Introduction

George Loewenstein, Daniel Read,
and Roy F. Baumeister

Psychology and economics have a classic love-hate relationship. Members of each discipline often express positive sentiments about the other in the abstract, and acknowledge complementarities between disciplines in methods, subject matter, and levels of analysis. Yet actual encounters often produce glassy eyes or, worse, overt hostility. Both disciplines set out to use scientific method to explain and describe human behavior. They differ, however, in the details of their respective paradigms. Psychology is mainly empirically based, embracing a great variety of theories (frequently incompatible, as some economists are not shy about pointing out). Psychologists usually start from phenomena, develop a local theory based on that phenomenon, then test this theory against further observations. Economics, in contrast, is more theory-based, with a single theoretical approach that can be tailored to a wide range of applications. Psychologists and economists converge on at least one point: both wish that members of the other discipline would be more like them. Psychologists wish economists were less theory-driven and more responsive to empirical data. Economists accuse psychologists of being too willing to interpret quirky behavior as evidence of irrationality, and wish that psychologists would subject their theories to the discipline of formal modeling. Each group tends to view the other's theories as invalid or trivial, with very few evaluations falling between these extremes.

A dispassionate mediator might think that both sides have merit, but might also propose that the two disciplines find some way of sorting out their differences and agree on a common ground that combines both their strengths into a greater whole. Now, more than ever before, there are grounds for optimism that such a reconciliation is beginning to occur. This book is a record of what may be the greatest success story in the unfolding interdisciplinary relationship: the development of our knowledge of intertemporal choice and its application to important economic and psychological problems.

Intertemporal choice is what we do when we make trade-offs between costs and benefits occurring at different points in time. We are always making intertemporal choices—when we choose between a hamburger now or a fine meal (or thinner body) later; between increasing our pension fund contribution or going to Hawaii; or between a sinful moment on earth or an eternity in heaven. Indeed, so broad is the domain of intertemporal choice that it is difficult to think of a consequential decision that is *not* an intertemporal choice. Given the importance of intertemporal choice, that it is a central theme in both psychology and economics is not surprising.

Basic research on intertemporal choice by contemporary economists has revolved mainly around testing the validity and implications of the discounted utility (DU) model, which posits that people have a single unitary rate of time preference that they use to discount the value of delayed events. This means, for instance, that if something is worth 10 percent less to you if it is delayed by one year, it should be worth a further 10 percent less when delayed by a second year, and so on. DU is a *normative* model because if we don't treat the future in this fashion, we won't be able to make plans that we can be sure we will implement in the future. Whether DU is a valid *descriptive* model of behavior is a matter of debate among economists (and the central focus of the following chapter).

Basic research by psychologists on intertemporal choice (a term psychologists rarely use) has tended to focus on different issues. Some have been concerned with measuring individual differences in the propensity to delay gratification, others with situational determinants of impulsivity, and still others with cognitive and emotional mechanisms underlying intertemporal choice. The chapters in this volume represent a sort of academic "busing" intended to speed up the integration of psychology and economics in the domain of intertemporal choice. These were papers first presented during a meeting at the Russell Sage Foundation in New York, where distinguished researchers with common interests (who in many cases had not heard of one another) shared their work. For many of us the meeting was revelatory. We differed in our terminology, disciplinary assumptions, and research methods, but serious efforts were made to speak a common language. Whenever possible, the economists stripped their talks of equations—and when they did not, were called to task by the psychologists. The psychologists took great pains to express theories in terms that economists could digest. In the end, all left the meeting with a somewhat changed perspective and certainly with an enhanced understanding of diverse aspects of intertemporal choice. We can only hope that these chapters will have a similar effect on the reader.

Completion of this volume occurs almost exactly ten years after the publication of a previous compendium on intertemporal choice (also published by the Russell Sage Foundation) titled *Choice Over Time*. Despite substantial overlap in authorship and subject matter between the two volumes, the differences are striking and underline the pace of progress that occurred in the intervening decade. Much of the original volume (fully six of fifteen chapters) was occupied with discussions of *hyperbolic time discounting*—that is, the observation that people tend to be more impatient toward trade-offs involving earlier rewards than those involving later rewards. These discussions focused mainly on its theoretical development and implications. Hyperbolic time discounting continues to be a central theme in the current volume (see chapters 7, 10, 14, and 18), but most chapters herein are oriented toward examining its implications rather than documenting the phenomenon, as in the earlier volume. Moreover, a variety of other important themes have emerged. For example, there is much more discussion of intra-individual variability in time discounting—why people sometimes behave as if there is no tomorrow and at other times seem obsessively focused on the future (chapters 1, 5, 6, 7, 8, 11, and 12); more discussion of the role played by emotions (chapters 1, 5, 13, 14, and 18); and a plethora of intriguing new theoretical perspectives (chapters 5 through 9). The applications have become far more sophisticated both in terms of underlying theory and quality of empirical research (chapters 13 through 18). Moreover, there is even research questioning the robustness of hyperbolic time discounting (chapters 1 and 10).

On top of the appeal of a round number, the decade gap between the two volumes seems close to ideal from the standpoint of scientific progress. Only a few years earlier, it might have been difficult to fill a volume with new research sufficiently distinct from that published in the original volume. Given the rapid pace of research in the field, doing justice to the breadth of the topic a few years from now may well require two volumes instead of one. As it is, we have inevitably neglected to include much important and high quality research, as well as a number of theoretical perspectives.

The remainder of the introduction provides a map to the thematic concerns in the chapters that follow. Part I looks at some basic issues: the philosophy of intertemporal choice, how it evolved, and what has evolved. Part II describes some theoretical contributions around the question of what is intertemporal choice. Most of these chapters focus on the problem of self-control, an issue that looms large throughout the book. Part III turns to specific patterns of time preference, looking at what we know about how people value future outcomes. Part IV,

the final section, deals with applications in the domains of health, drug addiction, dieting, marketing, savings, and public policy. The categorization is necessarily crude. Many of the more theoretical chapters are heavily informed by applied findings, and the applied chapters all derive novel basic theoretical insights from the specific domains they explore. Chapter 1 presents an overview of research on time discounting and time preference. The chapter focuses on the descriptive validity of the DU model, which remains the dominant economic theory of intertemporal choice. Frederick, Loewenstein, and O'Donoghue observe that virtually every assumption and implication of the DU model has been contradicted by empirical research. Moreover, they argue that the model cannot be salvaged by merely assuming a different—hyperbolic, for example—discount function. Rather, they argue, understanding intertemporal choice behavior requires an account of several distinct motives that can vary greatly across decisions—a theme that resounds throughout the book.

While Frederick and colleagues focus on the descriptive question of how people do discount the future, chapter 2, by Frederick, addresses the normative question of whether we *should* discount the future at all. Many argue that there is no good reason to weight utility in the near future more than utility in the more distant future—at least no reason other than the uncertainty of obtaining that future utility. Yet some philosophers, in particular Derek Parfit, have argued that our relation to future selves is fundamentally the same as our relation to other distinct individuals, that our connection to future selves diminishes over time, and that to give lesser weight to selves with whom we are less connected is reasonable—just as we would care less about strangers than about close relations. Along with examining both sides of this issue, Frederick presents a descriptive study that assesses whether diminishing identity might not only justify time preference, but also help explain it. In chapter 3, Kacelnik discusses the evolutionary basis of time discounting, and specifically provides an evolutionary account of anomalous patterns of discounting behavior observed both in humans and animals. Experiments with animals have sometimes been interpreted as showing that, contrary to evolutionary theory, animals do not maximize expected rewards. Kacelnik argues that the findings that support such conclusions are an artifact of research in which animals are placed in artificial environments unlike the natural environments they have evolved to deal with. Kacelnik also devotes considerable attention to drawing linkages between experimental results dealing with humans and other animals, and cautions against the temptation of automatically concluding that they are the product of similar underlying mechanisms and evolu-

tionary forces. Kacelnik's chapter underlines the crucial role that the ability to optimally delay gratification must have played in evolution. Not surprisingly, therefore, as reviewed in chapter 4 by Manuck and colleagues, a great deal of evidence has emerged showing that patience is influenced by specific brain structures and chemical processes. One such brain structure is the prefrontal cortex. People with damage to this area of the brain or to its connections with other brain structures (such as those involved in emotional experiences and motivation) do not take into account the future consequences of their actions, and so choose based on immediate rewards only (Damasio 1994). The ability to withstand impulses is also modulated by activity of the neurotransmitter serotonin. In laboratory animals, for example, administering drugs that reduce serotonin activity in the brain increases the likelihood of impulsive choice, whereas increasing serotonin activity (as by Prozac) makes animals more willing to wait. Variability in serotonergic activity also predicts variability in impulsive choices in humans, including suicide and aggression, as well as scores on a personality test of impulsivity. Manuck and colleagues also discuss genetic variation in a component of the serotonin system that is correlated, in humans, with both impulsivity and brain serotonergic activity.

The ability to wait is not only correlated with brain activity, but also with age and dispositional person variables. Mischel, Ayduk, and Mendoza-Denton describe Mischel's studies of delay ability using the classic *preschool delay of gratification paradigm*, which measures the time children wait for a preferred but delayed reward over a less desirable but immediately available reward. As children mature they become increasingly able to wait for longer periods, due to the development of cognitive-attentional strategies to prevent the immediate desirability of the small rewards from overwhelming their long-term interests. This development is seen as reflecting the maturation of a rational and far-sighted, cognitive, "cool" system (probably located in the frontal lobe and hippocampus) that gradually becomes able to moderate the impulses of the "hot" system (amygdala) and enables self-regulation. Independent of maturational level, stable individual differences also exist in the ability to access cool-system strategies in dealing with the frustration of the delay situation. This ability is predictive of social, emotional, and cognitive competencies throughout adolescence and adulthood. For instance, research showed that five-year-olds who resisted longer in the preschool delay paradigm also did better in college later in life. The authors summarize findings elucidating the attentional mechanisms that underlie this ability. They also highlight recent research exploring how delay ability protects against destructive behavior in interpersonal relationships.

Mischel and colleagues found that four-year-olds who waited longer in the preschool delay paradigm also did better in college later in life. We often speak of the ability to overcome temptation as involving the use of *willpower*. In chapter 6, Baumeister and Vohs report research suggesting that this term can be taken literally. They argue that people have a limited pool of resources that they can use to resist temptation. Successful resistance draws on this resource pool and makes people less able to resist subsequent (immediate) temptations. Indeed, they suggest that self-control is like a muscle that can be temporarily fatigued (and may also be able to be strengthened with judicious use). Baumeister and Vohs describe experiments in which people are required to exert willpower in one domain (resisting tempting food) and then called on to exert it in an unrelated domain (persistence on an unsolvable task). They find that there appears to be less willpower left at the second stage.

The contributions of Mischel, Ayduk, and Mendoza-Denton and Baumeister and Vohs focus on how difficult it can be to overcome the desire for immediate gratification. O'Donoghue and Rabin, in chapter 7, are concerned with the importance of people's awareness of their own self-control problems. They distinguish between two extreme states of awareness: at one extreme is sophistication, or full awareness of future self-control problems; at the other extreme is naivete, or full unawareness. O'Donoghue and Rabin show that to be sophisticated is sometimes better, and at times to be naive is better. A sophisticate will resist procrastinating too much because she knows she will have no more willpower tomorrow than today, while a naïf will (incorrectly) count on his future selves to do the job. Yet a naïf might persist despite inevitable (but unanticipated) future self-control failures, whereas a sophisticate, who predicts these failures, might never even try. Their chapter also describes the more realistic case of partial naivete: you know you will have trouble getting up when the alarm rings, but you underestimate how much. Indeed, one of the results of their analysis will be familiar to many readers—namely, that even a little naivete about the tendency to procrastinate can lead one to delay completing a task until the last minute.

O'Donoghue and Rabin discuss as well how our beliefs affect *when* we want to do things. Although their approach is different, this question is also at the heart of Liberman and Trope's *temporal construal theory* described in chapter 8. They argue that choice objects can be characterized in terms of their high-level and low-level features, and that our *construal*, or mental representations, of those objects will have more high-level features the more the object is delayed. Consequently, choices concerning delayed events will be based on higher-level con-

struals than will choices for immediate events. To illustrate, a high-level construal of writing a book chapter might be "building a career," while a low-level construal would be "sitting in front of the computer." The decision of whether to write a chapter in six months will be based on its effects on one's career, while the decision to work on it *tonight* will depend on whether typing is more fun than watching television. As Trope and Liberman show in numerous experiments, this seems to be precisely how we do decide. Theoretical approaches to self-control have largely focused on the conflict between immediate gratification and long-term objectives. In chapter 9, Prelec and Bodner note that many self-control problems also involve a problem of scale: success in the long run requires persistence and endurance over many smaller decisions, each of which has a negligible impact on the larger goal. To deal with this, they develop a self-signaling model of self-control, in which success "in the small" is a motivating signal of success in the long run. The model rests on a distinction between two types of reward: reward experienced directly from the consequences of a choice; and diagnostic reward, which is the moral pleasure or pain derived from learning something positive or negative about one's own disposition or future prospects. Anticipation of diagnostic reward or fear of diagnostic pain promotes self-control. In the model, whether a person is aware of the attempt to self-signal matters. Unawareness of self-signaling promotes good behavior and self-esteem. Awareness can trigger an excess of self-control, where good behavior is discounted for diagnostic motives, and being reasonably good is no longer "good enough." The fully aware self-signaler is on a self-control treadmill, engaged in perfect behavior with no commensurate improvement in self-esteem.

As previously mentioned, DU theory has been widely judged deficient. The most widely discussed deficiency—one that has led to the greatest degree of consensus between psychologists and economists—is the idea that the discount rate is not a single number but rather a function. As noted, a strong consensus has developed around the notion of hyperbolic time discounting. The experiments reported in Read's chapter 10 may therefore be met by some dismay by both economists and psychologists who, having finally acknowledged the limitations of exponential time discounting, were on the verge of throwing in their hat with hyperbolic discounting. These experiments show that much of the data adduced as evidence for a hyperbolic discount function are faulty. Read's studies show that the measured discount rate depends critically on the length of the interval being evaluated and may have little to do with when the interval occurs. This is called *subadditive intertemporal choice*, and Read questions the

strength of the evidence for hyperbolic time discounting once this subadditivity of time discounting is taken into account.

Although DU theory and its alternatives, such as hyperbolic discounting, are often studied with experiments comparing preferences for single outcomes, in practical terms their most important application is to sequences or streams of consequences, such as the value of a lifetime of paychecks or living with a chronic disease. According to DU theory, the value of such a sequence is the sum of the discounted values of each element in that sequence. These discounted values are assumed to be separable, meaning that the utility experienced at one period does not influence utility at a later period. Ariely and Carmon's review of the evidence (much from their own experimental research) in chapter 11 shows that this assumption is far from true. Rather, the value placed on a sequence of good or bad feelings is apparently based on the combination of a host of gestalt properties of that sequence. For instance, people weigh heavily on how good the end of a sequence is, its peak, and whether it is improving or getting worse (people like improvement). Ariely and Carmon's bottom line is that people encode or evaluate sequences based on abstract mental representations containing summary information (such as the mean, slope, and variance) about the pleasure of experiences, and not on representations of entire sequences. When people decide what *future* sequences they want to experience, they take these gestalt properties into account, thereby violating the assumption of separability. If the assumption of additivity—that the value of a sequence is the sum of the values of its parts—is unrealistic, an even more unrealistic (though all too common) assumption in modeling intertemporal choice is that tastes are fixed over time. Loewenstein and Angner, in chapter 12, underscore the artificiality of this assumption by enumerating some of the most important sources of preference change. They also review the burgeoning literature on predictions of future tastes, showing that, while many reasons exist for why people mispredict changes in their own tastes, these diverse reasons generally produce a common pattern of misestimation that Loewenstein, O'Donoghue, and Rabin (2002) call *projection bias*. Finally, Loewenstein and Angner challenge the universal applicability of a common assumption made in most research that does allow for changing preferences—namely, that people want to satisfy whatever preferences they expect to have in the future. This certainly is often true, but situations occur (which they discuss) in which people attempt to impose their current preferences on their own future selves.

One area in which time discounting plays an especially important role is that of decisions concerning health. For the consumer, virtually

all health-related decisions involve trade-offs between short-term and long-term gains. Taking medication, seeing the doctor, undergoing withdrawal—these are all behaviors that are good for us in the long term but can range from annoying to unbearable in the short term. Studies of discounting health outcomes are also theoretically important because health may be the only nonfungible domain of choice that we can study: whereas consumption of money is (in principle) independent of when it is received, health can *only* be consumed when it is received. An important question is whether qualitative findings reported in studies of monetary discounting are also found in studies of health; Chapman, in chapter 13, shows that some are and some aren't. For instance, people discount larger health effects at a lower rate than smaller ones, which is also true for monetary outcomes, but they sometimes would rather their health got worse rather than better (holding total health constant) when they usually want their income to increase (again, holding total income constant). A perhaps more important question is whether health discounting is correlated with discounting for other things; Chapman suggests it is not. This finding appears to challenge the common view that time discounting is a reliable individual difference.

One domain in which the future values of money and health, or at least health-related behavior, do seem to coincide is discussed by Bickel and Johnson in chapter 14. They find that addicts do indeed discount future money much more than do nonaddicts, suggesting that at least in this domain, discounting is related to health-related decisions. One might think that this is because people who discount the future more become addicts, but research suggests that perhaps the addiction creates the heavy discounting. The authors observe that heroin addicts and smokers discount the future heavily, but once they have fought off their addiction, their discount rates fall. Evidence that longtime ex-smokers discount the future the same way as nonsmokers also supports this notion.

Both Chapman's and Bickel and Johnson's chapters show how health-related behaviors are often quite shortsighted. Caplin's chapter 15 examines policy implications of this observation. He analyzes the common tendency to avoid learning about, or discovering if one has, a disease. Many people put off seeing the doctor because they fear what they might discover. For them, myopia arises not because they weigh too little on the future but because contemplating bad futures is frightening, and not making an appointment with the doctor is one way to avoid doing so. Many campaigns to increase the public take-up of testing for serious conditions (such as testicular cancer) use fear appeals that emphasize the consequences of the undiagnosed disease.

In line with current psychological theory, Caplin suggests that these fear appeals may have the paradoxical effect of making at-risk individuals less willing to undertake the preventive act. This is especially likely to happen in cases such as self-examination for testicular cancer, in which the preventive act itself triggers fear. In such cases the fear appeal may backfire by making the preventive act even more frightening, prompting greater efforts at avoidance. Caplin's work shows one way in which economic analyses can enhance psychological thinking: his model specifies circumstances under which fear appeals are likely to be effective and those when they are not. He also examines differences between economists and psychologists in how they judge the success of policy interventions, and proposes a hybrid approach that draws on both traditions.

Herman and Polivy, in chapter 16, provide an in-depth examination of a case of dynamic inconsistency that has actually served as the poster child for much writing about intertemporal choice. At the meeting in New York, almost everyone used failure to stick to a diet as the "perfect example" of dynamic inconsistency. The prototypical example is that we plan to eat healthily before dinner, but at dinnertime our good intentions fail and we take some frightfully fattening food. The cycle repeats itself dinner after dinner. Herman and Polivy, who have spent most of their careers studying dieters, pointed out that this is not exactly what happens; dieters do plan to diet when meals are still far away, but they do not necessarily fall prey when the moment for dessert arrives. Rather, dieters break down usually because some particular disruptive event (albeit often minor) disinhibits them enough to make them break their diet momentarily; then the fact that the diet has been broken often leads to a catastrophic binge. Herman and Polivy's skeptical conclusions are in line with many of the other chapters, pointing out that the road to understanding intertemporal choice is not through developing better discount functions but through understanding the variety of psychological processes that enter into future-based decision making. One such psychological process, investigated by Wertenbroch, is in the domain of consumer choice. In chapter 17, he starts with the well-established phenomenon that people seem to consume at a rate that is an increasing function of their immediately available (local) resources. What this means is intuitively compelling: we will eat more cheese if we have ten pounds in the fridge instead of five, and more if we have five pounds instead of one. Wertenbroch suggests that consumption decisions often are driven by lax local constraints (for example, what we have stockpiled in inventory) rather than by more stringent global constraints (such as long-term health concerns). Thus, for example, if we purchase beer by

the case (as is virtually mandated by Pennsylvania law), we are likely to consume more beer over the course of the month than if we buy by the six-pack. Wertenbroch presents wide-ranging evidence for the proposition that a large paycheck makes people overspend, a large box of chocolates makes people overeat, and the lack of an explicit deadline makes people squander time. He argues that consumers are partially sophisticated in dealing with this problem, and so respond by self-imposing rationing rules that tighten local constraints on consumption ("never borrow," "only $50 per month for clothes"). Such rationing strategies include, for example, self-imposing costly deadlines to prevent procrastination and *purchase quantity rationing* (buying small amounts at a time) to prevent overconsumption: many smokers, for instance, buy cigarettes one pack at a time rather than economizing by buying cartons. In chapter 18—the final chapter—the focus shifts from the individual to the entire population and asks, What are the consequences of a society of shortsighted decision makers? Angeletos, Laibson, Repetto, Tobacman, and Weinberg investigate the implications of hyperbolic discounting for predictions of national savings data. If consumers conform to the DU model, their rate of spending would be more or less even over their lifetime: early in life they would borrow on future income, then they would save in midlife and retire on their savings. In fact, consumers' spending seems to closely fit their current incomes. Early in life they spend little, in midlife they spend a lot, and they retire in relative poverty. Angeletos and colleagues conducted computer simulations of economies composed of either hyperbolic consumers (who spend too much too soon) or exponential consumers (who spend the right amount at the right time). The hyperbolic economies predicted real consumer behavior extraordinarily well.

As is clear from these synopses, the contributions to this book are remarkably wide ranging and comprehensive. The study of intertemporal choice is now a truly interdisciplinary project. We hope that you'll invest the time in familiarizing yourself with this exciting and active area of research. Please don't procrastinate; there's no time like the present.

References

Damasio, Antonio R. 1994. *Descartes' Error: Emotion, Reason, and the Human Brain.* New York: G.P. Putnam.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2000. "Projection Bias in the Prediction of Future Utility." Working paper, Department of Social Science and Decision Sciences, Carnegie Mellon University.