CHAPTER 13

A Neuroeconomist's Perspective on Thinking about the Future

Anna B. Konova
Paul W. Glimcher

What Would a Neuroeconomist Do?: Overview and Goals of This Chapter

Toward the end of the 1990s, a group of scientists from the disciplines of neuroscience, economics, and psychology came together, recognizing that they had a common goal: to understand how humans and other animals make decisions, express preferences, and trade off outcomes. Each discipline brings its unique contribution to this common goal. Speaking a bit too broadly, economics provides the theoretical framework for understanding choice, psychology the experimental tools for the study of choice behavior and constructs for capturing mental life, and neuroscience the measurement tools for elucidating the choice machinery at a reduced and mechanistic level. The hybrid approach that emerged from the conglomeration of each of these disciplines' contributions to the study of decision making is now known as *neuroeconomics*.

In this chapter, we describe how scientists who employ this neuroeconomic approach typically think about decisions that involve tradeoffs between outcomes that occur at different points in time (in the future vs. today, for example), also known as *intertemporal choices*. Our primary goal is to provide the reader with the context within which neuroeconomists think about this particular class of decision problems. To do this, we provide an introduction to more traditional (neoclassical) economic theories of choice as they influenced the development of neuroeconomics. We first review the history of modern notions of "utility" or "subjective value"—a subject-specific quantity associated with how valuable a given outcome is to a given individual—before turning our attention to a more detailed discussion of the intertemporal choice models commonly used to capture how time specifically influences subjective notions of value in *reduced-form* (mathematical–economic) theories. We then describe how these models are combined with neuroscientific tools by neuroeconomists to understand what happens in the brain as a person makes intertemporal choices. What happens in the brain, neuroeconomists typically argue, can in turn help us narrow down the set of plausible models of choice (be those models psychological or economic in origin), by excluding those which cannot be plausibly mapped to a neural hardware for implementation or inspire new models linked to biological constraints or features. We then summarize key empirical findings that show people *discount* (value less) outcomes set to occur in the future in highly predictable, albeit idiosyncratic, ways. We highlight a core set of brain regions whose coordinated activity is thought to underlie this behavior. We conclude by taking drug addiction as a case study of how a psychological pathology might reflect these brain and behavior interactions under conditions of ill health.

Neuroeconomic Tools

In the most general sense, a neuroeconomist: (1) observes people's *behavior*; (2) thinks about theory and tests which *models or classes of models of choice* explain this behavior (at various reductive levels of analysis); and (3) identifies what the *neurobiological processes* are that enable this behavior by linking the "best-fitting" model of choice to brain data. In this section, we review the central tools most neuroeconomists use to accomplish these goals before focusing on intertemporal choice as a representative area in which much progress has been made using the neuroeconomic approach.

Economic Models of Choice

Early Notions of Utility and What a "Neuroeconomic" Theory of Choice Might Look Like

To traditional economists, all of human behavior can be described as having the ultimate goal of maximizing a theoretical quantity called *utility*. Von Neumann and Morgenstern (1944) and Savage (1954) defined utility as the *implicit* value to a specific individual of any object or event in the outside world, be it the utility derived from wealth, love, social status, or another source. But it is important to remember that economists see this as a subjective experience. To make that clear, consider one of the earliest puzzles of the classical economic revolution of the 18th and 19th centuries: the *diamond–water paradox*. Water is essential for survival, and diamonds have limited use-value, but despite this fact, people are willing to pay far more for diamonds than for water. Why is the subjective value, or utility, of diamonds so high and the subjective value of water so low? The philosopher Adam Smith argued that this apparent contradiction stems from a person's subjective utility being fundamentally tied to scarcity. Because we are wealthy in water, Smith argued, we place very little value on a small increment in that "water-wealth," but the opposite is true of diamonds. The idea that the utility you experience from a single diamond or a single liter of water diminishes as you possess more and more (formally a *diminishing marginal utility*) was revolutionary and spurred what became known as the *marginal revolution.* Formal assumptions of the marginal revolution were that (1) people maximize utility as evidenced by their choices; (2) people experience utility from owning or consuming goods, and it is the behavior of those people which defines value (not the economist); and (3) the amount of utility experienced per unit of most goods follows a function that "diminishes at the margin," or slowly bends downward in slope (a curve relating happiness/satisfaction/utility—and hence price—to quantity; see Figure 13.1).

Interestingly, in some ways the paradox at the center of the marginal revolution echoed ideas from choice theory pioneered almost a century earlier by the probability theorist Blaise Pascal, who was working to understand people's decisions when gambling. He came to the notion of *expected value,* a quantity that is the product of the magnitude of a potential outcome (e.g., amount of money that could be won from a gamble) and the probability of that outcome occurring (e.g., the winning probability). Pascal argued that what one should do to maximize one's welfare (or net wealth) is to aim to maximize expected value. He proposed that, given two options, a chooser should compute each option's expected value and simply select the one with the higher of the two. But Pascal ignored a critical empirical observation that was important to the economists of the marginal revolution: There are significant individual differences in starting welfare (or "wealth"), which have implications for what one should do. What one should do depends on where one is on the curve—rich or poor in any particular quantity.

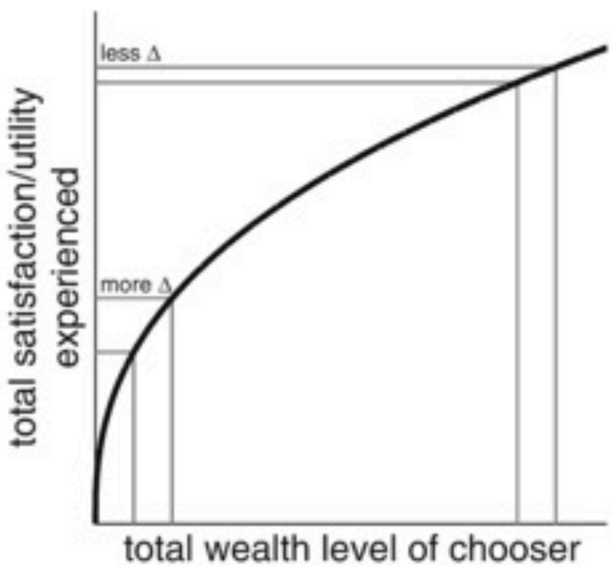

FIGURE 13.1. Diminishing marginal utility. During the marginal revolution, scholars working to develop a formal theory of choice realized that the "utility" a person experiences by owning things or by monetary gains depends both on the intrinsic value of those objects (which is itself very complex) and how much of them the person possesses. The function mapping objective value (or quantity) to utility is often assumed to be a power function with an exponent referred to as "alpha." Alpha captures individual differences in the rate of diminishing utility experienced for a given object by different people. Here we show such a function for money for an individual whose alpha takes on a value below 1. As the total number of dollars this chooser possesses increases, the amount of utility he or she experiences with each additional dollar increase diminishes. When he or she only has a couple of dollars, an additional dollar results in a greater increase in utility than when he or she has many dollars, as illustrated by the distance between the two sets of horizontal lines.

For a psychologist, reconciling the marginal revolution's idea of diminishing marginal utility with Pascal's expected value requires understanding the work of the mathematician Daniel Bernoulli, who suggested that, rather than maximizing expected value, choosers should be seen as maximizing instead a more subjective quantity: expected utility (a transformation of expected value). Expected utility differs critically from expected value because, unlike expected value, it is a more psychological and subjective, rather than an objective, measure. Bernoulli argued that the satisfaction or utility derived from a given unit of a good x is a nonlinear function of how much that unit of good x increases a chooser's total welfare, as determined by $u(x)$. For a chooser, which course of action offers higher utility (e.g., taking \$7,000 for sure or a 50% chance of \$20,000) depends on (1) what that chooser's $u(x)$ function looks like and (2) his or her total wealth. Bernoulli assumed that $u(x)$ was a logarithmic function for everyone, but it is now accepted that the form of $u(x)$ is more complicated and can be represented in many ways. For our purposes, we employ in this discussion a power function with an exponent (typically referred to as "alpha") that differs across individuals and determines the rate of diminishing utility for good x. That is, for Pascal:

$$\text{Utility} = \text{Expected Value and } (0.5 \times \$20{,}000) > (1 \times \$7{,}000).$$

The chooser should go for the gamble.

But for Bernoulli and modern utility theorists, for whom utility follows a nonlinear monotonic function,

$$\text{Utility} = \log(\$) \text{ or } \$^{\text{alpha}} \text{ and } [0.5 \times \log(\$20{,}000)] < [1 \times \log(\$7{,}000)]$$
$$\text{or } (0.5 \times \$20{,}000^{0.5}) < (1 \times \$7{,}000^{0.5}).$$

The chooser should go for the safe bet.

Notice that for Bernoulli all of the subjectivity of a chooser is embedded in the utility function. He assumed that there was no subjectivity in the perception of *probability*. It was not until the work of the economist Leonard Savage in the 1950s that the notion of "subjective probability" was introduced, a point that we return to later.

Revealed Preference

The end of the 19th century brought on a new revolution in economics—the ordinal revolution—which was spearheaded by the Italian economist Vilfredo Pareto. Although Bernoulli's and the marginal revolution's notions of utility significantly advanced our understanding of what a theory of choice might look like, both critically assumed that utility can be *measured* as a specific numerical function that allowed one to identify, for example, that one option was exactly twice as good as another. Pareto argued that such precise quantitative comparisons could not be made based on data about simple choice but, rather, argued that only statements about one option being better than another (not better by a certain fraction) could be made based on observations about what people prefer. He showed, through a famous mathematical proof, that asking people which of two options they preferred could only reveal a *rank order* for goods in terms of utility, never a numerical (or metric) distance. Take, for example, a chooser who prefers apples to oranges and oranges to grapes. Pareto showed that, for this chooser, one could represent her utilities as u(apples) = 3, u(oranges) = 2, and u(grapes) = 1, or that the utilities she assigned to these fruit could be any other ordered set of numbers that preserved this ranking and thus predicted the observed choice ordering! Pareto's demonstration that a unique value for utility could not be derived from simple choice experiments led to one of the most important tenets of early 20th-century economics: that utility must be considered a mathematically *ordinal* (just rank ordered) and not a cardinal (a truly numerical) quantity.

Recognizing the problem with measuring utility directly, the American economist Paul Samuelson proposed an advance to Pareto's approach that preserved most of its features while providing new tools for thinking about choice. The question Samuelson sought to answer was whether the choices we might observe in a given experiment could be proven to be inconsistent with all possible utility functions—whether there is a pattern of choice which fundamentally disproves utility theory. He sought to develop a simple way of describing patterns of choices that were fundamentally incompatible with the idea that the chooser assigned even simple rank-order utilities to options such as apples and oranges. To accomplish this, he developed a formal test for this kind of consistency, the weak axiom of revealed preference (WARP; now often confusingly called *rationality* [Samuelson, 1938]), a condition that must be satisfied by any pattern of choice that is consistent with maximizing utility—for any possible utility function (though with some minor restrictions on the craziness of the function). Simply put, Samuelson showed that, for WARP to be satisfied, the following must be true: (1) if A is chosen when both A and B are available, then A is said to be preferred to B ($A > B$, strict preference) or is at least as good as B ($A \gtrsim B$, weak preference); but (2) it cannot be the case that B is clearly preferred to A ($B > A$). This may seem trivial, but with just this assumption and a bit of math, Samuelson (1938) was able to prove that an individual who chooses apples over oranges and oranges over pears cannot also *prefer* pears over apples—if she or he has a stable underlying utility representation. It is worth noting that, although we can say that an individual obeying WARP in his or her choices shows behavior consistent with maximization of some utility function, because of WARP's ordinal nature, we cannot say anything about the shape of that utility function.

Extending Samuelson's (1938) approach further, Hendrik Houthakker (1950) developed what is now known as the generalized axiom of revealed preference (GARP). Whereas obeying WARP is a *necessary* condition for maximizing and representing utility, Houthakker sought to identify what might be a *necessary* and *sufficient* condition for a utility representation—the kind of choice behavior that indicates unambiguously that the individual was behaving *as if* she or he really did have a utility representation of some kind. What Houthakker did with GARP was to consider how an individual's choices order when several options are available with a subtle assumption about how increasing the quantity of a desirable good increased the utility of that good—a test to determine whether an individual's choices necessarily and sufficiently obey *transitivity* (Figure 13.2). A chooser who obeys GARP is one who is transitive and, roughly speaking, monotonic (he or she obeys the general rule: more is better, or at least not worse) in his or her preferences. Obeying GARP, in turn, means that there exists a monotonic utility function (or, more accurately, a family of such functions) that this chooser behaves as if he or she is maximizing. We *still* do not know from this test what that function(s) looks like, though now we do know that there is *at least one* function-shape that can account for the individual's choices and hence that the individual can, in principle, be viewed as a utility maximizer under the conditions being studied.

A chooser who violates GARP is said to be "irrational"—he or she has inconsistent preferences, and the claim that he or she is acting as if he or she is maximizing a monotonic utility function is falsified. To put that another way, a chooser who violates GARP cannot be described as having a stable internal representation of value guiding his or her choice, and when we can make this observation, we know this for sure. The notion of rationality causes a lot of confusion in the cross-talk between economists and other scholars. Rationality in the economic sense only means that a person's preferences are internally consistent—that they are transitive. That a person prefers a bag of chips over a lifetime supply of caviar is not irrational in the economic sense regardless of whether it seems crazy to a caviar-lover.

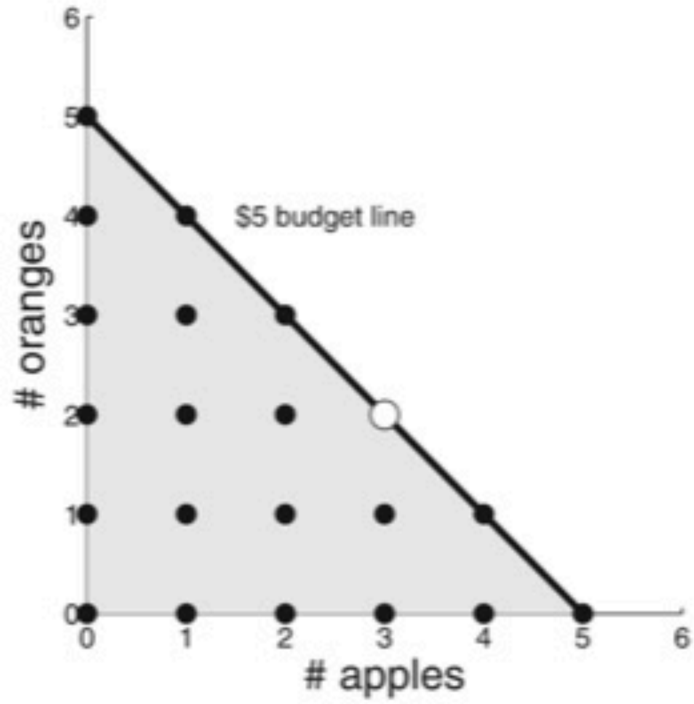

FIGURE 13.2. Economic rationality and a test of transitivity. Rationality in the economic sense means that a person's choices obey transitivity—that they are internally consistent. In the formal test of transitivity, using a budget set choice problem, choosers are asked to allocate a fixed amount of money to purchasing two goods. In this example, the chooser has a budget of $5 to spend on apples and oranges, which cost $1 apiece. Each point on the graph depicts a possible combination of apples and oranges the chooser can afford. Assuming the chooser must spend all of his or her money, he or she can choose any point (bundle) that lies on the $5 budget line, and in this example, the chooser selected the circled point as the preferred option (three apples and two oranges). What this tells us is that for this chooser the white circled point is *at least as good* as every other point on the budget line (because he or she selected it). It also tells us that every point in the shaded area below the line is *strictly worse* than at least one point on the line. Given this reasoning, we can conclude that anyone who also selects a point in the shaded area cannot be described as having a monotonic utility function that he or she is acting to maximize.

high-probability gains (they tend to *underweight* the likelihood of positive outcomes when the probability with which they are expected to occur is fairly *high,* for example, preferring option *A* over option *B* in our Allais paradox example); (2) risk averse for low-probability losses (they tend to *overweight* the likelihood of negative outcomes when the probability with which they are expected to occur is fairly *low,* such as when purchasing insurance); (3) risk seeking for low-probability gains (they tend to *overweight* the likelihood of positive outcomes when the probability with which they are expected to occur is fairly *low,* such as when going to a casino or purchasing lottery tickets); and (4) risk seeking for high-probability losses (they tend to *underweight* the likelihood of negative outcomes when the probability with which they are expected to occur is fairly *high,* preferring to avoid a sure loss at almost any cost). In prospect theory, risk attitudes are captured (though not in a unique way—many possible combinations of these parameters can capture each risk attitude) by (1) the curvature of the positive and negative value functions, (2) the magnitude of the loss aversion coefficient, and (3) the curvature of the subjective "probability weighting function."

Models of Temporal Discounting and How They Differ

Another factor (in addition to wealth and risk) that is consequential to how human choosers behave is how they weigh the present relative to the future, also known as *intertemporal choice*. Here we briefly review theoretical models of intertemporal choice before discussing the behavioral and neuroimaging evidence for this particular class of decision problems in the second part of this chapter.

It is no surprise that what we do today bears consequences for how much better or worse off we will be in the future. This is captured, for example, by our choice today of what we eat (e.g., high-fat foods), drink (e.g., massive amounts of alcohol), or choose more generally to do (e.g., exercise, smoke cigarettes). See Figure 13.4 for Walter Mischel's famous "marshmallow test," in which he showed that the ability of young children to delay gratification (waiting to get a better treat) was predictive of how financially and academically successful they became later in life and how likely they were to avoid going to jail or using drugs. See Mischel and Metzner (1962) for a description of the original studies and Mischel et al. (2011) for a complete survey of this work.

A consistent finding across all species tested with approaches like the marshmallow test is that choosers of all ages place less weight on outcomes the further in the future they are set to occur. This is the same (for an economist) as saying that there is a decrease in the utility of the delayed outcome as a function of time. For a given chooser, the rate at which the utility (or idiosyncratic subjective value) of an outcome decreases with time is referred to as this chooser's *discount rate,* a factor that shows considerable individual differences. The steepness of the discount rate describes a chooser's patience (with steeper discounting indicating more impatience). The discount rate, like most quantities stemming from economics, is estimated by observing people's choices, for example, when faced with sooner smaller and later larger outcomes ($50 today vs. $100 in 365 days).

The specific *shape of the discounting function,* however, is a point of ongoing debate that mirrors the tension between EU theory and prospect theory. The

FIGURE 13.4. The marshmallow test. In a famous series of studies in the 1960s conducted by then Stanford University psychologist Walter Mischel, kids were given a choice between having one marshmallow as soon as they wanted and waiting several minutes alone with the tempting single marshmallow until an experimenter returned to the room, at which point they would get two marshmallows. How willing kids were to wait (i.e., to delay gratification) was predictive of how successful they became later in life in multiple domains, including academics and financial stability. See Mischel and Metzner (1962) for a description of the original studies and Mischel et al. (2011) for a complete survey of this work. Photo by Kai Schwabe/Food Collection/Getty Images.

neoclassical economists Fishburn and Rubinstein (1982) and Strotz (1956) theorized that decision makers *should* (in order to be logically consistent) employ exponential discounting such that the utility of a given outcome (say monetary amount) should decrease at a fixed rate for each unit of time it is pushed into the future (Figure 13.5), given by:

$$\text{Discounted Utility } (A, t) = A \times e^{-\delta t}$$

Smaller values of the discount factor δ indicate shallower discounting. They were led to this conclusion when it became clear that any other shape of the discount function could lead to logically inconsistent choice behavior (Strotz, 1956). This is a critically important point that is often unclear. If a human chooser weighted the value of an outcome using anything other than an exponential function, then he or she would be *intransitive in time*, just as a person who prefers apples to oranges to pears but prefers pears to apples is intransitive now. And, of course, this would mean that no temporal-utility representation could account for his or her behavior.

Fishburn and Rubinstein (1982) completed what is now considered the standard proof of this fact using a modified form of the EU theory axioms with the addition of axioms about *impatience* and *stationarity*. Impatience refers to a preference for positive outcomes to occur sooner and negative outcomes to occur later. Stationarity refers to a consistency of behavior in time such that if a chooser is willing to wait 1 more day for a reward that could have been received in 364 days, he or she should also be willing to wait 1 more day for a reward that could have been received today.

Substantial empirical work, however, again shows that people violate the Fishburn and Rubinstein axioms under many conditions;[1] they do not have consistent preferences over time. Instead, most people's choices suggest that having to wait 1 day to receive a reward that could have been received today decreases the utility of that reward to a greater extent than having to wait 365 days to receive a reward that could have been received 364 days from today (Figure 13.5).

Mazur (1987) showed in seminal work that the discounting behavior of pigeons could best be explained by an alternative, hyperbolic or hyperbolic-like function, a form of discounting that has since received vast support across species, including in humans. This functional form assumes that choosers treat the near and distant future differently: the rate at which the utility of a reward decreases with time to its delivery is *faster* for the more proximal than the more distant future. Here, the

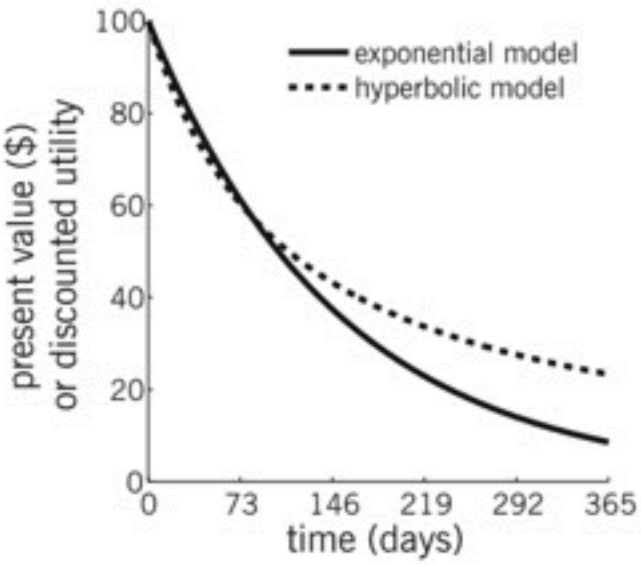

FIGURE 13.5. How the present value (or discounted utility) of $100 diminishes with time according to exponential discounting and hyperbolic discounting. The graph shows that having to wait 1 year to receive $100, for this moderately patient individual, is worth about the same as getting $8 today, according to exponential discounting, and $23 today, according to hyperbolic discounting. The rate of decrease in utility of the $100 is constant through time in the exponential model but is faster for the near future and slower for the distant future in the hyperbolic model. This difference in the rate of change through time in the hyperbolic model results in inconsistent preferences over time: having to wait 1 day from today to receive the $100 reduces its value by about $1, but having to wait 365 days instead of 364 days reduces its value by only $0.05.

[1]There is debate about whether they violate the stationarity axiom (Laibson, 1997) or the ordering axiom (Kable & Glimcher, 2010).

What is so important about revealed preference and rationality is that, with these tools in hand, a choice experiment reveals what an individual prefers. If Mary chose a cookie over an apple when she had the choice of both, then in that moment Mary must have preferred the cookie. That Mary feels regret that she went for the cookie instead of the apple has no bearing on what her preferences are, from an economist's perspective. That is, *unless her expressed regret has consequences for her behavior* (e.g., when faced with the choice between a cookie and an apple a second time, under equal circumstances, she chooses the apple), then—based on her behavior—an economist from the 1950s would conclude that Mary had made a mistake.

Expected Utility Theory

The groundwork reviewed above contributed to the origin of what is now known as expected utility (EU) theory (Savage, 1954; Von Neumann & Morgenstern, 1944). EU theory uses the neoclassical approach to describe choices under uncertainty developed by Samuelson (1938) but adds one new idea: that probabilistic outcomes can be used as a way to measure precisely the metrical shape of the utility representation (if it exists). To accomplish this, EU theory added additional axioms to the transitivity axiom (the axiom that requires a person's choices be internally consistent, as tested with GARP) that took advantage of *probability* as a ruler to determine the relationship in utility space between different choice options. These additional axioms are *completeness, continuity,* and *independence,* and their due attention is beyond the scope of this chapter. Very broadly speaking, completeness assumes that people have well-defined preferences, continuity implies that the utility function maximized is continuous, and independence guards against preference shifts in the presence of irrelevant choice alternatives (for a complete and accessible explanation of these axioms for this audience, see Glimcher, 2010). Critically, the objects of choice in EU theory are referred to as lotteries—defined by a probability and a prize (e.g., 25% chance of winning 100 pears). Taken together, if a chooser obeys GARP, there exists a utility function that the chooser is maximizing. If the chooser also obeys the other EU axioms, then we can further say that the expected utility of any option (the lottery) that the chooser is presented with is equal to the utility of the prize multiplied by the probability of receiving the prize.

RISK ATTITUDES

A critical insight here is that the curvature of the utility function tells us something about a person's risk attitude in the EU model—indeed, the curvature of the utility function fully specifies risk attitude for a consistent chooser (when one exists). In EU, a concave utility function over prize x (i.e., when the curvature of the utility function for x takes on a value below 1 and the function bends downward, as shown in Figure 13.1) is synonymous with saying that a person is risk averse in choices involving prize x. This person will, for example, prefer \$50 for sure to a 50% chance of \$100 because receiving \$50 carries a utility that is more than half the utility of receiving \$100—*consistent with diminishing marginal utility over total wealth.* The opposite is said for a person whose utility function is convex (i.e., when the curvature of the utility function for x takes on a value above 1 and the function bends discount rate κ determines the rate of decrease, and smaller values of κ indicate shallower discounting:

$$\text{Discounted Utility }(A, t) = \frac{A}{1 + kt}$$

An unsatisfying aspect of hyperbolic discounting for economists is that hyperbolic discounting implies that choosers are necessarily inconsistent in their preferences with respect to time. To account for the empirically observed departures from exponential discounting while preserving stationarity, Phelps and Pollak (1968), and later Laibson (1997), popularized the use of the quasi-hyperbolic discounting model. In this form of discounting, rewards are discounted exponentially (by a constant rate δ over time), but the utility of delayed rewards (i.e., when $t > 0$) is additionally reduced relative to immediate rewards by a second constant "bias" term β:

$$\text{Discounted Utility }(A, t) = \beta \times (A \times e^{-\delta t})$$

For economists, quasi-hyperbolic discounting is preferred to hyperbolic discounting, even if it may also not be very predictive under some conditions (e.g., Laibson, 1997), because it does not require a complete departure from exponential discounting, as the exponent is preserved and all of the inconsistency emerges uniquely from a "present bias." When the bias term β is 1, that is, when delayed rewards are not treated in some special way compared with immediate rewards, this functional form reduces to exponential discounting. When β approaches 0, delayed rewards are given less weight and immediate rewards prevail. The β parameter is thus an intuitive way to assess how far a chooser departs from normative, exponential discounting. Its usage has been central to mapping temporal discounting onto *dual process theories* of decision making in which there is a perceived battle between the hot–cool (Metcalfe & Mischel, 1999) or affective–deliberative (Bernheim & Rangel, 2004; Loewenstein & O'Donoghue, 2004) self. The distinctions among exponential, hyperbolic, and quasi-hyperbolic models make interesting predictions about what we might observe at the neural, implementation, level of intertemporal choice decision making. We come back to this point later, after we review the tools neuroeconomists use to get at these underlying mechanisms.

Before we do that, however, it is worth pausing for a moment to point out that, as the careful reader might have noticed, in the three prominent models of temporal discounting discussed, the delayed reward, A appears without an exponent in the equations. That is, these models assume an embedded *linear* utility function consistent with risk neutrality, such that there is no subjective transformation of the objective value of A. Yet we know that most people do not have linear utility functions; instead, most people behave as though they have nonlinear utility functions (usually concave, consistent with risk aversion). This assumption of linear utility might importantly affect the estimation of an individual's idiosyncratic discount rate, which is invariably influenced by that individual's idiosyncratic utility curvature. This issue has recently come to the forefront of research in temporal discounting, although there is still much work to be done. Nevertheless, it seems clear now, as Andreoni and Sprenger (2012) showed, that the presence versus absence of certainty can bias people's time preferences and that jointly eliciting risk and

Model-Based fMRI

In a typical model-based fMRI study (see Figure 13.6), a time series of values is derived from a computational model of a specific cognitive process (e.g., the discounted utility of the option offered on each trial of an intertemporal choice task as determined by a given participant's κ value) and then correlated against the fMRI data from that participant performing the task to determine which brain regions show a response profile consistent with the model. The goal is to map model-derived parameters onto their specific underlying brain circuits, to determine not only what regions are potentially responsible for generating this particular cognitive process but *how* this process is implemented in these brain regions. This is the key advantage of the model-based approach over more conventional analytical approaches to neuroimaging data (O'Doherty, Hampton, & Kim, 2007).

A special form of model-based fMRI is called a *neurometric* analysis (as opposed to a psychometric analysis, in which a model is fit to behavior or ratings provided by the participant). In a neurometric analysis, the model is fit directly to neural responses. This is difficult to do in associative brain areas (such as the prefrontal cortex) but is now common practice in the vision and audition sciences that focus on the sensory cortices in which signals are more closely tied to experimental inputs (visual stimuli or sounds as they evoke responses in visual or auditory cortices).

The Case of Intertemporal Choice: Putting It All Together

In this section, we focus on how neuroeconomists have worked to better understand why people often choose immediate rewards over rewards to be received in the future, even if the future rewards are bigger.

Quasi-hyperbolic (dual-system-style) discounting would predict the engagement of two neural systems in the decision to take the larger, later reward, an account we discussed earlier as having to do with the competition between two "selves." One system would push behavior toward less patience (i.e., toward the immediate reward as demonstrated by a lower beta term) and the other system would push behavior toward more patience. Initial evidence using fMRI in human participants indeed seemed to suggest, based on a fairly complex inference, that one set of brain regions, the ventromedial prefrontal cortex (VMPFC), ventral striatum, and posterior cingulate cortex (PCC), was more active when participants chose the immediate relative to the delayed options in an intertemporal choice experiment (McClure, Laibson, Loewenstein, & Cohen, 2004). Also another set of regions, the dorsolateral prefrontal cortex and posterior parietal cortex, was active when the participant made a more difficult choice compared with an easier choice. These two groups of regions could be construed as mapping onto beta and delta regions, respectively, because in the first case (when participants chose the immediate rewards) there was more activity in beta regions relative to delta regions, and vice versa when participants chose the delayed reward.

But putting this early evidence aside for a moment, it is, of course, also logically possible that a single hyperbolically discounting system is used to evaluate all rewards, with each reward's utility determined by how far away in time it is to be

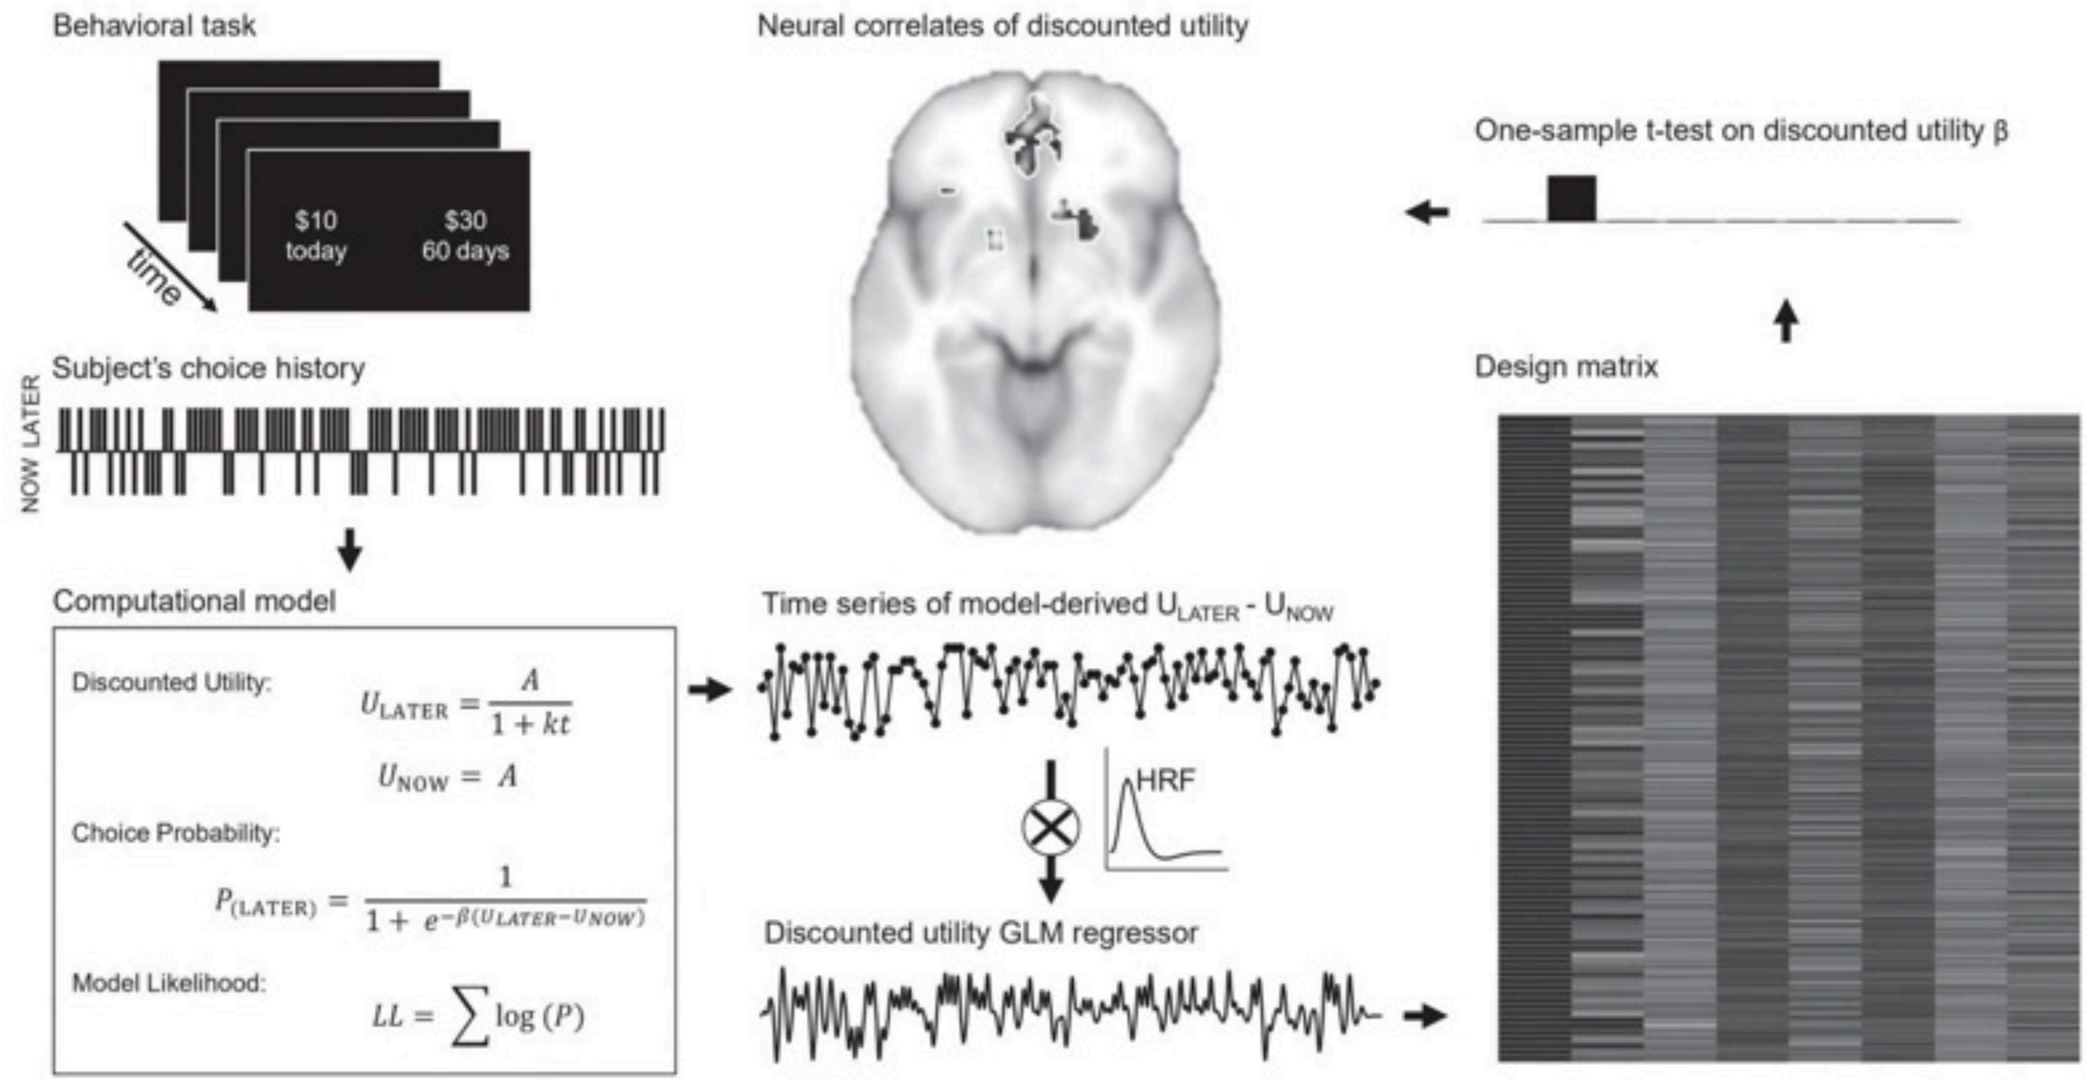

FIGURE 13.6. Steps involved in conducting a model-based functional magnetic resonance (fMRI) study. In this example of an intertemporal choice model-based fMRI study, participants first make binary choices between some amount of money to be delivered without delay (NOW) and a larger amount of money to be delivered with some delay (LATER). Based on the series of choices the participant makes, the experimenter estimates a discount rate parameter that minimizes the difference between the model predictions and the observed behavior. The best-fitting parameters are then used to generate a time series, for example, of discounted utility or difference in utility of the NOWversus LATER options seen on each trial of the experiment. This time series is then convolved with a basis function such as a canonical hemodynamic function (or HRF) to account for the lag in the BOLD response. This convolved time series of discounted utility can then be used as a predictor variable, along with other predictors, in a regression analysis against the simultaneously collected fMRI data. The experimenter can then visualize in which regions of the brain BOLD activity correlates with discounted utility. Adapted from O'Doherty, Hampton, and Kim (2007) and Glascher and O'Doherty (2010).

received by a simple hyperbolic calculation. Pitting these two competing accounts against each other (and avoiding some of the complex structural inference of the earlier McClure et al., 2004, study), Kable and Glimcher (2007) estimated what individual participants' discount rates were based on the choices they made in a long and detailed intertemporal choice experiment. The authors then used this estimate to construct what the discounted utility of each option on a given trial of the experiment would be for a given participant based on that participant's own discount rate. Using the model-based fMRI analysis approach we discussed in the previous sections, this study revealed that activity in a small set of brain regions, all part of the proposed beta regions, that is, the VMPFC, ventral striatum, and PCC, increased hyperbolically exactly as the utility of the offer increased (Figure 13.7)—a study now widely seen to demonstrate a single system for hyperbolically representing discounted value. In a follow-up study, Kable and Glimcher (2010) showed that this was also true in the context of deciding about two delayed rewards (in which there could be no "present bias") rather than about a delayed and an immediate reward. In this study, participants anchored their discounting behavior to the soonest possible reward, not simply to the present, and the magnitude of activation in the VMPFC, striatum, and PCC was not higher when an immediate reward could be chosen relative to when only a delayed reward could be chosen. Instead, activity in these regions encoded the subjective value of immediate and delayed rewards. Together, these two studies provided what is now typically seen as unequivocal evidence that the VMPFC, striatum, and PCC track discounted utility, or discounted subjective value, and not immediacy, impulsivity, or some "hot" response to a reward. Moreover, the shape of the discounting function that maximized the correspondence between the brain and behavioral data in both studies was hyperbolic, providing further evidence that human participants do not decide according to exponential discounting.

Although both perspectives have since found empirical support, the latter (unified system) view provides a more parsimonious account (from a neurobiological perspective) about how we might decide about the future. Further work continues to show that the VMPFC, striatum, and PCC are not active exclusively or disproportionately more for immediate rewards. Instead, these regions are now thought to form what is referred to as the brain's valuation system, a set of brain regions that track the value of choice options across a variety of decision contexts and reward types (Bartra, McGuire, & Kable, 2013).

Is Temporal Discounting a Loss of Self-Control?

An outstanding question pertains to whether steeper discounting reflects a loss of self-control, as the quasi-hyperbolic discounting model would predict. Many studies have linked steeper discount rates to maladaptive behaviors presumed to result from poor self-control and/or lack of forethought, including gambling, obesity, low achievement, and credit card debt (Story, Vlaev, Seymour, Darzi, & Dolan, 2014). Excessive discounting is also a feature of many psychiatric disorders that are characterized by problems with self-control, most notably impulse control disorders (American Psychiatric Association, 2013) and drug addiction (Baler & Volkow, 2006). Although these findings are highly suggestive of a link between self-control

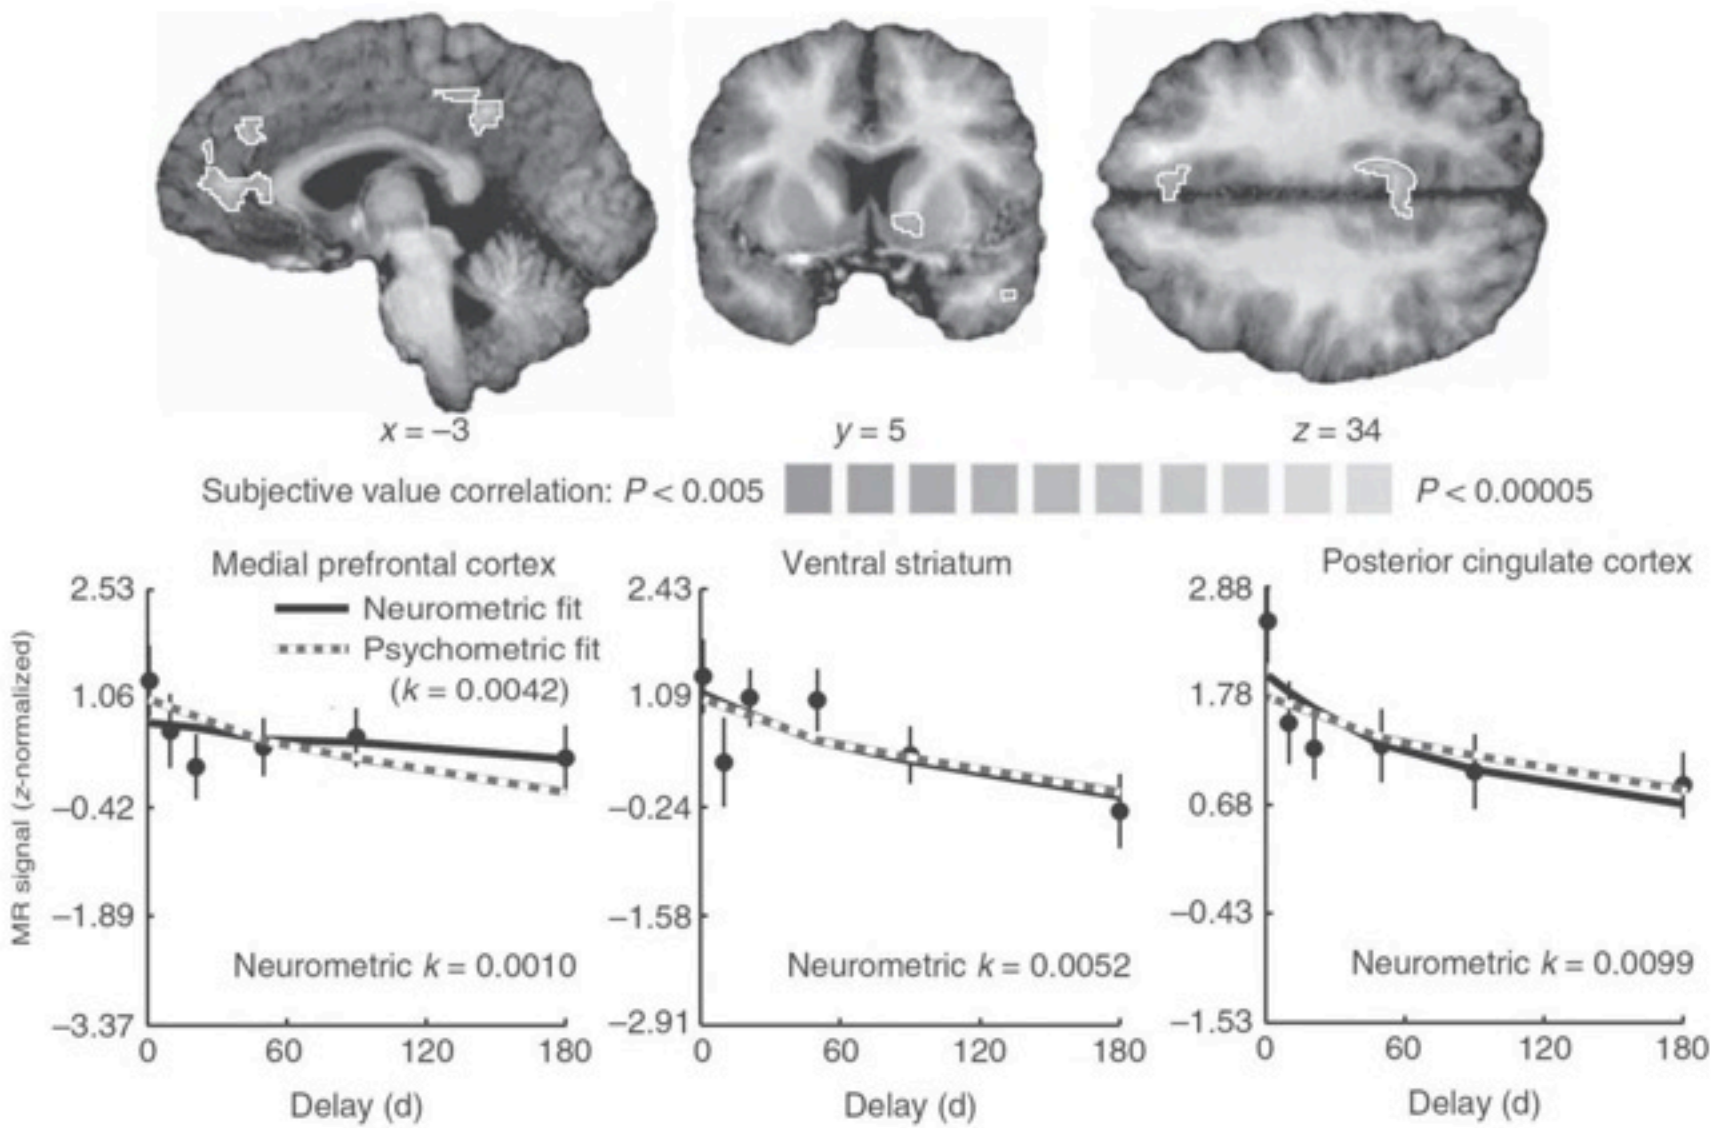

FIGURE 13.7. Brain areas that track individual discounted utilities. The top panel shows which regions' BOLD activity correlated with discounted utility for each individual based on his or her own discount rate parameter κ. These regions are the ventromedial prefrontal cortex (VMPFC), ventral striatum, and posterior cingulate cortex, which, based on meta-analytic data, are thought to form a unified system of valuation for all types of rewards and decision contexts. The bottom panel shows, for an example participant, a close match between the predicted response in each region as a function of delay given that participant's behaviorally derived discount rate and the brain's response to each delay (the neural discount rate). Together, these data show that these three regions are directly involved in the valuation of delayed outcomes to guide choice for those outcomes. Reprinted by permission of Macmillan Publishers from Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience, 10,* 1625–1633.

upward). This person is said to be risk seeking (or risk loving) and will prefer a 50% chance of \$100 to \$50 for sure as demonstrated by the same rationale.

Risky Choices as Captured by Prospect Theory

Economic models such as EU theory are built on logical statements about choice. Their strength is in their falsifiability. But this may also be their greatest weakness: Under many conditions we can show that people do not obey the EU axioms, and therefore their choices are inconsistent with EU theory. A series of observations made by experimental psychologists indeed highlighted this flaw, showing that the neoclassical theories did not capture many kinds of human choice behavior and may thus be insufficient as a basis for our understanding of choice. When faced with probabilities that are very small or very large, for example, people consistently behave as if they violate the independence axiom, as most famously demonstrated in the Allais paradox (Allais, 1953). In the Allais paradox experiments, people are asked to decide between two lotteries, *A* and *B,* and then between two other lotteries, *C* and *D*. In a typical example, *A* is a 100% chance of \$1 million, and *B* is an 89% chance of \$1 million, a 1% chance of nothing, and a 10% chance of \$5 million. In the second pair, *C* is an 89% chance of nothing and an 11% chance of \$1 million, and *D* is a 90% chance of nothing and a 10% chance of \$5 million. In reality, after removing any implied "common consequence" components of the lotteries (89% of \$1 million for *A* and *B,* and 89% of nothing for *C* and *D*), the two pairs of lotteries can be reduced to the same choice (11% chance of \$1 million vs. 10% chance of \$5 million). But despite this formal equivalence, most people choose *A* from the first pair and *D* from the second pair, which would imply that they both prefer an 11% chance of \$1 million over a 10% chance of \$5 million *and* a 10% chance of \$5 million over an 11% chance of \$1 million, violating the independence axiom. One implication of these findings is that any unified theory of choice needs to recognize that there are certain *constraints* (be it biological or other) that make us inconsistent and that influence our decision process and that perhaps these behavioral patterns should not be ignored but rather used to inform theory.

Aiming to allow choice prediction even when choosers were being inconsistent, when it could be proven that no fixed internal representation of value existed, Kahneman and Tversky (1979) suggested an alternative to EU theory, known as prospect theory. This made prospect theory fundamentally different from EU theory because it was, in essence, a fittable predictive model rather than a logical theory in the economic sense. Prospect theory permits fitted solutions that approximate behavior even when one can prove that the kind of representation prospect theory employs cannot, fundamentally, be correct in the economic sense. Prospect theory achieved this by modifying EU theory with three classes of "parameterizations." The first was the idea that, in addition to the transformation of prizes in the computation of utility, there was a transformation of probabilities as well that related objective probabilities to *subjective probabilities,* an extension of an idea first introduced by Savage (1954). A second change had to do with how utilities were defined. Instead of being referenced to a chooser's total wealth (the original notion of diminishing marginal utility), Kahneman and Tversky (1979) argued for the presence of a moving *reference point* against which utilities are compared (how much the time preferences can improve the accuracy of discount rate estimates (Andersen, Harrison, Lau, & Rutstrom, 2008). More recent work has focused on how to best capture these interdependencies in the experimental setting (Andreoni, Kuhn, & Sprenger, 2015).

Experimental Tools from Psychology and Neuroscience Used in Neuroeconomics

In this section, we review the basic tools that neuroeconomists use to study decision making in human choosers. We specifically focus on how these tools (mostly stemming from experimental psychology and neuroscience) can help researchers test some of the core predictions of the theoretical models (mostly stemming from economics) that we have reviewed in the preceding sections. We first describe the key features of a decision-making experiment and how researchers use the choice data generated to "fit" these models. We then describe how these experiments can be performed using functional brain imaging techniques, and how the neural data generated with those techniques can be combined with a study participant's choice data and the theoretical model under study.

The Notion of Incentive Compatibility

We can estimate a decision maker's risk preference and discount rate by asking him or her to participate in a laboratory experiment. In a typical neuroeconomic laboratory experiment, participants are asked to make a series of choices about goods (e.g., snacks) or money that they can *actually* receive with some probability or at a given delay. That participants make choices with real consequences (as compared with hypothetical choices with no direct consequences to them), referred to as *incentive compatibility,* is thought by many to be extremely important for eliciting an individual's "true preferences." Economic journals as a rule do not publish studies that lack real incentives (e.g., choosing a delayed reward of $100 in 365 days and actually getting the $100 in 365 days). This practice has been argued to contribute to the relatively high replication rates seen in experimental and behavioral economics studies relative to other fields, in which the use of hypothetical decisions and incentives is more common (Camerer et al., 2016).

Because in a typical experiment more than one choice needs to be made by participants and because paying according to each choice can become costly if meaningful incentives are used, task earnings are usually determined by a random draw of one or several of the series of choices the participant made. Paying from a subset of randomly determined choices after the experiment is complete also helps reduce strategy use—for example, stockpiling earnings to then gamble more in a later part of the experiment. Participants are instructed, without deception, that they should treat each of their choices as independent and important because only one or a subset of these choices will determine their earnings, and they do not know which at the time of choice.

Once enough choice observations are made, the experimenter typically picks a model (e.g., hyperbolic discounting) and, using minimization algorithms, tries to and temporal discounting, based on this evidence alone one might draw the conclusion that self-control as indexed by temporal discounting is a between-subjects phenomenon—a stable interindividual difference—rather than a capacity or resource that can be depleted or "lost."

However, there is also compelling evidence that temporal discounting is not (or at least is not only) a stable interindividual difference or trait. It can be affected, within an individual, by his or her current physiological needs or emotional state and by external, contextual influences (e.g., how the decision problem is framed; Lempert & Phelps, 2016). More recent work also shows that cognitively taxing activities (6 hours of working-memory and set-switching tasks) can make laboratory participants less patient, presumably by interfering with the neural structures that support self-control, such as the dorsolateral prefrontal cortex (Blain, Hollard, & Pessiglione, 2016). The problem here, though, is in the interpretation of an increase in discounting as self-control failure. There are many reasons why individuals should discount delayed rewards at different rates. One reason is that these delayed rewards might fail to materialize or come too late to satisfy the organism's *current* needs. There is an opportunity cost that comes with waiting (Story, Moutoussis, & Dolan, 2015), and certain states or external contexts might shift this cost. We return to this point in the next section, in which we describe how a neuroeconomist might think about excessive discounting behavior, such as that seen in drug addiction. Furthermore, without consideration for an individual's beliefs about the experimental paradigm, it is difficult to know whether these changes in the discount rate reflect a change in the certainty about receiving the prospective reward, a change in time perception, or a change in the rate at which future rewards are discounted.

Temporal Discounting and Psychopathology: The Case of Drug Addiction

Addiction is a chronic, relapsing disease characterized by continued drug taking despite many harmful health and social consequences (American Psychiatric Association, 2013). The addicted individual becomes preoccupied with the addictive substance (e.g., alcohol or heroin), despite efforts to stop or cut down drug taking, even when the positive experiences derived from the drug are markedly reduced. Addiction is considered a brain disease because drugs change the brain's chemistry, particularly in regions that form the brain's valuation system. At the core of this system are dopamine fibers that originate in the ventral tegmental area and terminate in the ventral striatum and in other regions, including the VMPFC. Although different drugs of abuse have different mechanisms of action, they have all been shown to produce a surge of dopamine in the ventral striatum in exerting their reinforcing effects (e.g., the feelings of "high" and "euphoria"; see Chen, Hopf, & Bonci, 2010; Laruelle et al., 1995; Sulzer, 2011; Volkow, Fowler, Wang, Ding, & Gatley, 2002; Volkow et al., 1997).

A large body of literature indicates that steeper discount rates are a hallmark feature of drug addiction (MacKillop et al., 2011). This is not surprising given that the brain circuits affected by addiction are primarily those that support intertemporal choice that we reviewed in the preceding sections (Volkow & Baler, 2015). The steepness of an individual's temporal discounting has even been proposed chooser is set to gain or lose relative to his or her reference point). The third change stemmed from the reference point idea and allowed "losses" to be considered as distinct from "gains" (a notion not present in EU theory) with their own mapping from prize (objective value) space to utility space. This is also referred to as *loss aversion*–the idea that people behave as if the shape of their utility function for losses is different (steeper) than that for gains (Figure 13.3).

A critical distinction between EU theory and prospect theory and, more broadly, a difference in the way economists and psychologists think about theories of choice is that, whereas the former describes how people *should* behave (in the sense that it describes consistent logical behavior), the latter makes it possible to basically describe with a utility-like model how people *do* behave, even when their behavior can be proven mathematically not to be describable with a utility-like functional representation. Put another way, EU theory is a *prescriptive* (or normative) theory of choice (and thus does a good job when people are being logical), whereas prospect theory is a *descriptive* theory of choice that can be used under nonlogical conditions.

To help illustrate this difference between EU theory and prospect theory more concretely, we return to the notion of risk attitudes. The specific formulation of utility in prospect theory is made with the forethought of *accommodating* empirical observations about how people behave in situations involving fully known (and hence explicitly stated) risk. These observations are: (1) people are risk averse for

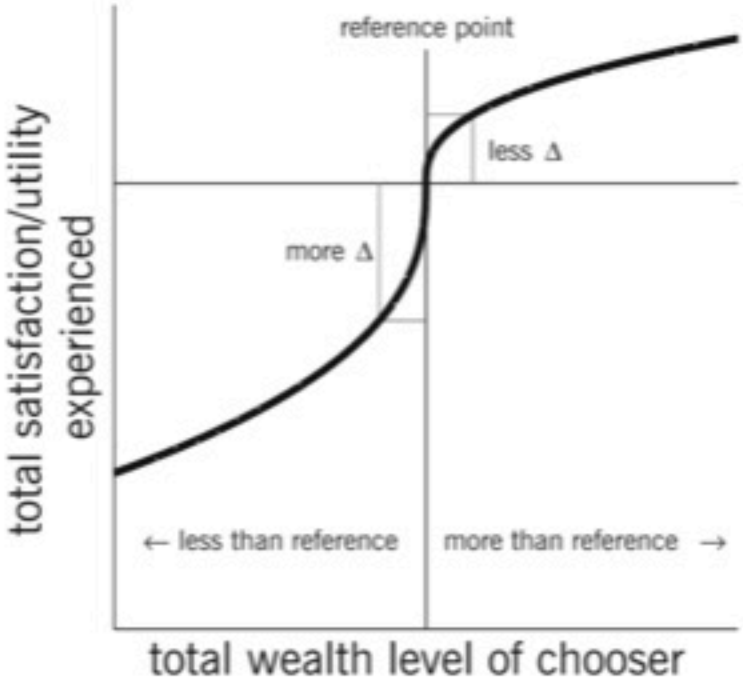

FIGURE 13.3. Shape of the "value" function according to prospect theory. The value function (Kahneman & Tversky [1979] never called it a utility function because they wanted the function to be used even when choice behavior proved that no utility representation was formally possible for a particular chooser) in prospect theory is centered on a subjective "reference point." The amount of gain or total wealth of a chooser is determined as a deviation from this reference point. The difference in the curvature of the utility function on the left versus the right side of the reference point accounts for the empirical observation that people are risk averse for gains (or relative gains) and risk seeking for losses (or relative losses). This effect, also referred to as loss aversion, is thought to occur because a given amount of gain produces smaller increases in satisfaction than the same magnitude of loss decreases satisfaction.

find the parameters of that model that best capture a participant's choices (e.g., her or his discount rate). This analytical approach is referred to as a "model-based" or "parametric" analysis. How good the model is at capturing a participant's preferences depends not only on the assumptions of the model but also on the quality of the experiment (e.g., if a sufficiently probative "choice set" was used and if enough observations were made).

Human Neuroscience Tools

The unique contribution neuroscience brings to neuroeconomics is the direct measurement of *subjective value* in the brain. Subjective values are related to utilities and value functions with the difference being that subjective values are directly observable (for a detailed discussion of this topic, see Glimcher, 2010, Section 3). Using a variety of neuroscientific tools, it is now possible to measure directly, or to infer quite precisely, the activity levels of neurons or populations of neurons encoding the subjective values of options in single-choice sets actually being presented experimentally to laboratory participants.

Measurement Techniques

The most widely used measurement tool by neuroeconomists to study how the brain represents utility or subjective value in humans is functional magnetic resonance imaging (fMRI). Although other measurement tools are also available to study human decision making—including scalp and intracranial electroencephalography (EEG), magnetoencephalography (MEG), and positron emission tomography (PET), fMRI is by far the most common. This is likely because fMRI is a noninvasive and fairly flexible technique with a good balance between spatial and temporal resolution. fMRI takes advantage of the fact that blood oxygen levels increase in areas of the brain where there is more metabolic activity due to neuronal activity. These increases in oxygenated blood—measured with the blood-oxygen-level-dependent (BOLD) contrast—can be used as an indirect readout of the amount of neuronal activity in the area of increase. For a comprehensive description of how fMRI and the BOLD signal work, the interested reader is referred to Huettel, Song, and McCarthy (2009).

In a typical neuroeconomic fMRI study, participants make choices on economic tasks such as the ones we described in the section above while fMRI data are collected. However, these tasks need to be adapted for use during fMRI, keeping a couple of limitations of the BOLD response in mind—mainly those pertaining to its "delay to peak," the observation that the BOLD signal is a sluggish representation of underlying brain activity (on the order of 5–15 seconds), and its troublingly low signal-to-noise ratio. These constraints require adequate temporal separation of the events of interest (in this case, the choices a participant is asked to make) and repetition to enhance detection power. To capture how the brain computes or represents utility and implements choice, data analysis typically focuses on the time window during which participants consider their options and prepare a response. A neuroeconomics experiment aims to correlate brain activity during this time window with economic model parameters or with a participant's choices or both.

use to self-regulate their discount rate, particularly in situations in which the more appealing immediate course of action is in conflict with one's long-term goals (e.g., smoking after having recently quit). One promising approach may be the use of mental contrasting (Oettingen, 2000, 2012; Oettingen, Pak, & Schnetter, 2001; Oettingen & Schwörer, 2013)—a process by which an individual compares his or her long-term goals (maintaining abstinence and overall good health) with the obstacles of current reality (relieving the stress of a work day with a cigarette) as a way to explicate tangible steps to overcome that reality and to meet his or her respective goal. Mental contrasting has also been combined with implementation intentions. These are if–then plans (Gollwitzer, 1999, 2014) that link goal-directed responses to specified situational cues. Mental contrasting with implementation intentions (MCII) turns out to be a particularly effective self-regulation tool to change behavior in multiple life domains (health, social, occupational; Oettingen & Schwörer, 2013). Even in laboratory studies of economic decision making, both MCII (e.g., involving social fairness; Bieleke, Gollwitzer, Oettingen, & Fischbacher, 2016) and implementation intentions (e.g., involving unknown or not fully known probabilistic outcomes; Doerflinger, Martiny-Huenger, & Gollwitzer, 2017) were effective self-regulation tools.

Mental contrasting and implementation intentions could thus also help to reduce the rate of temporal discounting. One psychological mechanism by which this strategy could work is by making the future seem more *concrete* and *certain*. Indeed, people consistently exhibit shallower discounting behavior when delayed outcomes are described in terms of a calendar date or event (e.g., December 25th, or Christmas) rather than in terms of a more abstract passage of time (e.g., 65 days from today), presumably because this framing manipulation helps one concretize his or her future (Read, Frederick, Orsel, & Rahman, 2005). In addition, a related process, episodic future-thinking, in which a person mentally simulates future experiences (what it would feel like to get $100 on Christmas), has been shown to reduce individual discount rates in laboratory studies (Peters & Buchel, 2010). Engaging in episodic future-thinking (and possibly mental contrasting) is likely inherent to choosing to wait for delayed rewards, as individuals with frontotemporal dementia, whose memory and imagination systems are atrophied, discount future outcomes *to a greater extent* than their neurotypical counterparts (Bertoux, de Souza, Zamith, Dubois, & Bourgeois-Gironde, 2015; Chiong et al., 2016; Lebreton et al., 2013).

Finally, strategies that focus on the choice process itself might also have some efficacy in biasing choice toward that consistent with one's long-term goals. For example, directing attention toward specific attributes of the choice options (the magnitude of reward vs. the delay to reward) can differentially promote immediate versus delayed choices (Lempert & Phelps, 2016). Other strategies include elaborating on the possible outcomes of a decision before making a choice (by first generating arguments favoring delayed consumption followed by arguments favoring immediate consumption; Weber et al., 2007) and focusing on choice options in the "explicit zero" format (as in a choice between $10 today and nothing in the future rather than simply $10 today; Magen, Dweck, & Gross, 2008; Magen, Kim, Dweck, Gross, & McClure, 2014).