1

Introduction

This book has a simple and straightforward message. The political and programmatic success of social programs requires improved target efficiency: directing resources where they do the most good. Although this fact is widely understood, it is seldom discussed, much less analyzed—and certainly not by the supporters of such programs. Our principal goal in writing this book is to make that discussion more coherent, better informed, and easier to conduct.

The public domain boasts many sound social programs. Some of these programs seek to allocate resources to individuals who are members of a legally defined target group—people whom politicians and policymakers have chosen to receive these resources. But many social programs are not nearly as well targeted as they could be, and a few are so poorly targeted as to call their social value into serious question. Public policy should improve the targeting of social programs so that they can accomplish more of their goals while using the same resources to assist the same needy populations.

We are particularly concerned with programs that seek to improve the conditions and opportunities of unfortunate, disadvantaged, usually low-income individuals, people whom we call *bad draws*.[1] We think of these bad draws as parties to a kind of

social contract for insurance against certain random misfortunes. This is why the government pays for medical care for the sick, unemployment benefits for those who lose their jobs, and food stamps for those who would otherwise be hungry or malnourished.

The category of bad draws, of course, is extremely broad, providing little guidance to policymakers who must allocate social welfare resources among the many bad draws who have plausible claims on society's solicitude. This book seeks to provide better guidance by focusing on two particular groups of bad draws who divert resources from the other bad draws for whom they were primarily intended. We call these two groups *bad bets* and *bad apples*. All of the individuals in these two groups are bad draws, but they are not well-targeted beneficiaries, either because they derive little benefit from a program (bad bets) or because their very participation imposes significant costs on other participants (bad apples).

Bad bets are individuals who are likely to benefit little from social resources relative to other bad draws. Our paradigmatic bad bet is a chronically ill nonagenarian who receives costly medical treatments at public expense, which predictably will yield little social benefit.

Bad apples are individuals whose irresponsible, immoral, or illegal behavior in the past—and predictably, in the future as well—marks them as unsuitable to receive the benefits of social programs.[2] We are concerned in this book with a subset of this category: those who interfere with the ability of deserving participants to benefit from a program. (Most bad apples also harm themselves, but our principal concern here is their adverse effect on *good apples* in the same programs.) An all-too-common example of a bad apple is the public school student who chronically disrupts class and thereby impairs the learning of others who desperately need a sound education. Another bad apple is the public housing tenant or homeless-shelter resident whose repeated misconduct debases his or her neighbors' quality of life. Bad apples are found in every segment of society; the category includes many who are relatively wealthy and advantaged (*good draws*).[3] Here, however, we focus on bad apples who are bad draws because they are the ones who consume the scarce resources available for important social welfare programs.

Bad bets and bad apples pose distinct challenges to policymakers, but we think it is useful to address them in a single book. We are keenly aware of the many hard issues raised by a serious effort to understand and address these two problems. First, merely defining these two categories of program beneficiaries is a profoundly difficult undertaking, necessitating tough line-drawing decisions. Second, an even more controversial challenge is the administrative task of assigning particular beneficiaries to these categories. This process in

effect labels some individuals as socially or programmatically undesirable (in the case of bad apples), and others as relatively unlikely to benefit from program resources (in the case of bad bets). Third, each of these judgments entails predicting future behavior or events on the basis of inevitably limited information, which makes some level of error inescapable. Implementing such judgments in the real world of program administration raises many challenges. Fourth, to make matters worse, we cannot precisely measure the social benefits of avoiding bad bets or removing bad apples, nor can we measure the likely costs of implementing policies that do so. Finally, the politics of dealing forthrightly and effectively with the problems posed by these two groups—or indeed, even candidly acknowledging these problems—are bound to be daunting. We strongly suspect that the sense of futility that many policymakers feel at the prospect of openly confronting these problems helps to explain why they have received relatively little attention. All the more reason, then, for academics like us to get the analytical ball rolling.

The book proceeds as follows. In chapter 2, we present the foundations of our analytical approach and introduce the subject of target efficiency, which is pivotal to this approach. We elaborate on the definitions and the positive, normative, and methodological assumptions that guide our discussion, and we highlight the most difficult issues raised by our analysis. In addition, we address a very hard question: what should social programs do with the bad bets they avoid or the bad apples they remove. This presentation, while longer than we would like, is essential to understanding what follows.

In chapter 3, we begin with an analysis of the stakes facing social policymakers, whom we urge to think more rigorously and courageously about bad bets and bad apples. We then explain the reluctance of politicians and bureaucrats to acknowledge and deal with the problems posed by bad apples and bad bets. This reluctance ends up harming precisely those bad draws who most deserve the help of social programs: individuals who are both good apples and good bets. We then present a taxonomy of bad policies—some handicapped by poor targeting, others exhibiting different flaws—and conclude by analyzing six pathologies that contribute to poor targeting across a wide range of social programs.

Citizens who want to maintain and expand social programs designed to promote the well-being of deserving bad draws have an important stake in avoiding bad bets and removing bad apples. We call such citizens *well-targeted redistributionists*, and we count ourselves among them. Well-targeted redistributionists should want to recapture the resources squandered on bad apples and bad bets so that they can be redirected to better purposes. Alas, as we explain, many well-targeted redistributionists impede this goal (some-

times naively) by pretending that bad apples and bad bets scarcely exist and ignoring the misallocation problems that these groups present. In this way, the redistributionists make it that much easier for politicians and bureaucrats to engage in the same neglect.

In chapters 4 and 5, respectively, we develop the analysis of bad bets and bad apples in considerable detail. These chapters examine several social program areas to help provide the focus and data needed for clearer thinking about improving targeting. For bad bets, health care examples get primary attention, largely because of the enormous amounts of money involved. For bad apples, our main examples are welfare programs, school classrooms, public housing, and homeless shelters.

As noted earlier, classifying an eligible beneficiary as a bad bet or a bad apple can require very difficult and controversial predictions of future events and behavior. In chapter 6, we examine the subject of target efficiency in more detail, focusing on three distinct targeting processes: sorting by authorities, sorting by recipients, and sorting with appeals. The difficulties of prediction and classification are discussed, as well as the procedural protections that are essential to making these decisions with acceptable levels of accuracy and fairness.

In chapter 7, we highlight the central themes of the book by considering the steps policymakers must take to improve program targeting. The first step is to obtain better information about individuals and programs. The second is to apply our methods to specific cases. We illustrate this process with several examples and conclude the chapter and the book by urging policymakers to employ both caution and urgency as they seek to improve target efficiency.

The basic goal of target efficiency—allocating resources to the individuals for whom and the purposes for which they will do the most good—is straightforward and should not be controversial. Nevertheless, this goal often proves elusive. Giving resources to A but not to B, or putting them toward goal X but not toward goal Y, immediately raises a host of challenges—some conceptual or analytical, but many of them purely political. For example, those who represent B will protest and may even sue. Even more vociferous will be the groups that deliver resources to B or sell services in support of Y. For example, organizations that purport to speak for low-income people often oppose removing bad apple tenants from public housing and bad apple students from traditional schools, while the good apple beneficiaries to whom these programs are targeted prefer that the disrupters be removed.

Conservatives are likely to be skeptical of any effort to improve targeting that they fear may be costly. Indeed, they may think that a poorly functioning

social welfare system better serves their efforts to discredit the welfare state and limit its expenditures. Liberals, for their part, seldom actively champion target efficiency. As we discuss in chapters 3 and 5, they tend to worry that acknowledging the problem of bad apples amounts to blaming the victim, maligning the poor, and undermining the legitimacy of cherished social programs. They exhibit a particular and seldom discussed (much less publicly defended) form of risk aversion: to draw on a cliché from criminal law, they would rather serve ten undeserving recipients than deny a single deserving one.

Expecting little sympathy for target efficiency from either end of the political spectrum, we look for support from well-targeted redistributionists, who cluster in the middle and who, we suspect, vastly outnumber those at the left and right extremes.[4] We expect that their willingness to devote resources to education, health care, low-income housing, rehabilitation of prisoners, and other social welfare programs would expand significantly if the target efficiency of such programs improved—that is, if they weeded out the chronically disruptive students, the patients getting little benefit from vast Medicaid and Medicare expenditures, the public housing and homeless shelter residents who spoil their neighbors' quality of life, and the recidivist criminals.

Liberals and conservatives should both want to spend money where it will accomplish the most good, without seriously undermining beneficiaries' incentives to improve their own status. We are encouraged in this expectation by the overwhelming public approval of the efforts to improve targeting in welfare programs that began in the mid-1990s—first with experimental state programs and then through the 1996 federal welfare reform law. The government now spends far more money per capita for support of the good apple poor than it did before the reforms were adopted, and politicians' attacks on the newer welfare programs are much reduced. For well-targeted redistributionists, this is good news indeed.

This book presents considerable empirical data to document the problems it explores. Little of that data is new, however, other than the tallies of media mentions of "welfare fraud" and "welfare recipient" in chapter 3 and the information, detailed in chapter 5, that we report from the agencies that deal with public schools, public housing, and homeless shelters in New York City. We recognize that these reports are neither rigorous nor systematic, but they do illustrate the nature and magnitude of the challenges raised by the bad apples problem and suggest some approaches that policymakers might take. Our main contribution, then, is not to adduce new data but rather to present a framework for thinking about social policy that has not come naturally to those engaged in policy debates or administering social programs. The data, examples, and analyses that we provide show how this

framework could and should apply in the design and implementation of social programs.

Some readers will criticize one or another example. Others will insist that even repeated misbehavior does not mean that one is a bad apple. Still others will maintain that it is wrong to exclude people of certain ages or medical conditions from public benefits simply because they are bad bets. Others will dispute our empirical claims—for example, that removing chronically disruptive students will improve the life chances of those who remain. These factual issues are highly relevant to sound policymaking, but we do not propose to debate, much less resolve, them here. Our goal, rather, is to establish sound policy precepts.

Policymakers should be eager to refine the specific examples that we cite and to learn about other examples that would help them to improve targeting in their programs. Again, our chief purpose is to establish sound principles and then provide examples that can help test their plausibility. Readers who find our framework useful should test it in their own policy domains. We imagine, for example, that thinking about antipoverty programs in terms of bad draws, bad bets, and bad apples could extend the lessons of the 1996 welfare reform, discussed in chapters 3 and 5, to a broader array of both actual and proposed social problems and interventions. By systematically developing the framework presented here, we hope to make the tough choices entailed by efforts to improve targeting in social policy more analytical, more accessible, and thus more thinkable.