

2

Psychology and Economics

What do people want? Do they even know? How do they make choices, big and small? Answers to questions like these—how individuals form preferences, how they make decisions—guide how economists think about the world. No matter how far removed the immediate questions of study in any particular field—macroeconomics, finance, trade—may seem from matters of individual choice, scratch the surface and the analysis nearly always depends in some part on assumptions or observations regarding how individuals choose and behave. They may be hidden or implicit, they may be ad hoc or unexamined, but there they are.

Public finance is no exception. While it is easy to think of public finance mainly in terms of more aggregate units of analysis—how markets fail, how they can be repaired—its conclusions are undergirded everywhere by a theory of individual choice. The occurrence and the consequences of market failures depend on elements of individual decisionmaking just as much as they do on the role of market structure. For example, the implication of negative externalities in the consumption of polluting forms of energy—like gasoline for cars—is a joint outcome of the failure of prices to reflect those external costs and the behavioral response of individuals to that pricing failure. Similarly, conclusions about whether and how the government should intervene in response to market outcomes turn on how we believe people will respond to those policies. A corrective tax on gasoline, for example, is presumed to be effective because of the way that it will enter the choice calculus of those who must pay it.

If the usual approach of public finance is to elide any serious discussion of how people form preferences and make choices, it is not, then, because such

questions are irrelevant. Rather, it is because they are treated, for most practical purposes, as settled. In the standard economic analysis, our answers primarily follow from the assumption that individuals optimize perfectly, by which we mean, roughly speaking, that individuals are good at choosing among the options that different market structures or different policy environments present them with. People know what makes them happy. They hold preferences that are complete, stable, and well specified. They make plans to maximize their well-being, and their choices reflect those plans. Many times we go further, making simplifying assumptions about what people want, as well—for example, that people are purely self-interested, or nearly so.

Even as public finance economists employ such assumptions, they recognize their limitations. Of course people make mistakes, of course they give in to temptations that they later regret, and of course they can be altruistic. Despite such violations of economic assumptions, the standard model remains a durable feature of public finance. Even if it is not completely true, it is useful. It gives clear guidance on how to set policies—everything from how to set a tax to correct an externality to how to structure social insurance in the face of adverse selection. In short, the standard model survives because it is presumed to be a reasonable approximation of reality for many problems and because there is no obvious alternative that is as concrete and useful.

Increasingly, however, the evidence suggests that deviations from the standard model are more the rule than the exception and that they have consequences in the aggregate and for policy responses. Psychology has demonstrated that violations of the standard economic assumptions about preference and choice are pervasive. Behavioral economics has identified a number of contexts in which deviations have consequences for market or policy outcomes. Centrally, that evidence suggests that when people deviate from the standard assumptions, they do so in predictable ways. Thus behavioral economics does not just question the validity of old assumptions; it replaces them with new ones.

In this chapter, we catalog evidence from psychology and behavioral economics of behaviors that are inconsistent with the standard assumptions, and we classify those specific findings into a set of broad behavioral tendencies. The core challenge for incorporating the results from psychology and behavioral economics into public finance is simply making sense of them. The literature on psychology is vast, providing deep and wide-ranging insights across a variety of phenomena of the mind. To economists, a first reaction to this literature can be to view it merely as a collection of isolated observations with an obscure taxonomy: Mental accounting. Asymmetric dominance. Choice overload. And so on.

By themselves these results seem both too specific and too diverse to be immediately useful for drawing general conclusions about how they relate to economic activity. From this large pool, however, only a small set of abstract insights about

behavior—insights that crudely capture many different psychological phenomena in a few broad behavioral tendencies—is sufficient for economic analysis. Obviously such a distillation results in a tremendous loss of rich insights for understanding the mind. But for understanding how to set policy and regulatory levers, that richness is neither necessary nor desirable. In general, the goals of economic policy are not to understand or correct the behavior of individuals but to affect behavior in the aggregate or on the margin.

So, for example, consider the specific findings mentioned above: mental accounting, asymmetric dominance, and choice overload. They refer to very particular and very different behaviors. Mental accounting refers, roughly speaking, to the tendency of individuals to fail to treat income or wealth as fungible across sources or uses. Asymmetric dominance refers to the tendency of individuals to allow their preferences to be swayed by the introduction of irrelevant alternatives. And choice overload describes the tendency of individuals to be put off from making a choice as alternatives proliferate. But all three can also be viewed as examples of a more general finding that people are not unbounded in their ability to consistently consider and respond to all of the features of complex choices. Therefore, these and similar findings can be categorized under a more general psychological tendency that might be referred to as limited computational capacity.

The other way in which public finance can abstract from psychology is to determine which results from psychology are important to incorporate into economics and which can be safely ignored. The test here is whether the psychological principle is likely to have much bearing on the validity of the standard assumptions about economic agents. So psychological results about decisionmaking errors (which bear on assumptions about the ability of economic agents to optimize perfectly) or other-regarding preferences (which relate directly to standard assumptions about the form of preferences) must be incorporated in some way. On the other hand, psychological results that do not relate directly to those assumptions—such as, say, the tendency to obey authority—though obviously not without economic consequences on some level, are largely beyond the scope of public finance.

We should be clear up front that while there are benefits to doing this type of selective aggregation and creating a broader taxonomy of behavioral tendencies, there also are costs. For example, while it is useful for many purposes to think of asymmetric dominance and mental accounting as manifestations of a broader tendency toward limited computational capacity, there will be instances in public finance in which the specific features of those behaviors and the distinctions between them remain important. In designing policy responses, for instance, asymmetric dominance may argue for restricting choice sets while mental accounting may argue for framing choices differently. Similarly, aspects of psychology that we will largely ignore, such as the tendency to obey authority,

will be in some instances important for public finance, and by excluding them we do lose some power for understanding and designing policy.

Below, we describe our classification scheme for results from psychology and behavioral economics, which organizes findings around three basic deviations from standard assumptions:

—*Imperfect optimization.* The classical model assumes that individuals are capable maximizers of their own utility—that is, that they know what they want and what will make them happy and that their choices and preferences are consistent. Behavioral economics, however, finds that individuals are imperfect in their ability to maximize their own welfare and that their choices are often inconsistent—that is, that individuals have more difficulty knowing what they want than the standard model assumes.

—*Bounded self-control.* Even when individuals accurately perceive their own interests, they can have difficulty realizing their intentions. The classical model allows for no such difficulty, and it assumes time consistency in preferences. Behavioral economics recognizes forces such as temptation and procrastination as real and meaningful phenomena—that is, that individuals have more difficulty doing what they want than the standard model assumes.

—*Nonstandard preferences.* Finally, the standard model also makes some weak assumptions about the shape of individual preferences. Behavioral economics finds two important cases in which those assumptions appear inaccurate: First, preferences appear to be set over changes in status rather than over end states. Second, the assumption of pure self-interest is often a bad assumption, in that individuals routinely hold preferences that are other-regarding—that is, that what people want is different from what we usually assume.

For each deviation, we describe some of the available empirical evidence. We relegate to appendix A a brief discussion of how to incorporate these deviations into economic models of choice and welfare in a slightly more formal way. Finally, note again that we make no attempt here to do a thorough review or complete summary of behavioral economics but seek instead to highlight features that will be important for public finance. Good reviews and summaries are available in the literature.¹

Imperfect Optimization

Economists famously assume that individuals are optimal decisionmakers. Technically, optimality in choice is a matter of adherence to a set of assumptions that impose both a degree of accuracy on choices—that choices reflect all of the relevant, available information, for example—and a logical consistency on choices—for example, that choices are independent of irrelevant alternatives. Less formally, optimality in choice amounts to an assumption that individuals

are basically good at making choices that maximize their own welfare: individuals know what they want, and they make choices that realize their desires.

However, psychology and behavioral economics have amassed a growing collection of findings suggesting that according to those criteria, individuals are, in practice, flawed decisionmakers. Consider one well-known case relating to choice behavior in the face of an increase in the number of alternatives. Standard assumptions of optimization imply that increasing the number of elements in a choice set should leave individuals at least as likely to choose from the set of increased choices as to choose from the original set. At worst, individuals will continue to select as they did from the original options, and at best they will make new selections from among the new options. But an experiment in which shoppers in a grocery store were given coupons for jam after being randomly offered samples of either a few selected varieties of jam or a wider assortment of jams found evidence that contradicted that prediction.² While 30 percent of subjects offered the smaller set of samples ultimately purchased jam, only 3 percent of subjects offered the wider set went on to purchase jam. Researchers interpreted that finding to mean that individuals offered the larger set of samples were actually put off from choosing by the difficulty of selecting from the greater number of options. Moreover, results of this type are not limited to psychology experiments. There is, for example, some evidence that individuals are less likely to participate in their employer's retirement plan as the number of investment alternatives increases.³

Such results are part of a large set of findings in the psychology of judgment and decisionmaking that suggest that, in fact, individuals are not always good at making choices.⁴ They make choices that appear to ignore or misconstrue available information or that exhibit the types of logical inconsistencies disallowed by full optimality. The full catalog of particular deviations is long, and it can be organized in different ways. For the purpose of working through their implications for public finance, we will group the deviations into three categories according to the general feature of decisionmaking that drives the deviation: *limited attention*, *limited computational capacity*, and *biased reasoning*. Limited attention captures deviations from optimality that appear to be due to the fact that there are limits to the bandwidth of the human brain in processing stimuli—that individuals cannot notice and attend to all of the features of choice simultaneously. Limited computational capacity captures deviations that are due to the limits of the processing power of the human brain—that even when individuals are capable of attending to the relevant features of a choice, making some choices simply is complex or otherwise intrinsically difficult. Biased reasoning captures deviations from optimality that are due to a set of persistent biases in the way that the human brain appears to subjectively evaluate alternatives, especially those involving probabilities or statistics.

Limited Attention

Optimal choice generally requires actively considering the properties of multiple alternatives along multiple dimensions. Unfortunately for human decisionmakers, psychologists have observed that individuals have a limited capacity to attend to multiple features of choice simultaneously.⁵ The mind appears able to attend to only a small fraction of the stimuli that it perceives, and that attention is focused in a way that is neither random nor entirely conscious. As a result, individuals can focus on, or attend to, only a few of the many features of their choice environment at once. As a result, choice becomes sensitive to the way in which attention is allocated or directed. That can cause individuals to ignore some features of choice and to be excessively sensitive to others, depending on the extent to which those features attract attention. It can also lead individuals to construe their choices in artificially narrow terms as they direct their attention across the features of choices, leading to locally rather than globally optimal choice. Following the psychology literature, we refer to this feature of decisionmaking as limited attention.

Limited attention is responsible for several features of observed choice behavior that are either broadly inconsistent with optimal choice or at least puzzling from the standard perspective. Two that are important for public finance are *salience effects* and *local construal*.

SALIENCE EFFECTS

Because individuals cannot attend to everything at once, salient features of their environment will command their attention and can influence behavior and choice. An illustrative finding from psychology is that while individuals are generally unable to simultaneously process a second set of words that they hear while paying attention to a first, an exception occurs when the second set includes a person's own name—literally an attention-grabbing word for most individuals.⁶ This result is sometimes labeled the cocktail party effect because of the way that it mirrors the familiar experience of overhearing, but not following, chatter at a cocktail party until someone speaks your name, which you immediately recognize. The relative salience of different features of choice directs attention in a similar way, and in doing so guides choice. Cues that direct attention toward or away from particular options or that highlight or conceal specific characteristics of alternatives can affect behavior even when the underlying choice set is preserved. For example, items in grocery stores sell better on shelves at eye level, where consumers' attention is focused by default, than at other shelf heights.⁷

In general, more salient features of choice get access to the limited attention of decisionmakers, while less salient features do not. That seems to be true in policy contexts, as well. For example, there is evidence that raising or lowering the

salience of taxes or fees, without changing their level, affects behavior.⁸ Another consequence of salience effects is that individuals can have trouble ignoring salient information even when they want to—told not to think of a white bear, many people will immediately conjure an image of a white bear.⁹ That can be economically significant when, for example, it would benefit individuals who have an informational advantage to be able to predict the behavior of those who do not—they may be unable to bring themselves to ignore their private information.¹⁰

LOCAL CONSTRUAL

The other set of effects of limited attention on choice comes about because individuals with limited attention can direct that attention. The ability of individuals to direct their limited attention is powerful. Psychological research that asks individuals to pay attention to one part of an image or video often finds that those individuals fail to notice even unusual or striking images when their attention is focused on the part that they were instructed to observe. In perhaps the most famous such example, when asked to count the number of passes made in a video of people playing with a basketball, many observers failed to notice an individual walking across the frame in a gorilla suit.¹¹ The result of focusing attention for choice is that it can lead to choice processes that result in what are local, rather than global, optimization patterns. For example, individuals may engage in elimination by aspects—whereby they consider aspects of available alternatives one at a time, eliminating options that are undesirable according to each aspect in sequence—or in similar choice behaviors.¹² Another type of local optimization that may be driven to some extent by limited attention is choice that narrowly construes not the choice set, but the hedonic consequences of choice. In particular, individuals may focus on immediate or salient outcomes rather than the full range and path of outcomes.¹³

Limited Computational Capacity

While limits to attention underlie many of the specific decisionmaking errors and biases that psychologists and others have observed, other anomalous behaviors appear to reflect a deeper set of cognitive limitations. Even when individuals are not constrained in terms of attention, they can find some choices hard to make because of the complexity of evaluating the alternatives and because they are not unbounded in their capacity to think and reason. Individuals appear to have difficulty thinking and reasoning accurately or consistently about choices and preferences. They hold subjective valuations that are inconsistent or arbitrary. They have difficulty penetrating opaque pricing schedules. They exhibit evidence of an inability to integrate decisionmaking across domains. And their decisions can be influenced by spurious features of the choice environment. In general, we group

findings of this type from the psychology and behavioral economics literatures as evidence of limited computational capacity.

The main implication of limited computational capacity for economic behavior is that optimization generally is only approximate, not accurate or precise. We collect a variety of specific classes of choice anomalies under this broad heading: *decisional conflict*, *inconsistent subjective valuation*, “*schmeduling*,” and *mental accounting and choice bracketing*.

DECISIONAL CONFLICT

The clearest demonstration of how individual choice is affected by the processing limitations of the human brain is the direct evidence that individuals appear to find the process of choosing itself to be difficult under some conditions. In the jam example above, the proliferation of choices appeared to inhibit choosing any alternative at all, a condition sometimes labeled choice overload.¹⁴ That is an extreme example of a more general finding that individuals sometimes seek to avoid making choices. More specifically, there is evidence that what individuals seek to avoid are difficult choices. When individuals face choices among options with no clearly dominant alternative, they are more likely to look for ways to avoid choosing, such as seeking additional alternatives or deferring choice, than when a dominant option is available.¹⁵

INCONSISTENT SUBJECTIVE VALUATION

An essential element of optimal choice is that it is based on an underlying set of consistent preferences. A range of evidence from behavioral economics suggests that individuals in fact have a difficult time forming consistent subjective valuations. Valuations instead appear malleable and arbitrary, as demonstrated in contexts in which alternatives have attributes that are not easily valued or that vary along multiple dimensions. For example, individuals often reverse their stated preferences when they are given choice attributes jointly instead of separately.¹⁶ Valuations of positive and negative attributes of alternatives differ depending on whether individuals are selecting or rejecting alternatives.¹⁷ And the attributes that individuals base their valuations on can be difficult to view as the result of perfect optimization. In one example, individuals tasting wines were found to peg their valuations of different wines—as indicated through brain imaging—to the price of the wine rather than the taste.¹⁸

Perhaps most dramatically, other results suggest that individuals’ preferences can be influenced by external cues that have no plausible connection to subjective value. For example, experiments have shown that reminding individuals of the last two digits of their Social Security number affects how they value goods—individuals with higher numbers will tend to value goods more highly than those with low numbers, even while being reminded of the arbitrariness of their Social

Security number.¹⁹ Finally, preferences appear to be very sensitive to the way in which choices are structured. The addition or subtraction of alternatives, even irrelevant alternatives, can also lead to preference reversals. For example, individuals often are influenced by the introduction of asymmetrically dominated alternatives, whereby adding a third alternative causes individuals to switch their preference over an initial pair.²⁰ In a similar manner, the existence of extreme alternatives can promote the selection of intermediate options.²¹

SCHMEDULING

If one cornerstone of optimal choice is the presumption that individuals can form and access consistent subjective valuations of choice alternatives, the other is that individuals correctly perceive their cost. Schmeduling is a label, coined by economists Jeffrey Liebman and Richard Zeckhauser, for behaviors that appear to be a result of difficulties that individuals may have with understanding price schedules—that is, with knowing what they are paying. It refers to a tendency of individuals to hold and act on only approximate mental representations of price schedules.²² Individuals are thought to be susceptible to two types of errors in particular, which are to incorrectly smooth price schedules, such as by mistaking average for marginal prices, and to incorrectly respond to local prices when the full schedule of prices is relevant for decisionmaking. Findings from psychology, such as those on the tendency of individuals to respond to incentives in a way that is attractive piecewise but suboptimal in the aggregate, suggest the susceptibility of individuals to such tendencies.²³

However, the bulk of the evidence for this difficulty comes from choices that individuals make in economic settings that are consistent with the hypothesis that individuals respond to complex price schedules in ways that are hard to square with perfect optimization. For example, the failure of incomes to bunch around the points in the income tax schedule where tax rates change discretely (kink points) is consistent with individuals responding to average rather than marginal tax rates.²⁴ Individuals are thought to face similar difficulties when the price schedules of consumer goods are complex.²⁵ Evidence suggests that individuals choosing prescription drug plans in Medicare Part D had difficulty choosing the least costly plan, a result due in part to the complexity of the price schedules involved.²⁶

MENTAL ACCOUNTING AND CHOICE BRACKETING

A further important set of deviations from perfect optimization is captured by the concepts of mental accounting and choice bracketing. Mental accounting is the tendency of individuals to evaluate choices with respect to discrete, notional accounts rather than general measures of financial status, such as overall wealth, total income, or total spending.²⁷ The clearest cases of mental accounting come

from household budgeting behavior, wherein individuals routinely treat income from different sources as flowing to, or spending for different purposes as drawing down, distinct mental accounts. For example, individuals spend money differently—usually more frivolously—when they perceive it as having been won or found rather than having been earned.²⁸ Individuals can also be more or less willing to take actions depending on the mental account to which they post—credits or debits.²⁹ In one well-known study, individuals were more willing to drive twenty minutes in order to save \$5 on a \$15 calculator than to save \$5 on a \$125 jacket, apparently because they evaluated the \$5 discount relative not to their overall wealth but to a mental account out of which they were spending \$15 in one case and \$125 in another.³⁰

Such behavior may explain similar relationships between the source and disposition of funds that we observe in policy contexts. For example, tax benefits seem to be more likely to lead to increased spending on children simply by virtue of being labeled child credits, possibly by affecting the mental accounting of the benefit.³¹ Mental accounting can also affect how individuals make choices about the time path of consumption, payment, and debt, depending on how individuals form the respective accounts.³² The process of structuring and balancing mental accounts is closely related to another dimension, sometimes known as choice bracketing, in which limits to computational capacity can affect choice.³³ Individuals can choose to use broader or narrower brackets, and the bracket used will have an impact on choice. For example, individuals may be more inclined to commit to making small, recurring payments, such as to a charity, because they bracket the choice narrowly—comparing the payments to other small, frequent expenses rather than considering the aggregate expense.³⁴

Biased Reasoning

A final category of behavioral tendencies that are broadly inconsistent with optimal choice has to do with statistical reasoning and judgments of probability. These deviations take the form of biases that individuals exhibit when assessing the probabilities associated with risky choice or when making judgments about their own place in the distribution of possible outcomes. These deviations are slightly different from those above in that they appear to reflect not limits to processing capacity but a set of persistent biases in the way that the human brain processes probabilities. Put another way, if limited attention and limited computational capacity are for the most part a result of the imprecision of the human brain as a decisionmaking organ, these deviations are about its inaccuracy. We group this set of behaviors into a category we call biased reasoning. Biased reasoning of this sort is manifested in two broad categories, *probabilistic reasoning* and *motivational biases*.

PROBABILISTIC REASONING

Individuals appear to have difficulty making correct or consistent decisions under uncertainty.³⁵ For example, individuals appear to employ an availability heuristic, in which they assess an event or outcome as more or less likely depending on how easily it can be thought of or imagined.³⁶ Similarly, individuals appear to employ a representativeness heuristic, in which they tend to ignore the relative frequency of alternatives in probability judgments.³⁷ In addition, individuals appear to systematically overweight low probabilities and underweight high probabilities in decisionmaking.³⁸ That is, they make decisions treating very unlikely events as more likely to occur than they are and likely events as less likely to occur than they are. They mistake randomness for patterns.³⁹ These biases also are consistent with behavior outside the lab. The same availability heuristic that leads to deviations in controlled settings is seen in the behavior of individuals who purchase flood insurance after being hit by a flood rather than before.⁴⁰ The underlying probability of being flooded in any given year is unlikely to change; what changes is that the flood itself causes individuals to appreciate the risk differently.

MOTIVATIONAL BIASES

Related but slightly different are biases in probability assessments related to individuals' chances of success in their own endeavors, what are sometimes referred to as motivational biases. One such result is overconfidence. Individuals are found to be routinely overconfident about their own abilities.⁴¹ For instance, the majority of drivers believe themselves to be better-than-average drivers.⁴² Overconfidence also appears to be related to some economic behaviors, like risk taking by entrepreneurs.⁴³ A related but distinct bias is a tendency toward over-optimism, of which there also is evidence.⁴⁴ For example, unemployed workers appear to be excessively optimistic about their chances of finding work, which appears to affect their search effort.⁴⁵ Individuals also appear to possess a self-serving bias, a tendency to consider their own self-interested judgments as fair; that tendency has been shown to lead to difficulties in negotiations.⁴⁶

Summary

Taken together, limited attention, limited computational capacity, and biased reasoning have broadly similar consequences in that they allow for the possibility that individuals make systematic errors in attempting to maximize their own utility. They lead individuals to make decisions based on heuristics and biases. They suggest that rather than the kind of deliberate choice that the standard model envisions, individuals more often use shortcuts or crude rules of thumb that can be incorrect. They each, in their own way, show how making good decisions

is hard—much harder than the standard model emphasizes. They are, broadly speaking, manifestations of imperfect optimization.

Finally, it is worth noting before moving on that none of these limitations—limited attention, limited computational capacity, or biased reasoning—is tantamount to limited intelligence. They reflect decisionmaking as practiced, not capability. Moreover, they appear to be nearly universal features of decisionmaking. Students at MIT give intuitive but incorrect answers to questions designed to prey on cognitive biases, just as anyone else does.⁴⁷ Even professionals such as physicians demonstrate these biases within their areas of expertise.⁴⁸ So these limitations are not related to intelligence, or even expertise. They simply appear to be properties of the way that the human brain forms judgments and makes decisions most of the time. There is even some reason to think that such errors and biases may be adaptive. Some evidence indicates that in certain contexts people are more satisfied with their choices when they neglect conscious deliberation and rely on heuristics instead.⁴⁹

Bounded Self-Control

In addition to assuming that individuals are good at knowing what to choose, economists further assume that individuals are good at implementing their choices—in particular, that they possess what can be broadly termed self-control, that they do not face any tension between what they intend to do and how they act. Slightly more formally, we might say that the standard economic model is one in which choices display time consistency. When choices are time consistent, consumption patterns observed *ex post* are consistent with consumption plans made *ex ante*.

But just as psychologists and economists have found that individuals can have difficulty knowing what they want, in the case of imperfect optimization, behavioral economists have uncovered evidence that individuals have difficulty doing what they want. In addition to failures of reason, individuals also often suffer from failures of self-control. Individuals choose and act in ways that are time-inconsistent, and they often display a bias for present over future consumption. Consider, for example, the finding that individuals' preferences often are inconsistent over delayed rewards.⁵⁰ Individuals often prefer to receive a larger delayed reward later in the future, but a smaller more immediate reward today: for example, they prefer to receive \$110 in thirty-one days rather than \$100 in thirty days, but they prefer to receive \$100 today rather than \$110 tomorrow. This behavior is time inconsistent: if both choices were executed, on day thirty the individual would find that she had committed to a path that she no longer found optimal.

Similar conflicts are evident in real-world behaviors. For example, individuals have been found to face a similar choice when selecting pricing plans for gym membership. Expensive monthly plans make sense only if individuals intend to go to the gym a sufficient number of times over the month to make the average daily cost of the monthly plan less than the price of a day pass. Research finds that in fact individuals who choose the monthly pass attend the gym too few times during the month to make it worthwhile.⁵¹ Time-inconsistent individuals decide months or weeks in advance that on some fraction of future days they will want to go to the gym but then, when many of those days actually arrive, they decide that they would rather not go.

The failure of individuals to display time-consistent preferences is an example of a general tendency that we label bounded self-control. A multitude of findings from psychology and behavioral economics indicate that individuals make choices over time in ways that are broadly inconsistent with the standard model.⁵² In general, translating intention into action seems to involve difficulties that the standard model does not allow for and results in behaviors that it does not predict and cannot easily accommodate. People sometimes do things that they really do not want to do or fail to do things that they wish they had done. They can be influenced toward or away from actions by minor inducements or inconveniences. And so on. Among the many manifestations of this general tendency, important classes of behavior include *procrastination and temptation, channel factors, state and affect, and addiction*.

PROCRASTINATION AND TEMPTATION

One major consequence of bounded self-control is the gap that it can create between intention and action. As in the case of gym membership and exercising, individuals may engage in procrastination, failing to take actions that they intended to take. Conversely, it can lead individuals to succumb to temptation, taking actions from which they intended to refrain. For example, when choosing for immediate consumption, individuals prefer junk food and trashy movies, even while stating a preference for healthy food and high-brow films when making plans for later.⁵³ One important source of evidence that individuals are subject to temptation is their demand for and behavior in the presence of commitment devices—for example, deadlines can be shown to improve student outcomes.⁵⁴ Similarly, the willingness of individuals to engage in illiquid forms of savings even in the absence of higher expected returns, apparently to avoid the temptation of consuming out of savings, is another consistent piece of evidence.⁵⁵ The role of commitment devices also serves to highlight the fact that while individuals have only bounded self-control, that does not imply that they are necessarily naïve about their lack of self-control.⁵⁶

CHANNEL FACTORS

One of the most striking results in psychology is that allowing for a gap between intention and action, research finds that only very minor features of choice can serve to widen or narrow that gap. Psychologists have labeled those features of choice channel factors.⁵⁷ Channel factors can explain the tendency of individuals to be steered toward or away from choices by ostensibly quite minor barriers or inducements. One classic study in psychology finds, for example, that a message to receive an immunization was many times more likely to result in an individual following up and receiving the shot if the initial message was accompanied by a map to the health clinic and a request that the person decide on a time to get the shot.⁵⁸ Moreover, individuals often can be unaware of the influence of these factors on their own intentions—in one study in which minor cues such as reminders were found to influence behavior, those cues did not affect self-predictions about behavior.⁵⁹ Effects consistent with channel factors also are observed in many real-world contexts. The dramatic results of automatic or simplified enrollment procedures in social programs, such as college financial aid programs or employment benefit programs such as 401(k) plans, are likely due in part to channel factors.⁶⁰

STATE AND AFFECT

Another important aspect of bounded self-control is that the ability of individuals to exhibit self-control depends not just on the context of choice but also on the state of the decisionmaker. There are at least two ways in which state and affect can influence the ability of individuals to take actions that match their intentions. First, when individuals find it difficult to exert self-control, other aspects of their mental state can modulate their ability to overcome that difficulty. For example, stress and cognitive load may cause individuals to act impatiently—something as simple as asking people to hold a long string of digits in their head can make them more likely to select a more tempting, less healthy snack.⁶¹ Outside the lab, episodes of stress have been found to have a similar effect on the ability of quitters to refrain from smoking.⁶² Similar effects may result from other visceral states, such as hunger or fear.⁶³

Second, state and affect can play a role in time-inconsistent behavior to the extent that the inconsistency comes about because of the difficulty that individuals have in predicting their hedonic state, or forecasting their affect, at the time of forming their intentions. In particular, individuals tend to display what has been labeled projection bias—a tendency to project their current preferences onto their future selves.⁶⁴ So, for example, individuals will elect to receive more or less healthy snacks one week from now depending on whether they are hungry now.⁶⁵ Similar effects have been observed with catalog orders, when individuals were

more likely to return orders for cold weather gear when orders were placed on unusually cold days, suggesting that at the time of placing the order, individuals were projecting a desire for such gear that did not persist when the order arrived and the weather had improved.⁶⁶

ADDICTION

Finally, at the extreme, individuals may lose self-control outright due to addiction, which is a behavioral tendency worth separating from the others. Tobacco use and smoking—and substance use more generally—is the common case, but other behaviors, such as gambling, also have properties of addiction. Standard economic analysis can accommodate even models of addictive behavior.⁶⁷ But evidence suggests that characterizing addiction as a process by which individuals lose the ability to maintain self-control may fit the data better.⁶⁸ Recent models of addiction, which are more grounded in the physiology of the brain, argue that addiction is a matter of substances or behaviors leading to a direct interference with the ability of the brain to forecast hedonic states.⁶⁹ Those models capture common features of addiction, such as a stated preference for quitting in the face of the oftentimes practical inability to do so.

Summary

Before moving on it is worth commenting on a feature of our aggregation and classification of psychological findings into behavioral tendencies that is especially evident in our discussion of bounded self-control. We are categorizing behaviors according to how they operate and in a way that will be useful for thinking about their consequences for public finance. Consequently, from the perspective of psychology or behavioral economics, the categorizations may be somewhat loose with respect to both the underlying nature of the behaviors and some related terminology. With respect to the underlying psychology of failures of self-control, for example, there are alternative models and hypotheses that we subsume in our discussion. Failures of self-control can be thought of as a result of present-biased preferences due to quasi-hyperbolic discounting.⁷⁰ They can also be thought of as a result of conflict between the mental processes by which individuals plan and those by which they act.⁷¹ Alternatively, they might be thought of as a result of a decisionmaking process in which self-control demands willpower and willpower is costly to exercise.⁷² They also can be thought of as a result of individuals construing the time dimension of choice in some nonstandard way.⁷³ In part as a result of the variety of processes that might in fact generate such behaviors, it is somewhat imprecise to label all of the behaviors described here as features of bounded self-control. That loss of precision at this stage is deliberate—it is the cost of having a convenient shorthand for referring to classes

Box 2-1. *Bounded Rationality*

The integration of findings from psychological research and economic analysis described here owes much to the behavioral economics literature of recent decades that follows the work of Daniel Kahneman and Amos Tversky.^a This literature stresses the way that the human brain ordinarily approaches decisionmaking. It highlights the shortcuts, the heuristics and biases, that individuals commonly employ and the ways in which they lead to decisionmaking patterns that are at odds with the patterns that a model of fully optimal decisionmaking would predict.

Before applying a behavioral approach to public finance, it is worth pausing to note that there is an older strain of research, going back to the work of Herbert Simon, that describes a slightly different approach to thinking about the implications of psychological realities for decisionmaking and economic analysis.^b This alternative line of research also acknowledges the limits of the human brain as a decisionmaking organ and recognizes that there are costs to thinking and deciding. But this approach preserves the possibility of what is referred to as bounded rationality: that individuals make optimal decisions subject to those constraints. That is, rather than limits to attention or computational capacity leading directly to imperfect decisionmaking, in this model individuals optimally allocate attention and computational capacity. Individuals remain limited in their capacity to choose optimally and consistently, but they can be savvy about how they manage those limits.

Which model is correct, in the sense that it best describes how individuals make decisions and why we observe behavior that violates standard assumptions, is ultimately an empirical question. The answer is surely a mix of both. Moreover, it is important not to fashion out of the distinction a false dichotomy: the two approaches are clearly related. That said, we tend to take the approach described in this chapter—of stressing imperfections in decisionmaking rather than the bounds on rationality—and we do so for several reasons.

of behavioral tendencies when we move on to our central goal of reviewing public finance through this lens.

That said, our approach to considering these behavioral tendencies does make one substantial but implicit assumption about the decisionmaking process that generates these tendencies. In particular, we follow recent developments in behavioral economics in taking the view here that both imperfect optimization and bounded self-control, in general, derive from an underlying psychology of judgment and decisionmaking that leads individuals to act in ways that are inconsistent or even erroneous, not from a considered judgment on the part of individuals about how to deploy limited cognitive resources or manage limited reserves of willpower. For more on this distinction, see box 2-1.

First, this is the direction in which psychological research has trended in recent decades. Many of the pieces of individual evidence described here are difficult to view as consistent with bounded rationality. For example, the evidence on the importance of context and situational factors, whereby extremely minor or apparently irrelevant features of the choice environment affect behavior, seems more of a piece with a model in which individuals are simply imperfect decisionmakers.

Second, in other contexts that suggest direct tests of imperfect decisionmaking models against cost-of-thinking models, we view the evidence as favoring imperfect optimization or bounded self-control. The payoffs to decisions such as electing to participate in 401(k) plans or choosing prescription drug plans optimally in Medicare Part D are so large compared with the costs that they are a poor fit for models of bounded rationality. That is, a decisionmaker who was optimally allocating her scarce attention or computational resources would almost surely have devoted it to making such high-return decisions in an optimal manner.

Finally, to the extent that the different models matter for policy design, the imperfect decisionmaking model has in some ways done a better job at identifying opportunities *ex ante* than models of bounded rationality. Even if, *ex post*, costs of thinking or processing information can explain outcomes such as the failure of qualified individuals to apply for college financial aid, the hypothesis that application assistance could be important came out of an imperfect decisionmaking model. Even without settling the question, for the narrow purposes of this book, that is a substantial practical advantage of the imperfect decisionmaking approach.

a. Daniel Kahneman and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, vol. 47 (March 1979), pp. 263–91.

b. Herbert Simon, "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, vol. 69, no. 1 (1955), pp. 99–118.

Nonstandard Preferences

Usual economic assumptions about choice include some weak assumptions about the shape and content of preferences. Two, in particular, are relevant here. First, economists typically assume that individual utility is a function of end states—that is, how individuals value an outcome usually does not depend on the path taken to realize it or on the position of the outcome relative to other possible outcomes, but simply on the outcome itself. Second, economists commonly assume that individuals are purely self-interested. It should be noted that in neither case are those assumptions essential features of the standard model of choice; they are instead standard simplifying assumptions.

Psychology and behavioral economics have produced findings that suggest that in many cases those assumptions are a poor fit with the preferences exhibited by many decisionmakers. Take, for example, the results of a study investigating how individuals form valuations in markets.⁷⁴ In that experiment, the authors first randomly distributed coffee mugs to half of the participants, leaving the remaining participants empty handed. They then asked each participant with a mug to indicate the price at which he or she would be willing to sell the mug and each participant without a mug to indicate the price that he or she would be willing to pay to acquire a mug. Using that information, the researchers set a market-clearing price and executed the trades indicated at that price. Because the mugs had been allocated randomly, the expectation was that about half of the participants would trade mugs in that market arrangement. But, in fact, very few trades occurred—only about one-tenth of the mugs were traded. What happened? Against expectations that valuations would be randomly distributed throughout the class, the mugs came to be systematically valued more highly by those to whom they were initially allocated than by those to whom they were not. Those given mugs were, on average, unwilling to sell them for less than about \$5, while those who initially did not receive mugs were, on average, willing to offer only about \$3 for a mug. Merely assigning ownership of the mug appeared to affect how much participants came to value it. Preferences were formed with respect to the initial allocation—people put a value on giving up the mug or on acquiring the mug—not with respect to an abstract valuation of the mug. Similar results have been observed in the field. For example, homeowners and homebuyers sometimes display similar preferences.⁷⁵

This result, an example of what is known as the endowment effect, is a consequence of reference-dependent preferences, a violation of the assumption that individuals' preferences are over end states. Reference-dependent preferences are one type of nonstandard preferences that is especially important for public finance. Other-regarding preferences—a violation of the assumption of pure, or nearly pure, self-interest—is another.

Reference-Dependent Preferences

Choice theory in economics typically assumes for the sake of simplicity that goods enter individual utility functions in absolute terms. That is, goods have an intrinsic value that does not depend on how they compare with alternatives. In many instances however, individuals appear to evaluate many choices in relative terms, in particular in comparison with some reference point.⁷⁶ Preferences over alternatives might depend on whether an alternative represents a gain or a loss relative to expectations or to prior experiences. They may depend on whether individuals are valuing a good to sell it or to buy it. Or they may depend on their

relationship to the status quo. These results are manifestations of what behavioral economists refer to generally as reference-dependent preferences.

Among the catalog of choice anomalies observed by behavioral economists, several important examples are arguably a result, at least in part, of reference-dependent preferences. They include the *endowment effect*, *loss aversion*, and *status quo bias*.

ENDOWMENT EFFECT

The endowment effect is described above in the coffee mug experiment. The general finding is that where individuals start from, in terms of their endowment, matters for choice because it creates a reference point that affects how they value outcomes. The authors interpret the findings in the coffee mug example as evidence that individuals think of, and subjectively value, the experience of acquiring a good differently from the experience of giving one up. Parting with an item that individuals think of as their own seems to hurt them more than acquiring the same item benefits those who do not own it. As a consequence, owners required more compensation to give up the mugs than non-owners were willing to pay to obtain them. More generally, the endowment effect may result from individuals setting reference points around expectations—evaluating outcomes relative to those expectations.⁷⁷ That interpretation can also incorporate apparent rejections of the endowment effect, as exceptions that prove the rule: where experimental conditions mitigate expectations of continued ownership, the endowment effect will not manifest.⁷⁸ One important consequence of the endowment effect is that willingness-to-pay valuations may not match willingness-to-accept valuations.

LOSS AVERSION

Another reference point around which individuals tend to form preferences is zero; that is, individuals do not value or experience losses and gains symmetrically. This result is known as loss aversion, because of the consistent finding that individuals perceive losses more intensely than gains.⁷⁹ That is, to give someone with loss aversion some amount of money and then take it back would not leave the individual's welfare unchanged, as in the standard model—rather, the individual would feel worse off, because paying the money back would reduce his or her welfare by more than the original gift increased it. That effect has been demonstrated in a number of contexts, including policy-relevant ones. For example, experimental work suggests that the tendency of individuals to spend out of a tax cut might be sensitive to whether the cut is framed as a tax rebate or a bonus because one is perceived as repayment of a loss and the other as a gain.⁸⁰

One important consequence of loss aversion is that individuals can express what appear to be odd preferences with respect to risk. In particular, individuals can make choices that reveal an extreme aversion to risk.⁸¹ What looks like

extreme aversion to risk may be driven in part by an aversion to the possibility of loss (combined with the tendency of individuals to overweight small probabilities, as described above). That has some real-world consequences. Loss aversion may be behind why individuals insure against very small risks or choose to have very low deductibles.⁸² Loss aversion also has the effect of making individuals risk averse with respect to gains, as standard assumptions predict, but possibly risk loving with respect to losses—that is, individuals may be willing to accept risk in order to avoid greater losses. One important demonstration and implication of the effects of loss aversion is that whether choices are framed as losses or gains can have a measurable effect on choice.⁸³

STATUS QUO BIAS

Another consequence of reference-dependence is status quo bias, the tendency of people to stick with what they have. It was originally noted in the tendency of individuals to stick with their health insurance plan and retirement options over time.⁸⁴ Similar results have been found in other contexts as well, such as individuals' preferences with respect to service options from utility providers.⁸⁵ This effect operates at least partly in conjunction with other tendencies—such as procrastination—but it also seems to be partly a function of using one's current situation as a reference point in evaluating alternatives. The effectiveness of defaults in promoting enrollment in employment benefits and social programs, noted above as consistent with the effect of channel factors, is also reinforced by status quo bias.

Other-Regarding Preferences

One final assumption of the standard model that leads economists and policy-makers astray is the assumption that people are purely self-interested. While it is only a simplifying assumption on the part of the standard model, it is central to a number of specific results, including results in public finance. Findings from psychology and behavioral economics suggest that preferences and choices are interdependent in a wide variety of ways. People care about the outcomes realized by others, or at least they act as if they do. They care about the outcomes for groups and how those outcomes are generated. They care about how their choices compare with those of others and how they are viewed by others. And so on. In general, we categorize the ways in which individual preferences are related to the choices and outcomes of others as demonstrations of what we label other-regarding preferences. There are several facets to other-regarding preferences that are relevant to the economics of the public sector. They include *altruism*, *fairness*, *social norms*, and *interpersonal preferences*.

ALTRUISM

Evidence from multiple domains supports a view of human nature that is less dim than what economists typically suppose. Rather than pursuing narrow conceptions of self-interest alone, people frequently act as though they care about the outcomes of others, either individually or as a group.⁸⁶ Many results from laboratory experiments, for example, are inconsistent with strong forms of self-interest. Consider the results of the “ultimatum game,” which has been repeated many times in many different contexts: There are two players, one of whom proposes a certain split of a pool of money. The other player can either accept the proposed split—in which case it is realized—or reject it, in which case neither player receives a payout. Were the game played by purely selfish individuals, a bare minimum offer should be made and accepted. In fact, when this game is played, offers tend to be around 40 percent of the pool.⁸⁷ In addition, the consistency of this behavior suggests that rather than being an artifact that might result from imperfect optimization, findings of this sort reveal preferences for altruism.⁸⁸

One important consequence of this behavior for public finance is that individuals may engage in what amounts to voluntary redistribution. Indeed, an important piece of real-world evidence for altruism comes from donations to charities, which are substantial and difficult to explain if individuals are purely selfish.⁸⁹ Another important consequence for public finance is that this behavior may lead to voluntary contributions to public goods. Evidence from laboratory experiments suggests that individuals tend to contribute to public goods in excess of what an assumption of pure self-interest would predict.⁹⁰ There also is some evidence from the field that voluntary contributions are made to public goods such as public radio stations and schools that are difficult to reconcile with pure self-interest.⁹¹

FAIRNESS

A related finding but one that has distinct consequences is that individuals have preferences with respect to the process that generates outcomes, as well as the outcomes themselves.⁹² That is, there is evidence that individuals have preferences for fairness. Survey responses, for example, indicate that individuals value fairness in price and wage setting.⁹³ Behaviors in experimental settings also are consistent with preferences for fairness. For example, individuals appear to value cooperation and more generally to act in accordance with reciprocity.⁹⁴ Those behaviors are in many ways more consistent with a taste for fairness than strategic behavior.⁹⁵

SOCIAL NORMS

Individuals are influenced by the behavior of others and by the way that others expect them to act to an extent that is surprising in the standard model.

Individuals will often behave in a way that conforms to community norms.⁹⁶ For instance, results suggest that in addition to an intrinsic preference for fairness, a desire to be perceived as fair by others—in particular, by following the norm of splitting rewards evenly—partly drives such results in laboratory games.⁹⁷ Field experiments confirm that social norms influence behavior. For example, individuals given a flyer are less likely to dispose of it by littering in environments that have been manipulated to be relatively free of litter.⁹⁸ Direct messages that indicate to people that most other people behave a certain way have been found to promote conformity to that behavior.⁹⁹ In one striking set of findings, showing individuals how their consumption of residential energy use compared with that of their neighbors and framing above-average energy use as undesirable was found to reduce energy consumption.¹⁰⁰

INTERPERSONAL PREFERENCES

A final set of interrelationships among the choices that individuals make arises from the fact that people care how they are viewed by and how they are positioned relative to others. For instance, in a set of results that combines reference-dependent preferences and other-regarding preferences, individuals often are found to have positional preferences.¹⁰¹ That is, their utility is a function of their outcomes relative to the outcomes of others. For example, in one study, individuals were found to be less happy when their neighbors had higher earnings.¹⁰² Another result of interpersonal preferences is that individual choices have been found to depend to some extent on how individuals identify socially.¹⁰³ For instance, individuals exhibit preferences that depend on which of their multiple social identities—for example, student, employee, spouse, American—is salient at the time of choice.¹⁰⁴

Summary

What these findings on reference-dependent and other-regarding preferences indicate, ultimately, is that results in public finance that are sensitive to assumptions about the form and content of preferences require review. For example, as noted above, the levels of private contributions to public goods, such as public schools, are sometimes higher than predicted by the standard model, a result that can be explained in part by the existence and nature of other-regarding preferences. Conclusions about the efficient level of public provision of public goods such as these, derived under assumptions of perfect self-interest, will need to be revisited in light of empirical evidence with respect to other-regarding preferences.

What the results on nonstandard preferences should not be interpreted to mean is that they have taken down a straw man. Pure selfishness and reference-

independent preferences are not central features of the standard model of decisionmaking; they are merely simplifying assumptions that have been convenient to use in economic analysis. The issue becomes one of whether, for example, self-interest or altruism is the more appropriate operating assumption for how economic agents interact with the world. Are expressions of altruism unimportant or rare deviations, or are they common, widespread, and important for understanding behavior? Questions such as these are an empirical matter. A similar caveat, in fact, applies not just to the findings on nonstandard preferences but to all of the findings described in this chapter—the standard model is a set of assumptions, and the issue is not whether those assumptions ever fail, but whether they fail so regularly that they are worse operating assumptions than the alternatives.

Ultimately, the important question for public finance, in all cases, is whether allowing for these key behavioral deviations from the usual assumptions—imperfect optimization, bounded self-control, and nonstandard preferences—matters. They seem on their face to create the possibility that results in public finance may change if revised assumptions about individual decisionmaking are incorporated. But do they? Do choice errors matter in the aggregate or in equilibrium? Do individual decisionmaking failures matter for market failures? And so on.

The answers are not obvious. For one, it may be that individuals exhibit these behaviors in experimental settings but not in real life, where the stakes are higher and the influence of experience and learning may be more substantial.¹⁰⁵ There is some evidence that behavioral tendencies that stand out in the lab can be attenuated in the field or in agents with greater experience or strong financial incentives. For example, the endowment effect can be less pronounced among individuals with more experience in relevant markets.¹⁰⁶ However, given evidence from the field, this is clearly not always true, as in the case of automatic enrollment in retirement saving plans. Separately, it may be the case that markets operate in a way that neutralizes the effects of individual decisionmaking errors on aggregate outcomes. That is theoretically possible;¹⁰⁷ however, so is the converse.¹⁰⁸ Tests under market conditions, including the evidence from the field and simulated markets, suggest that markets are not always sufficient to enforce the outcomes predicted by the standard model.¹⁰⁹

To understand the true implications of behavioral economics for public finance, then, we cannot simply apply findings from psychology directly to issues in public finance piecemeal. It is necessary instead to integrate findings on behavioral tendencies into the economic framework of public finance—into the analysis of externalities and asymmetries of information, and so on—and work through their implications for the role of government and for the design of public policy.