1

Think Again, Again

Every year, 9 million children die before their fifth birthday.[1] A woman in sub-Saharan Africa has a one-in-thirty chance of dying while giving birth—in the developed world, the chance is one in 5,600. There are at least twenty-five countries, most of them in sub-Saharan Africa, where the average person is expected to live no more than fifty-five years. In India alone, more than 50 million school-going children cannot read a very simple text.[2]

This is the kind of paragraph that might make you want to shut this book and, ideally, forget about this whole business of world poverty: The problem seems too big, too intractable. Our goal with this book is to persuade you not to.

A recent experiment at the University of Pennsylvania illustrates well how easily we can feel overwhelmed by the magnitude of the problem.[3] Researchers gave students $5 to fill out a short survey. They then showed them a flyer and asked them to make a donation to Save the Children, one of the world's leading charities. There were two different flyers. Some (randomly selected) students were shown this:

> Food shortages in Malawi are affecting more than 3 million children;
> In Zambia, severe rainfall deficits have resulted in a 42% drop in maize

> production from 2000. As a result, an estimated 3 million Zambians face hunger; Four million Angolans—one third of the population—have been forced to flee their homes; More than 11 million people in Ethiopia need immediate food assistance.

Other students were shown a flyer featuring a picture of a young girl and these words:

> Rokia, a 7-year-old girl from Mali, Africa, is desperately poor and faces a threat of severe hunger or even starvation. Her life will be changed for the better as a result of your financial gift. With your support, and the support of other caring sponsors, Save the Children will work with Rokia's family and other members of the community to help feed her, provide her with education, as well as basic medical care and hygiene education.

The first flyer raised an average of $1.16 from each student. The second flyer, in which the plight of millions became the plight of one, raised $2.83. The students, it seems, were willing to take some responsibility for helping Rokia, but when faced with the scale of the global problem, they felt discouraged.

Some other students, also chosen at random, were shown the same two flyers after being told that people are more likely to donate money to an identifiable victim than when presented with general information. Those shown the first flyer, for Zambia, Angola, and Mali, gave more or less what that flyer had raised without the warning—$1.26. Those shown the second flyer, for Rokia, after this warning gave only $1.36, less than half of what their colleagues had committed without it. Encouraging students to think again prompted them to be less generous to Rokia, but not more generous to everyone else in Mali.

The students' reaction is typical of how most of us feel when we are confronted with problems like poverty. Our first instinct is to be generous, especially when facing an imperiled seven-year-old girl. But, like the Penn students, our second thought is often that there is really no point: Our contribution would be a drop in the bucket, and the bucket probably leaks. This book is an invitation to think again, *again*: to turn away from the feeling that the fight against poverty is too overwhelming, and to start to think of the challenge as a set of concrete problems that, once properly identified and understood, can be solved one at a time.

Unfortunately, this is not how the debates on poverty are usually framed. Instead of discussing how best to fight diarrhea or dengue, many of the most vocal experts tend to be fixated on the "big questions": What is the ultimate cause of poverty? How much faith should we place in free markets? Is democracy good for the poor? Does foreign aid have a role to play? And so on.

Jeffrey Sachs, adviser to the United Nations, director of the Earth Institute at Columbia University in New York City, and one such expert, has an answer to all these questions: Poor countries are poor because they are hot, infertile, malaria infested, often landlocked; this makes it hard for them to be productive without an initial large investment to help them deal with these endemic problems. But they cannot pay for the investments precisely because they are poor—they are in what economists call a "poverty trap." Until something is done about these problems, neither free markets nor democracy will do very much for them. This is why foreign aid is key: It can kick-start a virtuous cycle by helping poor countries invest in these critical areas and make them more productive. The resulting higher incomes will generate further investments; the beneficial spiral will continue. In his best-selling 2005 book, *The End of Poverty,*[4] Sachs argues that if the rich world had committed $195 billion in foreign aid per year between 2005 and 2025, poverty could have been entirely eliminated by the end of this period.

But then there are others, equally vocal, who believe that all of Sachs's answers are wrong. William Easterly, who battles Sachs from New York University at the other end of Manhattan, has become one of the most influential anti-aid public figures, following the publication of two books, *The Elusive Quest for Growth* and *The White Man's Burden*.[5] Dambisa Moyo, an economist who previously worked at Goldman Sachs and at the World Bank, has joined her voice to Easterly's with her recent book, *Dead Aid*.[6] Both argue that aid does more bad than good: It prevents people from searching for their own solutions, while corrupting and undermining local institutions and creating a self-perpetuating lobby of aid agencies. The best bet for poor countries is to rely on one simple idea: When markets are free and the incentives are right, people can find ways to solve their problems. They do not need handouts, from foreigners or from their own governments. In this sense, the aid pessimists are actually quite optimistic about the way the world works. According to Easterly, there are no such things as poverty traps.

Whom should we believe? Those who tell us that aid can solve the problem? Or those who say that it makes things worse? The debate cannot be solved in the abstract: We need evidence. But unfortunately, the kind of data usually used to answer the big questions does not inspire confidence. There is never a shortage of compelling anecdotes, and it is always possible to find at least one to support any position. Rwanda, for example, received a lot of aid money in the years immediately after the genocide, and prospered. Now that the economy is thriving, President Paul Kagame has started to wean the country off aid. Should we count Rwanda as an example of the good that aid can do (as Sachs suggests), or as a poster child for self-reliance (as Moyo presents it)? Or both?

Because individual examples like Rwanda cannot be pinned down, most researchers trying to answer the big philosophical questions prefer multicountry comparisons. For example, the data on a couple of hundred countries in the world show that those that received more aid did not grow faster than the rest. This is often interpreted as evidence that aid does not work, but in fact, it could also mean the opposite. Perhaps the aid helped them avoid a major disaster, and things would have been much worse without it. We simply do not know; we are just speculating on a grand scale.

But if there is really no evidence for or against aid, what are we supposed to do—give up on the poor? Fortunately, we don't need to be quite so defeatist. There are in fact answers—indeed, this whole book is in the form of an extended answer. It is just that they are not the kind of sweeping answers that Sachs and Easterly favor. This book will not tell you whether aid is good or bad, but it will say whether particular instances of aid did some good or not. We cannot pronounce on the efficacy of democracy, but we do have something to say about whether democracy could be made more effective in rural Indonesia by changing the way it is organized on the ground, and so on.

In any case, it is not clear that answering some of these big questions, like whether foreign aid works, is as important as we are sometimes led to believe. Aid looms large for those in London, Paris, or Washington, DC, who are passionate about helping the poor (and those less passionate, who resent paying for it). But in truth, aid is only a very small part of the money that is spent on the poor every year. Most programs targeted at the world's poor are funded out of their country's own resources. India, for example, receives essentially no aid. In 2004–2005, it spent half a trillion rupees ($31 billion USD purchasing power parity [PPP])[7] just on primary-education programs for the poor. Even in Africa, where foreign aid has a much more important role, it represented only 5.7 percent of total government budgets in 2003 (12 percent if we exclude Nigeria and South Africa, two big countries that receive very little aid).[8]

More important, the endless debates about the rights and wrongs of aid often obscure what really matters: not so much where the money comes from, but where it goes. This is a matter of choosing the right kind of project to fund—should it be food for the indigent, pensions for the elderly, or clinics for the ailing?—and then figuring out how best to run it. Clinics, for example, can be run and staffed in many different ways.

No one in the aid debate really disagrees with the basic premise that we should help the poor when we can. This is no surprise. The philosopher Peter Singer has written about the moral imperative to save the lives of those we don't know. He observes that most people would willingly sacrifice a $1,000 suit to rescue a child seen drowning in a pond[9] and argues that there should be no difference between that drowning child and the 9 million children who, every year, die before their fifth birthday. Many people would also agree with Amartya Sen, the economist-philosopher and Nobel Prize Laureate, that poverty leads to an intolerable waste of talent. As he puts it, poverty is not just a lack of money; it is not having the capability to realize one's full potential as a human being.[10] A poor girl from Africa will probably go to school for at most a few years even if she is brilliant, and most likely won't get the nutrition to be the world-class athlete she might have been, or the funds to start a business if she has a great idea.

It is true that this wasted life probably does not directly affect people in the developed world, but it is not impossible that it might: She might end up as an HIV-positive prostitute who infects a traveling American who then brings the disease home, or she might develop a strain of antibiotic-resistant TB that will eventually find its way to Europe. Had she gone to school, she might have turned out to be the person who invented the cure for Alzheimer's. Or perhaps, like Dai Manju, a Chinese teenager who got to go to school because of a clerical error at a bank, she would end up as a business tycoon employing thousands of others (Nicholas Kristof and Sheryl WuDunn tell her story in their book *Half the Sky*).[11] And even if she doesn't, what could justify not giving her a chance?

The main disagreement shows up when we turn to the question, "Do we know of effective ways to help the poor?" Implicit in Singer's argument for helping others is the idea that you know how to do it: The moral imperative to ruin your suit is much less compelling if you do not know how to swim. This is why, in *The Life You Can Save*, Singer takes the trouble to offer his readers a list of concrete examples of things that they should support, regularly updated on his Web site.[12] Kristof and WuDunn do the same. The point is simple: Talking about the problems of the world without talking about some accessible solutions is the way to paralysis rather than progress.

This is why it is really helpful to think in terms of concrete problems which can have specific answers, rather than foreign assistance in general: "aid" rather than "Aid." To take an example, according to the World Health Organization (WHO), malaria caused almost 1 million deaths in 2008, mostly among African children.[13] One thing we know is that sleeping under insecticide-treated bed nets can help save many of these lives. Studies have shown that in areas where malaria infection is common, sleeping under an insecticide-treated bed net reduces the incidence of malaria by half.[14] What, then, is the *best* way to make sure that children sleep under bed nets?

For approximately $10, you can deliver an insecticide-treated net to a family and teach the household how to use it. Should the government or an NGO give parents free bed nets, or ask them to buy their own, perhaps at a subsidized price? Or should we let them buy it in the market at full price? These questions can be answered, but the answers are by no means obvious. Yet many "experts" take strong positions on them that have little to do with evidence.

Because malaria is contagious, if Mary sleeps under a bed net, John is less likely to get malaria—if at least half the population sleeps under a net, then even those who do not have much less risk of getting infected.[15] The problem is that fewer than one-fourth of kids at risk sleep under a net:[16] It looks like the $10 cost is too much for many families in Mali or Kenya. Given the benefits both to the user and others in the neighborhood, selling the nets at a discount or even giving them away would seem to be a good idea. Indeed, free bed-net distribution is one thing that Jeffrey Sachs advocates. Easterly and Moyo object, arguing that people will not value (and hence will not use) the nets if they get them for free. And even if they do, they may become used to handouts and refuse to buy more nets in the future, when they are not free, or refuse to buy other things that they need unless these are also subsidized. This could wreck well-functioning markets. Moyo tells the story of how a bed-net supplier was ruined by a free bed-net distribution program. When free distribution stopped, there was no one to supply bed nets at any price.

To shed light on this debate, we need to answer three questions. First, if people must pay full price (or at least a significant fraction of the price) for a bed net, will they prefer to go without? Second, if bed nets are given to them free or at some subsidized price, will people use them, or will they be wasted? Third, after getting the net at subsidized price once, will they become more or less willing to pay for the next one if the subsidies are reduced in the future?

To answer these questions, we would need to observe the behavior of comparable groups of people facing different levels of subsidy. The key word here is "comparable." People who pay for bed nets and people who get them for free are usually not going to be alike: It is possible that those who paid for their nets will be richer and better educated, and have a better understanding of why they need a bed net; those who got them for free might have been chosen by an NGO precisely because they were poor. But there could also be the opposite pattern: Those who got them for free are the well connected, whereas the poor and isolated had to pay full price. Either way, we cannot draw any conclusion from the way they used their net.

For this reason, the cleanest way to answer such questions is to mimic the randomized trials that are used in medicine to evaluate the effectiveness of new drugs. Pascaline Dupas, of the University of California at Los Angeles, carried out such an experiment in Kenya, and others followed suit with similar experiments in Uganda and Madagascar.[17] In Dupas's experiment, individuals were randomly selected to receive different levels of subsidy to purchase bed nets. By comparing the behavior of randomly selected equivalent groups that were offered a net at different prices, she was able to answer all three of our questions, at least in the context in which the experiment was carried out.

In Chapter 3 of this book, we will have a lot to say about what she found. Although open questions remain (the experiments do not yet tell us about whether the distribution of subsidized imported bed nets hurt local producers, for example), these findings did a lot to move this debate and influenced both the discourse and the direction of policy.

The shift from broad general questions to much narrower ones has another advantage. When we learn about whether poor people are willing to pay money for bed nets, and whether they use them if they get them for free, we learn about much more than the best way to distribute bed nets: We start to understand how poor people make decisions. For example, what stands in the way of more widespread bed net adoption? It could be a lack of information about their benefits, or the fact that poor people cannot afford them. It could also be that the poor are so absorbed by the problems of the present that they don't have the mental space to worry about the future, or there could be something entirely different going on. Answering these questions, we get to understand what, if anything, is special about the poor: Do they just live like everyone else, except with less money, or is there something fundamentally different about life under extreme poverty? And if it is something special, is it something that could keep the poor trapped in poverty?

TRAPPED IN POVERTY?

It is no accident that Sachs and Easterly have radically opposite views on whether bed nets should be sold or given away. The positions that most rich-country experts take on issues related to development aid or poverty tend to be colored by their specific worldviews even when there seem to be, as with the price of the bed nets, concrete questions that should have precise answers. To caricature ever so slightly, on the left of the political spectrum, Jeff Sachs (along with the UN, the World Health Organization, and a good part of the aid establishment) wants to spend more on aid, and generally believes that things (fertilizer, bed nets, computers in school, and so on) should be given away and that poor people should be enticed to do what we (or Sachs, or the UN) think is good for them: For example, children should be given meals at school to encourage their parents to send them to school regularly. On the right, Easterly, along with Moyo, the American Enterprise Institute, and many others, oppose aid, not only because it corrupts governments but also because at a more basic level, they believe that we should respect people's freedom—if they don't want something, there is no point in forcing it upon them: If children do not want to go to school it must be because there is no point in getting educated.

These positions are not just knee-jerk ideological reactions. Sachs and Easterly are both economists, and their differences, to a large extent, stem from a different answer to an economic question: Is it possible to get trapped in poverty? Sachs, we know, believes that some countries, because of geography or bad luck, are trapped in poverty: They are poor because they are poor. They have the potential to become rich but they need to be dislodged from where they are stuck and set on the way to prosperity, hence Sachs's emphasis on one big push. Easterly, by contrast, points out that many countries that used to be poor are now rich, and vice versa. If the condition of poverty is not permanent, he argues, then the idea of a poverty trap that inexorably ensnares poor countries is bogus.

The same question could also be asked about individuals. Can people be trapped in poverty? If this were the case, a onetime infusion of aid could make a huge difference to a person's life, setting her on a new trajectory. This is the underlying philosophy behind Jeffrey Sachs's Millennium Villages Project. The villagers in the fortunate villages get free fertilizer, school meals, working health clinics, computers in their school, and much more. Total cost: half a million dollars a year per village. The hope, according to the project's Web site, is that "Millennium Village economies can transition over a period from subsistence farming to self-sustaining commercial activity."[18]

On a video they produced for MTV, Jeffrey Sachs and actress Angelina Jolie visited Sauri, in Kenya, one of the oldest millennium villages. There they met Kennedy, a young farmer. He was given free fertilizer, and as a result, the harvest from his field was twenty times what it had been in previous years. With the savings from that harvest, the video concluded, he would be able to support himself forever. The implicit argument was that Kennedy was in a poverty trap in which he could not afford fertilizer: The gift of fertilizer freed him. It was the only way he could escape from the trap.

But, skeptics could object that if fertilizer is really so profitable, why could Kennedy not have bought just a little bit of it and put it on the most suitable part of his field? This would have raised the yield, and with the extra money generated, he could have bought more fertilizer the following year, and so on. Little by little, he would have become rich enough to be able to put fertilizer on his entire field.

So is Kennedy trapped in poverty, or is he not?

The answer depends on whether the strategy is feasible: Buy just a little to start with, make a little extra money, and then reinvest the proceeds, to make even more money, and repeat. But maybe fertilizer is not easy to buy in small quantities. Or perhaps it takes several tries before you can get it to work. Or there are problems with reinvesting the gains. One could think of many reasons why a farmer might find it difficult to get started on his own.

We will postpone trying to get to the heart of Kennedy's story until Chapter 8. But this discussion helps us see a general principle. There will be a poverty trap whenever the scope for *growing income or wealth at a very fast rate* is limited for those who have too little to invest, but expands dramatically for those who can invest a bit more. On the other hand, if the potential for fast growth is high among the poor, and then tapers off as one gets richer, there is no poverty trap.

Economists love simple (some would say simplistic) theories, and they like to represent them in diagrams. We are no exception: There are two diagrams shown below that we think are helpful illustrations of this debate about the nature of poverty. The most important thing to remember from them is the shape of the curves: We will return to these shapes a number of times in the book.

For those who believe in poverty traps, the world looks like Figure 1. Your income today influences what your income will be in the future (the future could be tomorrow, next month, or even the next generation): What you have today determines how much you eat, how much you have to spend on medicine or on the education of your children, whether or not you can buy fertilizer or improved seeds for your farm, and all this determines what you will have tomorrow.

The shape of the curve is key: It is very flat at the beginning, and then rises rapidly, before flattening out again. We will call it, with some apologies to the English alphabet, the *S-shape curve*.

The S-shape of this curve is the source of the poverty trap. On the diagonal line, income today is equal to income tomorrow. For the very poor who are in the *poverty trap zone*, income in the future is lower than income today: The curve is below the diagonal line. This means that over time, those in this zone become poorer and poorer, and they will eventually end up trapped in poverty, at point N. The arrows starting at point A1 represent a possible trajectory: from A1, move to A2, and then A3, and so forth. For those who start outside of the poverty trap zone, income tomorrow is higher than income today: Over time they become richer and richer, at least up to a point. This more cheerful destiny is represented by the arrow starting at point B1, moving to B2 and B3, and so forth.

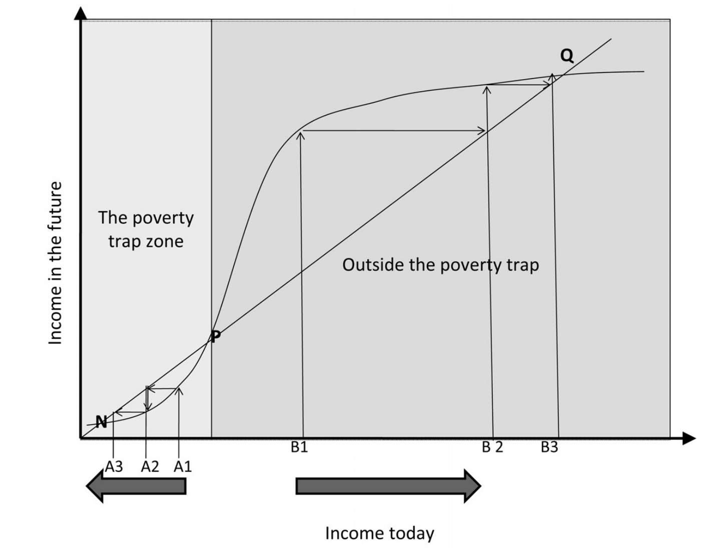

Figure 1: The S-Shape Curve and the Poverty Trap

Many economists (a majority, perhaps) believe, however, that the world usually looks more like Figure 2.

Figure 2 looks a bit like the right-hand side of Figure 1, but without the flat left side. The curve goes up fastest at the beginning, then slower and slower. There is no poverty trap in this world: Because the poorest people earn more than the income they started with, they become richer over time, until eventually their incomes stop growing (the arrows going from A1 to A2 to A3 depict a possible trajectory). This income may not be very high, but the point is that there is relatively little we need or can do to help the poor. A onetime gift in this world (say, giving someone enough income that, instead of starting with A1 today, he or she start with A2) will not boost anyone's income permanently. At best, it can just help them move up a little bit faster, but it cannot change where they are eventually headed.

So which of these diagrams best represents the world of Kennedy, the young Kenyan farmer? To know the answer to this question we need

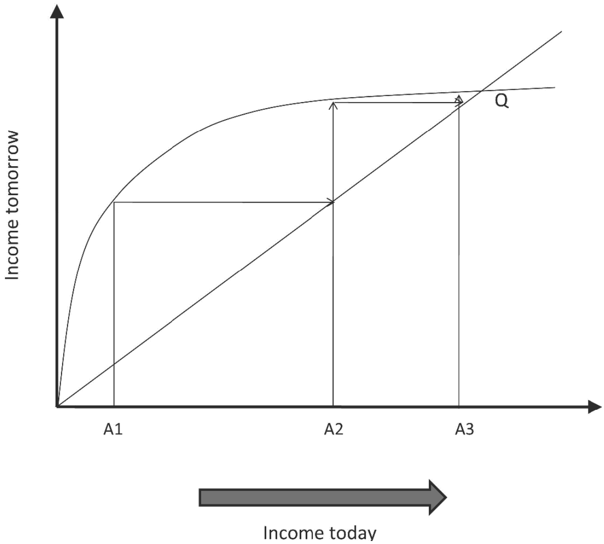

Figure 2: The Inverted L-Shape: No Poverty Trap

to find out a set of simple facts, such as: Can one buy fertilizer in small quantities? Is there something that makes it hard to save between planting seasons, so that even if Kennedy can make money in one season, he cannot turn it into further investment? The most important message from the theory embedded in the simple diagrams is thus that theory is not enough: To really answer the question of whether there are poverty traps, we need to know whether the real world is better represented by one graph, or by the other. And we need to make this assessment case by case: If our story is based on fertilizer, we need to know some facts about the market for fertilizer. If it is about savings, we need to know how the poor save. If the issue is nutrition and health, then we need to study those. The lack of a grand universal answer might sound vaguely disappointing, but in fact it is exactly what a policy maker should want to know—not that there are a million ways that the poor are trapped but that there are a few key factors that create the trap, and that alleviating those particular problems could set them free and point them toward a virtuous cycle of increasing wealth and investment.

This radical shift in perspective, away from the universal answers, required us to step out of the office and look more carefully at the world. In doing so, we were following a long tradition of development economists who have emphasized the importance of collecting the right data to be able to say anything useful about the world. However, we had two advantages over the previous generations: First, there are now high-quality data from a number of poor countries that were not available before. Second, we have a new, powerful tool: randomized controlled trials (RCTs), which give researchers, working with a local partner, a chance to implement large-scale experiments designed to test their theories. In an RCT, as in the studies on bed nets, individuals or communities are randomly assigned to different "treatments"—different programs or different versions of the same program. Since the individuals assigned to different treatments are exactly comparable (because they were chosen at random), any difference between them is the effect of the treatment.

A single experiment does not provide a final answer on whether a program would universally "work." But we can conduct a series of experiments, differing in either the kind of location in which they are conducted or the exact intervention being tested (or both). Together, this allows us to both verify the robustness of our conclusions (Does what works in Kenya also work in Madagascar?) and narrow the set of possible theories that can explain the data (What is stopping Kennedy? Is it the price of fertilizer or the difficulty of saving money?). The new theory can help us design interventions and new experiments, and help us make sense of previous results that may have been puzzling before. Progressively, we obtain a fuller picture of how the poor really live their lives, where they need help, and where they don't.

In 2003, we founded the Poverty Action Lab (which later became the Abdul Latif Jameel Poverty Action Lab, or J-PAL) to encourage and support other researchers, governments, and nongovernmental organizations to work together on this new way of doing economics, and to help diffuse what they have learned among policy makers. The response has been overwhelming. By 2010, J-PAL researchers had completed or were engaged in over 240 experiments in forty countries around the world, and very large numbers of organizations, researchers, and policy makers have embraced the idea of randomized trials.

The response to J-PAL's work suggests that there are many who share our basic premise—that it is possible to make very significant progress against the biggest problem in the world through the accumulation of a set of small steps, each well thought out, carefully tested, and judiciously implemented. This might seem self-evident, but as we will argue throughout the book, it is not how policy usually gets made. The practice of development policy, as well as the accompanying debates, seems to be premised on the impossibility of relying on evidence: Verifiable evidence is a chimera, at best a distant fantasy, at worst a distraction. "We have to get on with the work, while you indulge yourselves in the pursuit of evidence," is what hardheaded policy makers and their even harder-headed advisers often told us when we started down this path. Even today, there are many who hold this view. But there are also many people who have always felt disempowered by this unreasoned urgency. They feel, as we do, that the best anyone can do is to understand deeply the specific problems that afflict the poor and to try to identify the most effective ways to intervene. In some instances, no doubt, the best option will be to do nothing, but there is no general rule here, just as there is no general principle that spending money always works. It is the body of knowledge that grows out of each specific answer and the understanding that goes into those answers that give us the best shot at, one day, ending poverty.

This book builds on that body of knowledge. A lot of the material that we will talk about comes from RCTs conducted by us and others, but we also make use of many other types of evidence: qualitative and quantitative descriptions of how the poor live, investigations of how specific institutions function, and a variety of evidence on which policies have worked and which have not. In the companion Web site for the book, www.pooreconomics.com, we provide links to all the studies we cite, photographic essays that illustrate each chapter, and extracts and charts from a data set on key aspects of the lives of those who live on less than 99 cents per person per day in eighteen countries, which we will refer to many times in the book.

The studies we use have in common a high level of scientific rigor, openness to accepting the verdict of the data, and a focus on specific, concrete questions of relevance to the lives of the poor. One of the questions that we will use these data to answer is when and where we should worry about poverty traps; we will find them in some areas, but not in others. In order to design effective policy, it is crucial that we get answers to such questions right. We will see many instances in the chapters that follow where the wrong policy was chosen, not out of bad intentions or corruption, but simply because the policy makers had the wrong model of the world in mind: They thought there was a poverty trap somewhere and there was none, or they were ignoring another one that was right in front of them.

The message of this book, however, goes well beyond poverty traps. As we will see, ideology, ignorance, and inertia—the three Is—on the part of the expert, the aid worker, or the local policy maker, often explain why policies fail and why aid does not have the effect it should. It is possible to make the world a better place—probably not tomorrow, but in some future that is within our reach—but we cannot get there with lazy thinking. We hope to persuade you that our patient, step-by-step approach is not only a more effective way to fight poverty, but also one that makes the world a more interesting place.