

1

Introduction

When should the government intervene in the economy? When do markets fail? How do we craft policies that maximize social welfare? How do we design policies to minimize unintended consequences? Traditional public finance provides a powerful framework to tackle those questions. This framework, however, relies on an overly simple model of human behavior. This book revisits the core questions of public finance but with a psychologically richer perspective on human behavior. We do not merely apply psychology to economic problems; instead, we explore how psychological factors reshape core public finance concepts such as moral hazard, deadweight loss, and incidence.

The Promise

To build our case, we construct a single analytical framework that encompasses both traditional policy levers—taxes and subsidies—and psychologically informed ones—such as defaults and framing. Three examples—health insurance, taxes, and externalities—illustrate how this approach alters our understanding of basic policy problems.

Health Insurance

Models of health insurance emphasize moral hazard. Individuals choose care by comparing the price of care with its benefits. Since under insurance the price of care is often below its actual cost, people may overuse it. For example, because a consumer pays only a fraction of the full cost of an MRI, he or she may decide

to get one even if it provides minor benefits. Insurance design seeks to balance the benefits of insurance against inefficient overuse, such as through copayments, health savings accounts, or consumer-directed health plans.¹ For our purposes, notice how the logic of overuse relies crucially on individuals making choices in a narrow, calculating fashion: it occurs because consumers make a trade-off between the price of care and its true benefit.

Medical studies, however, suggest that health care choices are significantly more complex.² Take the case of a diabetic who is prescribed medication. The cost-benefit calculus for taking the medication is clear cut. Diabetes is a serious disease, and insulin provides an important tool to manage it: the long-term health benefits drastically outweigh the monetary and “hassle” costs of buying and taking the medication. Human psychology can short-circuit that calculus. A patient focused on day-to-day concerns may simply forget to take his medication; another patient may simply “feel good” and decide that taking the medication is not worth it; and still another may decide to skip a dose simply because the benefits are in the future and not salient right now. Missing a single dose may not feel especially costly relative to the salient hassle costs (“I really don’t feel like experiencing the pain of an injection right now”). Medication use by diabetics is not a unique example. Psychology affects decisions about nearly all types of medical care. In other words, the “psychic” cost-benefit calculus may be very different from the economic calculus.

For our purposes, we are particularly interested in how such deviations interact with traditional economic concepts, in this case moral hazard. We must now look beyond *overuse* of care. We must also consider the possibility of *underuse*: care that patients fail to use even when their benefits exceed the cost.³ That has important implications for policy design. Take the case of copayments—the payments made by an insured person each time he or she uses a medical service. The usual policy logic dictates that we can use elasticity of demand for a category of care to set copayments. A high demand elasticity means that the care is of low value. If small changes in price (which bring it closer to true cost) dissuade many people, the value of that care must not have been very high: a high demand elasticity signals overuse. As a result, copayments should increase with the elasticity of demand.

That logic fails in a behaviorally augmented model. A high elasticity of demand no longer indicates overuse. When a copay increase reduces demand, we can no longer infer that the care is actually of low *social* value; perhaps people were underusing it and we are worsening the problem. When individuals do not choose optimally, a change in demand tells us only that people choose *as if* they do not value the care. Return to the case of insulin treatment for diabetes. A patient who was non-adhering on some days because he feels that medication is optional on days when he “feels good” will show price sensitivity: he will

skip more doses on those days if prices are high. In effect, he feels that the care is optional. Increasing copays for such a patient on the basis of that elasticity would, however, be worsening a behavioral bias. In effect, psychology forces us to reinterpret empirical data on demand. Empirical studies can no longer simply use demand elasticities to measure moral hazard. We must understand more about the category of care where the elasticity appears. The demand elasticity is no longer sufficient for setting policy. The optimal amount of a copayment must be based on both knowledge of elasticity and an external assessment of the value of the treatment. In some cases, optimal copayments may even be negative—for example, in cases in which it is worthwhile to pay people to take their medication because of the positive spillover effects of doing so.

We can also examine nonprice levers. Consider the provision of “nudges,” the label given by Richard Thaler and Cass Sunstein to psychologically astute interventions that influence behavior.⁴ In the case of drug adherence, an example would be simple reminders to take medications. Once we recognize that nudges can affect use of care, we must examine insurance design more broadly. When will insurers nudge patients to use care? When will they nudge patients to reduce use of care? The answer depends on how profits align with health outcomes. Take again the case of drug adherence. Patients’ failure to adhere to treatment regimes has long-term costs: hospital admissions, for example, will be higher. A long-term insurer will bear those costs; as a result, a profit-maximizing long-term private insurer will have incentives to devise and implement nudges to increase adherence. Investments in disease management—which many companies increasingly make—can be understood from that perspective. In contrast, a short-term insurer bears none of the costs of patients’ non-adherence. They not only have zero incentives to provide nudges to improve adherence, they also have perverse incentives to find nudges that discourage use, even when use has high long-term benefits for the patient. For example, the short-term insurer can create costs by making it a hassle to schedule a doctor’s appointment or to refill a prescription. The psychological perspective therefore can add to our understanding of why health insurance is structured in certain ways when provided in a private market.

A fuller integration of behavioral economics and public finance allows us to go beyond just suggesting specific psychologically astute policies to experiment with. It provides a different framework for understanding such traditional public finance levers as copayments and market structure.

Taxes

Governments must raise revenues to provide services. Traditional public finance has a well-developed framework for determining how to set taxes optimally. Models of incidence help us understand who bears the burden of taxes; models

of efficiency help us understand how taxation can hinder economic activity. Together they offer practical insights for designing policy for taxes of all stripes: income, sales, and so on. For example, one broad insight is that efficient taxes are those that minimally distort consumer choices. Since individuals were choosing optimally in the absence of taxes, a change in the choices that they make represents a welfare cost. Concretely, one should raise revenues by, for example, taxing low-elasticity goods—taxes on, for example, cigarettes are often justified in part for this reason.

Behavioral economics complicates that logic. One recent study finds, for example, that individuals may fail to perceive sales taxes that are not included in the prices posted on store shelves but are computed at the register.⁵ People may simply fail to attend to them—they are not salient at the time of choice. Applying traditional logic, tax non-salience represents an opportunity for governments: they can raise revenues without distorting behavior. That logic, however, is incomplete. Lack of response to a non-salient tax is not the same as lack of response to a salient tax. When people fail to respond to a non-salient tax, there is an error: they make consumption choices as if an item costs $\$X$, but in purchasing the item they actually spend $\$X + \Y . As a result, they have $\$Y$ less to spend in the future than they had planned.

How that affects all other consumption must now enter the welfare calculation. Consider two polar cases. The lost money could be treated as a pure income effect: individuals see that they have $\$Y$ less to spend on all other goods and adjust accordingly. That would, in effect, turn the non-salient tax into a lump-sum tax, and governments therefore should use non-salient taxes heavily. Alternatively, suppose that the $\$Y$ is taken out of a narrower mental account. For example, rather than thinking of their overall budget as depleted by $\$Y$, individuals think of $\$Y$ as depleting their grocery budget specifically, and they may spend $\$Y$ less on their next trip to the grocery store. Or they may never change consumption and instead simply end up saving less. In such cases, the low demand response to non-salient taxes is misleading: though it does not generate distortions in the demand for the good being taxed, it is creating possibly higher distortions elsewhere. As a result, governments would need to take into account other potential distortions before using non-salient taxes.

In this case we also see that it is impossible to think about the implications of a nudge on tax salience—for example, excluding taxes from posted prices—in isolation from the public finance framework. The simple application of traditional logic suggests that one should always use nudges to reduce tax salience. In an integrated framework, that is no longer the case. The effects of reduced salience must include all the demand responses that it elicits.

Externalities

One of the triumphs of public finance is to provide a clear understanding of how to deal with external costs (externalities). Take the case of carbon emissions, which contribute to global warming, a typical negative externality. Individuals and firms do many things that affect carbon emissions, from driving automobiles to retrofitting factories; in making their choices, they impose costs on society. Traditional public finance provides an elegant solution to ensuring that those externalities are internalized in choice: individuals and firms must face the full costs of their carbon-emitting activities. The prices that they pay must include not only the marginal cost of the goods that they consume but also the cost of the carbon emissions that those goods produce. Put simply, we can achieve economic efficiency by placing a carbon tax on goods that is equal to the social cost of the carbon emissions produced by those goods. There are technical and political challenges in implementing such a tax, but the conceptual solution is clear.

As with the other examples, decisions involving carbon emissions may not be made in accordance with standard assumptions. For example, psychological studies suggest that social comparisons can drive behavior. Being told, for example, that “you used x kilowatt hours last month, but your neighbors used y kilowatt hours” can reduce a person’s consumption of electricity. Based on that insight, a company called OPOWER has implemented a large-scale program that charges utilities to send social comparison reports to consumers. In randomized, controlled trials with hundreds of thousands of utility customers across the United States, the reports have been shown to reduce electricity consumption in the average household by about 2 percent.⁶ Notice several interesting aspects of this example. First, even with a traditionally efficient carbon tax, there may be inefficiency if consumers do not choose their energy consumption levels optimally. Second, in addition to the role the prices play in affecting behavior, nudges or other interventions can play a powerful role. Third, and most important, in this case the private sector has generated a nudge—social comparison reporting—in order to affect energy consumption.

The last point is especially interesting because it suggests that policy levers besides carbon taxes and government-imposed nudges can be devised. Can the government somehow induce firms to nudge effectively? They have levers that can be used for that purpose. Consider decoupling for utilities, under which the profits of electricity retailers are no longer directly related only to the volume of electricity sold; they also receive revenues for *reducing* consumption. That type of lever, if it encourages utilities to nudge consumers toward reducing energy use (as it has in the case of OPOWER), is a powerful tool.

The logic here involves both economics and psychology. Psychology recognizes the power of nudges; economics recognizes the power (and peril) of markets. Firms may have nudges available to them that the government does not. So the government can do better than just implement its own nudges; it can look for policy levers such as decoupling that encourage firms to create and use nudges to improve consumer well-being.

These examples suggest that, first, psychological insights must be applied more deeply and broadly in public policy. They can be more than an added-on tweak at the end of a predetermined economic policy—they can alter the basic policy framework, from deadweight loss to moral hazard. Second, many of the policy suggestions based on behavioral insights are not especially behavioral. Decoupling is a traditional economic policy lever, but the behavioral approach enriches our understanding of its impacts. Finally, these examples illustrate that the law of unintended consequences continues to be important: policy changes (nudges or otherwise) must continue to be analyzed within the broader system in which they operate.

The Pitfalls

While integrating the psychological insights of behavioral economics into public finance policy holds great promise, as illustrated above, doing so also introduces a set of potential stumbling blocks for analysis and policy design. The approach described in this book overcomes or at least alleviates two of the major challenges.

Can the number of potential psychological factors be made manageable?

Psychology is, naturally, a very rich discipline, full of insights. That richness creates an overload of information. For any policy problem, it seems that an endless array of psychological phenomena could be relevant. The length of unemployment spells could be influenced by cognitive dissonance, hyperbolic discounting, anchoring, overconfidence, and loss aversion, to cite just a few examples. How do we handle such a vast array of possibilities?

We believe that the answer lies in abstraction. Knowing the specific psychological factor that drives a behavior is important in designing nudges. For example, job seekers may procrastinate in searching for jobs for a variety of reasons. Some activities (going out with friends, watching TV) may be enjoyable and therefore hard to resist. On the other hand, unemployment can sap a person's motivation, making it hard to exercise the self-control needed to engage in day-to-day activities such as sending out resumes. Those factors suggest different interventions: should we reduce procrastination by offering people a chance to commit themselves to searching for work in the future, or by finding a way to remotivate the unemployed?

Our insight is that despite the differences in those two examples, they have much in common: in both, the unemployed individuals recognize the future benefits of searching for a job; in both, they would like to and plan to search for a job; in both, they are unable to implement their desires because, at the moment, something (a tempting activity, lack of motivation) intervenes. We can lump those factors and other phenomena into a particular category labeled *bounded self-control*—the category of psychological factors that reflect a general tendency whereby people would like to take an action with future benefits but fail to do so. Categorizing helps us to craft policy principles. For example, when bounded self-control is a problem, we would argue that one must be very careful about the structure of incentives. Giving a person a bonus to leave unemployment will have weak effects if the benefits are realized far in the future. Those with bounded self-control already recognize and would like to capitalize on the future benefits of searching for a job. Their problem is implementing their desires; adding a modest bonus to those future benefits will not help much.

More generally, we create three categories of deviations from the standard economic model of decisionmaking: imperfect optimization, bounded self-control, and nonstandard preferences. These categories capture much of the psychological evidence that is both robust (supported by a vast majority of evidence) and important across a broad range of policy applications. Different psychological factors are considered similar if they call for the use of similar kinds of public policy levers: taxes, eligibility rules, and so on. Even when the focus is on the creation of nudges, categorization helps us see the general psychological force on which a nudge ought to operate. Our categorization is by no means perfect; there inevitably will be important psychological factors that are hard to categorize. Nor is it a magic bullet. But we do feel that categorization greatly simplifies addressing policy problems and in several important cases allows us to make significant progress with little reference to specific psychological factors.

The first category, *imperfect optimization*, captures errors: mistakes that people make in choosing among alternatives. For example, overconfidence or misunderstanding risks could lead people to under-demand insurance. Imperfect optimization means that people may have desires that do not match hedonic utility and may make choices that do not match their desires. The second category is *bounded self-control*, discussed above. Bounded self-control means that people, even when they are accurate in what they want, often are unable to implement their wants. The third category is *nonstandard preferences*: what individuals want is not what we presume. Even when people are accurate in their wants and they are able to choose in accordance with their wants, those wants may be different from the standard model. Their preferences include components or take a shape that the standard model usually assumes away.

Does imperfect optimization make welfare economics impossible?

Even if the myriad psychological phenomena can be made generally accessible to policymakers and economists, can we implement welfare economics if we incorporate them? Welfare analysis is built on the assumption that choice reveals preference, with social welfare reflecting an aggregation of the utility functions thus revealed. In the behavioral model, choices no longer reveal preference. Take the case of cigarettes. A behavioral approach emphasizes a conflict of preferences here. On one hand, people do not want to consume them; they would like to quit. On the other hand, they would like to quit in the future; right now, they would like a cigarette. That generates a preference inconsistency.

If we take the stated desire to quit seriously, we might use taxes to make it harder for people to smoke. Research in fact finds that cigarette taxes can be shown to assist individuals who have problems with self-control to do better for themselves.⁷ But should we have such taxes? Ultimately, who is to say that individuals should smoke fewer cigarettes? Notice that we can no longer assume that choices reveal people's preferences because people may reveal multiple preferences. The failure of revealed preference deprives public sector economics of a clean analytical foundation for assessing the welfare impact of policies. When individuals behave in inconsistent ways, what actions should public finance economists take to reflect welfare?

We focus on two complementary ways to solve that problem. First, we observe a practical reality. In the vast majority of cases, public finance economists (behavioral or otherwise) are not asked to make such judgment calls. Instead, policymakers and societies more broadly typically make those judgments. Policies already reflect a decision to discourage smoking, encourage saving, and ensure adherence to some drug treatment regimes. Instead, the role of most public finance economists is to design policies that take such welfare functions as given. That is similar in some ways to how, in traditional public finance, we do not expect economic theory to resolve *interpersonal* preference conflicts. When, for example, economic policies will have distributional consequences, public finance does not in general offer a way to compare the losses of one group against the benefits to another. We take as given the weights that the social planner gives to different people.

Leaving it to society to resolve *intrapersonal* preference conflicts is not too different. Of course, for economics as a field, it is important to make progress on the fundamental question of inferring hedonics in a world of behavioral agents. The most complete work to date on this fundamental question has been done by Douglas Bernheim and Antonio Rangel, who rigorously draw out the serious challenges of making such inferences.⁸ Overcoming this problem will be a key challenge for behavioral public finance. In this book, we sidestep the question by examining the design of policy when the policymaker has already made such inferences.

Second, we observe a misleading aspect of the cigarette example. In that example, the only reason for government to intervene is to solve a *behavioral* problem—to “fix” smokers’ mistake in smoking. In sharp contrast, most policies aim to solve *nonbehavioral* problems. Social programs aim to redistribute social benefits, taxes aim to raise revenues or address externalities, Social Security and health insurance policies aim to solve market failures, and so on. With those problems, policies already have been implemented to solve other market failures. As a result, the issue of welfare, while not disappearing, becomes secondary. In those cases, there are first-order consequences of individuals’ behavior to *society*, independent of the consequences for their own welfare. For example, if decision-making biases lead individuals to systematically disfavor fuel-efficient vehicles, one can debate what their true utility function is. But the carbon externality that they impose in making their choice is clear. In most of what we do, we focus on how behavioral economics changes policies in areas in which government already plays a traditional role.

The Payoff

Integrating behavioral economics into public finance results in a new set of principles for both understanding the role of government in the economy and informing policy design. This approach to public finance reveals deep insights for policymakers—for how incentives operate, for how markets work and fail, and for the role of information—which yield a variety of results.

Perceived prices drive behavior.

Standard public finance emphasizes the use of price changes, through taxes and subsidies, for example, to attain efficiency. Behavioral public finance recognizes that psychology mediates consumers’ responses to prices. For example, individuals with limited attention and limited computational capacity respond not to actual prices but to the prices that they *perceive*. Similarly, responses to prices may not reflect *intended* responses because of an individual’s imperfect capacity for self-control. As a result, prices will not always be effective levers for changing behavior, especially when prices are not salient or when the targeted behaviors already are the result of imperfectly optimal behavior.

For example, complicated subsidies may prove ineffective. Take the case of the Saver’s Credit, a policy that subsidizes retirement saving. In part because the credit is somewhat obscure and difficult to understand, its effectiveness in actually increasing retirement saving among targeted individuals appears to be limited. The evidence suggests that, dollar for dollar, a subsidy structured in a more straightforward way, such as a savings match, might have a greater impact on saving behavior.⁹ That is emblematic of a behavioral policy error: the presumption

that the objective price (the extent of the Saver's Credit) matches the subjective price (the perceived subsidy in the Saver's Credit).

Nudges have social as well as private effects.

The success of automatic enrollment in increasing the contribution rates to savings plans begs its application to other contexts. Many have suggested that if the application process for means-tested programs such as Temporary Assistance to Needy Families (TANF) or Medicaid were simplified so that benefits were easier to claim, the benefits would reach more qualifying individuals. There is some evidence that simplifying application procedures also works, for example, for college financial aid.¹⁰ But is that a good idea? Public finance demands that we integrate the psychological approach closely with the original rationale for intervention. In these cases, government is attempting to redistribute income or assist those with low incomes. But while we want everyone to do at least some saving for retirement, we do not want to redistribute equally to everyone. In fact, creating hurdles to claiming public benefits could screen out those who need them least. Automatic enrollment could subvert a screening process that is actually economically efficient. The behavioral public finance framework suggests that it is necessary to answer empirically the question of who is screened out of programs by enrollment procedures in order to understand the impact of simplification of these programs on social welfare.

Take another example, the Medicare prescription drug program, also known as Medicare Part D. Medicare Part D provides prescription drug insurance for seniors, who must choose among private plans. Evidence demonstrates that the choice among plans is difficult for individuals and that they make mistakes in choosing.¹¹ Intelligent assignment, automatically enrolling individuals in low-cost plans, is one possible way to structure this policy—in fact, some states did so for their low-income participants—and one that would be suggested by the automatic enrollment experience. But allowing for individuals to make those mistakes—or randomly assigning low-income participants to plans, as other states did—might have had beneficial effects in terms of risk pooling.

Nudges cannot be assessed by whether they help individuals. One must understand how nudges interact with the market failures that motivated the nudge policy and evaluate them within the broader social welfare function.

The social welfare function has psychological aspects.

Without looking through the behavioral lens, we may also misunderstand the social welfare function. Take Social Security. To understand its role, economists look for a market failure. As we age, we face the risk of outliving our resources. While annuities could solve that problem, adverse selection makes annuities very expensive or unavailable for some. Social Security exists to solve that market

failure. This story has some truth—longevity risk is a genuine problem—but it seems incomplete. Surely Social Security was motivated in part by a belief that people will fail to save effectively for retirement. Bounded self-control in the face of day-to-day consumption demands and temptations makes it hard to implement one's saving preferences. Imperfect optimization makes retirement planning difficult and error prone (How much to save? Where?). From a behavioral perspective, one of the primary *purposes* of Social Security is to reduce the demands on willpower and the complexity of saving for retirement. That understanding affects the form and design of Social Security policy. It also reinforces the earlier point about how the welfare problem is solved in practice: policymakers and society have, as in other cases, already adjudicated intrapersonal conflicts. They have sided with the self that wants to save more over the one that fails to save. They have decided that some choices, such as paying high fees for an index fund, are simply errors.

Unintended behavioral responses to policies do not necessarily represent moral hazard. Economics often uses behavioral responses to make important inferences, but psychological factors can change what those inferences can or should be. We saw this in the example of health insurance, taxes, and externalities, but it operates more generally. Consider unemployment. We might attribute a person's disinclination to look for a job to moral hazard: knowing that they get unemployment benefits, people enjoy their leisure until they exhaust those benefits. Alternatively, unemployment may undermine the willpower needed to search for work. Misunderstanding the original problem can lead to faulty policies. Long-term incentives work if the behavior is driven by moral hazard, but they work poorly if it is driven by lack of willpower or procrastination. That may help us understand, for example, why some experiments with creating incentives to counteract moral hazard have proven disappointing.¹²

Selection effects reflect both incentives and psychology.

Much of public finance emphasizes the role that prices, incentives, and information play in screening or generating selection effects. But behavioral economics emphasizes that individuals can respond to incentives in nonstandard ways that can undo or reverse selection effects. That might be true both in markets with asymmetric information, where the standard approach might identify adverse selection, and in cases where public policy wants screening in order to generate efficient outcomes.

An example of a screening problem in which behavioral tendencies may pose a design challenge is when the government seeks to induce efficient screening, which arises, as noted above, in targeted transfer programs. Traditional economic logic suggests that barriers to program take-up, such as application cost or

waiting time, can serve as an effective way to screen the needy from those who simply seek to exploit the program. However, if people fail to participate because of human frailties—procrastinating in filing the application form, being put off by the tediousness or hassle of completing it, or failing to understand program rules—screening may not be efficient. Nonparticipants then are not those who value the program the least but those who understand the rules the least or who have the biggest procrastination problem. In some cases, such as transfer programs, those individuals might be the very population targeted by the program.

Similar forces might operate to affect outcomes in markets with asymmetric information. For example, in health insurance markets, individuals are thought to have an informational advantage (asymmetric information) with respect to their own health status, which is believed to lead to adverse selection, which in turns undermines the efficient operation of such markets. However, the extent to which individuals correctly perceive and act on any such information may be mediated by psychological factors that affect their demand for health insurance.

Government intervention is more effective when attuned to the market's choice architecture.

A final insight involves markets. Regulations, taxes, or subsidies that better align firm profits with *true utility* mean that markets can be used to solve behavioral biases, as in the case of OPOWER described previously. That mirrors one of the innovations of modern public finance: even in cases of market failure, clever policies can harness market forces to resolve the original market failures (as in the case of tradable pollution permits). Similarly here, careful policy can harness market forces to resolve behavioral biases.

That also means that when creating markets, governments must be careful to minimize choice errors. A recent example of this lesson is the case of Medicare Part D, which was designed as a marketplace in which seniors could choose subsidized coverage from private providers. The hoped-for gains from competition, however, may have been dissipated by choice errors. Part D choice is rife with complexity: participants choose from dozens of plans that are differentiated in ways that make it hard to value them. For example, each plan has a unique schedule of benefits, so different drugs are covered differently by each plan. And subsequent empirical research has shown that the program's complexity has had quantitatively large consequences: seniors make errors in plan choice that, on average, cost them hundreds of dollars a year.¹³ Similar difficulties have been observed in markets established by policymakers to provide, for example, education.¹⁴ When individuals choose badly, firms compete to cater to their bad choices, leaving little hope for maximizing welfare.

Organization of the Book

In the chapters that follow, we develop our insights more systematically and in greater detail. The rest of the book proceeds in two parts.

In part 1, we set the stage for integrating behavioral economics into public finance by interpreting the evidence from psychological studies and developing a framework for applying it to questions in public finance. Chapter 2 presents and organizes the evidence from the psychology and behavioral economics literatures, abstracting from the specific results in a manner that will make the results useful for economic analysis. Chapter 3 introduces and develops our framework for integrating behavioral economics into public finance on a conceptual level.

In part 2, we apply that framework to topics in public finance. Chapter 4 considers problems and policies stemming from asymmetries of information, with an emphasis on social insurance, including old-age insurance, health insurance, and unemployment insurance. Chapter 5 treats externalities and public goods, with a focus on applications to environmental externalities, public health externalities, and education. Chapter 6 applies behavioral insights to issues related to income support and redistribution. Chapter 7 explores the behavioral dimensions of the economics of taxation and revenue.