, carrying a positive balance constitutes an invasion of the future income account. Because of the entry fee, individuals who invade their asset accounts take discrete bites from the account balance. Although they might take a small bite, they do not nibble. Shefrin and Nicols (2014) find that a sizable proportion of cardholders take major bites from their future income accounts. In this regard, 25% of cardholders make only the minimum required payment every month. This is roughly one half of all cardholders who carry a credit card balance from month to month.

Using credit cards to access future income accounts is particularly problematic for individuals who run up maximum balances, and as mentioned above fall into delinquency.[73] Cutting up credit cards, a standard recommended self-control technique, corresponds to shifting from invasion to non-invasion, and in so doing leads to reduced current consumption: hence, the negative marginal propensity to consume from income.[74] Consider the following excerpt from *Fortune* magazine:[75]

> **Why it's important to get out of credit card debt**
>
> It's important to avoid credit card debt for one big reason: Carrying a balance can be costly; credit cards are one of the most expensive forms of credit, with higher interest rates than car loans or personal loans.[76]
>
> According to the Federal Reserve, the average annual percentage rate (APR) for cards that assessed interest was 22.16%.

Cutting up credit cards is one way to shift from invasion to non-invasion of the future income account, or the asset account for that matter. In general, the shift from invasion

[73]Survey findings from 2024 reported by Bankrate indicate the following: 35% of those "carrying a balance on their credit cards have more credit card debt" in August 2024 "than they did at the beginning of 2022." In addition, 24% indicate that "they feel less confident in their ability to get out of credit card debt ... than they did at the beginning of 2022." In this regard, 17% "who carry a balance on their credit cards worry they might not be able to make their minimum credit card payments at some point in the next six months." Source: *Bankrate's 2024 Credit Card Debt Report*, https://www.bankrate.com/credit-cards/news/credit-card-debt-report/#credit-card-rewards

[74]In the relevant case described in Theorem 8, an increase in income leads the planner to shift from invasion to non-invasion, as opposed to invoking a rule such as "cut up the credit cards."

[75]See Tretina, Kat, 2024. "Ditch the plastic: 8 strategies for crushing your credit card debt," *Fortune*, https://fortune.com/recommends/credit-cards/how-to-get-out-of-credit-card-debt/.

[76]Fast and frugal heuristics, a technique discussed in Gigerenzer and Todd (1991) can be applied to help cardholders decide whether or not to seek professional help in managing their credit card debt. See Shefrin and Nicols (2014).

to non-invasion constitutes one of the cases in Theorem 8. Recall that this case pertains to individuals who find themselves at the point of indifference between invasion and non-invasion. The significance of this tipping point is that such individuals have a negative marginal propensity to consume from income.

Consider two types of individuals who have low assets and invade their future income accounts by using credit card debt, but at date t are at the margin between invasion and non-invasion. The first type has a high value of $\varphi_{1,F,t}$ and therefore incurs high willpower costs. As a result, the first type chooses high consumption from the future income account. This corresponds to the case of borrowing near or at the credit limit. For this type, a small increase in income at t leads to a sharp drop in consumption at t. This is the type of person who would benefit from cutting up their credit cards, in case no such increase in income materializes, to induce them to avoid invading the future income account. The second type of individual has a low value of $\varphi_{1,F,t}$ and therefore incurs low willpower costs. As a result, the first type chooses low consumption from the future income account. This corresponds to the case of minimal borrowing. For this type, a small increase in income at t leads to a small drop in consumption at t, and therefore a negative marginal propensity to consume that is low in absolute value.

In respect to the aggregate MPCs from income and wealth, a key issue is the contribution of individuals at the margin, and the degree to which they are close to each type. Keep in mind that individuals who invade their future income accounts, but are not at the margin, do not have negative MPCs from income.

The statistics on credit card utilization rates, meaning credit card usage relative to credit limits, do not suggest that credit card users are mostly nibblers, or that the tipping point case comprises a significant proportion of cardholders. About half of cardholders do not use their cards to borrow; the 50th percentile uses an average of about 9% of their available credit limit. The members of the top 25th percentile use an average of 56% of their available credit limit. Finally, members of the top 10th percentile nearly max out their available credit limit, at an average of 93%.[77] The cumulative distribution function

[77] The source for the percentile data is the Federal Reserve (FRED): https://fredblog.stlouisfed.org/2022/12/credit-card-balances-utilization-rates/.

(cdf) of the credit utilization rate features a mass, about 50%, at zero. To the right of zero, the cdf features an inverse S-shape, which is near linear except in the regions near 0 and 1. The cdf characterizes the degree of heterogeneity among cardholders.

Willpower and mental accounting are subtle. Account invasion is subtle. While the BLC framework highlights three accounts, in practice there are many more, especially in the multicommodity setting. Lottery winnings might sit in a different asset account than emergency savings. In terms of the model, the relevant φ variables might be different across accounts. People might view winnings from gambling more as "house money" than hard earned savings.

Subsection 13.4 provides statistics about retirement preparedness in the United States. The BLC approach focuses on the heterogeneous nature of how people save. Some use Christmas clubs, which are internally enforced rules. Some join firms that offer bonuses paid in large lump sum amounts. Some choose excessive income tax payments from their wages and salaries in order to generate lump sum tax refunds; this simulates the lump sum bonus. Some people are successful savers, and some are not. Some are disciplined about dipping into assets, and some are not. There is heterogeneity in respect to the degree to which people engage in rational saving behavior. The BLC framework seeks to model the range, including the use of internal self-nudges such as Christmas clubs and external programs such as 401(k) plans.

Overwithholding of taxes is an example of changing the shape of the income profile. Recall case 1 above, where the individual's income profile coincides with the first best consumption plan. Choosing an income profile is part of choosing a rule. That is a key point associated with the four cases. Cases 2, 3, and 4 feature undersaving stemming from present bias, and the associated need to employ costly willpower. At the same time, a case featuring high assets in old age might feature underconsumption at this stage of life, because of the psychological difficulty of breaking a rule that was well suited to accumulating assets earlier in life. This is the point about $\varphi_{4,A,t}$ in the second part of Theorem 8.

13.7 Windfalls

Windfalls provide a major opportunity for individuals to engage in overconsumption. The BLC treats windfalls somewhat differently than the traditional framework. In a traditional expected utility framework, an individual spreads a windfall over time in order to equate marginal expected utility, with the usual caveat about liquidity constraints. However, present bias will induce current consumption out of the windfall to be excessively high relative to the traditional case.

In the BLC, a windfall typically arrives through the income account. All else being the same, the higher account balance generates higher marginal doer utility for each consumption level funded by the income account. Temptation m also increases for the income account, and this results in greater willpower being required to induce any interior consumption level funded by the income account. The net effect of a windfall will be for the planner to choose to increase consumption from the income account; and whatever is not consumed will be saved and deposited into an asset account. Notably, the higher asset account balances result in higher future temptation for this account. This combination shifts the calculus for deciding whether or not to invade the asset account, and if so, how much to consume from that account.

Shefrin and Thaler (1988) hypothesized that the marginal propensity to consume (MPC) would differ across the three categories of accounts, being highest for the income account, lower for the asset account, and zero for the future income account. They tested the hypothesis using survey evidence from part time MBA students, all of whom were working professionals. The survey presented subjects with a \$2,400 windfall, received in three different ways. First, individuals received a workplace bonus consisting of 12 monthly payments consisting of \$200 each. Second, individuals received workplace bonus consisting of a lump sum payment of \$2,400. In the first two cases, the announcement of the bonus occurred in the month they received their first bonus payment. Third, individuals received an inheritance with a present value of \$2,400, to be paid to them as a lump sum five years hence.

For each scenario, subjects were asked by how much they would increase their consumer purchases, first in the month they received the announcement about the windfall, and second during the first year. The median responses for the 12-equal payments scenario were 4% and 50%. The median responses for the lump sum scenario were 17% and 33%. The median responses for the inheritance scenario were 0% and 0%.

Fagereng, Holm, and Natvik (2021) present evidence about the MPC-behavior of windfalls received by lottery winners in Norway. In their sample, all lottery winners receive their prizes within a few weeks of the lottery draws. Fagereng et al. report that winners of small prizes, less than $2,000, tend to spend everything within the year of winning. They note that these winners tend to spend even more by combining the money won with other sources of financing. For larger prizes, above $8,000, winners spend less than half of their winnings within a year of winning. In this case, spending peaks in the year of winning and reverts to normal within five years.

The survey MPC-results from Shefrin and Thaler (1988) are somewhat lower than those reported by Fagereng et al., but roughly consistent. The $2,400 windfall in the former would correspond to $3,640 in 2000 dollars, where the year 2000 is the midpoint of the Norwegian lottery data time frame (1994-2006). This would make $2,400 an intermediate size windfall. The second scenario from the survey would be most similar to the Norwegian lottery case, as that scenario involves a lump sum payoff.

Fagereng et al. develop a discrete time model in which an individual with power utility maximizes lifetime expected utility over a finite number of periods. The individual faces idiosyncratic income risk, a life-cycle income profile, and a borrowing constraint. The individual can save in two assets: a liquid asset paying no interest, and a high return illiquid asset having associated transaction costs. The model implicitly assumes no present bias.

Clearly, the empirical findings in Fegereng et al. imply a differential MPC effect. Their model implies that the one-month MPC is 0.37, the quarterly MPC is 0.47, the half-year MPC is 0.54, and the one-year MPC is 0.63. Because future consumption from windfalls is funded from savings (an asset account), there are differential MPCs between accounts. The one-year MPC is certainly less than 12 times the one-month MPC.

Fagereng et al. note that although their model generally accounts for the estimated MPC heterogeneity and time profile, it "struggles" to account for the MPC level. Specifically, the model fits the systematic patterns, but in the year of winning misses a "base" response of approximately 0.35. To fill the gap, Fagereng et al. discuss three possible enhancements, to be added on top of their model. The first involves "temptation shocks," the second involves "near rational behavior," and the third involves "consumption durables" (which the authors find less plausible than the other two).

Fagereng et al. do not cite any of the behavioral literature, and in personal correspondence have stated a lack of familiarity with this literature. Yet, in the end, they implicitly acknowledge that it is the behavioral element, temptation, which is missing from their analysis. Also missing from their formal approach is mental accounting, the basis of the BLC architecture.

With Theorem 8 in mind, consider the finding that small windfalls are consumed very quickly but large windfalls are consumed over five years. An individual who consumes the full amount in their income account, but nothing more, will tend quickly to consume 100% of a small lump sum windfall. This is because anything saved from the windfall will flow to the asset account and not be available before retirement, without invading the asset account and incurring the entry fee. On the other hand, if the lump sum windfall is large, the individual will typically find it beneficial to use willpower to consume less than 100% immediately and transfer the remainder to the asset account. The additional balance in the asset account will increase the temptation variable m associated with the account, and if sufficiently large will induce the individual to incur the entry fee and consume from that account.[78]

Boehm, Fize, and Jaravel (2025) also analyze MPCs (from income) associated with windfall effects, using an experiment in which some participants receive a MasterCard debit card that is linked to a new transactions account having an initial balance of €300. They present five findings from their experiment, which are as follows:

[78] Some individuals will set up a separate asset account for windfall winnings. Doing so allows the individual to spend down the windfall at a different rate than other assets. This feature is consistent with MPC estimates reported by Fagereng et al. being higher than those obtained by others discussed below, namely Levin (1998) and Kaplan and Violante (2014, 2022).

1. The four-week MPC associated with the windfall is 23%. This value is similar to the one-month 17% figure reported by Shefrin and Thaler (1988) for a lump sum bonus. This contrasts with the 4% value for the MPC that Shefrin and Thaler (1988) report for the same size bonus when it is distributed uniformly over the course of a year. While the 4% pertains to the total annual bonus, the associated MPC out of cash received during the month is actually 50%. As discussed below in the subsection about dividends, the MPC from dividend income is 49%, and dividend income also tends to have a uniform profile during the course of a year.

2. A much higher MPC of 61% applies when the terms of the card are such that the any unused funds expire after three weeks. Boehm et al. note that this finding is inconsistent with money fungibility. This finding provides support for a proposal in Shefrin (2009b) recommending the use of vouchers with expiration dates for generating fiscal stimulus.

3. The consumption response is concentrated in the first three weeks. This finding is similar to that reported by Fagereng et al. for small lottery prizes, and as mentioned above, can be explained in a BLC framework.

4. MPCs vary with household characteristics but are high even for households having significant amounts of liquid wealth. This finding is consistent with the numerical BLC example discussed above. Households whose desired consumption profiles are close to their income profiles, perhaps by design, will use mental accounting to protect their tempting liquid assets from invasion. However, incremental income is unprotected from a weakness of will, and individuals will respond to incremental income by consuming more than the associated "annuity" value.

5. The unconditional MPC distribution exhibits large variation. As discussed above, heterogeneity in respect to willpower costs, along with demographics, produces substantial variation in MPCs (from income). Boehm et al. also report striking demographic effects. Notably, the MPCs associated with men are about twice those of women. MPCs also increase with age, a finding common to both the BLC and

neoclassical models. Boehm also find little evidence of substantial negative MPCs. Their finding contrasts that of Misra and Surico (2014) who do report substantial negative MPCs in the consumer population.[79] As discussed above, Theorem 8 identifies conditions under which the MPC from income is negative. Notably, these conditions correspond to a set of measure zero in parameter space, and appear to be unrepresentative of the economic situations faced by many households.

Boehm et al. state that their "results are difficult to reconcile with agents being rational and treating money as fungible." Nonfungability is a key feature of the BLC. In this regard, Boehm et al. develop a mental accounting model to describe some of their findings. Separately, they also discuss present bias, but not in conjunction with mental accounting; nor do they mention temptation, a key feature of the BLC. However, they do point out that the $\beta\delta$ framework, which does not include temptation or mental accounting, cannot explain framing effects along the lines of their second finding described above. What is worth noting is that Boehm et al. discuss how their approach relates to a two-system TF&S-framework, one of the few papers discussing MPC issues to do so.

As was just mentioned, the analysis in Boehm et al. includes three key features: present bias, mental accounting, and two-systems. While these are discussed separately, it is vitally important to build economic models which integrate all three, as does the BLC. It is also important to incorporate the concept of temptation, a critical feature of present bias; and temptation is one of the key issues formally introduced by Shefrin and Thaler (1988) into the present bias approach. In the BLC, the MasterCard account balance can be placed in its own mental account, without a pecking order restriction.[80] Significantly, it can also feature a different degree of temptation than other accounts. This recognition, along with matrix mental accounting, offers insights into several of the findings reported by Boehm et al.

[79] Boehm suggests a possible econometric problem in Misra and Surico (2014) which is associated with their interpretation of two variances. The variances in question pertains to consumption changes, one for the receivers of incremental income (transfers) and the other with the not-yet or previously-received. The former has a larger variance than the latter, an inequality which forms the basis of Misra and Surico's conclusion about the very large relative proportion of MPCs that are negative.

[80] See the discussion earlier in this subsection, in the previous subsection, and the subsection about optimization nuances.

13.8 Dividends: Consumption, Withdrawals, and Clienteles

The matching of expenditures to sources of wealth is a critical feature of the BLC, and leads to differential MPC behavior. This is particularly the case with income received from cash dividends [Shefrin and Statman (1984), Shefrin (2009a)].

As discussed above, many individuals use the adage "don't dip into capital" to describe their reluctance to fund consumption from their portfolios, instead preferring to use dividend income for that purpose. In other words, individuals prefer to fund consumption using dividends instead of capital because dividends flow through the income account, and capital resides in the asset account. Not dipping into capital is a way to exercise self-control and avoid overconsumption. Not dipping into capital is a self-imposed liquidity constraint, which the BLC predicts will have differential MPC effects.

Baker, Nagel, and Wurgler (2007) provide the strongest evidence for a connection between consumption and dividends. Their analysis makes use of two types of data sets: the Consumer Expenditure Survey (CEX) from the Bureau of Labor Statistics and brokerage account data. The CEX provides panel data on household consumption, income, wealth, dividends, and demographic information. The brokerage data provide information about how individual investors behave when they receive dividends.

The CEX data used by Baker et al. (2007) cover the period 1988–2001 and involve 3,272 household-year observations. Total expenditure including durables is about $50,000 (measured in December 2001 dollars). Total income, which includes dividends but not capital gains, has a mean of $56,789. Financial wealth is typically around a third of total wealth. For the mean household, interest income is $1,207 and dividend income is $891. On average, interest and dividends account for 4% and 2% of total income, respectively. The distribution is skewed, with the median household reporting zero dividend income.

As discussed above, the BLC implies that the marginal propensity to consume from the income account exceeds the marginal propensity to consume from the asset account. Shefrin and Statman (1984) suggest that dividends flow through the income account and that capital gains flow through the asset account. This implies that the marginal propensity to consume from dividend income is about the same as the marginal propensity to consume from total income, which in turn exceeds the marginal propensity to consume from capital gains.

Baker et al. (2007) estimate that the marginal propensity to consume from dividend income is 0.49, the same as the marginal propensity to consume out of total income, whereas the marginal propensity to consume from total current-year returns is close to zero. This estimate is virtually identical to the value reported by Shefriin and Thaler (1988) for the case when incremental income is distributed evenly over the course of a year.

Baker et al. analyze a data set developed by Barber and Odean (2000). These data, which cover household portfolios, contain monthly position statements and trading activity for a sample of 78,000 households that had accounts at a large discount brokerage firm. The analysis covers the period January 1991 to December 1996.

The mean account value in the Barber-Odean data set is \$54,400 and the median is \$28,400. For the mean household, about 83% of this value stems from holdings of common stock and 14% stems from holdings of mutual funds. The average total monthly return is 1.11% . Notably, the average dividend income per month is 0.20% of the beginning-of-month portfolio value. In about one-half of household-months, dividend income is positive. Here, 78% of dividend income is from ordinary dividends, with the remainder from mutual funds. Special dividends are infrequent but typically large. Households using dividend income to finance consumption must first withdraw those dividends from their brokerage accounts.

The Barber-Odean (2000) data provide an opportunity to investigate the connection between ordinary dividends and net withdrawals. Baker et al. (2007) report a positive relationship between net withdrawals and ordinary dividends that is close to one-for-one for small dividends. This statement does not apply to mutual fund dividends, which appear to be automatically invested.

Notably, households only withdraw a portion of large dividends. In this regard, the BLC stipulates that households require much less willpower to save large windfalls that find their way into the income account or for that matter the asset account. Therefore, they might increase consumption in the face of a windfall (e.g., a special dividend) but still be able to save most of it. Baker et al. (2007) report that households reinvest very large special dividends, suggesting that households allocate these amounts to the asset account, effectively treating them as capital.

Consistent with the BLC, Baker et al. (2007) report a much smaller effect for capital gains than for dividends. Specifically, regardless of the level of capital gains, median contemporaneous net withdrawals are close to zero. On average, the marginal propensity to withdraw contemporaneous dividends is 0.35 and the marginal propensity to withdraw ordinary dividends is 0.9. In contrast, the marginal propensity to withdraw capital gains is 0.02. Still, the size of capital gains is sufficiently large that the aggregate effect of that 0.02 is fairly sizable.

Baker et al. (2007) note that investors withdraw mutual fund dividends at a much lower rate, presumably because mutual fund investors use automatic reinvestment policies. The authors also note that investors withdraw small special dividends at roughly the same rate as investors withdraw ordinary dividends, but they mostly reinvest large special dividends.

A key feature of the Shefrin-Statman (1984) framework is that households view ordinary dividends as predictable, similar in nature to pension income, Social Security payments, interest payments, and labor income. This is because households choose portfolios with the intent to consume ordinary dividends. According to Baker et al. (2007), ordinary dividends, scaled by beginning-of-period portfolio wealth, are predictable. They find that dividends lagged by one year explain 57% of the variation in ordinary dividends. Moreover, the combination of dividends lagged one year and dividends lagged three months explain 81% of the variation. Baker et al. also report that mutual fund dividends are less predictable than ordinary stock dividends, and that special dividends are unpredictable.

Dividends are not completely predictable. For this reason, the question is worth asking whether households attempt to smooth consumption or whether they consume the unpredictable component as well as the predictable component. Baker et al. (2007) investigate this issue by adding the 12-month lag of dividends to their analysis of withdrawals, to control for the predictable component. They find that lagged dividends explain a portion of withdrawals but that their inclusion has no discernable effect on the impact of contemporaneous dividends. Similar statements apply to the CEX data, although dividends are less predictable in that data set than in the Barber and Odean (2000) data set.

The work of Rantapuska and Kaustia (2007) provides additional information about reinvestment behavior associated with ordinary dividends, special dividends, and tender offers. Their study uses Finnish data from the Finnish Central Securities Depositary (FCSD). A virtue of these data is that the FCSD records every transaction involving Finnish securities on an account basis. Rantapuska and Kaustia use transactions between 1995 and 2002. Their data include transactions by households and various nonhousehold investors.

The main finding from the Finnish data is that investors reinvest less than 1% of dividend payments within ten days of receiving payment. Some of this behavior can be attributed to the dividends being small, which does not justify the transaction costs of reinvestment. However, investors are also averse to reinvesting larger ordinary dividends, where transaction costs are not a major issue. Here reinvestment rates lie in the range of 4% to 8%. Reinvestment rates are larger for special dividends and tender offers. Surprisingly, the associated reinvestment rates for special dividends and tender offers are not large. For special dividends, the reinvestment rate is about 8% and for tender offers it is about 13% . In fact, 57% of households choose not to reinvest either special dividends or tender offer proceeds within ten days of receipt.

The behavioral theory of dividends predicts how both ordinary dividends and special dividends will affect consumption. The theory also predicts how the demand for dividends will vary over the household life cycle, depending on a household's needs to replace labor income. In particular, behavioral theory predicts that older, retired households will find dividends more attractive than younger households that are still in the workforce. Moreover, households that have few options for replacing labor income after retirement will find dividends especially attractive. Low-income households are similar to older, retired households. For both groups, the shadow price of current income is high, and yet both are trying to protect the balances in their asset accounts. In this respect, demand for dividends will be negatively related to income.

Behavioral dividend theory predicts that clienteles will form around age, retirement status, and low income. Graham and Kumar (2006) document how the demand for dividends depends on age, retirement status, and income. As with Baker et al. (2007), Graham and Kumar analyze the data used in Barber and Odean (2000).

Demographic information including age, income, occupation, marital status, and sex is available for a subset of 31,260 of the 77,995 households in the sample. Graham and Kumar (2006) focus on the behavior of older and low-income investors. They partition households by age, income, and occupation. For age, the partition boundaries are below 45 (younger), between 45 and 65, and above 65 (older). For income, the partition boundaries are annual household income below $40,000 (low income), between $40,000 and $75,000, and above $75,000 (high income). For occupation, the partition boundaries are professional, nonprofessional, and retired.

About 15% of investors in their sample are at least 65 years of age and about 17% have annual income below $40,000. Older investors hold larger portfolios than their younger counterparts ($55,685 versus $24,402, respectively), and their portfolios contain more stocks (six versus four, respectively). Older investors also turn their portfolios over at a lower rate than younger investors (monthly portfolio turnover of 4.46% versus 5.36% respectively).

Graham and Kumar (2006) measure the preference for dividends using portfolio dividend yield. They report that irrespective of income, older households prefer dividend-paying stocks over non-dividend-paying stocks. For younger households, those with low income exhibit a somewhat stronger preference for dividend-paying stocks than households with higher income. Moreover, the preference for dividends is relatively stable, with dividend yields of individual investors' portfolios changing slowly over time. These findings support the predictions of behavioral dividend theory about the impact of age and income.

In terms of stock selection, Graham and Kumar (2006) report that older, retired investors hold about 80% of their portfolios in dividend-paying stocks, while younger investors hold about 65% of their portfolios in dividend-paying stocks.[81] In particular, older investors with low incomes attach more weight to utility stocks than do younger investors with high incomes. The reverse is true for the stocks of firms in the computers and business services industries. This pattern is consistent with the idea that older, low-income investors focus more on industries with a reputation for paying higher dividends.

Bräuer, Hackethal, and Hanspal (2022) provide a thorough analysis of dividend consumption behavior. Their results are consistent with the studies described above. Most importantly, they report that the marginal propensity to consume from dividends is strongest for those who are disciplined planners, and greater than for those who appear to be more impulsive spenders. Bräuer et al. suggest that the latter finding is contrary to the theory presented in Shefrin and Statman (1984), which they describe as positing "that investors with self-control problems can 'safeguard their wealth against compulsion and immediate gratification by employing a rule' where they are to consume only from dividends rather than the portfolio capital." This description is accurate. However, Bräuer et al. then continue, stating: "It follows that those with self-control problems would use dividends as a commitment device and tie their consumption to this source of income." This last statement does not follow from the description they provide. A more accurate statement would involve adding the term "some of" so that the phrase reads *some of those with self-control problems ...*

The point is that not everyone with a potential self-control problem is able to structure a disciplined rule to deal with that problem. Those that do will be perfectly comfortable planning to consume dividends when they arrive as anticipated. Indeed, Bräuer et al. report that dividend income is consumed quickly, in line with the BLC. They state: "[P]rudent investors have higher levels of wealth, tilt their portfolios towards larger and more frequently occurring dividends, and consume out of them. Their average MPCs

[81]Consistent with the predictions of behavioral dividend theory, retirement status increases the demand for dividends. Graham and Kumar (2006) regress portfolio yield on a series of variables including age and a retirement status dummy. They report that the age variable has a positive and significant coefficient estimate (0.038 with a t-statistic of 4.22), while the dummy for retirement status has a coefficient estimate of 0.010 with a t-statistic of 2.03.

are nearly 30% in the week following dividend receipt, compared to 1.7% found among investors classified as imprudent."

In the subsection above about differential marginal propensities to consume, there is a discussion about why individuals having the discipline to follow rules such as "don't dip into assets" have higher marginal propensities to consume, and more saving than individuals lacking the discipline to do so. This theoretical feature is consistent with the findings in Bräuer et al., which overall are are highly supportive of the behavioral approach.

13.9 MPC-Behavior in a Multicommodity Setting During Late Middle Age

Levin (1998) tests the BLC for people in late middle age, using multicommodity data from the Longitudinal Retirement History Survey (RHS). These data pertain to 11,000 household heads whose ages spanned 57 to 62 at the beginning of the survey. This age group is of particular interest for testing the BLC, because it has higher asset balances than its younger counterparts. The survey was repeated every two years, for ten years, and asks participants to indicate whether or not they are retired.[82]

The consumption data in the RHS is broken down into eight consumption categories such as food consumed at home (groceries), restaurant meals, transportation, entertainment, gifts, and dues. This enables Levin to analyze the category-based MPCs associated with income and different types of assets. Examples of liquid assets are checking and savings accounts, stocks and bonds. Notably, stocks pay dividends and bonds pay interest, which are both forms of income. An example of an illiquid asset is property, with rents being a form of income. Levin's main findings are as follows.

1. MPCs from income are dramatically different than MPCs from wealth: a $ 1 increase in income has as much effect on consumption as a $26 increase in wealth.

[82] The data included detailed information about family status, wealth, social security, income and expenditures. The data were collected between 1969 and 1979.

2. Households are much more willing to spend out of liquid assets than they are out of other forms of wealth. In three of the eight spending categories, the difference between the coefficient of liquid wealth and property is positive and statistically significant. In cases where the difference is negative, the t-statistics are always small. In the aggregate, the effect of property on aggregate consumption is one-third the size of liquid wealth's effect. Changes in housing wealth have no detectable effect on consumption. For this group, future income largely consists of payments from Social Security. The effect of future wealth on aggregate consumption is approximately 75% of the effect of liquid wealth, largely concentrated on grocery spending.[83]

3. Households with little or no emergency (rainy day) funds are liquidity constrained. Levin partitions his sample into two groups, one unconstrained and the other constrained. The unconstrained group have liquid wealth that is at least three months of household income in at least five of the six sample years. The constrained group consists of everyone else. Levin states: "Individuals do not appear to finance consumption out of less tempting or illiquid assets unless their more liquid assets are exhausted. Thus liquidity constraints seem to be internally not externally imposed... [T]he constrained group has already broken the mental barrier against spending out of the illiquid assets. Therefore, their spending will be sensitive to changes in those assets, whereas the unconstrained group's spending will not be very sensitive. [T]he propensity to consume out of any form of illiquid wealth is smaller for the unconstrained group than it is for the constrained group. [T]he unconstrained subsample contains relatively successful savers who may have adopted the behavioral rule of not consuming out of less tempting forms of wealth..." (p. 79)

   For every extra dollar of future wealth, aggregate consumption for the constrained group increases by 2.16 cents. In contrast, for the unconstrained group, the increase is 0.80 cents. For the unconstrained group, the MPC for property wealth is 0.25 cents per dollar, and 3.26 cents for the constrained group, (13 times more). For housing

[83] Levin states: "One possible reason for this is that future wealth is mainly the value of social security, which increased rapidly in this period. Perhaps an increase in one period increased people's confidence that it would continue to increase in the future."

wealth, the unconstrained group's MPC is generally negative (for six of the eight consumption categories) and has an aggregate value of -0.59 cents. Notably, the constrained group's MPC from housing wealth is 5.42 cents. Generally, as Levin notes, "the unconstrained group's propensity to consume out of illiquid forms of wealth is significantly lower than the constrained sample – just the reverse of what the external liquidity constraint hypothesis would have predicted but exactly in line with the prediction of the behavioral life-cycle." (p. 79)

4. The unconstrained group's MPC out of income is approximately 0.4, and for the constrained group is 0.25. The lower value for the constrained group is consistent with this group having invaded its asset accounts more than the unconstrained group. The discussion above provides a theoretical explanation, based on Theorem 8 and the principle of "bang for the buck," for why the MPC from income for the constrained group is less than the MPC for the unconstrained group.[84]

5. By and large, the unconstrained group spends little of their illiquid assets, such as the value of their homes, thereby increasing the size of wealth to be bequeathed. Levin states: "A large increase in home value will be entirely 'saved' for a bequest

[84]The constrained group is more apt to invade the asset account than the unconstrained group. In the numerical example presented in subsection 11.5, young individuals with low incomes (below 1.1 in the example) borrow against future income. As their incomes increase over time, they pay down debt, which leads to lower marginal propensities to consume from income than those who have higher incomes and are less constrained. Relatedly, Theorem 8 establishes that individuals who are indifferent between invading and not invading, but choose to invade, have negative marginal propensities to consume from income. For example, if they are carrying credit card debt, the extra income induces them to avoid taking on additional debt. In the aggregate, the indifferent-subgroup will reduce the average marginal propensity to consume for the constrained group. The above discussion about credit card debt suggests that the constrained group is a mix. Many are young and have low income. While some incur low costs from exercising willpower, others have high costs. The discussion above about "bang for the buck" notes that marginal propensity to consume reflects how marginal income impacts relative willpower costs. If the planner judges that at the margin, a low marginal propensity to consume leads to a sharp decline in marginal willpower cost, then it will choose a low marginal propensity to consume from income, even though the marginal unit of consumption is from the asset account. Sharp declines in marginal utility tend to occur near values of $m_{i,t}$, which for asset accounts can be credit limits. See Figure 7 in this regard. If marginal willpower costs for the constrained do not decline, relatively, when the allocation from marginal income is low, then the planner will prefer that the marginal propensity to consume be high. This might be the case for the constrained when their income accounts are active, and they are engaged in discretionary saving. Pertinent to this last point is the magnitude of the respective φ-variables, and the extent to which these variables depend on changing values of temptation.

while an equivalent rise in a more liquid asset will not." (p. 79)[85]

6. In respect to budgeting behavior in the multicommodity setting, Levin presents evidence for matrix mental accounting, meaning how consumption categories are funded from different assets. As an example, he mentions the possibility that "an increase in housing equity coupled with an equivalent decrease in another asset may cause vacation spending to rise while grocery spending falls..." (p. 80).[86] To test for this type of phenomenon, he focuses on each consumption category, and computes the ratios of the MPC-wealth coefficients divided by the corresponding MPC-income coefficient. In a neoclassical setting, these ratios should all be equal between any two consumption pairs. The evidence indicates that this is not the case for many such pairs. For example, charitable donations are funded very differently from vacations, transportation, and restaurant meals. Groceries too are funded differently from vacations and transportation. Levin states: "Vacation spending was most likely to reject the equal ratio test ... perhaps because vacations are more like a durable good than the other seven [consumption categories]." (p. 81)[87]

13.10 Hand-to-Mouth Behavior, Heterogeneity, and MPCs

There is a macroeconomics literature analyzing the degree to which the MPC from income varies across three particular groups of households. For example, one study reports that the elasticity of consumption of nondurables with respect to income is 0.21. [Ganong,

[85]Graham and McDowall develop a mental accounting model to explain U.S. consumption spending responses to predictable income. They report that "even for households with large liquid asset balances, we find no spending in anticipation of income receipt, substantial spending following receipt, and significant front-loading with respect to date of receipt." This is consistent with the MPC from future income being zero, and the MPC from income being positive and large.

[86]Boehm et al. (2025) also report effects consistent with matrix mental accounting, noting that households which receive the €300 incremental income spend a disproportionate share on clothing and household equipment.

[87]Data from a 2024 Discretionary Spending Survey by Bankrate provides support for Levin's findings about mental accounting relationships. These data indicate that 38% of adults in the U.S. would have been willing to go into debt to travel, dine out, or see live entertainment. Of those willing to go into debt, 27% would have been willing to take on debt to travel. The corresponding percentages for the other two categories are 14% for dining out and 13% for live entertainment. See Gillespie, Lane, 2024. "Bankrate's 2024 Credit Card Debt Report," edited by Tori Rubloff, https://www.bankrate.com/credit-cards/news/credit-card-debt-report/#credit-card-rewards.

Jones, Noel, Greig, Farrell, and Wheat (2023)]. Ganong et al. also report that the elasticity declines monotonically as liquid assets increase, ranging from 0.50 for the lowest-asset households to 0.08 for the highest-asset households.[88]

The MPC from income is a central variable in macroeconomic analysis. Kaplan and Violante (2022) point out that the MPC determines the size of fiscal multipliers, the transmission mechanism of monetary policy, the amplification of aggregate shocks, and portfolio choice between risky and safe assets. Macroeconomic models which explain average MPC as an aggregate of MPCs from heterogeneous subgroups of households are known as "heterogeneous agent neo-Keynesian (HANK) models [Violante (2021)].[89]

Kaplan and Violante (2022) analyze the average MPC by identifying how the MPC for a $500 increment in income varies across heterogeneous groups. In this regard, they identify three groups. They call the first group "poor hand-to-mouth" (HtM) households. They state that this group has "zero wealth" and are "borrowers for whom credit constraints bind. Similar to what Levin (1998) reports, these types of households consume the majority of any extra liquidity they receive." (p. 749) The second group they call "wealthy HtM households." They note that members of this group holds "the bulk of their wealth in illiquid assets. These households' consumption responds to small windfalls, similarly to the consumption of poor HtM households." Kaplan and Violante call the third group "non-HtM households."

Kaplan and Violante's classification of households is similar to the four credit card user clusters discussed above. Two of the clusters correspond to poor HtM households, one to wealthy HtM households, and one to non-HtM households. Poor HtM households

[88]The data for this study comes from bank account data from the JPMorgan Chase Institute (JPMCI). These data feature monthly observations on 20 million households having a Chase checking account, for the period January 2018 through November 2022. The analysis focuses on three variables from these data, namely income, spending, and assets.

[89]Gul and Pesendorfer (2004) provide a self-control based framework for consumption-savings problems, stating: "Our goal is to reconcile the experimental evidence with tractable, dynamically consistent preferences and to apply the resulting model to the analysis of problems in macroeconomics and finance." They do not cite works such as Shefrin and Statman (1984, 1985) which apply self-control to finance or Shefrin and Thaler (1988) which applies self-control to macroeconomics. Gul and Pesendorfer develop a model with a penalty function-structured valuation equation similar to (55) above, which they describe as "main theoretical result of this paper (Theorem 1)..." There is no discussion in their article about marginal propensity to consume.

share important traits with cluster 2, whose members relative to other clusters are more concentrated in the youngest age group, are the most inclined to carry a credit card balance, and who rely the most on their credit cards to finance emergencies. In contrast, wealthy HtM households share important traits with cluster 3, whose members are more concentrated in the oldest age group, use multiple credit cards mainly for everyday purchases, and are most inclined to pay the full balance every month. Members of cluster 4 feel a sense of confidence and control about managing their household finances, sharing important traits with non-HtM households.

Heterogeneity applies to willpower costs as well as income and wealth. In this respect, the members of cluster 1 are similar in many respects to those in cluster 2, as poor HtM, but engage in more impulse buying and feel much less in control of their finances.[90]

Kaplan and Violante (2014) partition assets into two classes, liquid assets such as cash and checking accounts and illiquid assets such as housing and retirement accounts.[91] Kaplan and Violante (2022) use data from the 2019 Survey of Consumer Finances (SCF). There are several possibilities for operationalizing the notion of HtM, based on the SCF. One definition proposed by Kaplan and Violante is households with liquid wealth less than half of monthly income. By this definition, 14% of households are poor HtM, 27% are wealthy, implying that 41% of US households are HtM. Because MPCs vary across the three groups, the aggregate MPC is a weighted average of the MPCs across the three groups. Kaplan and Violante state the following: "Multiplying the average MPC for poor HtM households (24%) by their share (13%), the average MPC for wealthy HtM households (30%) by their share (27%), and the average MPC for non-HtM households (7%) by their share (60%), and adding up these three components, we obtain an average

[90]Recall from the above discussion that 36% of U.S. adults have more credit card debt than emergency savings.

[91]Kaplan and Violantge (2014) state: "Our definition of liquid assets comprises: money market, checking, savings and call accounts plus directly held mutual funds, stocks, bonds, and T-Bills net of revolving debt on credit card balances. All other wealth, with the exception of equity held in private businesses, is included in our measure of illiquid assets. It comprises housing net of mortgages and home equity loans, vehicles net of installment loans, retirement accounts (e.g., IRA, 401K), life insurance policies, CDs, and saving bonds.
As expected, the bulk of household wealth is held in (what we call) illiquid assets, notably housing, vehicles, and retirement accounts."

MPC close to 16%."[92]

Kaplan and Violante (2014) evaluate the MPC associated with a $480 tax rebate connected to the 2001 recession. They report two MPCs. For HtM households, the MPC is 45%. This is close to the 50% figure reported in Shefrin and Thaler (1988) for working professional MBA students, most of whom were liquidity constrained. For non-HtM households, Kaplan and Violante report an MPC which is near zero. They note that "most of the households in the model behave exactly as predicted by the PIH [permanent income hypothesis] and have zero MPC." Kaplan and Violante note that there are high earnings households at both ends of the distribution, made up of both wealthy HtM households and unconstrained non-HtM households. Unconstrained non-HtM households are income-rich with low expected earnings growth. In contrast, wealthy HtM households are income-rich with high expected earnings growth, a combination which induces a boundary consumption choice ($c_t = y_t$). As a result, with the tax rebate being a lump sum, the income-richest have the highest MPC among constrained agents. Kaplan and Violante also note the possibility of negative MPCs, stating: "When transaction costs are low enough, upon receiving the rebate, some agents choose to anticipate the adjustment decision and save the rebate, together with their current holdings of cash, into the illiquid account. As a result, they save more than the rebate amount (which explains the negative MPC ...)"

Kaplan, Violante, and Weidner (2014) discuss additional insights about the characteristics of HtM households.

1. The majority of the HtM group consist of wealthy HtMs, and therefore tend to be missed by measurements of HtM behavior that are based on net worth.

2. Over the life cycle, households that qualify as poor HtM do so at young ages. In contrast, the age profile of the wealthy HtM-group is hump-shaped; and it peaks in the vicinity of age 40.

[92]Levin (1998) reports that 38% of his sample are constrained/HtM, which is close to the 41% figure mentioned above. However, Levin reports higher values for MPC from income than in the discussion above. In Levin's sample, the average MPC is 34%, very close to the MPC of 33% reported for the bonus case in Shefrin and Thaler (1988).

3. The wealthy HtM typically hold sizable amounts of illiquid wealth: The median at age 40 is about $50,000.

4. Wealthy HtM households appear similar to the unconstrained non-HtM in the age profiles of their income and their shares of illiquid wealth held in housing and retirement accounts.

The general findings from the various studies discussed above are consistent with the BLC. Recall that the numerical example discussed in Section 11 describes a series of cases associated with varying incomes for young households. The households that are poorest while young have low incomes and invade their future income accounts by borrowing to fund consumption. Households that are slightly wealthier do not invade their future income accounts, but consume 100% of their take-home pay. Those at the margin between the first two groups feature negative MPCs out of income, as the marginal income induces them to stop borrowing to fund consumption. These first two groups qualify as HtM households. A third group with higher income, also young, manage to save, and therefore its members are not constrained. However, their consumption might be close to their take-home pay, as temptation renders it difficult for them to save a lot. In the numerical example, members of this group anticipate significantly higher income in middle age, but choose to behave in HtM fashion and not to invade their asset accounts. They do so to protect against the possibility that they will overconsume if they invade these accounts. The middle aged households in the example are all unconstrained. Because of willpower costs and temptation, they tend to undersave for retirement. Households with low willpower costs do save adequately for retirement, but once retired can become miserly and engage in underconsumption: see the discussion in the next subsection.

Kaplan and Violate develop a series of models to fit joint distributions of MPC, income and wealth. Some of these models are neoclassical, but others are behavioral, incorporating temptation and self-control in a two-asset framework.[93] Of the two assets, the first is liquid, paying a low return while the second is illiquid, paying a high return. The household can fund consumption directly from income and the liquid asset, but not from

[93] Kaplan and Violante (2022) also develop a $\beta\delta$ version of their model.

the illiquid asset.[94] In this respect, funding consumption indirectly from the illiquid asset requires the household to pay a fixed transaction fee (κ) in order to shift wealth from the illiquid asset in its portfolio to the liquid asset.

A key empirical feature is that wealthy HtM households hold low amounts of liquid assets. Kaplan and Violante explain this feature by assuming that the illiquid asset pays a higher return than the liquid asset. Attanasio, Kovacs, and Moran (2024) suggest that this feature renders the Kaplan-Violante explanation fragile, in that assuming that the two assets offer the same returns would lead wealthy HtM households to hold a larger proportion of their portfolios in the liquid asset.[95]

Attanasio et al. propose a temptation-based model to explain why wealthy HtM households choose to hold low levels of the liquid asset. Their model is similar in nature to the BLC.[96] In the two-account numerical BLC-example, the planner's willingness to avoid low consumption during the retirement phase depends on the shape of the income profile. A planner that can constrain consumption during the middle years, without incurring high willpower costs, will be able to shift consumption from that period to the retirement period.[97] If there is a fixed cost guardrail associated with the illiquid asset,[98] then the planner can potentially protect that asset during the middle years, just as if

[94]This last condition imposes a borrowing constraint.

[95]See also Kovacs, Low, and Moran (2021).

[96]Attanasio et al. base their model on Gul and Pesensdorfer (2001), and do not cite Thaler and Shefrin (1981) or Shefrin and Thaler (1988). Recall Levin's finding that illiquid assets are concentrated in housing or future Social Security payments. Attansio et al. develop a model in which the illiquid asset is housing, which provides both direct utility as well as serving as a store of value. They note that homeowners might use mortgage payments as devices to help them save, in the form of acquiring home equity. This aspect of mortgage payments is similar to Christmas clubs, savings programs, and whole life insurance, all of which facilitate asset accumulation.

[97]In the model developed by Attanasio et al., asset accumulation leads to higher temptation in future years, and therefore a higher penalty for consuming in those years.

[98]Attonasio et al. do not associate a fixed cost with the liquid asset. In contrast, the three-account version of the BLC features endogenous entry fees for both the (liquid) asset account and the future income account.

income were lower during those years.[99]

In a neoclassical setting without behavioral elements, Kaplan and Violante report that illiquid assets need to earn about 8% more than liquid assets: 8% is quite high. Kaplan and Violante (2022) adapt the temptation framework developed in Attanasio et al. to reduce the 8% premium by roughly half; and this with an average MPC of 21%.

The valuation function which Kaplan and Violante employ to model temptation and self-control in a discrete time one-asset version of their framework has the form:

$$U_t = u(c_t) + \lambda E\{U_{t+1}\} - \varphi[u(m_{i,t}) - u(c_t)] \tag{59}$$

Here U_t is a function of y_t and A_t, and $\varphi[u(m_{i,t}) - u(c_t)]$ can be interpreted as a penalty function $\Pi(c_t, m_{i,t})$, analogous to the role played by Π in the BLC valuation function (55). The penalty function in (59) conforms to the rightmost term in (42) from Gul and Pesendorf (2001). Notably, the utility functions in Gul and Pesendorfer's framework are linear. However, Kaplan and Violante do not impose linearity; and as Fudenberg and Levine forcefully argue, linearity is an unrealistic property for these types of issues. In the discussion of the BLC above, the utility function is logarithmic. If $\Pi = \varphi[u(m_{i,t}) - u(c_t)]$, then the reduced form doer utility function Z_R has the form

$$Z_R = u(c) - \varphi[u(m) - u(c)] = (1 + \varphi)ln(c) - \varphi ln(m) \tag{60}$$

Equation (59) applies to a one-asset discrete time model where the single asset is liquid. Kaplan and Violante extend this model to continuous time with two assets, one liquid and one illiquid, with a fixed transaction fee for shifting funds from the illiquid asset to the liquid asset. Notably, the valuation function for the two-asset model features a comparison of two local maxima, one where the transaction fee is paid and a fund

[99] In the BLC, individuals seek to control the shapes of their income profiles, for example by selecting stocks with particular dividend payout ratios. As discussed above, they do so as part of a willpower-based rule stipulating "don't dip into capital." Attanasio et al. make no mention of willpower, but do have a penalty function structure for utility. They report that their model "generates patterns of consumption behaviour in response to changes in income that are in line with the available evidence ... that the consumption response to income shocks declines only slowly with shock size." They also "summarize the implications of temptation and commitment for fiscal stimulus targeting."

transfer takes place, and the other where no fund transfer takes place.

Kaplan and Violante describe φ as a temptation parameter. They note that adding even a small amount of temptation substantially raises the value of the average MPC. For example, with a temptation parameter of $\varphi = 0.01$ instead of 0, the average quarterly MPC increases from 16% to 24%, and with a temptation parameter of $\varphi = 0.05$, the MPC increases to 47%. In respect to these magnitudes, keep in mind that the MPC estimates discussed in earlier sections, in Shefrin and Thaler (1988), Levin (1998), and Fagereng et al. (2021), are all significantly higher than the values that appear in the work of Kaplan and Violante.

Consider Kaplan and Violante's temptation model viewed through the lens of the planner-doer two-system framework.[100] To do so, begin with Figure 5 which depicts three functions: the planner function $u = ln(c_t)$, the doer function $Z_t(c_t, .)$ for fixed θ_t and $m_{i,t}$, and the reduced form doer utility function $Z_R(c_t, .)$ for fixed $m_{i,t}$. The penalty function Π_t is given by $\Pi_t = u(c_t) - Z_R(c_t, .)$. As was mentioned above, in the BLC, Π_t is a power function with argument $m_{i,t} - c_t$, whereas in Kaplan and Violante, Π_t is equal to $\varphi[u(m_{i,t}) - u(c_t)]$. Casting the Kaplan-Violante model as a planner-doer model[101] is accomplished by substituting $\Pi_t = \varphi[u(m_{i,t}) - u(c_t)]$ for (41) in (32).[102]

Although Kaplan and Violante do not explicilty incorporate mental accounts, their two asset model with temptation is very similar in structure to a two-asset BLC. Kaplan and Violante's model features a liquid asset and an illiquid asset. The BLC features a liquid asset account and a future income account which is treated as being illiquid. However, the analogy is not exact, because in Kaplan-Violante, the household does not face any direct costs for financing consumption from the liquid asset account. It only faces

[100]This discussion focuses on the discrete time setting.

[101]A similar argument establishes that the framework in Fudenberg and Levine (2006) can be described as a planner-doer model.

[102]Kaplan and Violante (2022) assume that $u(c_t)$ is logarithmic for what they call their baseline parameterization (p.751), but also analyze Epstein-Zin preferences. Theorem 7 above is developed for log-utility, but is easily extended to other utility functions and their corresponding derivatives. For the case of $u()$ being log-utility, $u(m_{i,t}) - u(c_t$ is the percentage difference between $m_{i,t}$ and c_t, as opposed to the additive difference used in the BLC. As a result, the penalty function is proportional to the percentage decline in consumption below the temptation level. To be clear, although the specific functional form of the penalty function in the above discussion involves a power function, the framework developed in Shefrin and Thaler (1988) also permits a logarithmic penalty function.

a fixed fee κ for invading the illiquid asset account in order to transfer funds to the liquid asset account. For this reason, the mental accounting structure in Kaplan-Violante has the form [I+A,F], whereas in the baseline BLC model it is [I,A,F]. Notably, the structure [I+A,F] implies that for the active account I+A, the temptation variable $m_{i,t}$ is the sum $y_t + A_t$, the respective balances in the two accounts. As for account F, in the BLC, the endogenous (fixed) entry fee provides the guardrail to discourage invasion. In contrast, Kaplan and Violante use a transaction cost κ for this purpose. With the embedding of the Kaplan-Violante temptation model into the planner-doer framework, it becomes possible to interpret κ as a psychological guardrail which the planner uses to mitigate the tendency of the doer(s) to overconsume by invading the future income account.[103]

Because Kaplan and Violante effectively assume the mental account structure [I+A, F] instead of [I,A,F] as in the BLC, in their framework the MPCs from income and liquid assets will be equal to each other. However, as in the case of the BLC, the MPC from income will typically exceed the MPC from illiquid assets. In this regard, Kaplan and Violante (2022) state: "In our model, the average quarterly MPC out of an illiquid windfall of \$500 is 1.4%, and the MPC out of an illiquid windfall of \$5,000 is 2.4%. As in the data, the average illiquid MPC is roughly an order of magnitude smaller than the liquid MPC." This finding is generally consistent with those reported by Levin (1998), and Kaplan and Violante also report that these estimates are consistent with the general literature.[104]

Figure 5 can be applied to the Kaplan-Violante temptation model, where Z_R is given by (60). Figure 6 also applies, but with two accounts, not three, and the discontinuity in Z_R corresponding to an exogenous fixed cost rather than an endogenous entry fee. As a result, in the Kaplan-Violante model, the left and right limits of marginal utility are

[103] Similar remarks apply to the model in Attanasio et al. The budget constraints in their model do not allow for borrowing against future income. The penalty function is not account-specific. Instead the counterpart to the temptation variable is the choice of a consumption-housing pair that would be achieved by allocating all wealth held at a given date on consumption and housing at that date.

[104] They state: "To put our analysis in the context of recent empirical contributions, we note that Carroll et al. (2011) estimate an average quarterly MPC out of housing wealth of around 2%. Mian et al. (2013) estimate it at 1.5%, and they uncover a great deal of heterogeneity with respect to income and leverage. Di Maggio et al. (2020) analyze unrealized capital gains in stock market wealth and conclude that the literature, including their own research, finds values between 0 and 2.5% for the average quarterly MPC."

equal at the mental account borders, whereas in the BLC the right hand limit is higher than the left. See Theorem 7 above.[105]

Kaplan and Violante's two-asset temptation model rationalizes high observed MPCs in the data with empirical income and wealth distributions. As was mentioned above, the introduction of a temptation parameter(φ) enables this to be accomplished with a rate of return premium on the illiquid asset that is well below 8%.[106] In Kaplan-Violante, φ and κ are independent. However, as implied by Theorem 7, in the BLC, the entry fees depend on φ_1, the analogue of φ.

13.11 Miserliness in Retirement

Consider Levin's finding about the relationship between bequests and illiquid assets. Thaler and Shefrin (1981) suggest that some people who develop strong saving habits during the wealth accumulation phases of their lives will find it difficult to break the saving habit during their later years. There is compelling empirical evidence of high saving rates among the elderly [De Nardi, French, and Jones (2016), Lockwood (2018), Fella, Holm, and Pugh (2024)].

Fella et al. state: "The average retired household dissaves at a much lower rate than predicted by the standard life-cycle model." They note that the literature offers two potential explanations for this phenomenon. The first pertains to hedging the risks associated with expenses for medical care and long-term healthcare. The second pertains to the bequest motive. Their study uses Norwegian data, which is important because in Norway, medical and nursing home expenses are effectively fully insured. This feature allows them to focus on the second explanation, bequest behavior.

Fella et al. report that at age 85, what they describe as a "bequest motive" comprises approximately 75% of aggregate wealth. To be sure, some of this wealth might be intended to fulfill a true need to leave a bequest. However, in the context of the BLC, it might also

[105]Another difference pertains to prudence, defined as the third derivative of the utility function being negative. In the BLC, the third derivative of Z_R is negative in a neighborhood of $m_{i,t}$ whereas in Kaplan-Violante, it is positive throughout.

[106]As discussed above, Fagereng et al. also introduce behavioral issues to explain MPCs from income.

reflect staus quo bias relative to an wealth accumulation rule. In this regard, Fella et al. find that their "estimate of the utility of residual wealth is very similar for households with offspring as those without." This finding, they interpret "as strong, prima facie, evidence that the utility of residual wealth is driven by forces beyond an actual bequest motive." The BLC offers rule-status quo bias as constituting the main such force; and it is strong enough to overcome present bias and result in underconsumption during retirement, not overconsumption.[107]

The findings about saving behavior during retirement indicate that we live in a Goldilocks world. For some the porridge is too hot, and for others it is too cold. Many people overconsume/undersave, but not all. Others underconsume/oversave. Present bias might be a factor for members of both groups, but not the only factor. Human nature is complex, indeed more complex than optimization models alone can adequately explain.[108]

14 Risk and Dynamic Inconsistency

There is a close formal analogy between intertemporal choice and decision making involving risky prospects. Notably, a profile of state–contingent outcomes is similar to a profile of time–contingent outcomes, and discount weighting is similar to probability weighting. The following extension of (1) illustrates a framework which incorporates both intertemporal and risk elements.

$$U(c) = \sum_{t,i} \delta^{t-1} p_i u(c_{t,i}), \tag{61}$$

where $c_{t,i}$ denotes consumption at date t conditional on the occurence of event i, and p_i denotes the probability attached to event i. Observe that in (61) utility $u(c_{t,i})$ is

[107]Neurologically, this is an issue about inhibitory behavior. Chambers et al. (2009) state: "Neural mechanisms of cognitive control enable us to initiate, coordinate and update behaviour. Central to successful control is the ability to suppress actions that are no longer relevant or required." Underconsumption during retirement reflects the inability to update behavior.

[108]Neurologically, rule enforcement has a physical neural network structure, and in some cases neural circuitry can be rewired without physical intervention. Rewiring is what occurs during some therapy sessions. Sometimes, permission from a trusted authority figure can be sufficient to rewire a neurological guardrail.

discounted by $\delta^{t-1}p_i$. An individual with time invariant preferences represented by (61) will exhibit dynamically consistent behavior.

Despite the formal similarities between the intertemporal and uncertainty elements of choice, there are important and subtle differences. Consequently, I begin with some examples to fix ideas. Consider yourself to be in a situation where you are presented with three pairs of risky prospects, and asked to express your preference in respect to each pair. For each pair, imagine that you have an opportunity to choose to play exactly one member from each pair on a one–time only basis. The first two pairs $(1A, 1B)$ and $(2A, 2B)$ are as follows:

$1A$. A 90% chance of winning $2000, and a 10% chance of zero.

$1B$. A 45% chance of winning $4000, and a 55% chance of zero.

$2A$. A 0.2% chance of winning $2000, and a 99.8% chance of zero.

$2B$. A 0.1% chance of winning $4000, and a 99.9% chance of zero.

The third choice pair involves compound prospects consisting of two stages. In the first stage, the probability of winning a prize is 2/900 (that is, 2.22%). Now imagine that you have won at the first stage, and must choose among two prospects to be played at the second stage:

$3A$. A ticket for 1A (described above).

$3B$. A ticket for 1B (described above).

The three binary choices just presented are taken from Kahneman and Tversky (1979). In an experimental setting, Kahneman and Tversky elicited individual preferences for all three. They report that there is a systematic tendency among their subjects to choose in a particular way. Specifically, $1A$ tends to be favored over $1B$, $2B$ tends to be favored over $2A$, and $3A$ tends to be favored over $3B$.

Imagine someone who expresses the preceding preference pattern. McClennen (1990) states that this pattern reflects dynamic inconsistency. Notably, the compound prospect associated with $3A$ is equivalent to $2A$, and the compound prospect associated with $3B$ is equivalent to $2B$. Consequently, at the first stage the individual acts as if he prefers $2B$ over $2A$, but reverses this preference at the second stage. To complete the analogy with the certainty case, consider the following modification to the third choice pair.

Imagine that at the first stage of the compound prospect you face a choice. You may play the compound prospect described above, or you may play the following simple propsect:

$2A+$. A 0.2% chance of winning \$2001, and a 99.8% chance of zero.

Observe that $2A+$ stochastically dominates $2A$ in that it offers a prize which is \$1 dollar more. We view \$1 as an ϵ which is sufficiently small so that $2B$ is preferred to $2A+$ as well as to $2A$.

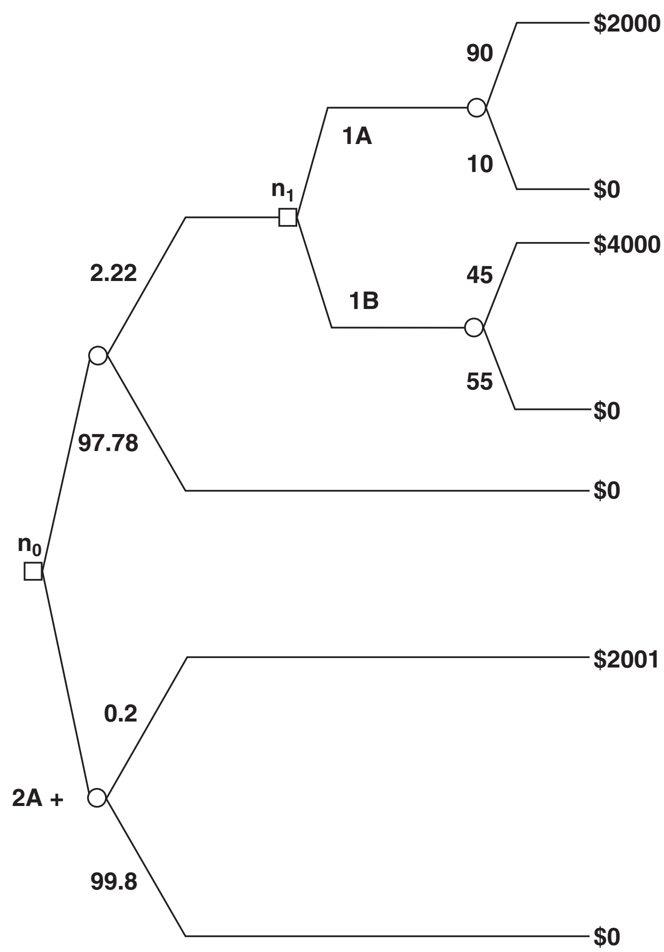

Figure 8 : Uncertainty Tree

Figure 8 portrays the choice situation. Consider the analogue to naïve choice. At the first stage the agent favors $2B$, which emanates from the top branch at stage 1, over $2A+$, which emanates from the bottom branch at stage 1. Naïve choice implies that during the first stage, the agent believes that if he selects the top branch and subsequently wins at the first stage, then he will choose to play $1B$ over $1A$ at the second stage. Of course, if the agent manages to end up facing this choice at the second stage, he will choose $1A$, not $1B$. Consequently, naïve choice leads the individual to actually choose to play the compound prospect $2A$ over $2A+$.

Recall that $2A+$ actually stochastically dominates $2A$. Consequently, sophisticated choice would involve the selection of $2A+$ at the first stage, since the sophisticated agent makes his choice by taking proper account of his future self's choice behavior.

Keep in mind that the intertemporal certainty framework has been extended to accommodate uncertainty. This extension essentially involves the recasting of time from the single variable t to the ordered pair (t, s), where s is a state of nature. Hence the root of dynamic inconsistency in the intertemporal setting will carry forward to the uncertainty setting. Moreover, the additively separable intertemporal utility function has expected utility as its uncertainty analogue: both involve weighted sums of sub–utilities.

In the certainty model with intertemporally additive utility, a nonexponential discount function lies at the root of the divergence between naïve choice and precommitted choice. Consider the same question in connection with the uncertainty framework.[109] In other words, if we take the discount function to be exponential, then what is the root cause of dynamic inconsistency in choices over lotteries? Hammond has forcefully argued that in a consequentialist setting, it is the independence axiom associated with expected utility which lies at the root of the issue.

In order to illustrate the key aspect underlying the connection between expected utility and dynamic inconsistency, consider the example discussed above. Let the agent's

[109] Von Auer's treatment of the subtleties associated with indifference extends to the uncertainty case as well.

preferences satisfy the axioms of expected utility. Then the agent possesses a utility function u on the set of consequences. For simplicity, we normalize u and set $u(0) = 0$ and $u(4000) = 1$. Then preferences over prospects are represented by expected utility. Since in the example $1A$ is preferred to $1B$, $1A$ must have a higher expected utility than $1B$. This holds if and only if $u(2000) > 0.5$. However, this last inequality implies that $2A$ must be preferred to $2B$, in contrast to the preference pattern assumed in the example.

The independence axiom underlies the representation of preferences over lotteries by means of an expected utility functional. It should come as no surprise that the independence axiom is central to the issue of dynamic consistency. The independence axiom is a statement about compound prospects. Let A, B, and C be lotteries, with A being at least as good as B. The independence axiom pertains to two compound prospects: AC, which is A with probability q and C probability $1 - q$, and BC which is B with probability q and C probability $1 - q$. The independence axiom asserts that since A is at least as good as B, then AC must be at least as good as BC. As McClennen's example makes clear, the prospects offered at any node in the uncertainty tree are comprised of compound prospects based upon prospects attached to future nodes. Therefore the independence axiom effectively imposes dynamic consistency across the tree.

A common feature of the two choices {1A, 1B} and {2A, 2B} is that the probability ratios of a positive gain are the same: $0.9/0.45 = 0.002/0.001$. This feature is known as the *common ratio effect* and in a consequentialist expected utility setting it induces the same A vs. B choice patterns in options 1 and 2 above, as well as inducing dynamic consistency.[110] If $2A$ is preferred to $2B$, then there is no conflict between the agent's preference ordering at the two stages. Prospect $2A+$ would be selected, regardless of whether the choice mechanism were naïve, sophisticated, or precommitted.

Prospects 2A and 2B are lotteries, in the sense of featuring small probabilities of receiving large prizes. The experimental evidence described above not only suggests that

[110] In addition to the common ratio effect, Kahneman and Tversky (1979) introduce a property known as the *common consequence effect*. This effect involves probability mixtures of two risky choices, such as A and B, with a third risk C using the same mixing probabilities. If A is preferred to B, then under expected utility theory, so will be their associated mixtures. In practice, people often violate the common consequence effect just as they violate the common ratio effect.

people find lotteries attractive, but that they also typically display dynamic inconsistency in choices involving lotteries. There is also evidence linking such behavior to temptation, self-control, and as discussed below, mental accounting. Hamilton, Liu, Miranda-Pinto, and Sainsbury (2024) report that during the Covid pandemic, Australia permitted the withdrawal of $20,000 (Australian) from mandatory retirement accounts. Generally, these accounts cannot be accessed before retirement. They found that about 16% of those who could do so, made withdrawals, with the total amount comprising about 2% of GDP. Hamiton et al. argue that this behavior reflects present bias, not just impatience. In addition, they estimate MPCs associated with the withdrawals to be in the range 0.43–0.48 within eight weeks, with gambling being the largest spending category.[111] This finding is the mirror image of the windfall effect discussed above, which focused on MPCs associated with lottery prizes. Here, the MPCs pertain to invading the future income account to spend on lottery-like prospects.

15 Prospect Theory

Experimental evidence suggests that dynamic inconsistency is prevalent. If so, then it is important to understand why this is the case. Economists and psychologists have a common interest in this question. Recall the passages quoted in Section 4 from Strotz (1956) which dealt with the general shape of the discount function. Subsequent to the appearance of his article, psychologists proposed that most individuals exhibit dynamic inconsistency because they discount the future using hyperbolic functions. In this section, I discuss the contributions of psychologists Kahneman and Tversky to the analogous question of why individuals exhibit dynamic inconsistency when facing uncertainty. In addition, I describe how their work has been used to provide additional insight into dynamically inconsistent intertemporal choice.

Kahneman and Tversky (1979) have developed a framework to describe the character

[111]Hamilton et al. state: "Poor financial health and gambling strongly predict withdrawal and spending. We [demonstrate] ... that while impatience under liquidity constraints reconciles the observed withdrawal behavior, only present bias reconciles the magnitude and frequency of the observed spending response."

of the heuristics which lead many people to exhibit preferences such as those described in the examples discussed in the previous section. They call their framework *prospect theory*. Like standard expected utility theory, prospect theory is a one-system framework.

Recall that in the examples described in the previous section, dynamic inconsistency leads to the selection of a stochastically dominated lottery under naïve choice. Indeed it is common for prospect theoretic individuals to choose stochastically dominated lotteries. Prospect theory seeks to describe why this can occur, but hardly suggests that such behavior is normatively desirable. Nevertheless, Kahneman and Tversky's findings suggest that dynamic inconsistency among individuals is prevalent.[112]

15.1 Structure of Prospect Theory

Prospect theory involves two stages. In the first stage, a decision problem is decomposed into a series of subproblems. Kahneman and Tversky refer to this stage as *editing*. The editing phase pertains to the establishment of *mental accounts*. These record the outcomes associated with the various subproblems. In the second stage, the *evaluation stage*, the alternatives in each mental account are evaluated by means of a weighted sum of utilities. The actual decision is assumed to maximize the value of the weighted sum. I note that although the weights are functions of the probability distribution, they are not identical to probabilities. In particular, Kahneman and Tversky suggest that individuals tend to overweight low probability events relative to high probability events. The nonuniform weighting function lies at the heart of the breakdown of expected utility across the first and second pairs of lotteries in the preceding example.

The formal structure of prospect theory involves the specification of the mental accounting structure, reference points, utility function, and probability weighting function. Define a general prospect (x, p) in terms of the two vectors x and p. Here p_i denotes the

[112] This feature leads to question of whether all those choosing stochastically dominated lotteries can be subjected to a money pump which leads them to lose all their wealth in short order. An affirmative answer does not follow automatically, and it may not follow at all. There are many complex issues which can interfere with the operation of money pumps, such as: transaction costs, customizing the money pump to individual preference patterns, whether individuals learn when they recognize a sequence of losses, and stop participating without adjusting their preferences, institutional structure and securities market regulation.

probability attached to the occurrence of event i, and x_i represents the outcome attached to this event. Consider a decomposition of x into J subvectors x^j, where:

$$x = \sum_{j=1}^{J} x^j, \tag{62}$$

Consider the prospect (x^j, p). A decision maker who faces concurrent prospects $(x^j, p), j = 1, ...J$, effectively faces the overall prospect (x, p). A mental account consists of a prospect together with a reference point ρ^j, and hence can be expressed as a triple (x^j, p, ρ^j). If event i occurs then this mental account registers a net gain of $x_i^j - \rho^j$. The prospect theoretic counterpart to expected utility is a function:

$$\sum_i \psi_i(p) v(x_i^j - \rho^j), \tag{63}$$

where v is a utility function over the domain of net gains. Kahneman and Tversky call v a *value function*, and call $\psi_i(p)$ a probability weighting function.

By defining v as a function of $x_i^j - \rho^j$, utility is taken to be defined over the domain of gains and losses, rather than final outcome. Notably, Kahneman and Tversky postulate that utility is concave in gains but convex in losses. That is, $v(x_i^j - \rho^j)$ is concave when $x_i^j - \rho^j$ is positive, and convex when $x_i^j - \rho^j$ is negative. This property gives rise to an S-shaped utility function, which Kahneman and Tversky call a *value function*.

Kahneman and Tversky postulate that the value function is nondifferentiable at the origin. In particular, it is steeper in losses than in gains. They refer to this property as *loss aversion*. The convexity of the function in losses can lead to risk seeking behavior for prospects involving losses, especially when one choice involves a sure loss. Kahneman calls such behavior *aversion to accepting a sure loss*.

To see the implications of the Kahneman-Tversky assumption, consider the following example. Consider the following two pairwise choices. First, choose between:

$4A$. A 100% chance of winning \$1800; and

$4B$. A 45% chance of winning \$4000, and a 55% chance of zero.

Second, choose between:

$5A$. A 100% chance of losing \$1800; and

$5B$. A 45% chance of losing \$4000, and a 55% chance of losing zero.

Because $4A$ and $4B$ have the same expected payoff, none of the probabilities lie at the extreme ends of the spectrum, and the lotteries are defined in the domain of gains, a prospect theoretic individual would select $4A$. The same is true for $5A$ and $5B$, except that here the two lotteries are defined in terms of losses. Consequently, a prospect theoretic individual would select $5B$ over $5A$, whereas a risk averse individual would select $5A$.

Expected utility requires that $\psi_i(p) = p_i$. In the Kahneman–Tversky (1979) treatment of prospect theory, the function $\psi_i(p)$ is a function $\pi(p_i)$. In addition, ψ satisfies $\psi_i(p) > p_i$ for positive p_i near zero, and $\psi_i(p) < p_i$ for p_i near unity. This specification has some attendant difficulties with violation of stochastic dominance. The issue motivated Quiggin (1982) to suggest a *rank dependent* structure for ψ based upon the cumulative density function.

Consider a specific mental account (x^j, p, ρ^j), and order the events i so that x_i^j is lowest in event 1 and monotonically increases over events $2, 3, ...$ Rank dependent utility involves the cumulative probability distribution C, where $C(y)$ is the probability (given p) that the random payoff x^j to mental account j is y or less. Let $\Psi(C(x))$ be a transformed cumulative function, with $[\psi_i]$ being the associated density function. If ψ_i is obtained in this way, then ψ_k and ψ_l need not be equal, even if $p_k = p_l$. In their (1992) version of prospect theory, called cumulative prospect theory, Kahneman and Tversky adopt the rank dependent formulation with a weighting function that has an inverse-S shape.[113]

Probabilistic-based dynamic inconsistency is built into the rank dependent structure. This is because the weight attached to the probability of a specific event depends on the rank of that event in any gamble within which it appears. As a result, conditional probabilities do not remain invariant over time, when the rank of specific events changes with realized outcomes.

[113]In cumulative prospect theory, there are two transformation functions, one for gains and the other for losses.

Barberis (2012) analyzes the impact of the cumulative prospect theory weighting function on casino gambling decisions. He states: "[T]he probability weighting embedded in prospect theory leads to a plausible time inconsistency: at the moment he enters a casino, the agent plans to follow one particular gambling strategy; but after he starts playing, he wants to switch to a different strategy. The model therefore predicts [that] how a gambler behaves depends on whether he is aware of the time inconsistency; and, if he is aware of it, on whether he can commit in advance to his initial plan of action." Relatedly, Heimer, Iliewa, Imas, and Weber (2023) use a brokerage dataset to study decisions by stock market investors. In a series of experiments, they report the following: "We compare people's initial risk-taking plans to their subsequent decisions. Across settings, people accept risk as part of a 'loss-exit" strategy—planning to continue taking risk after gains and stopping after losses. Actual behavior deviates from initial strategies by cutting gains early and chasing losses." The latter behavior is the *disposition effect* [Shefrin and Statman (1985)].

In general, there may be several ways to decompose or frame a particular decision problem. Consequently, the reference point which is used as a basis for computing gains and losses may fail to be unique. This feature can give rise to *framing* effects, in which the decision taken is sensitive to the way that the decision problem is framed. For example, an individual who receives a coincident payment of $2000 when facing pair 5, may factor this payment into his evaluation of the two lotteries. If he does, then he frames his choice as a gamble involving gains, and this would lead him to select $5A$. If he views the coincident payment of $2000 to be independent of the gamble, and therefore ignores it when making his choice, then he frames his choice in the domain of losses and chooses $5B$. Therefore, prospect theoretic individuals are vulnerable to *static inconsistency* as well as dynamic inconsistency. Notice that the core issue is not so much the violation of expected utility, as the absence of a consequentialist framework. The domain over which utility (or value v) is defined concerns lottery outcomes rather than final consequences.

15.2 Intertemporal Issues

The terminology of framing and reference points explicitly links prospect theory to the notion of moving frames of reference. In effect, the reference point serves as a degenerate (certainty) reference lottery, against which the alternatives are compared.[114]

Loewenstein and Prelec (1992) use prospect theory to develop an intertemporal choice model without uncertainty. Their model is designed to accomodate a variety of empirical regularities which have been observed about the way that people discount future outcomes. Notably, such discounting is a function of the way that outcomes are transformed into utility as well as the extent to which future utility is discounted to the present. Consider the *equivalent variation* E associated with a reward of X at time t. That is, E satisfies:

$$u(E) = \lambda(t)u(X).$$

Notice that the ratio E/X can be considered the discount factor applied to X being received at t.[115] If u were linear, then E/X would just be $\lambda(t)$. However if u is strictly concave, then it is easily demonstated that $E/X < \lambda(t)$.

In addition to the common difference effect described in sections 4 and 5, Loewenstein and Prelec discuss three other effects which are inconsistent with the discounted utility model. They are:

1. The gain–loss asymmetry where gains are discounted at higher rates than losses; and

2. The absolute magnitude effect where small amounts are discounted at a higher rates than large amounts;[116]

3. The delay–speedup asymmetry in which the valuation placed upon a change in timing of receipt depends upon whether it is framed as a delay or as a speedup.

[114] The reference lottery need not be degenerate.

[115] We could also have used the compensating variation in place of the equivalent variation.

[116] This property is actually consistent with the discounted utility model, but becomes anomalous in combination with the gain–loss asymmetry.

The following example illustrates some of the features with which Loewenstein and Prelec are concerned. Two groups of subjects are presented with different descriptions of the *same* decision problem involving the installment purchase of a television set. The first group is told that they must choose between:

6A. A payment this week of \$160 and a second payment of \$110 in six months; and

6B. A payment this week of \$115 and a second payment of \$160 in six months.

The second group is told that the plan normally calls for two equal payments of \$200, one this week and one in six months. However, because of a promotional sale, there are rebates which will be applied to both payments, and timed to coincide with the payments themselves. The second group can choose between two rebate schedules, namely:

6C. A rebate of \$40 on the initial payment and a rebate of \$90 on the later payment.

6D. A rebate of \$85 on the initial payment and a rebate of \$40 on the later payment.

Notice that alternatives *6A* and *6C* are equivalent to each other as are *6B* and *6D*. Nevertheless in experimental settings, there is a systematic preference for *6A* over *6B*, but *6D* over *6C*. Loewenstein and Prelec explain this finding in terms of their model as follows. The participants in the first group frame the decision problem as a choice between two loss schedules, where as the participants in the second group frame the problem in terms of gains. Because losses are discounted less than gains, *6A* is selected on the basis of a lower total payment. However, the second frame features *both* gains *and* smaller magnitudes. Therefore the later rebate is discounted highly, and this leads to the preference for *6D* over *6C*.

There is a well documented phenomenon known as *preference reversal* which was discovered in connection with decision making under uncertainty but also applies to intertemporal choice. Tversky and Thaler (1990) survey the literature on this phenomenon. Imagine that an individual is asked for two pieces of information about a pair of lotteries, such as $1A$ and $1B$ in section 14. The first piece of information is how he ranks the pair.

The second piece of information is the value of each certainty equivalent $CE(1A)$ and $CE(1B)$.[117] If an individual claims to favor $1A$ over $1B$, but attaches a higher certainty equivalent to $1B$ than to $1A$, then he is said to commit preference reversal. This leads to a preference cycle since $CE(1A)$ is indifferent to $1A$, $1A$ is preferred to $1B$, $1B$ is indifferent to $CE(1B)$, and $CE(1B)$ is preferred to $CE(1A)$.

Tversky and Thaler suggest three possible explanations of preference reversal: 1) intransitive preferences; 2) violation of the independence axiom; and 3) cognitive mispricing. Although the evidence suggests that there is more than one single explanation for what causes the phenomenon, Tversky and Thaler indicate that mispricing appears to the most prevalent. Specifically, many individuals are prone to attach too high a certainty equivalent to a lottery such as $1B$. They do this because in their computation of $CE(1B)$, they overweight the importance of the \$4000 payoff relative to the probability of its occurence. Tversky and Thaler indicate that such errors are not confined to uncertainty. Similar reversals occur when delayed rewards are being compared, and their equivalent variations computed. The phenomenon of preference reversal raises the question of whether individuals can even be said to have intact preference orderings. Tversky and Thaler suggest that what they have instead is a set of basic value principles, which they use to construct preferences each time that a choice needs to be made. If so, then preference reversal may reflect systematic cognitive errors in reconstructing the relevant potion of the preference ordering from the underlying value primitives.

15.3 Embedding Prospect Theory Within a Two-System Framework: A Direct Approach

Kahneman (2011) states that in developing prospect theory, he and Tversky "were not working with the two-systems model of the mind" (p. 211). However, he later came to recognize that three of prospect theory's cognitive features are "common to many automatic processes of perception, judgment, and emotion [which] should be seen as

[117]Let the individual own the lottery ticket, so that the certainty equivalent refers to the lowest amount they would accept in order to voluntarily sell the ticket.

operating characteristics of System 1." The three features are

- evaluation of consequences relative to a reference point;
- diminishing sensitivity to changes in wealth; and
- loss aversion.

The remainder of this subsection, and the next discuss the formal embedding of prospect theory into a two-system framework. This subsection describes a direct approach, and the subsequent subsection describes an indirect approach. To highlight some of the key features in this discussion, consider the original Allais paradox and the way that prospect theory explains the paradoxical behavior.

In its original form, the Allais paradox asks respondents to make two choices. The first choice is between

7A. $1 million with probability 100%
and

7B. $0 with probability 1%
$1,000,000 with probability 89%
$5,000,000 with probability 10%

The second choice is between

7C. $0 with probability 89%
$1,000,000 with probability 11%
and

7D. $0 with probability 90%
$5,000,000 with probability 10%

Most people choose 7A over 7B and 7D over 7C. On the face of it, the choice of 7D over 7C seems plausible. The expected value of 7C is \$110,000 and of 7D is \$500,000, a difference of \$390,000. The \$390,000 premium serves as compensation for facing an additional 1% probability of incurring a payoff of \$0. In contrast, the expected value of 7A is \$1,000,000 and of 7B is \$1,390,000, which also involves a premium of \$390,000, as compensation for facing an an additional 1% probability of incurring a payoff of \$0. Apparently, the \$390,000 premium offered by 7B is insufficient to overcome giving up the certainty of a sure \$1,000,000 in 7A, whereas it is sufficient when there is no sure loss involved, as in 7C and 7D.

Standard expected utility theory implies that a preference for 7A over 7B requires a preference for 7C over 7D. To see why, let $u = [u_0, u_1, u_5]$ be a vector of utiltiies and $p_j = [p_{j,0}, p_{j,1}, p_{j,5}]$ be a vector of probabilities for $j = A, B, C, D$. Notably p_C and p_D have the forms $p_C = p_A + \Delta p$ and $p_D = p_B + \Delta p$. Therefore

$$p_A - p_B = p_C - p_D$$

which implies that for the dot product $u \cdot p$,

$$u \cdot (p_A - p_B) = u \cdot (p_C - p_D)$$

That is, the expected utility differences for {A,B} and {C, D} are equal; and it is expected utility differences which dictate which is the preferred risk.

Prospect theory provides an explanation for the Allais paradox, through the weighting function. For sake of exposition, let the units be in millions. The prospect theory value of 7A, $v(1)$ exceeds the weighted average $\psi(0.1)v(0) + \psi(0.89)v(1) + \psi(0.1)v(5)$ associated with choice 7B. Likewise, the choice of 7D over 7C implies that $\psi(0.1)v(5) > \psi(0.11)v(1)$. Given that $v(0) = 0$, these two inequalities imply that:

$$1 - \psi(0.89) > \psi(0.11)$$

In turn, the above inequality implies that

$$1 > \psi(0.11) + \psi(0.89)$$

so that probability weights do not sum to unity, as do probabilities. Kahneman and Tversky (1979) call this subadditivity-condition *subcertainty*. Subcertainty is not necessary for the result. Rather, the key issue is that either ψ sufficiently overweights probabilities that are low, or underweights probabilities that are high, or both. In this regard, Tversky and Kahneman (1992) develop an extension of prospect theory called *cumulative prospect theory* which uses rank dependence in order to generate decision weights. Rank dependence does not produce subcertainty, but can produce mis-weighting of probabilities.[118]

Fudenberg and Levine (2011) provide an explanation for the Allais paradox which focuses, not on the weighting function, but on the value function.[119] Intuitively, temptation associated with the sure outcome in 7A leads to this choice carrying much more weight than 7B. In contrast, there is no sure outcome in the choice between 7C and 7D. In this respect, the expected prospect theory value computation is consequentialist, in that the argument of value is the payoff and only the payoff. However, in the two-system planner-doer framework, value (or utility) has other arguments besides consumption, as discussed below.

Notably, the two-system explanation provided by Fudenberg and Levine focuses on the the penalty function being strictly convex. As discussed above, reducing the probability attached to a reward from 100% reduces the temptation variable m, and with it the cost of willpower. As a result, the impact of the penalty function is to reduce risk aversion. The argument in reverse is that the 100% probability associated with a certain outcome

[118]Rank dependence generates decision weights by applying an inverted S-shaped weighting function to (cumulative) probability distribution functions. The inverse S-shape implies overweighting of small probabilities attached to the lowest ranked outcomes and also to the highest ranked outcomes. As a result, the decision weights associated with a specific probability can vary from one prospect to another. In the case of the Allais paradox, the probability attached to outcome $0 in risk 7B can be overweighted, thereby favoring the selection of 7A over 7B. In the case of risks 7C and 7D, the weighting function can be relatively insensitive to intermediate probabilities such as 0.10 and 0.11, which can cause 7D to be favored over 7C.

[119]See also Noor and Takeoka (2010).

increases risk aversion. Observe that choice 7A features certainty, and high risk aversion will induce a preference for 7A over 7B. In contrast, there is no certain outcome in the choices 7C and 7D.[120]

In their conclusion, Fudenberg and Levine state: "One class of alternative models that can potentially be used to explain a wide range of phenomena are probability weighting models, including ... prospect theory and models of ambiguity aversion." (p. 55). The remainder of this section discusses a two-system planner-doer approach to prospect theory. Fudenberg and Levine also state: "These models are typically static... However, there are recent efforts to extend these models to examine dynamic phenomena." The analysis below pertains to dynamic models.

As was mentioned above, Shefrin and Statman (1985) discuss prospect theory as an element of the disposition effect. In this respect, the disposition effect explicitly embodies additional psychological elements such as mental accounting and self-control. This suggests embedding prospect theory into the two-system planner-doer framework.[121]

The planner-doer framework is inherently a reference point based theory, with temptation variables serving as reference points, and mental accounts serving as the structure for prospect theory's editing function. The doer utility function can incorporate impulses relating to gains and losses, not just present bias. Log utility is the fittest utility function, and can form the basis for a notion of the System 2 planner as an expected log-utility maximizer, who needs to take account of System 1 impulses. The interaction between the two systems reflects a conflict between log-utility risk aversion and System 1 reference point-based loss aversion.

Shefrin and Statman (1985) describe several self-control options which disposition effect-prone investors can use to counter their tendencies to hold onto losers too long. Chief among these is the use of stop loss orders, an external device. They also discuss internal rules such as developing habits to sell a stock whenever its price drops by some prespecified percentage below the original purchase price. These are self-nudge techniques.

[120]Fudenberg and Levine (2006) also discuss how the two-system approach can explain the paradox presented by Rabin (2000), which features high risk aversion for small stake gambles.

[121]In his Nobel lecture, Kahneman discussed prospect theory broadly, and later expanded that discussion to two-system and two-selves in his 2011 book *Thinking, Fast and Slow.*

What makes them important from a theoretical perspective is that they make clear that investors are working to exercise self-control and overcome tendencies such as aversion to accepting a sure loss. That is, their planners are seeking rules to help overcome the impulses generated by their doers.

There are at least two ways to embed prospect theory into a two-system planner-doer framework. The first way is to treat a prospect as a commodity in the multicommodity planner-doer system, with risk. In this regard, two different frames of the same fundamental risk correspond to two different commodities. Recall that in the multicommodity planner-doer framework, the planner needs to use willpower (θ) to impact the doers' choice of consumption. With this embedding of prospect theory, doer utility Z is the decision-weighted prospect theory value attached to the prospect in question. The doers simply correspond to agents with prospect theory preferences.

Prospect-commodities might be index funds, government bonds, sports bets, meme stocks, and lottery tickets. The convex portion of the prospect theory value function might lead a doer to favor imprudent risks when facing the prospect of a sure loss. The loss aversion property might induce a doer to favor taking little or no risk, when the planner judges otherwise. The self-control issues with prospects are akin to self-control issues with physical commodities. Taking insufficient risk is akin to choosing to engage in insufficient exercise. Taking excessive risk is akin to consuming too much alcohol. Recall that the multicommodity version of the BLC is structured to address self-control issues in which consumption of some commodities is too low, and for others too high.

The risks attached to prospect-commodities are fully specified, allowing the planner to compute the values of the arguments in its log-utility function. If the planner wishes to achieve an outcome that differs from one that is doer-driven, as in the case of the disposition effect, it will need to use willpower or rules or a combination of both. If willpower is excessively costly, or the planner does not have access to feasible rules that are effective, then the individual will behave in accordance with prospect theory.

The planner-doer approach to prospect theory just described has some of the features in Kösegi and Rabin (2006). They develop a reference point-based multicommodity theory in which utility is the sum of a neoclassical component and a gain-loss component. This formulation can be interpreted as utility being a mix of a planner component (neoclassical utility) and a doer component (prospect theory value function, probability weighting function, and editing structure), but with no self-control features.

15.4 Embedding Prospect Theory Within a Two-System Framework: An Indirect Approach

The second way to embed prospect theory features into a two-system planner-doer framework is to treat individual commodities as state contingent payoffs, not prospects. Under this approach, doer utility is interpreted as the value of a realized outcome, meaning experienced utility. Unlike the first way, which explicitly incorporates the structure of prospect theory into the two-system framework, the second way focuses on incorporating only the key features from the experimental evidence described above, but not the S-shaped value function and associated weighting function.

The multicommodity planner-doer framework involves two types of doers, those for whom willpower is required to reduce the consumption of a specific commodity, and those for whom willpower is required to increase consumption of a specific commodity. The extended planner-doer model being discussed here incorporates a third type of doer. This third type is depicted in panel B of Figure 9, which also depicts the first two types of doer utility functions and the logarithmic planner utility function. In Figure 9, the three types of doer utility functions are termed Z_{high}, Z_{mid}, and Z_{low} respectively. These serve as a *basis* for doer utility functions.

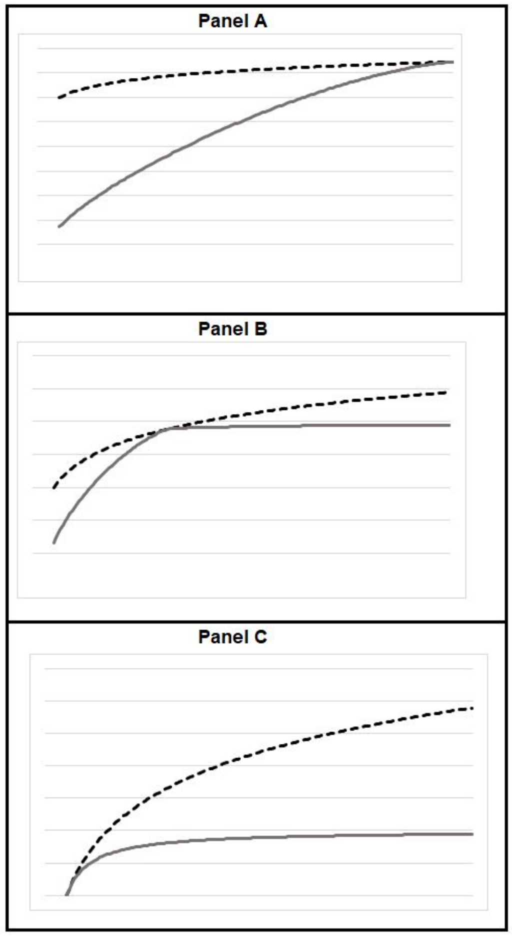

Figure 9: Doer utility function basis for the embedding of prospect theory within a two-system planner-doer framework. The top curve in all three panels is log-utility (dotted black curve). The gray solid curve in each of the three panels is doer utility. Panel A depicts Z_{high}, panel B depicts Z_{mid}, and panel C depicts Z_{low}.

The doer utility function basis provides the structure for alternative mental states. In this regard, activation of the region of the brain known as the nucleus accumbens (NACC) takes place during the anticipation of diverse rewards associated with potential risks [Knutson, Wimmer, Kuhnen, and Winkielman (2008)]. The nature of this anticipation depends on the nature of the risks being faced, and can give rise to alternative mental states.[122]

To model alternative mental states, consider that the three doer utility functions discussed above differ in respect to reference point. The reference point for Z_{high} is

[122]Dopamine and testosterone also play important roles in human risk taking [Coates (2012), Mikella, Crawford, Kuhnen, Samanez-Larkinand, and Seaman, (2023)].

the maximum value of the mental account housing a prospect. The reference point for Z_{low} is the minimum value of any commodity in the mental account housing a prospect. The reference point for Z_{mid} is the value of any commodity in the mental account when the resources allocated to that account are evenly distributed across commodities in the account. For example, in Figure 9, there are three commodities, and so the location of the reference point for Z_{mid} is one third of the mental account balance corresponding to the reference point for Z_{high}. Because the commodities in this framework are state-contingent payoffs, the equal distribution allocation is a *certain* outcome.

Consider a "standard case" in which all doers have $Z_{mid}-$utility functions, and the certain (risk-free) outcome is the status quo position associated with $\theta = 0$. In consequence, doer utility Z achieves a local maximum at the risk-free value. This is akin to Z_{low} which achieves a local maximum at $c = \epsilon = 1$. If the planner prefers that the individual take some risk in the mental account, it will have to use willpower (or a rule) to induce the account doers to change their choice of c. If willpower is costly, which is the situation depicted in panel B of Figure 9, then the planner will find it inconvenient to shift the account allocation far from the risk-free point.

The consequence of high willpower cost results in the same feature as loss aversion in prospect theory. In particular, if willpower costs are higher to the left of the reference point than to the right, then losses will be accorded more weight than gains of the same magnitude. This is what loss aversion entails; and by choosing parameters for Z_{mid} for which the left derivative at the reference point is 2.25 times the size of the right derivative, the magnitude of loss aversion (for small gains and losses) will conform to the findings in Tversky and Kahneman (1992).

Choice within a mental account involves a priority ordering among doers. In the single commodity BLC, the priority ordering is based on time: doers at early dates make choices before doers at later dates. Present bias typically leads the advantage conferred on the doers at early dates to result in overconsumption at early dates. Similar features apply in the mental accounting setting discussed here. In this regard, imagine that the mental account priority ordering provides an advantage to doers associated with high probability events and low ranked (unfavorable) outcomes. This can be accomplished by using a linear scoring priority rule, with a positive coefficient on state probability and a negative weight on ranking. In this situation, a top priority doer will have the first opportunity to assign the mental accounting budget to itself. Doers down the line will need to make do with the residual, just as happens in the single commodity BLC. This pecking order feature has the effect of according more weight to high probability low ranking outcomes. If the priority ordering favored outcomes with low probability and low ranking, the asymmetric weighting would be similar to the impact associated with the inverse S-shaped weighting function in cumulative prospect theory.

Mental states involve a series of regions besides the nucleus accumbens. Notably, activation of the region known as the anterior insular induces fear and anxiety [Shi, Feng, Wei, and Zhou (2020)]. In this regard, Lopes (1987) proposes a psychological model of risk featuring fear, hope, and aspiration as the key elements driving decisions about risk. Knutson et al. (2008) state that "[a]nticipation of shifting to the high-risk option versus shifting to the low-risk option [is] correlated with activation in the bilateral NAcc ... as well as deactivation of the right anterior insula." (p. 511). This means that the evaluation of high risk options involves mental states associated with decreased fear. Knutson et al. also report that "Of the model covariates, ... only losses on the previous trial significantly predicted shifts to the high-risk option." (p. 512.)

From the perspective of the two-system model, prior losses induce a change in basis for doer types, from Z_{mid} to Z_{high} for some outcome states, and from Z_{mid} to Z_{low} for other outcome states. In this case, the status quo consumption profile for the account is not risk-free, but high risk, with consumption outcomes at one boundary or the other of the mental account. If the planner prefers a less risky prospect than the status quo account prospect, it will have to employ willpower or a rule to induce the doers to different decisions. In particular, it might have to employ willpower at both ends of the spectrum, reducing high state-contingent payoffs, and increasing low state-contingent payoffs. If willpower costs are high, achieving the risk-free outcome, a sure loss, might be excessively expensive. Such a situation would conform to the feature of prospect theory in which individuals are averse to accepting a sure loss.

One way to see how the second version of prospect theory works is by analogy to the BLC with a single mental account. For simplicity, consider the case of $T = 3$, when the subjective rate of time preference and market interest rate are both zero Recall that the doer at $t = 1$ has the first shot at consuming lifetime wealth. At this date, the balance in the (single) mental account is at its peak. Therefore so is temptation, and so is the cost of exercising willpower. As a result, the planner will allocate relatively more consumption to $t = 1$ than at the other two dates. Similarly, the planner will choose consumption at date 2 to be higher than at date 3.

The doer utility function in the BLC framework is displayed in Panel A of Figure 9. Consider what would happen in the above case if the doer utility function corresponded to Panel B rather than Panel A. Now, willpower is required to move consumption away from the equal-consumption allocation. If the planner prefers not to do, the end result will be the equal-consumption allocation.

Next, suppose that the market interest rate is positive, instead of zero. Now, there is value to deferring consumption, as reflected by the market discount factor associated with any date, that being the price of consumption relative to date 1. A positive interest implies that consumption is most expensive at date 1 and least expensive at date 3. Typically, this feature will induce consumption to be monotone increasing in time. However, the kink in doer preferences will create friction, which might induce consumption to remain at the equal-consumption allocation. Only if the interest rate is sufficiently high, will the planner choose to shift away from the equal-consumption allocation towards a monotone increasing profile.

The example just described can be reinterpreted in terms of states of nature (risk) instead of time. In the risk version, there are three states (and a single period). State prices, meaning the prices of state-contingent claims, are monotone decreasing, as above, with state 1 being the most expensive. Doer utility is given by Panel B of Figure 9. Instead of prices reflecting market discount rates, prices reflect market supply of consumption. The kink(s) in the doer utility function(s) will induce a risk-free allocation, unless the differences in market prices are large enough to shift the planner away from the kink(s).

State probabilities are akin to discount rates associated with subjective time preference. Uniform probabilities are akin to a zero rate of time preference. In the BLC, a high subjective rate of time preference induces a preference for earlier consumption over later consumption, all else being the same. Similarly, there will be a preference for consumption in high probability states over consumption in low probability states.

In the BLC, there is a natural order in which doers make claims on wealth, the order being dictated by time t. In the risk interpretation, the order is not as natural. For sake of discussion, let state 2 be associated with the highest probability, and let the doer priority-order be [2,1,3]. Now, movement away from equal-consumption will be determined by the doers' sense of probabilities, state prices, and willpower costs associated with doer utility functions.

Viewing the Allais paradox through a planner-doer lens provides insight into how the two-system approach works in respect to dynamic inconsistency. By assumption, the planner is an expected utility maximizer; and standard utility maximizing behavior does not produce the Allais paradox. To see why, consider the two Allais choices, described using a model with four states of nature. States of nature are described below, in terms of their respective probabilities and state prices (in \$millions). In the description, the symbol i denotes the one-period rate of interest.

1. state 1 featuring probability 1% and state price $0.4/(1+i)$;

2. state 2 featuring probability 1% and state price $0.4/(1+i)$;

3. state 3 featuring probability 10% and state price $0.1/(1+i)$; and

4. state 3 featuring probability 88% and state price $0.1/(1+i)$.

The four states are assumed to be mutually exclusive and exhaustive: the four probabilities sum to unity. The pricing kernel, defined as the ratio of state prices-to probabilities is weakly monotone declining in proceeding from state 1 to state 4. The pricing kernel indicates which states are more expensive than others. Here states 1 and 2 are expensive and states 3 and 4 are cheap.

To form risk 7A in Allais choice 1, the individual purchases a one-unit contingent claim for each state, the cost of which is easily seen to be $1/(1+i)$. Notice that this combination pays one unit at the subsequent date, no matter which of the four states occurs.

To form risk 7B, the individual purchases one S2-contingent claim, one S4-contingent claim, and five S3-contingent claims. The combination pays $1 with probability 89% = 1% + 88%, and $5 with probability 10%. The cost of this combination is $1/(1+i)$.

To form risk 7C, the individual purchases one S2-contingent claim and one S3-contingent claim. The combination pays $1 with probability 11% = 1% + 10%. The cost of this combination is $0.5/(1+i)$.

To form risk 7D, the individual purchases five S3-contingent claims. The combination pays $5 with probability 10%. The cost of the claims is $0.5/(1+i)$.

An individual who holds claims to risk 7A can transform that risk into risk 7B by reducing the probability of receiving $1 by 11%, increasing the probability of receiving $0 by 1% and increasing the probability of receiving $5 by 10%. Because risks 7A and 7B have the same market value, $1/(1+i)$, the trades required to make the transformation have a net value of $0.

Likewise, an individual who holds claims to risk 7C can transform that risk into risk 7B by reducing the probability of receiving $1 by 11%, increasing the probability of receiving $0 by 1% and increasing the probability of receiving $5 by 10%. Because risks 7A and 7B have the same market value, $0.5/(1+i)$, the trades required to make the transformation have a net value of $0.

Notice that the two transformation trades are identical. This feature lies at the heart of what makes the typical Allais choice pattern paradoxical. In a standard expected utility framework, the change in expected utility associated with both transformation trades is

$$.01[Z(0) - Z(1)] + .1[Z(5) - Z(1)] \tag{64}$$

Because the expected utility function is linear in probabilities, (64) is independent of the initial risk being transformed. This why an expected utility maximizer does not exhibit Allais paradox behavior.

Individuals who exhibit typical Allais paradox behavior appear to focus heavily on the 1% probability of receiving \$0 in 7B, relative to the 0% probability of receiving \$0 in 7A. Doing so in a planner-doer setting can be modeled as having a lower value for doer utility $Z(0)$ in risk 7B than its counterpart value in risk 7D. Given that the function $Z()$ is the difference between logarithmic planner utility and a penalty function, such a feature would come about because risk 7B involves a higher penalty value $\Pi(0)$ than does risk 7D, in the region to the left of the sure gain-reference point. Therefore, in Panel B of Figure 9, the portion of Z to the left of the sure gain would be more steeply sloped when considering 7B than when considering 7D. Let $\Pi_L()$ denote the incremental penalty associated with the pain of choosing risk 7B and having the resulting outcome be less than the sure gain.

Consider the planner-doer explanation for why an individual would choose 7B from the choice set {7A, 7B}, but 7D from the choice set {7C, 7D}. The key point is that the presence of the sure gain 7A in the choice set {7A, 7B} involves the incremental penalty function $\Pi_L()$ in 7B, but not in 7C or in 7D. To understand this feature in the context of the planner-doer model, think about zero-willpower condition as an initial condition. In respect to Z_{mid}, zero-willpower leads to the choice of the risk-free outcome.[123]

In the case of 7C and 7D, the value of the risk-free outcome is $0.5/(1+i)$. Call this risk-free outcome 7E. To transform risk 7E into risk 7C, the individual needs to sell 0.5 units each of S1- and S4-contingent claims, and purchase 0.5 units each of S2- and S3-contingent claims. Doing so reduces the payoff in states S1 and S4, which occur with probability 89%, from 0.5 to 0 and increases the payoffs in states S2 and S3, which occur with probability 11%, to 1.

In the planner-doer model, choosing the risky 7C over the sure gain 7E involves the transformation of 7E into 7C. It is this transformation that triggers the presence of the incremental penalty $\Pi_L()$. However, 7E is not a member of the choice set {7C, 7D}. Therefore, the transformation is entirely theoretical, meaning that it has no impact on realized utility. Choosing 7D over 7C involves making a set of trades to transform 7C

[123] In the intertemporal choice framework based on Z_{high}, zero-willpower leads to the date t doer consuming the entire mental account balance at t.

into 7D. Because 7C is not risk-free, the incremental penalty $\Pi_L()$ is not triggered by the transformation. This feature allows for 7D being preferred to 7C, even as 7A is preferred to 7B.[124]

To conclude this section, consider the following theoretical points, when comparing the direct prospect theory two-system formulation to the indirect formulation.

- Concave utility functions: The basic structure of prospect theory mimics that of expected utility theory, with a value function in place of a utility function. Notably, the value function has both concave and convex regions, which contrasts with expected utility theory, which has only a concave region. The convex region is prone to generate extreme boundary solutions, which can be problematic for modeling purposes. In contrast, the second version of the planner-doer approach features only concave utility functions, both for the planner and the doers. As Figure 6 illustrates, the doer utility function might be much more concave in the portion of the domain associated with the income account than the portions associated with the asset account and future income account. From the planner's perspective, this suggests being more risk averse for risks associated with outcomes in the income account than the other two accounts. In this regard, keep in mind that facing risk typically impacts the size of these accounts, with attendant consequences for the amount of temptation to be faced.[125]

- Temptation: In the intertermporal framework, because of temptation, the doer utility function depends on choice set parameters (especially mental account balances). In the risk framework, the doer utility function also depends on choice set parameters (such as prior gains and losses, probability distributions, and presence of a risk-free option).

- Probability: The basic structure of prospect theory mimics that of expected utility theory, with a decision weighting function in place of a probability distribution.

[124]The planner-doer explanation of the Allais paradox involves doer utility functions that vary with the choice set. Analogously, in cumulative prospect theory, the decision weights vary with the ranking on prospect outcomes that occur with positive probability.

[125]That is, the m-variables are stochastic.

Notably, the decision weighting functions in prospect theory are rank dependent, with inverse S-shapes, that distort probabilities in the computation of weighted value. This compares to expected utility theory which uses undistorted probabilities. In contrast, the second version of the planner-doer approach features no direct distortion of probabilities, as the features associated with decision weights enter through the structure of mental accounts, and associated doer utility functions. Choice architecture is an important element, in respect to the priority ordering in which doers make selections.[126]

- Fourfold pattern: Prospect theory identifies four behaviors: general risk aversion in the domain of gains, general risk seeking in the domain of losses, risk seeking in the domain of gains when the probabilities of favorable events is small, and risk aversion in the domain of losses when the probabilities of unfavorable events is small. The planner-doer framework model features two modes, risk avoidance due to loss aversion driven by doer utility function characterized by Panel B in Figure 9 and risk seeking driven by doer utility functions characterized by Panel A and Panel C in Figure 9. In respect to the fourfold pattern, either prior losses or low probabilities of highly favorable gains induce the risk seeking mode. As discussed above, risk seeking behavior is related to a reduction in fear associated with activation of the anterior insula. Notably, the favoring of $2B$ over $2A$ in the low probability risk discussed above, while the favoring of $1A$ over $1B$ in the companion high probability risk is consistent with this feature.[127]

- Reference points and framing: The reference point in the above example is the sure outcome, \$1m, received at date $t+1$. The cost of purchasing this outcome in the market at date t is $1/(1+i)$, the discounted value. The gain in wealth between the two periods is i, which is taken to be the reference point. If $i=0$, then

[126]The need for an ordering reflects the fact that attention is a scare resource, which can limit the planner's ability to optimize.

[127]Choice architecture is especially important for the risk seeking mode, as doers concentrate first on purchasing claims to states featuring high payoffs. These states can involve strong temptation effects leading to skewed choices similar to what occurs in intertemporal choice with single mental accounts and strong self-control challenges. For high payoff states, Z_{high} is the operative doer utility function.

the reference point corresponds to the intuitive interpretation of gains and losses. Framing effects in prospect theory involve changes in reference point. For example, the individual might receive risk 7A as a windfall instead of having to pay for it. This situation leads to a framing effect associated with "house money" [Thaler and Johnson (1990)]. The house money effect leads to a shift in reference point from the sure outcome ($1m) to $0. In addition, significant framing effects in the second version of the planner-doer approach arise through mental states that impact the configuration of doer utility functions drawn from the basis, that is the combination of Panels A and C, instead of Panel B.

- Rules: The BLC associates several types of choices to the planner, including the selection of rules. A stop loss order is an example of a precommitment device to restrict imprudent choices by doers facing losses. Structurally, the planner-doer approach to modeling present bias is very similar to the planner-doer approach to modeling the behavioral features emphasized by prospect theory which are at odds with the features associated with expected utility theory.

- Planner rationality: The assumption that the System 2 planner is an expected utility maximizer is done for analytical convenience, so that in the model departures from rational behavior stem from the costs of controlling System 1. Of course, this is an idealization. As Kahneman (2011) notes, System 2 might be far from perfect.

- Integration: The integration of the risk-based dimension and the intertemporal dimension can be accomplished by nesting the former into the latter, using sub-accounts of mental accounts. Notably, the intertemporal analogue to consumption variables are expenditure allocations to mental accounts, and the analogue to doer utilities is the expected utilities associated with these accounts and subaccounts. Notably, two-system issues are used to model risk choices related to mental sub-accounts, along with present bias issues associated with expenditures on mental accounts. Each has their own combination of control variables (θ) and penalty functions (Π).

16 Concluding Remarks

This chapter articulates the relative merits of transitioning from a one-system neoclassical based economics framework to a two-system TF&S-framework, where TF&S-economics stands for "thinking, fast and slow." Making the case for this transition comprises the overarching theme of the chapter, and applies to three types of decision problems in which people make dynamically inconsistent choices. The three ways are: intertemporal substitution associated with present bias, the impact of self-control challenges on coincident substitution between different commodities, and the effect of changing mental states in respect to substitution between alternative risky prospects.

TF&S models feature two mental systems, a deliberative system and an automatic system. The first formal TF&S economic models were proposed in Thaler and Shefrin (1981) and Shefrin and Thaler (1988), where the deliberative system was termed the planner and the automatic system was termed the doer. Shefrin and Thaler (1988) introduced two important features to the planner-doer framework, temptation and mental accounting. The discussion in the chapter emphasizes the significance of both features for the three types of decision problems.

For the most part, economic models are one-system models. Certainly neoclassical models are structured as constrained optimization problems, where rational behavior corresponds to optimizing solutions. However, some behavioral models also conform to one-system optimizing models, especially those which feature intrapersonal game theoretic frameworks. As discussed in the chapter, a broad academic literature supports the view that two-system planner doer models involving temptation and mental accounting better accommodate the way people make economic decisions than one-system models.

Consider some of the highlights from the broad academic literature.

- The neuroeconomics literature on temptation and willpower is consistent with the planner-doer perspective about self-control being a simultaneous conflict involving different regions of the brain. The literature on nature, nurture, and financial literacy provides no support for the contention that most people are capable of solving the kind of complex dynamic lifetime optimization problems associated with the neoclassical approach and the intrapersonal game theoretic approach. Instead the literature on these topics finds that most people's financial decision skills are limited, and that they are better able to deal with simple rules with limited optimization.

- The literature on behavioral nudging supports the idea that people can be nudged into precommitment behavior which can help them successfully reduce overconsumption. In contrast, the game theoretic approach suggests that people avoid precommitment because liquidity constraints are inflexible. As for the neoclassical approach, there is no need for nudging as people are assumed to make fully rational choices on their own.

- The literature on household budgeting supports the notion that people budget by using matrix mental accounting, and that doing so enables them to avoid the problem of "many margins" for overconsuming. In contrast, the game theoretic approach ignores mental accounting, and suggests that because of "many margins," people will find ways to defeat self-selected liquidity constraints. Mental accounting is a bounded rationality feature, which plays no germane role in neoclassical frameworks that assume full rationality.

- The literature on windfalls is consistent with the differential MPC hypothesis implied by BLC theory. In contrast, the neoclassical approach ignores present bias, and finds the need for a "temptation component" to be added to the model in order to achieve an acceptable fit between theory and evidence.

- The literature on cash dividends suggests that households use mental accounting to separate the value of their stocks into dividend income and capital, for the purpose of consuming dividends but not "dipping into capital." In contrast, the neoclassical approach treats dividends and capital as being equivalent, but for taxes and transaction costs.

- The literature testing the BLC, for households headed by people in late middle age, provides strong evidence of a general differential MPC mental accounting effect for both general consumption and for matrix mental accounting in respect to specific consumption categories. In contrast, the neoclassical approach posits no such differential property for MPCs.

- The findings from the HANK macroeconomic literature fully support the BLC differential MPC hypothesis. An important finding from this literature is the need to include temptation in order for HANK models to fit the empirical data.
- The literature on wealth accumulation shows that, in contrast to present bias, retired individuals oversave. This finding is consistent with miserliness, which emerges in the two-system framework as a consequence of status quo bias in respect to wealth accumulation rules.
- The literature on prospect theory finds evidence of a fourfold pattern with a tendency to be risk-seeking in the domain of losses and to favor prospects with tempting lottery-like features. There is evidence that investor view risk-seeing in the domain of losses as a self-control problem, and seek commitment devices that will help them with self-control. There is also evidence that relaxing restrictions on accessing wealth from illiquid assets associated with self-control safeguards leads a significant number of people to use that wealth in order to gamble. These patterns strongly relate to self-control, temptation, and mental accounting.

The planner-doer framework was the first neuroeconomics model, and its structure captures key features of the manner in which people make decisions, and for that matter the decisions themselves. BLC-hypotheses about behavior relating to savings, differential MPCs, windfalls, and miserliness were advanced in the 1980s, and are supported by decades of data. For example, the BLC offers a natural explanation for why people overconsume during their accumulating years, but underconsume during retirement. For one-system models, this duality finding is a puzzle.

While the planner-doer framework is cast as an optimization, the intent is not to suggest that most people behave as if they are capable of solving complex dynamic maximization problems. Rather, the intent is to identify plausible heuristics that people use to mimic full optimization.

During the twenty-first century, the economics literature formally modeled temptation very similarly to the approach advanced by Shefrin and Thaler (1988). Indeed, most of the recent approaches can be embedded within a planner-doer framework, and therefore within a two-system framework. There is good reason to do so. The TF&S approach, with temptation and mental accounting, is strongly consistent with key findings in the neuroeconomics literature and with evidence about behavior in the applied economics literature.

References

Aczel, J., 1966. *Lectures on Functional Equations and Their Applications.*, Academic Press, New York.

Adriaanse, M. A., F. M. Kroese, M. Gillebaart, M., and D. T.D. De Ridder, 2014. "Effortless Inhibition: Habit Mediates the Relation Between Self- control and Unhealthy Snack Consumption," *Frontiers of Psychology*, Vol. 5, 1–6.

Ainslie, G., 1974. "Impulse Control in Pigeons," *Journal of the Experimental Analysis of Behavior*, 4, pp.485-89.

Ainslie, G., 1975. "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control," *Psychological Bulletin*, pp. 463-96.

Ainslie, G., 1985. "Beyond Microeconomics. Conflict Among Interests in a Multiple Self as a Determinant of Value," in J. Elster (ed.) *The Multiple Self*, Cambridge University Press, Cambridge, U.K.

Ainslie, G., 1992. *Picoeconomics: The Strategic Interaction of Successive Motivational States Within the Person*, Cambridge University Press.

Ainslie, G. and N. Haslam, 1992. "Self-Control," in *Choice over Time*, edited by George Lowenstein and Jon Elster, Russell Sage Foundation, New York.

Allais, M., 1952. "Fondements d'une Théorie Positive des Choix Comportant un Risque eE Critique des Postulats et Axiomes de l'école Américaine," *Econométrie*, 257-331.

Amador, M., M. Werning, and G. Angeletos, 2006. "Commitment vs. Flexibility," *Econometrica*, 74 (2), 365-396.

Ambuehl, S., B. D. Bernheim, and A. Lusardi, 2014. "The Effect of Financial Education on the Quality of Decision Making," National Bureau of Economic Research, Working Paper 20618, http://www.nber.org/papers/w20618.

Attanasio, O. A. Kovacs, and P. Moran, 2024. "Temptation and Commitment: a Model of Hand-to-Mouth Behavior," *Journal of the European Economic Association*, Vol. 22(4), Issue 4, 2025–2073.

Baker, M., S. Nagel, and J. Wurgler. 2007. "The Effect of Dividends on Consumption."" *Brookings Papers on Economic Activity*, 277–291.

Barber, B., and T. Odean. 2000. "Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors," *Journal of Finance* 55:2, 773–806.

Barberis, N., 2012. "A Model of Casino Gambling," *Management Science*, vol. 58, issue 1, 35-5.

Becker, G. and K. Murphy, 1988. "A Theory of Rational Addiction," *Journal of Political Economy*, 96, 675-700.

Becker G., M. Grossman, and K. Murphy, 1992. "Rational Addiction and the Effect of Price on Consumption," in *Choice over Time*, edited by George Lowenstein and Jon Elster, Russell Sage Foundation, New York.

Benzion, U. A. Rapoport, and J. Yagil, 1989. "Discount Rates Inferred from Decisions: An Experimental Study,', *Management Science*, 35, pp. 270-84. 2015.

Bernheim, B. D., D. Ray, and S. Yeltekin, 2015. "Poverty and Self-Control," *Econometrica*, 83 (5), 1877-1911.

Bernheim, B. D., and Dmitry Taubinsky, 2018. "Behavioral Public Economics," in *Handbook of Behavioral Economics: Applications and Foundations, 1,* Vol. 1, Elsevier, 381-516.

Beshears, J., J. Hurwitz, D. Laibson, and B. C. Madrian, 2015. "Liquidity in Retirement Savings Systems: An International Comparison," *American Economic Review*, 105 (5), 420-25.

Boehm, J. E. Fize, and X. Jaravel, 2025. "Five Facts about MPCs: Evidence from a Randomized Experiment," *American Economic Review*, Vol115(1), 1–42.

Böhm-Bawerk, E., 1889. *Capital and Interest*, South Holland, IL: Libertarian Press 1889, 1970.

Bräuer, Konstantin, Andreas Hackethal, and Tobin Hanspal. 2022. "Consuming Dividends," *Review of Financial Studies*, 35 (10), 4802-4857.

Chambers, C. D., H. Garavan, and M. A. Bellgrove, 2009. "Insights into the Neural Basis of Response Inhibition from Cognitive and Clinical Neuroscience," *Neuroscience & Biobehavioral Reviews*, Volume 33, Issue 5, 631-646.

Coates, J., 2012. *The Hour Between Dog and Wolf: Risk-taking, Gut Feelings and the Biology of Boom and Bust*, New York : Penguin Press.

Constantinides, G., 1990. “Habit Formation: A Resolution of the Equity Premium Puzzle”, *Journal of Political Economy*, 98, 3, pp. 519-43.

Cronqvist, H. and S. Siegel (2015). “The Origins of Savings Behavior,” *Journal of Political Economy*, Vol. 123, No. 1 (February 2015), 123-169.

DellaVigna, S. and U. Malmendier, 2004. “Contract Design and Self-Control: Theory and Evidence,” *Quarterly Journal of Economics,* Vol. 119, 2, pp. 353–402.

DellaVigna, S. and U. Malmendier, 2006. “Paying Not to Go to the Gym,” *American Economic Review*, vol. 96, no. 3, 694–719.

De Nardi, M., E. French, E. and J. B. Jones, J. B., 2016, “Saving After Retirement: A Survey,” *Annual Review of Economics* 8, 177–204.

Epstein, L. and S Zin, 1989. “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, Vol. 57, No. 4, 937-969.

Fagereng, A., M. B. Holm, and G. J. Natvik, 2021. “MPC Heterogeneity and Household Balance Sheets,” *American Economic Journal: Macroeconomics*, Vol. 13, No. 2, 1-54.

Fisher, I., 1930. *The Theory of Interest*, MacMillan, New York.

Fisher, F.M., 1977. “On Donor Sovereignty and United Charities,” *American Economic Review*, Vol. 67.

Frank, R., 1992. “The Role of Moral Sentiments in the Theory of Intertemporal Choice,” in *Choice over Time*, edited by George Lowenstein and Jon Elster, Russell Sage Foundation, New York.

Fudenberg, D. and D. Levine, 2006. “A Dual-Self Model of Impulse Control,” *American Economic Review*, Vol. 96, no. 5, 1449–1476.

Fudenberg, D. and D. Levine, 2011. “Risk, Delay, and Convex Self-Control Costs,” *American Economic Journal: Microeconomics*, Vol. 3, 34–68.

Fudenberg, D. and D. Levine, 2011. “Timing and Self-Control,” *Econometrica*, Vol. 80, No. 1, 1–42.

Simone Galperti, S. 2019. “A Theory of Personal Budgeting,” *Theoretical Economics*, Vol. 14, 173–210.

Ganong, P., D. Jones, P. J. Noel, F. E. Greig, D. Farrell, C. Wheat, 2023. "Wealth, Race, and Consumption Smoothing of Typical Income Shocks," National Bureau of Economic Research, Working Paper 27552.

Gathergood, John, Neale Mahoney, Neil Stewart, and Joerg Weber, 2019. "How Do Individuals Repay Their Debt? The Balance-Matching Heuristic," *American Economic Review*, 109(3), 844-875.

Gigerenzer, G., Todd, P., & The ABC Research Group, 1999. *Simple Heuristics that Make Us Smart.* New York: Oxford University Press.

Gigerenzer, G., 2007. *Gut Feelings: The Intelligence of the Unconscious.* New York: Viking Press.

Gimeno-Ribes, J. 2024. "A Model of Consumption with Mental Accounting and Incomplete Markets," Working paper: University of Valencia.

Giulio Fella, G. M. B. Holm and T. M. Pugh, 2024. "Saving after Retirement and Preferences for Residual Wealth," Staff Working Papers 24-21, Bank of Canada.

Goldman, S., 1979. "Intertemporally Inconsistent Preferences and the Rate of Consumption," *Econometrica*, 47(3), 621-626.

Goldman, S., 1980. "Consistent Plans," *Review of Economic Studies*, 47, 3, 533-537.

Gorman, W.M., 1967. "Tastes, Habits, and Choices," *International Economic Review*, Volume 8, pp. 218-222.

Graham, J., and A. Kumar. 2006. "Dividend Preference of Retail Investors: Do Dividend Clienteles Exist?" *Journal of Finance* 6:3, 1305–1336.

Graham, J. and R. McDowall, 2024. "Mental Accounts and Consumption Sensitivity Across the Distribution of Liquid Assets." Working paper, University of Sydney.

Gul, F. and and W. Pesendorfer, 2001."Temptation and Self-Control," *Econometrica*, Vol. 69, No. 6, 1403-1435.

Gul, F., and W. Pesendorfer, 2004. "Self-control and the Theory of Consumption," *Econometrica* 72, 119–158.

Gul, F., and W. Pesendorfer, 2007. "Harmful Addiction," *Review of Economic Studies,* 74, 147–172.

Hamilton, S. G. Liu, J. Miranda-Pinto, and T. Sainsbury, 2024. " A $100,000 Marshmallow Experiment: Withdrawal and Spending Responses to Early Retirement-Savings Access." Working paper: George Washington University.

Hammond, P., 1976a. "Changing Tastes and Coherent Dynamic Choice," *Review of Economic Studies*, XLIII, 133, pp. 159-73

Hammond, P., 1976b. "Endogenous Tastes and Stable Long-Run Choice," *Journal of Economic Theory*, Volume 13, Number 2, pp. 329-340.

Harris, C. and D. Laibson,"Instantaneous Gratification," *The Quarterly Journal of Economics*, 2013, 128 (1), 205.

Hayashi, T. and N. Takeoka, 2022."Habit formation, Self-deception, and Self-control," *Economic Theory*, Springer; Society for the Advancement of Economic Theory (SAET), Vol. 74(2), 547-592.

Heimer, R. Z. Iliewa, A. Imas, and M. Weber, 2023. "Dynamic Inconsistency in Risky Choice: Evidence from the Lab and Field," forthcoming *American Economic Review*.

Herrnstein, R., 1961. "Relative and Absolute Strengths of Response as a Function of Frequency of Reinforcement," *Journal of the Experimental Analysis of Behavior*, 4, pp. 267-272.

Horowitz, J.K., 1988. "Discounting Money Payoffs: An Experimental Analysis," working paper Department of Agricultural and Resource Economics, University of Maryland.

Jevons, W.S., 1871. *Theory of Political Economy*, Macmillan, London.

Kahneman, D., 2011. *Thinking, Fast and Slow*, Farrar, Straus, and Giroux, New York.

Kahneman, D. and A. Tversky, 1979. "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, Volume 47, No. 2, 1979, pp. 263-292.

Kahneman, D. and A. Tversky, 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, pp. 297-323.

Kaplan, G., and G. L. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, Vol. 82 (4), 1199-1239.

Kaplan, G., G. L. Violante, and J. Weidner, 2014. "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, Spring, 77-138.

Kaplan, G., and G. L. Violante, 2022. "The Marginal Propensity to Consume in Heterogeneous Agent Models," *Annual Review of Economics*, 14, 747-775.

Knutson B., , G. E. Wimmer, C. M. Kuhnen, and P. Winkielman, 2008. "Nucleus Accumbens Activation Mediates the Influence of Reward Cues onFinancial Risk Taking," *Brain Imaging*, Vol 19 No 5, 509-513.

Koopmans, T., 1960. "Stationary Ordinal Utility and Impatience", *Econometrica*, 28, pp. 287-309.

Köszegi, B. and M. Rabin. 2006. "A Model of Reference-Dependent Preferences", *Quarterly Journal of Economics,* 121 (4), 1133-1165.

Kovacs, A., H. Low, and P. Moran, 2021. "Estimating Temptation and Commitment Over the Life-Cycle." *International Economic Review*, Vol. 62 (1), 101–139.

Kuhnen, C., 2015. "Asymmetric Learning from Financial Information," *Journal of Finance*, Vol 70, Issue 5, 2029-2062.

Laibson, D., 1995. "Hyperbolic Discount Functions, Undersaving, and Savings Policy," working paper, Harvard University.

Laibson, D., 1997. "Golden Eggs and Hyperbolic Discounting," *The Quarterly Journal of Economics*, 112 (2), 443-478.

Laibson, D., 1998. "Life-Cycle Consumption and Hyperbolic Discount Functions," *European Economic Review*, 42 (3-5), 861-871.

Laibson, D., 2015. "Why Don't Present-Biased Agents Make Commitments?," *American Economic Review*, 105 (5), 267-72.

Levin, L., 1998. "Are Assets Fungible? Testing the Behavioral Theory of Life-cycle Savings," *Journal of Economic Behavior & Organization*, Vol. 36(1), 59-83.

Liao, C., 2021a. "ADHD Symptoms and Financial Distress," *Review of Finance*, Vol. 25, No. 4, 1129–1210.

Liao, C., 2021b. "Supplemental notes to Liao (2021a)." https://papers.ssrn.com/sol3/papers.cfm?abstrac

Lockwood, L. 2018."Incidental Bequests: Bequest Motives and the Choice to Self-Insure Late-Life Risks," *American Economic Review*, 108(9), 2513–2550.

Loewenstein, G, 1992. "The Fall and Rise of Psychological Explanantions in Intertemporal

Choice", in *Choice over Time*, edited by George Loewenstein and Jon Elster, Russell Sage Foundation, New York.

Loewenstein, G, and D. Prelec, 1992. "Anomalies in Intertemporal Choice: Evidence and Interpretation", in *Choice over Time*, edited by George Loewenstein and Jon Elster, Russell Sage Foundation, New York.

Lopes, L., 1987. "Between Hope and Fear: The Psychology of Risk," *Advances in Experimental Social Psychology*, Volume 20, 255-295.

Lusardi, A., 2010. "Americans' Financial Capability," Report Prepared for the Financial Crisis Inquiry Commission.

Lusardi, A., and O. Mitchell, 2014. ""The Economic Importance of Financial Literacy: Theory and Evidence," *Journal of Economic Literature*, 52(1): 5-44.

Maxted, P., 2024. "Present Bias Unconstrained: Consumption, Welfare, and the Present-Bias Dilemma," Working paper University of California, Berkeley, Haas School of Business.

McClennen, E., 1990. *Rationality and Dynamic Choice: Foundational Explorations*, Cambridge University Press, Cambridge, United Kingdom.

McClure, S. M., D. I. Laibson, G. Loewenstein, and J.D. Cohen, 2004. "Separate Neural Systems Value Immediate and Delayed Monetary Rewards," *Science*, 306(5695), 503–507.

Mikella, A., J. Crawford, C. Kuhnen, G. R. Samanez-Larkinand, and K. Seaman, 2023. "Multivariate Associations Between Dopamine Receptor Availability and Risky Investment Decision-Making Across Adulthood," *Cerebral Cortex Communications*, Vol. 4, 1–8.

Mischel, W., Y. Shoda, and M. Rodriguez, 1992. "Delay of Gratification in Children," in *Choice over Time*, edited by George Loewenstein and Jon Elster, Russell Sage Foundation, New York.

Misra, K. and P. Surico. 2014. "Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs." *American Economic Journal: Macroeconomics* Vol. 6(4), 84–106.

Muraven, M., M.D. Tice, and R. Baumeister, 1998. "Self-Control as a Limited Resource," *Journal of Personality and Social Psychology*, 74, 774–789.

Noor, J. and N. Takeoka, 2010. "Uphill Self-Control." *Theoretical Economics*, Vol.5(2), 127–58.

Peston, M. 1967. "Changing Utility Functions," in Martin Shubik, *Essays in Mathematical Economics, Essays in Honor of Oskar Morgenstern*, Princeton University Press, Princeton, NJ, pp. 233-236.

Phelps, E.S. and R.A. Pollak, 1967. "On Second-Best National Saving and Game-Equilibrium Growth," *Review of Economic Studies*.

Pollak, R.A., 1968. "Consistent Planning," *Review of Economic Studies*, 102.

Pollak, R., 1971. "Habit Formation and Dynamic Demand Functions," *Journal of Political Economy*, 78, 4, pp. 745-63.

Pollak, R., 1971. "Habit Formation and Long-Run Utility Functions," *Journal of Economic Theory*, Volume 13, Number 2, pp. 272-297.

Prelec, D., 1989. "Decreasing Impatience: Definition and Consequences." Harvard Business School working paper.

Quiggin, J., 1982. "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization*, December 1982, 3, pp. 323-43.

Rabin, M., 1998. "Psychology and Economics," *Journal of Economic Literature*, 36, 11-46.

Rabin, M., 2000. "Risk Aversion and Expected-Utility Theory: A Calibration Theorem." *Econometrica*, 68(5): 1281–92.

Rae, J., 1834. *The Sociological Theory of Capital*, reprint of original 1834 edition, MacMillan, London.

Rantapuska, E., and M. Kaustia. 2007. "Do Investors Reinvest Dividends and Tender Offer Proceeds?" Working Paper, Aalto Univerity, Finland.

Samuelson, P., 1937. "A Note on Measurement of Utility," *Review of Economic Studies*, 4, pp.155-61.

Schelling, T., 1992. "Self-Command: A New Discipline," in *Choice over Time*, edited by George Loewenstein and Jon Elster, Russell Sage Foundation, New York.

Sen, A., 1971. *Collective Choice and Social Welfare*, Oliver and Boyd, Edinburgh.

Senior, N.W. 1836. *An Outline of the Science of Political Economy*, Clowes and Son, London.

Shefrin, H., 2009a. "Behavioral Explanations of Dividends," in *Dividends and Dividend Policy*, edited by H. Kent Baker. Wiley.

Shefrin, H., 2009b. "One Way to Compel Taxpayers to Spend," *Boston Globe*, February 15.

Shefrin, H. , 2020. "Unfinished Business: a Multicommodity Intertemporal Planner–Doer Framework," *Review of Behavioral Finance*, Vol. 12(1), 35-68.

Shefrin, H., 2021. "Broadening the Scope of Financial Literacy to Incorporate Self-Control, Budgeting and Heuristics," in: R. Viale & U. Filotto, B. Alemanni and S. Mousavi (ed.), *Financial Education and Risk Literacy*, chapter 1, 5-26, Edward Elgar Publishing.

Shefrin, H., 2023. *The Behavioral Economics and Politics of Global Warming: Unsettling Behaviors.* Cambridge University Press, Cambridge.

Shefrin, H. and Statman, M. 1984. "Explaining Investor Preference for Cash Dividends," *Journal of Financial Economics* 13 (June), pp. 253-82.

Shefrin, H. and Statman, M. 1985. "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence," *The Journal of Finance*, Vol. 40, No. 3, 777-790.

Shefrin, H. and R. Thaler, 1988. "The Behavioral Life Cycle Hypothesis," *Economic Inquiry*, Vol. XXVI, pp. 609-43.

Shefrin, H. and Statman, M. 1992. *Ethics, Fairness, Efficiency, and Financial Markets*, Research Foundation of Institute of Chartered Financial Analysts, Charlottesville, Virginia.

Shefrin, H. and C. Nicols, 2014. "Credit Card Behavior, Financial Styles, and Heuristics," *Journal of Business Research*, 67: 1679–1687.

Shi, T., S. Feng, M. Wei, and W. Zhou, 2020. "Role of the Anterior Agranular Insular Cortex in the Modulation of Fear and Anxiety, *Brain Research Bulletin*, Volume 155, 174-183.

Slovic, P., M. L. Finucane, E. Peters, D. G. and MacGregor, 2007. "The Affect Heuristic," *European Journal of Operational Research*, Vol. 177, 1333–1352.

Smith, A., 1776. *The Wealth of Nations*, George Bell and Sons, London.

Stone, R., 1954. "Linear Expenditure Systems and Demand Analysis," *Economic Journal*, Volume 64, pp. 511-527.

Strotz, R.H., 1956. "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 23.

Thaler, R., 1981. "Some Empirical Evidence on Dynamic Inconsistency," *Economics Letters*, 8, pp. 201-207.

Thaler, R. , 1985. "Mental Accounting and Consumer Choice," *Marketing Science*, 27(1): 15-25.

Thaler, R., 2015. *Misbehaving*, W. W. Norton, New York.

Thaler, R. and E. J. Johson, 1990. "Gambling With the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice," *Management Science*, Vol. 36(6), 643-660.

Thaler, R., and H. Shefrin, 1981. "An Economic Theory of Self Control," *Journal of Political Economy*, 89(2): 392-406.

Thaler, R. and S. Benartzi, 2004. "Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving," *Journal of Political Economy*, vol. 112, no. 1, pt. 2: S164-187.

Tversky, A. and R.H. Thaler, 1990. "Preference Reversals," *The Journal of Economic Perspectives*, Vol. 4, No. 2, pp. 201-211.

Thaler, R. and C. Sunstein, 2008. *Nudge:The Final Edition.* Penguin Books.

Vanguard, 2024. *How America Saves 2024.* https://institutional.vanguard.com/insights-and-research/report/how-america-saves.html

Viale, R., 2022. *Nudging*, London: The MIT Press.

Violante, G., 2021. "What Have We Learned From HANK Models, Thus Far?" *Proceedings of the ECB Forum on Central Banking.*

von Auer, L. 1999. "Dynamic Choice Mechanisms," *Theory and Decision*, 46, 295–312.

von Auer, L. 2004. "Revealed Dynamic Preferences under Uncertainty," *Theory and Decision* 56 (3):269-290.

von Weizsäcker, 1971. "Notes on Endogenous Changes of Tastes," *Journal of Economic Theory*, Volume 3, pp. 345-372.

Yarri, M., 1977. "Consistent Utilization of an Exhaustible Resource or How to Eat an Appetite-Arousing Cake," Research Memorandum No. 26, Hebrew University, Jerusalem.

Zhang, C. Y., A. B. Sussman, N. Wang-Ly, and J. K. Lyu , 2022. "How Consumers Budget," *Journal of Economic Behavior & Organization*, Vol. 204, 69-88.

17 Appendix

Proof of Theorem 1: The proof of the theorem follows from the optimality condition for the U_τ optimization problem:

$$\lambda_\tau(t)u'(c_t^1) = \lambda_\tau(s)u'(c_s^1), \tag{65}$$

which is equivalent to:

$$u'(c_t^1)/u'(c_s^1) = \lambda_\tau(s)/\lambda_\tau(t). \tag{66}$$

The optimality of c^1 under U_τ, for all τ, requires that condition (66) hold for all τ. However, for $\tau = 1, 2$, and $t, s = 2, 3$, with $\lambda_1(1) = 1$, it follows that:

$$\lambda_1(3) = \lambda_1(2) \times \lambda_2(3).$$

Moreover, the shifting discount function condition (4) implies that $\lambda_1(2) = \lambda_2(3)$. Therefore

$$\lambda_1(3) = (\lambda_1(2))^2$$

which, applying the argument recursively implies that $\lambda_1()$ is exponential. **QED**

Proof of Theorem 2: For any t', and $\gamma t + (1-\gamma)t'$, (11) implies that

$$u(x)\lambda(t') = u(y)\lambda(kt' + s) \tag{67}$$

and

$$u(x)\lambda(\gamma t + (1-\gamma)t')$$

$$= u(y)\lambda(k(\gamma t + (1-\gamma)t') + s)$$

$$= u(y)\lambda(\gamma(kt + s) + (1-\gamma)(kt' + s)$$

$$= u(y)\lambda(\gamma\lambda^{-1}(u(x)\lambda(t)/u(y)) + (1-\gamma)\lambda^{-1}(u(x)\lambda(t')/u(y)) \tag{68}$$

Define $r = u(x)/u(y)$, $w = \lambda(t)$, $z = \lambda(t')$, and $v = \lambda^{-1}$. Substitution of these variables into (68) leads to the functional equation:

$$rv^{-1}(\gamma v(w) + (1-\gamma)v(z)) = v^{-1}(\gamma v(rw) + (1-\gamma)v(rz)). \tag{69}$$

Equation (69) is a functional equation in v. From Aczel (1966, p. 152, equation (18)), the only solutions to this equation are:

$$v(t) = c\log(t) + d$$

and

$$v(t) = ct^{\mu} + d.$$

Since $\lambda(t) = v^{-1}(t)$, the discount function must be either hyperbolic or exponential. In particular, it must have the form (10). **QED**

Proof of Theorem 3: Let $\tilde{c}_2$ be defined by the expression

$$u'(\tilde{c}_2) - \lambda_2(3)u'(0) = 0.$$

If $u'(0) = \infty$, then $\tilde{c}_2$ is undefined. Recall that $q_1 = q - c_1$. Observe that if $q_1 < \tilde{c}_2$, then $c_2^*(q_1) = q_1$ and $c_3^*(q_1) = 0$. Let $\tilde{\lambda}_2$ be defined as $1/\lambda_2(3)$. Since $u()$ is assumed to be concave, $\tilde{c}_2$ is an increasing function of $\tilde{\lambda}_2$. Consider the case in which $q_1 > \tilde{c}_2$. Then a consistent plan satisfies:

$$\tilde{\lambda}_2 u'(c_2^*) = u'(q_1 - c_2^*) \tag{70}$$

so that

$$\frac{dc_2}{dq_1} = \frac{u''(q_1 - c_2)}{\tilde{\lambda}_2 u''(c_2) + u''(q_1 - c_2)} \tag{71}$$

and

$$\frac{dc_3}{dq_1} = \frac{\tilde{\lambda}_2 u''(c_2)}{\tilde{\lambda}_2 u''(c_2) + u''(q_1 - c_2)} \tag{72}$$

Recall that $c_3 = q_1 - c_2$ so that $dc_3/dq_1 = 1 - dc_2/dq_1$. Observe from (71) that dc_2/dq_1 is decreasing in $\tilde{\lambda}_2$ and dc_3/dq_1 is increasing in $\tilde{\lambda}_2$. However for fixed q_1, (70) implies that $dc_2/d\tilde{\lambda}_2 > 0$.

Define $\tilde{\lambda}_1 = \lambda_1(2)/\lambda_1(3)$. Because the discount function is hyperbolic, $\tilde{\lambda}_2 > \tilde{\lambda}_1$, which by (70) and because $c_3^* = q_1 - c_2^*$, implies:

$$\tilde{\lambda}_1 u'(c_2^*) < u'(c_3^*) \tag{73}$$

This inequality indicates that date 1 preferences would prefer a reallocation of the bequest q_1 from date 2 consumption to date 3. Since this is not possible, the next question concerns the impact on the selection of c_1. Observe that the left hand derivative $V_C'^-(\tilde{c}_2)$ is equal to $\lambda_1(2)u'(\tilde{c}_2)$. The right hand derivative $V_C'^+(\tilde{c}_2)$ is given by

$$\lambda_1(2)u'(\tilde{c}_2)\frac{dc_2}{dq_1} + \lambda_1(3)u'(0)\frac{dc_3}{dq_1}. \tag{74}$$

It follows that $V_C'^+(\tilde{c}_2) > V_C'^-(\tilde{c}_2)$. Furthermore, $u'' < 0$ implies that V_C' is decreasing in q_1 except at $\tilde{c}_2$. Hence, V_C' has the shape depicted in Figure 10.

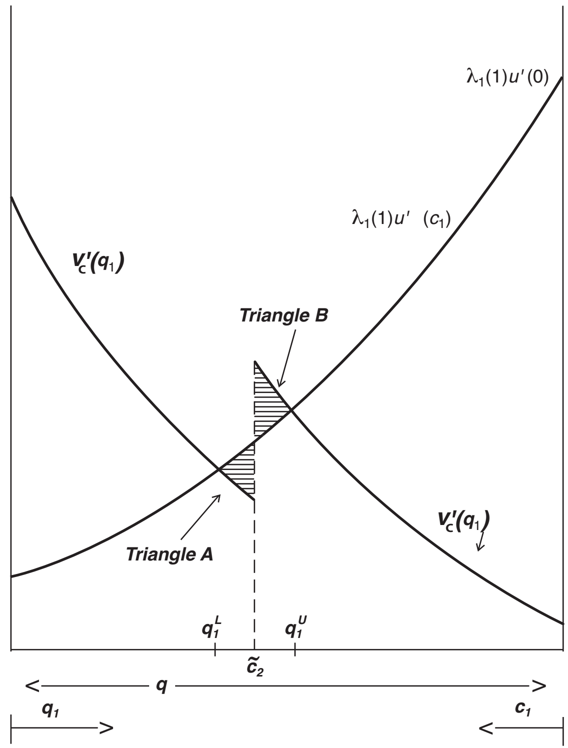

Figure 10: This graph depicts marginal utility V_C' of the value of a bequest q_1 against the marginal utility $\lambda_1(1)u'(c_1)$ of consumption, both along a consistent path. The axis for V_C' is at the left of the figure, while the axis for $\lambda_1(1)u'(c_1)$ is at the right of the figure. Note that the latter graph is depicted as a mirror image in order to reflect the fact that the sum of consumption c_1 and the bequest q_1 at date 1 sum to total wealth q.

Figure 10 illustrates how the optimal value of q_1 is selected along the consistent path. The right hand origin is the zero point for c_1 and the left is the zero point for q_1. Since $c_1 + q_1 = q$, the total distance between the origins is q. The optimal value of q_1 must satisfy $u'(c_1) = V_C'(q_1)$. If the u' curve intersects the V_1' curve at only one such point, then that point is c_1^*. However, there may be two such points, on account of the discontinuity in V_C'. This is the case illustrated in Figure 10, where the two points are denoted q_1^L and q_1^U. Date 1 preferences select q_1^U over q_1^L if the area of triangle B in the figure exceeds that of triangle A.

Recall that (3) defines $U_\tau()$ as the intertemporal utility function at τ associated with precommitment. Focus attention on Figure 11 which compares two marginal utility functions V_C' and V_P', where V_P' is the marginal utility function associated with the precommitment plan.

Define $\bar{c}_2$ such that: $\lambda_1(2)u'(\bar{c}_2) = \lambda_1(3)u'(0)$. Observe that $\bar{c}_2 < \tilde{c}_2$. Therefore when $\bar{c}_2 < q_1 < \tilde{c}_2$ we have $V_C'(q_1) < V_P'(q_1)$, meaning that the curve V_C' is steeper than the curve V_P'.

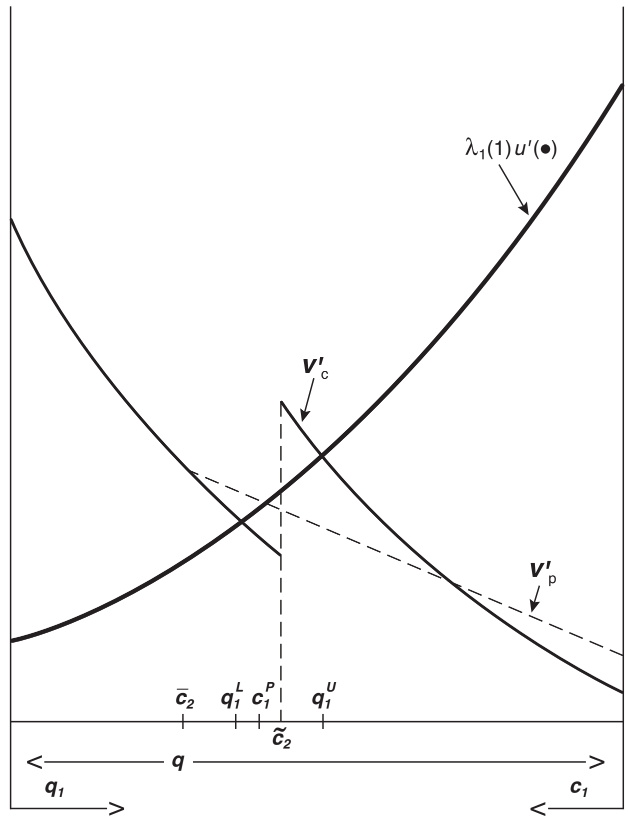

Figure 11: This figure compares V_C' and V_P', and displays the locations of key values for consumption and bequest variables at $t = 1$.

Figure 11 displays both V_C' and V_P'. Notice that V_P' is continuous but nondifferentiable at $\bar{c}_2$. It is continuous because of equation (70) (when evaluated at the tipping point $\bar{c}_2$ where $c_3 = 0$), and the fact that the precommitment plan corresponds to the case $\tilde{\lambda}_1 = \tilde{\lambda}_2$. It is nondifferentiable at $\bar{c}_2$ because to the left of $\bar{c}_2$, $dc_3/dq_1 = 0$ while to to the right of $\bar{c}_2$, $dc_3/dq_1 > 0$. These two derivatives are the weights in a weighted average of marginal discounted utilities for $t = 2$ and $t = 3$. Differentiability pertains to the derivative of marginal utility, meaning the second derivative of utility. The second derivative is assumed to be negative, and because $u' < 0$ throughout, u'' will not be constant: Quadratic utility has a constant second derivative, but quadratic utility is ruled out by the assumption that

$u' < 0$ throughout. It follows that the nonconstancy of u'' means that the right and left derivatives of V_P' at $\bar{c}_2$ are not generally equal to each other.

At $q_1 = \tilde{c}_2$, Figure 11 portrays the value of V_C' jumping up above the value of V_P'. To understand how this inequality can come about, observe that when $q_1^* = q_1^L$, $c_1^* > c_1^1$, but when $q_1^* = q_1^L$, $c_1^* < c_1^1$. This situation provides us with the case required to establish the theorem.[128]

It remains to demonstrate that V_C' can lie above V_P' in a region of the discontinuity point $\tilde{c}_2$ of V_C'. Using (70), (71), (72), and the definition of V_C, obtain the following equation for V_C'':

$$V_C''(q_1) = \frac{u'(c_2)(\tilde{y}_2/y_1)\tilde{\lambda}_2 u''(c_3)}{\tilde{\lambda}_1(\tilde{\lambda}_2 u''(c_2)) + u''(c_3)} \tag{75}$$

At the point of discontinuity $q_1 = \tilde{c}_2$, we have $u'(\tilde{c}_2) = u'(0)/\tilde{\lambda}_2$. Therefore equation (75) for V_1'' becomes:

$$V_C''(\tilde{c}_2) = \frac{u'(0)}{\tilde{\lambda}_2\tilde{\lambda}_1} \frac{(\tilde{\lambda}_2/\tilde{\lambda}_1)\tilde{\lambda}_2 u''(\tilde{c}_2) + u''(0)}{\tilde{\lambda}_2 u''(c(\tilde{c}_2)) + u''(0)} \tag{76}$$

Let $c_1^1(q_1)$ be the optimal precommited value of c_2 as a function of the date 1 bequest q_1. Observe that:

$$V_P'(q_1) = u'(c_2(q_1))/\tilde{\lambda}_1.$$

When $q_1 = \bar{c}_2$, $\tilde{\lambda}_1 u'(c_2) = u'(0)$. Recall that when q_1 lies between $\bar{c}_2$ and $\tilde{c}_2$, $c_2^1(q_1) > \bar{c}_2$. Since u is concave, $u'(0) > \tilde{\lambda}_1 u'(c_2^1(q_1))$ in this interval of q_1. It follows that:

$$u'(0)/\tilde{\lambda}_1^2 > u'(c_2^1(q_1))/\tilde{\lambda}_1 = V_P'(q_1).$$

Observe that when $\tilde{\lambda}_1 = \tilde{\lambda}_2$, the value of (76) reduces to $u'(0)/\tilde{\lambda}_1^2$. Since the right hand side of (75) is continuous in $\tilde{\lambda}_2$ and q_1, the preceding paragraph implies that we can find a value of $\tilde{\lambda}_2 > \tilde{\lambda}_1$ such that $V_C'(q_1) > V_P'(q_1)$ in a region of $\tilde{c}_2$. **QED**

[128] The reason why V_C' can feature a gap, while V_P' does not is that in the case of exponential discounting, the decisions at every date coincide with the optimal plan under date 1 preferences. The gap in Figures 3 and 4 occurs because date 2 preferences allocate bequests from date 1 differently than what date 1 preferences would favor.

PROOF OF THEOREM 4: At date t, the individual observes the history c^{t-1} defined as $(c_1, \cdots, c_{t-1})$, and chooses c_t. In turn, this produces the history $c^t = (c^{t-1}, c_t)$. Define $X^t(c^t)$ as the set of full $T-$length consumption paths c which can emerge, given history c^t and subsequent optimizing behavior. Because of indifference, $X^t(c^t)$ need not be a singleton. Observe that for given c^{t-1}, each choice of c_t in $[0, q_{t-1}]$ has an associated $X^t(c^t)$. Define Y^t as $\cup X^t(c^t)$, where the union is over all c^t such that $q_t \geq 0$.

The claim is that both $X^t(c^t)$ and Y^t are compact and nonempty. The argument for this claim is based upon backward induction. For $t = T$, the case is clear. $X^T(c^T)$ can only be c^T, and Y^T must be the set of all budget feasible consumption paths. Next proceed inductively from $t+1$ to t, assuming that X^{t+1} and Y^{t+1} are both nonempty and compact.

Since $X^{t+1}(c^{t+1})$ is compact, its elements can be ordered by date $t+1$ preferences. Define $W^{t+1}(c^{t+1})$ as the worst consumption paths in $X^{t+1}(c^{t+1})$. These are the worst consumption paths which might result after c_{t+1} has been chosen at date $t+1$. With c^t fixed, vary c_{t+1} to find the best consumption path in the closure of the union $\cup W^{t+1}(c^{t+1})$. Call this the maximin set $B^{t+1}(c^t)$.

Formally, $X^t(c^t)$ may be defined as the set of consumption paths faced by the date $t+1$ self which are at least as good as those in $B^{t+1}(c_t)$. Since the upper contour set to $B^{t+1}(c_t)$ and $\cup X^{t+1}(c^{t+1})$ are both closed and the intersection is nonempty, it follows that $X^t(c^t)$ is nonempty; and being a closed subset of the compact set $[0, q]^T$, it is also compact.

Clearly Y^t is nonempty. To see that it is also compact, consider a sequence $\{{}^n c\}$, where ${}^n c \in X^t({}^n c^t)$, such that ${}^n c$ converges to $\bar{c}$. The claim is that $\bar{c}$ belongs to Y^t. Recall that ${}^n c \in X^t({}^n c^t)$ implies that date $t+1$ preferences rank ${}^n c$ at least as highly as the elements in the union of all worst point sets which could be generated by the budget feasible choices at date $t+1$. Therefore date $t+1$ preferences rank $\bar{c}$ at least as highly as the elements in the set of cluster points attached to sequences selected from the worst point sets $W^{t+1}(\bar{c}^t, c_{t+1})$ where c_{t+1} is budget feasible. Recall that budget feasible means that $c_{t+1} \in [0, q_t]$. It follows that $\bar{c}$ is at least as good as any maximin path in $B^{t+1}(\bar{c}^t)$.

Hence $\bar{c}$ belongs to Y^t, which establishes compactness.

Apply the induction process to obtain X^1, which is compact and nonempty. The claim is that any member c^* of X^0 is a consumption path for some Strotz–Pollak equilibrium. To prove the claim, we need to construct reaction functions which constitute a perfect Nash equilibrium for the game. Let the reaction function for date 1 specify the choice c_1^*. Now assign each budget feasible choice of c_1 to a full path $s^1(c_1)$ in Y^1, such that s^1 maps c_1^* to c^* and all other budget feasible elements to the dominated paths $W^1(c_1)$.

Define the reaction function for the date 2 self as the second component of s^1. This means that if the date 1 self chooses c_1^*, then the date 2 self responds with c_2^*. However, suppose that the date 1 self chooses something other than c_1^*. Then if the date 2 self can both maximize its own utility, and simultaneously punish the date 1 self, it will do so. In other words, if because of indifference, the date 2 self has multiple optimal responses to the date 1 choice $c_1 \neq c_1^*$, then it will choose an optimal response which punishes date 1.

Proceeding inductively, we obtain a sequence of reaction functions which together form a perfect Nash equilibrium. The strategy is Nash because given these reaction functions, at no date t does the individual perceive himself to benefit by deviating from the strategy in question. If he does, he anticipates that his future selves will punish him if they can by choosing a path in $W^t(c^t)$. The equilibrium is subgame perfect because the X^t sets were constructed using backward recursion. **QED**

Proof of Theorem 5: We provide the proof for the case $T = 3$. The general case easily follows. Recall that $q_1 = q - c_1$ and observe that along a consistent path, c_1 is selected to maximize:

$$u(c_1) + \lambda_1(2)u(c_2(q_1)) + \lambda_1(3)u(c_3(q_1)), \tag{77}$$

so that

$$u'(c_1) = \lambda_1(2)u'(c_2)\frac{dc_2}{dq_1} + \lambda_1(3)u'(c_3)\frac{dc_3}{dq_1} \tag{78}$$

along c^*.

Now (73) and (78) imply that:

$$\lambda_1(2)u'(c_2^*) < u'(c_1^*) < \lambda_1(3)u'(c_3^*) \tag{79}$$

Recall that $c_3 = q_1 - c_2$. Hence, $dc_3/dq_1 = 1 - dc_2/dq_1$. Thus, (78) becomes

$$(dc_2/dq_1)[\lambda_1(2)u'(c_2) - \lambda_1(3)u'(c_3)] + \lambda_1(3)u'(c_3)$$

The term in square brackets is negative, by inequality (73) and the definition of $\tilde{\lambda}_1$. This yields the left-hand inequality in (79).

The second inequality above implies that a transfer from date 1 to date 3 is regarded as an improvement from the perspective of the date 1 utility function. This transfer would also be regarded as an improvement as judged by the utility functions U_2 and U_3. This establishes that c^* is Pareto–inefficient. Finally notice that there is some sufficiently small transfer from date 2 to date 3 which, when combined with the transfer from date 2 to date 3, leaves the individual better off from the perspective of U_2. **QED**

Proof of Theorem 6: Begin by considering the case of hyperbolic discounting with no endogenous appetite-arousal. The first part of the proof focuses on a comparison of T_P and T_C.

1. The first order condition associated with the precommitment plan is:

$$u'(c_1^P) \geq \lambda_1^P(t)u'(c_t^P) \tag{80}$$

   for all t with strict inequality holding if $c_t = 0$. Let c_1^P be associated with an optimal solution. Because the discount factors $\{\lambda_1^P(t)\}$ decline monotonically to zero, there is some finite value of t for which (80) holds, and moreover, the strict inequality holds for all $\tau > t$, but for no $\tau < t$. This establishes that T_P is finite.

2. Consider the case when $c^P_{T_P+1} = 0$ and

$$u'(c^P_1) = \lambda^P_1(t)u'(c^P_{T_P+1})$$

In this case, the precommitment plan features a tipping point such that a marginal increase in q leads the cake to be consumed in positive quantity at $T_P + 1$.

3. Suppose that $T_P = 1$. Then the $t = 1$ marginal utility comparison of consumption levels at $t = 1$ and $t = 2$ is the same for both the precommitment plan and the consistent plan, because $\lambda_2()$ plays no role in the decision criterion at $t = 1$. Hence $T_C = 1$. The same argument implies that if $T_P = 2$, then $T_C = 2$.

4. The proof of Theorem 3 provides the structure for completing the proof of this theorem. Recall that Theorem 3 focuses on the case $T = 3$ with hyperbolic discounting, but not endogenous appetite arousal. Notably, the proof of Theorem 3 features the case $T_P = 3$. Figure 4 portrays two marginal utility functions, V'_C for the consistent plan and V'_P for the precommitment plan. Notice that at a tipping point for V'_P, V'_P features a kink, while at a tipping point for V'_C, V'_C features a discontinuity. Because of this discontinuity, there can be two local maxima for the consistent plan: one features positive consumption at all three dates, while the other features positive consumption only at $t = 1$ and $t = 2$. In the first local maxima, $T_C = 3$, while in the second local maxima, $T_C = 2$. There is only one maximum corresponding to V'_P, and it features $T_P = 3$. Therefore, in this case, $T_P \geq T_C$.

5. The consistent plan must take into account that the cake is appetite-arousing. This means that at date t, the maximization needs to account for the fact that an increase in c_t will lead to higher future discount rates which lower the present value of consumption, for any fixed consumption profile. In respect to Figure 11, date 1 preferences need to account for the fact that an increase in c_1 will alter the manner in which date 2 preferences allocate the bequest from $t = 1$, to the right of the discontinuity. The reallocation will favor $t = 2$ over $t = 3$, and in consequence,

the marginal utility curve V_C' associated with the bequest will rise, to the right of the discontinuity, relative to the exogenous case (with the same value of $\tilde{\lambda}_2$). This implies that endogenous appetite arousal involves an offset to exogenous hyperbolic discounting. Relative to the exogenous case, endogenous This will shift the bequest q_1^U to the right and increase the area of triangle B. If an increase in c_1 increases the area of triangle B from being less than the area of triangle A to being greater, then T_C will increase from 2 to 3. Therefore, endogenous appetite arousal reduces the tendency for $T_P > T_C$, relative to $[\lambda_t(\tau)]$ being exogenous.

6. The issues discussed in the previous point, about the impact of endogenous appetite-arousal on allocation, can be seen formally in the case when $T_C = 3$. It follows from (71) and (72) that:

$$V_C'(q_1) = (\lambda_1(2)u'(c_2^*)u''(q_1 - c_2^*)$$

$$+\lambda_1(3)u'(q_1 - c_2^*)u''(c_2^*)\tilde{\lambda}_2(c_1^*)$$

$$-\frac{[\lambda_1(3)u'(q_1 - c_2^*) - \lambda_1(2)u'(c_2^*)][y_2'(c_1^*u'(c_2^*)]}{[\tilde{\lambda}_2(c_1)u''(c_2^*) + u'(q_1 - c_2^*)]} \tag{81}$$

Notice that this equation is the sum of three terms. By (73), the third term being subtracted in the numerator is positive in the endogenous case, but zero in the exogenous case. Since the denominator is negative, endogeneity serves to make $V_C'(q_1)$ larger than it would be in the case where y_t is exogenous. Ceteris paribus, this makes the bequest q_1 more attractive from the perspective of date 1 preferences, and therefore consumption c_1 less attractive.

7. The argument for $T = 3$ places a finite restriction on both T_P and T_C. This restriction only has an impact when q is sufficiently large, meaning that augmenting the problem to $T = 4$ would lead to positive consumption at $t = 4$. Therefore, to make the argument for $T = 3$ general, for fixed q, let T be sufficiently large that point 1 implies that consumption will be zero at $t = T$. Consider Figure 11 once

again, which features three marginal utility functions: one for $t = 1$, one for V_P, and one for V_C. The marginal utility function for $t = 1$ does not depend on whether $T > 3$. The other two functions do. Specifically, at any tipping point, where the Euler condition associated with zero consumption at $T_P + 1$ or $T_C + 1$ holds with equality, rather than inequality, V_P' will feature a new kink, and V_C' will feature a new discontinuity. That is, these two marginal utility functions will feature multiple kinks and discontinuities respectively. The location of the maxima will still feature intersections of the latter two marginal utility functions with the marginal utility function for $t = 1$. Therefore, the calculus underlying the maximization for the case $T = 3$, as depicted in Figure 11, carries over to general T. In respect to T_P and T_C, the key issue in Figure 411relates to

$$\bar{c}_t < \tilde{c}_t$$

which states that the value of the minimum bequest q_1 which would produce a tipping point at date t under the precommitment plan is less than the corresponding bequest associated with the consistent plan; and this feature follows from the assumption of hyperbolic discounting.

8. The variable $\tilde{\lambda}_2'$ in (81) represents the strength of appetite-arousal. If $\tilde{\lambda}_2'$ is sufficiently small, then its impact on the marginal utility function V_C' will not interfere with the inequality $T_P \geq T_C$.

9. In respect to the comparison between T_P and T_N, consider point 7 above (in the proof). The aspect of Figure 11 which is germane to this comparison involves the structure of V_P', which is concave, with kinks at tipping points. Consider what happens to the marginal utility function V_P' when the rate of time discount increases. In this respect, keep in mind that date 1 preferences are exponential. An increase in the rate of time preference reduces the marginal utility $\lambda_1(\tau)u'(c_\tau)$ at every c_τ. Hence, an increase in the discount rate, which reduces the discount factor, lowers the marginal utility function $V_P'(q_1)$. It follows from (70), (71), and (72) that to the

right of tipping points, the higher discount rate also makes V_P' steeper, because the ability to buffer a decline in future marginal utility by allocating a bequest across multiple dates is less impactful when utility from those future states are discounted more heavily.

10. For fixed q, the intersection between V_P' and $\lambda_1 u'(c_1)$ provides the maximizing value of c_1. An increase in the discount rate leaves the latter curve unchanged but reduces the former curve, as described in point 9. In consequence, the intersection point moves to the left, implying that the value of c_1 is higher as a result. In turn, this implies that the ratio c_1/q increases with the discount rate.

11. When the discount rate for date 2 preferences is higher than for date 1 preferences, point 10 implies that a naïve individual will observe that at $t = 2$, the value of the ratio c_2/q_1 will be higher than its anticipated value at $t = 1$. The cake will be consumed more quickly than the individual had earlier anticipated. Because the cake is consumed in finite time, this implies that $T_N \leq T_P$.

12. The argument just made is not impacted by whether the degree of appetite-arousal is endogenous or exogenous. All that matters is whether the value of past cumulative consumption leads to higher future discount rates. **QED**

Proof of Theorem 7:

1. For (29), when $c_t \geq 1/\theta_t$, $Z_t(c_t, \theta_t, m_{i,t})$ is quadratic in c_t, for fixed θ_t and $m_{i,t}$. For $\zeta_{1,i,t}$ and $\zeta_{2,i,t}$ strictly positive, Z_t achieves a maximum for c_t at

$$c_t^* = \frac{\zeta_{1,i,t}}{2\zeta_{2,i,t}}$$

This equation, together with (30) and (31), imply statement 2 of the theorem.

2. Statement 3 of the theorem is based on the equation for c_t^* and (32), as the latter can be written as:

$$\zeta_{0,i,t}+\zeta_{1,i,t}(\frac{1+\theta_t}{\theta_t})-\zeta_{2,i,t}(\frac{1+\theta_t}{\theta_t})^2 = ln(\frac{1+\theta_t)}{\theta_t})-[(\varphi_{1,i,t}(\varphi_{2,i,t}(m_{i,t}-(\frac{1+\theta_t}{\theta_t})))^{\varphi_{3,i,t}})+\varphi_{4,i,t}]$$

The left-hand-side of this equation is the value of the quadratic portion of Z_t, evaluated at c_t^*, for fixed θ_t and $m_{i,t}$. The right-hand-side provides the hedonic cost of employing willpower, expressed as the reduction in utility from $ln(c_t^*)$. Statement 3 of the theorem will follow, once statement 1 is established, as there is a need to ensure that c_t^* maximizes the reduced form function Z_t^R, not just its quadratic portion.

3. The continuity portion of statement 1 of the theorem follows from (33), which implies

$$\zeta_{0,i,t} + \zeta_{1,i,t}(\frac{1}{\theta_t}) - \zeta_{2,i,t}(\frac{1}{\theta_t})^2 = ln(\frac{1}{\theta_t}) - \varphi_{0,i,t}$$

This equation establishes that the right-hand quadratic component of Z_t and left-hand logarithmic component are equal at $c_t = 1/\theta_t$. The differentiability condition follows by establishing that the first derivatives of the two components are also equal at $c_t = 1/\theta_t$. the first derivative of the quadratic component is given by

$$\zeta_{1,i,t} - 2\zeta_{2,i,t}(\frac{1}{\theta_t})$$

and the first derivative of the logarithmic component is given by

$$\frac{1}{\theta_t}$$

Use the fact that

$$\frac{1}{\theta_t} = c_t^* - 1$$

along with equations (30) and (31) to establish that the values of the two derivatives are equal at $c_t^* - 1$. This establishes statement 1 of the theorem. Statement 3 of the theorem follows because the logarithmic function is monotone increasing, as is the quadratic component for $c_t < (1+\theta_t)/\theta_t$.

4. Statements 4 and 7 are established through mechanical computation. The associated inequalities in statement 4 are easily verified by inspection of (32) and (33). The inequality in (33) holds because $E_{i,t}$ is computed by subtracting an increasing nonnegative function from $ln(c_t)$, which is zero at the upper boundary of the mental account range.

5. To establish statement 5, observe that

$$\frac{\partial Z_t^R(c, c_t(\theta_t, m_{i,t}), m_{i,t})}{\partial c_t} = \frac{\partial Z_t^R}{\partial c} + \frac{\partial Z_t^R}{\partial \theta_t}\frac{\partial \theta_t}{\partial c_t}$$

Statement 5 follows because when willpower is employed, $\frac{\partial Z_t^R}{\partial c} = 0$.

6. Statement 6 follows from statement 5 and (34), which implies that

$$\frac{\partial \theta_t}{\partial c_t} = -\frac{1}{(c_t+1)^2} < 0$$

Statements 8 and 9 follow from statement 4. Statement 10 of the theorem can be verified by computation. **QED**

Proof of Theorem 8: 1. If $\varphi_{1,i,t} = \varphi_{4,i,t} = 0$ for all i and t, then willpower is costless. With no external liquidity constraints, the marginal propensity to consume out of income and marginal propensity to consume out of wealth will be equal. The rest of the proof focuses on the case where $\varphi_{j,i,t} > 0$ for all i and $j \leq 3$.

2. Consider the case in which $c_t < y_t$, prior to the marginal increase. If the individual were to maintain consumption at c_t, then the increase in income will lead to an increase in the temptation level $m_{I,t}$. In turn, this will increase marginal utility $\frac{\partial U}{\partial Z_t^R}\frac{\partial Z_t^R}{\partial c_t}$ at the original value of c_t, but have no impact on marginal utilities at other dates (for the original solution). The corresponding increase in wealth does not increase the temptation vari-

able $m_{I,t}$, and therefore leaves marginal utility $\frac{\partial U}{\partial Z_t^R}\frac{\partial Z_t^R}{\partial c_t}$ unchanged at the original solution.

However, the increase in y_τ increases temptation associated with the income account at date τ; and an increase in y_τ for $\tau > t$ gives rise to an increase in wealth at t. If the original planner solution features $c_\tau < y_\tau$, then marginal utility $\partial Z_\tau^R/\partial c_\tau$ at c_τ will increase, because of the higher temptation level. Hence, for this situation, the roles of t and τ are reversed, and so the increase in wealth leads to an increase in marginal utility at τ but not at t.

In the revised solution involving consumption responses to the two respective marginal increases, a marginal increase in income at t favors consumption at t over τ, whereas a marginal increase in wealth (at t) does the reverse. Hence, in this case, the marginal propensity to consume from income exceeds the marginal propensity to consume from wealth.

Suppose that the original planner solution calls for the asset account to be invaded at some date between t and τ. In this case, the increase in wealth will lead to an increase in temptation at the intermediate date. Increased temptation generates higher marginal utility of consumption from the asset account, but not temptation from the date t income account. Therefore, the planner will respond to the revised marginal utilities by favoring consumption at the intermediate date over consumption at t. That is, at date t, the marginal propensity to consume from wealth is muted, relative to the marginal propensity to consume from income, because the increase in wealth does not impact the temptation variable for the income account at date t.

3. Consider the case in which $c_t = y_t$, prior to the marginal increase. If the condition featuring equal planner marginal utility holds across all dates, then the same argument applies as above. If not, then at the optimum, the marginal planner utility for t will exceed the marginal utility at some other dates. Suppose that marginal utility is maxi-

mized at date t, and the planner finds it suboptimal to invade the asset account. Then the planner will choose to consume the entire marginal increase in y_t. For a marginal increase in wealth, not income, the planner will still prefer not to invade the asset account, as doing so is more costly for any given consumption level associated with consuming from the asset account. Therefore, in this case, at date t, the marginal propensity to consume out of wealth is zero while the marginal propensity to consume from income is unity.

If at the optimal planner solution, marginal utility at some date τ exceeds marginal utility at date t, then the planner will find it too costly to transfer purchasing power from t to τ. Therefore, in this case, the planner will prefer that the date t-doer consume the entire marginal increase in income. A marginal increase in wealth, not income, does not impact marginal utility from the income account, and will not induce the planner to invade the asset account. In this respect, the incremental wealth reduces the utility of consumption in the asset account, for each level of c_t. Therefore, in this case, the marginal propensity to consume from wealth at date t is zero.

4. Consider the case in which the planner chooses to invade the asset account at $t = 1$. The decision to invade the asset account involves a comparison of two local optima. The first local optimum is associated with non-invasion, and the second with invasion. Non-invasion features consumption at $t = 1$ being lower ($c_{t,I} = y_t, c_{t,A} = 0$) than with invasion ($c_{t,I} = y_t, c_{t,A} > 0$). Total planner utility is the sum

$$U = \sum_{\tau=1}^{T} \lambda(t) Z_{\tau}^{R} = Z_{1}^{R} + \sum_{\tau=2}^{T} \lambda(t) Z_{\tau}^{R} \tag{82}$$

Here the first term on the right-hand-side is higher for the invasion-local optimum than the non-invasion optimum, while the reverse is true for the second term. The second term reflects total consumption of $W - y_1$ for non-invasion, but the lower value of $W - y_1 - c_{1,A}$ for invasion. Suppose that the planner is indifferent between the two local optima, so that the difference in first terms is equal and opposite to the difference in second terms.

The key issue is how marginal increases in income and wealth impact the decision about whether or not to invade the asset account. With non-invasion, a marginal increase in income at $t = 1$ will be consumed, generating marginal (logarithmic) utility of $1/y_1$, and not impacting the value of the second term in (82). With invasion, and an internal solution, a marginal increase in income will be partly consumed at $t = 1$ and partly consumed after $t = 1$. At the original value of $c_{t,A}$, the associated marginal value of willpower will be unchanged because the amount of saving will be unchanged, even as total consumption rises. This is because $m_{t,A}$ increases by the marginal increase in income; and the marginal utility associated with Z_t^R corresponds to the marginal cost of willpower. This implies that the difference in first terms of (82), for invasion minus non-invasion, decreases. The impact on the difference in second terms of (82) reflects the transfer of a portion of the marginal increase in income to dates after $t = 1$. As a result, a marginal increase in income favors the non-invasion choice over the invasion choice. At a point of indifference between the two, the planner will respond to marginal income at $t = 1$ by shifting from invasion to non-invasion. Doing so will reduce total consumption, as $c_{t,A}$ falls from a positive value to zero.

A marginal increase in wealth, instead of income, operates on temptation in the wealth account, but not the income account. If the planner were to choose non-invasion, then there will be no impact on consumption at $t = 1$ from marginal wealth: Temptation is the same, and income is the same.

If the planner were to choose invasion, then achieving the same total consumption at $t = 1$ as in the original solution requires greater willpower, as $m_{t,A}$ has increased. Therefore, $c_{t,A}$ will increase in response, to reflect the higher marginal utility at the original value. However, the increase in temptation level $m_{t,A}$ will be the same, regardless of whether the source of the higher temptation level is marginal income or marginal wealth. Therefore, in this case, at date t, the marginal propensity to consume from income and from wealth will be the same. In consequence the second term in (82) will decline, as total consumption after $t = 1$ declines.

In the situation just described, the individual will not consume the entire marginal increase, be it from income or from wealth. To see why, observe that doing so would involve no change in the values of marginal utilities, in order to preserve optimality. Now the value of the derivative of the penalty function would remain the same when consumption c_t increased by a marginal unit, because the temptation variable also increases by a marginal unit, and so the difference between the two remains the same. However, marginal logarithmic utility declines, and therefore so does marginal doer utility. This decline will lead the increase in c_t to be smaller than the marginal increase in temptation.
As for switching from invasion to non-invasion, as a result of the marginal increase in wealth or income, this will depend on the comparison of the two differences discussed above. The difference in first terms, for non-invasion minus invasion, is negative. The difference in second terms is positive. Together, these imply that a marginal increase in wealth favors maintaining invasion.

If the planner solution for invading the asset account occurs at the boundary, the argument is only strengthened, as a marginal increase in wealth will be consumed entirely at $t = 1$, just as the case with the non-invasion choice.

It follows that in the case of indifference, in which the choice is made to invade the asset account, a marginal increase in income induces a reduction in total consumption, while a marginal increase in wealth induces an increase in consumption.

5. The second statement of the theorem follows because increasing $\varphi_{1,A,t}$ increases the asset account entry fee to a level that negates the benefits of any positive consumption financed from that account. Equation (40) implies that the planner will not invade the asset account, at any date t, if $E_{i,t} > ln(m_{i,t}) - ln(m_{i-1,t})$; and a sufficiently high value of $\varphi_{1,A,t}$ will ensure that this inequality holds. **QED**