

Introduction

In June of 2013, the Obama administration released its *Climate Action Plan*. The plan starts by making a case for action:¹

Climate change is no longer a distant threat—we are already feeling its impacts across the country and the world. Last year was the warmest year ever in the contiguous United States and about one-third of all Americans experienced 10 days or more of 100-degree heat. The 12 hottest years on record have all come in the last 15 years. Asthma rates have doubled in the past 30 years and our children will suffer more asthma attacks as air pollution gets worse. And increasing floods, heat waves, and droughts have put farmers out of business, which is already raising food prices dramatically.

These changes come with far-reaching consequences and real economic costs. Last year alone, there were 11 different weather and climate disaster events with estimated losses exceeding \$1 billion each across the United States. Taken together, these 11 events resulted in over \$110 billion in estimated damages, which would make it the second-costliest year on record.

The plan proceeds with a list of many proposals for regulations and incentives to reduce carbon emissions and mitigate climate change. These include (among many others) upgrading the electric grid, strengthening regulatory standards for automobile fuel economy and power plant carbon pollution, increasing funding for clean energy, advanced transportation technologies, energy efficient construction, and accelerating clean energy permitting.

There is a fairly simple diagnosis of the source of climate change that is agreed to quite broadly by social scientists. Carbon emissions—from factories, passenger vehicles, or what have you—cause climate change. Each of us who consumes fossil fuels bears only a tiny fraction of the costs associated with our personal emissions. I don't suffer much from the increased probability that you get asthma caused by my car. As a consequence, each of us emits too much carbon relative to what would be socially desirable. Put differently, if we could

¹Executive Office of the President. June 2013. *The President's Climate Action Plan*. <https://www.whitehouse.gov/sites/default/files/image/president27sclimateactionplan.pdf>

all agree to emit less, we'd all be better off. But each of us, as individuals, has an incentive to keep emitting.

There is also a pretty simple, and widely agreed upon, way to address this kind of problem. Increase the price of carbon so that people's individual costs reflect something closer to the true social costs of emissions. This price-based approach has at least two virtues. First, an increase in prices works directly on individuals' incentives. If the price of carbon is higher, people will use less of it all on their own. They don't need to be monitored or further regulated. Second, a price increase incentivizes people and businesses to reduce carbon consumption in as cost-effective a way as possible. Even among existing technologies, there is uncertainty about whether the most cost-effective way to reduce emissions is through increased fuel economy, greater investment in alternative energy sources, cleaner power plants, or what have you. Moreover, who knows what new approaches to reducing emissions will emerge if people have stronger incentives to innovate? Programs of the sort proposed in the President's Climate Action Plan incentivize particular approaches. A price increase on carbon incentivizes reduced emissions, but is agnostic as to how this should be achieved. This gives people and firms the flexibility to choose the most cost-effective strategies and encourages them to innovate.

As we will discuss throughout this book, there are a variety of political reasons that, despite its merits, a price-based approach might lose out. Policies that work through targeted subsidies and regulations do a worse job at mitigating carbon emissions. But, by creating specific winners from the policy process, they allow political leaders to build a coalition of support. This coalition building is often the critical step in achieving policy change. So perhaps we should not lament this outcome too much.

Suppose, however, that we wanted to think seriously about a price-based approach. Broadly speaking, there are two largely equivalent ways of directly increasing the price of carbon through policy. First, the government could impose a carbon tax. Mankiw (2009) describes studies suggesting the optimal such tax is somewhere between \$30 per ton (roughly 8 cents per gallon of gasoline) to \$300 per ton (80 cents per gallon). Second, the government could directly cap the level of carbon emissions—issuing permits for emissions and allowing companies to trade those permits. Such a system is typically called *cap-and-trade*.

Textbook policy analysis prefers the carbon tax to cap-and-trade (see Mankiw, 2009, for a clear articulation of this position). Here are two reasons. First, the carbon tax is more flexible. If there is significant fluctuation in the demand for carbon emissions over time, cap-and-trade might impose an inefficiently low level of emissions in a high-demand year. By contrast, a carbon tax allows firms the flexibility to use more carbon, if it is worth it to them, in such years. Second, a carbon tax generates government revenue. Such revenue

could be used to offset other, less efficient sources of government revenue. A cap-and-trade system does not generate such revenue unless the permits are auctioned off at the outset. Thus, standard policy analysis suggests that a carbon tax is preferable to cap-and-trade and that if we must do cap-and-trade, we should auction the permits, using the revenues to reduce other distortional taxes.

These arguments have a lot of merit. But, in my view, they are incomplete. If we are interested in actually reducing carbon emissions and mitigating climate change, we have to take the political incentives just as seriously as the economic incentives.

I would argue that the politics flip the standard policy advice on its head. If you want to achieve an increase in carbon prices and reduction in carbon emissions, you should prefer cap-and-trade without a permit auction (i.e., permits distributed to current emitters) over cap-and-trade with an auction, which you should prefer to a carbon tax. Why do I say that?

First, let's think about adoption. The winners from a carbon tax are the broad public who benefit from mitigating the risks of climate change. But the broad public is a diffuse and unorganized interest. It can exert very little political power. The losers from a carbon tax include oil and gas companies, automobile manufacturers, and emitting industries. These are highly concentrated and organized interests. They can exert significant political power to block a carbon tax. What about cap-and-trade? If cap-and-trade is coupled with an auction for emission permits, the same analysis holds. Powerful, concentrated, and well-organized interests are being asked to pay for emissions that they previously made for free. They have every incentive to block such a policy. But what about if permits are given for free? Now the policy has costs and benefits for emitters. On the one hand, they are being forced to reduce emissions, a cost. On the other hand, they are being given control over a tradable asset of considerable value. Current emitters who believe they can reduce emissions relatively cost-effectively will be able to sell those permits for a profit. This might create the sort of organized and powerful interest needed to move the policy through the political process.

Second, let's think about sustainability. Even if a carbon tax were somehow implemented, the analysis above holds. The supporters of a carbon tax are diffuse, unorganized, and relatively weak. The opponents of such a tax are concentrated, organized, and strong. It would require remarkable vigilance to keep a carbon tax on the books. Cap-and-trade, with or without an auction, is just the opposite. Once permits are issued, they become a valuable financial asset. In addition to the owners of the permits themselves, once a market for such permits emerges, there are brokers, investment bankers, and a variety of financial services providers with a stake in the market. An organized and

powerful set of interests will fight to sustain the policy. Again, then, the politics favor cap-and-trade.²

I hope you find these arguments provocative. My goal, at this point, is not to convince you about the right way to think about environmental policy. Rather, I want you to see that, while traditional policy analysis is an essential input to thinking carefully about policy, it is not the end of the story. Policy is made in political environments. If you want to understand policy or effect policy change, you must take the politics of the policymaking process seriously. So let's get going.

Three Goals

In service of the broad aim of understanding the politics of public policy, this book pursues three interrelated goals.

The first goal—tackled in Part I—is to spend some time thinking about the normative foundations of policymaking. That is, to ask questions like, “What are the proper goals of public policy?” In so doing, we will discover how difficult it is to come up with a coherent, unified, normative framework to motivate policy decisions. We will arrive at one limited version of such a framework, but we will also see that the ambitious program of formulating a broad conceptualization of the public interest that everyone can endorse might be a fool's errand. Nonetheless, along the way we will discover some useful and provocative ideas that will make our thinking about these issues more structured and nuanced.

The second goal is to think through some fundamental aspects of social life that create opportunities for public policy to make the world a better place. I call these aspects of social life—which are developed in Part II—*social dilemmas*. These social dilemmas are ubiquitous features of human interaction that give rise to predictably regrettable outcomes. We will discuss ways in which policy might improve those outcomes. The objective, in this part of the book, is to develop some habits of mind that let you see the world through a few analytical lenses. I believe you will find that once you get used to viewing the world through the lens of these social dilemmas, you will start to see them everywhere. This may be a bit depressing, but it is also useful in identifying opportunities for policy to do good.

The third goal is to think seriously and conceptually about the politics. As I said in the preface, I'm not going to do this by analyzing the politics of policymaking institution by institution. Instead, in Part III, we look at two classes of explanations for why governments might not always achieve good policy outcomes. The first focuses on *technological constraints*—facts about the

²This argument is based on the discussion of the success of cap-and-trade in curbing sulphur dioxide emissions in Patashnik (2008, Chapter 8).

policymaking process that limit the government's ability to achieve good policy outcomes even when policymakers are genuinely motivated to do so. The second focuses on *incentive constraints*. Policymakers are people with their own interests and concerns. We analyze how some fundamental facts of politics—most especially leaders' desire to remain in power—interact with political institutions to determine when policymakers have better or worse incentives to pursue good policy. We end this part of the book by looking at a bit of evidence on how the organization of government affects policy outcomes—including an extended discussion of foreign aid policy.

The Role of Models

The main analytic tool we will use in this book is theoretical models—mostly mathematical models. As such, it is important to think a little bit about how to relate to or evaluate a model. What is a model, and what is one good for?

The world—be it physical, biological, or social—is too complicated for us to think through coherently all at once. To make sense of the world, we need to break it up into bite-size pieces that we can think about carefully. The hope is that, in so doing, we will figure out some general principles that help guide our thinking. Models are an attempt at doing this.

Given this, the goal of a model is emphatically not to describe the world in all its richness. A model is an abstraction and a simplification. Its purpose is to isolate some aspect of a problem or phenomenon, so that you can think it through carefully and without distraction. If a model captured all the richness of the world, you would be just as perplexed thinking about your model as you are thinking about the world. That would be a lousy model.

Let me be a little more concrete. A classic example of a model is a map. A map of a city, for instance, is a model of a particular part of the earth. It abstracts away from almost everything that is interesting about that part of the earth—the people who inhabit it, the quality of the restaurants, the crime rate, school districts, in some cases the topography, and so on. To be a successful model, a map must ignore this richness. A map of place and street names isolates precisely those bits of the world that you need to navigate the city. More information would distract from the task at hand—trying to get around. Less information would leave the map insufficiently rich to be useful. It's a delicate balance.

The analogy to a map points to an important fact about models. What a model should and should not include depends critically on what you intend to do with it. Consider two maps—one a street map and the other a topographical map. One cannot say, in the abstract, which of these is the better model of a bit of the earth. If one's goal is to drive, then the street map is the

right model. If one's goal is to hike without trails, then your map needs topographical information.

The same holds true with models of social interaction. It is tempting, when presented with a model of human beings, to start listing things that you consider important but that are not included in the model. To react that way, however, is to miss the point of a model—just like criticizing a street map for failing to include restaurant information is to miss the point of a street map. If you evaluate a model by asking whether it captures everything that is interesting about some question or situation, you will be perpetually dissatisfied. The world is an infinitely interesting and wondrous place. Our poor powers to abstract and to generalize are not up to the task of capturing all its richness. So, yes, all our models will leave you wanting, if you expect them to be full representations of the world, or even of some little slice of the world. But this doesn't mean our models are bad. It just means that expecting models to capture everything that is interesting about some aspect of the world is the wrong goal. What, then, is the right goal?

First, for a model to be useful, it cannot be a purely abstract object. We must be able to relate it to some aspect of the world in which we are interested. As the philosopher of science Ronald Giere puts it, models are “artful specifications . . . designed so that elements of the model can be identified with (or coordinated with) features of the real world” (Giere, 2006, p. 63). Notice the key point here. Some element of the model must correspond with some feature of the world. The model need not capture *everything* that is interesting or important about the bit of reality you are interested in. But it must capture *something* that is interesting or important about it.

Second, the purpose of a model is to help you think through some aspect of the world that is too complicated to think through in its totality. As such, the model should teach you something that you didn't see before you broke the situation up into bite-sized pieces. Otherwise, the model wasn't terribly useful.

In my view, the right way to evaluate a model involves asking something like the following two questions:

1. Is something in the model like something I am interested in out there in the world?
2. Did the analysis of the model teach me something about that aspect of the world that I didn't know before?

If the answer to both questions is yes, then the model has done its job. It has laid out an abstraction that you are able to relate to the world in some way. And the abstraction did in fact help you learn about that aspect of the world. These are the goals of a model.

If this is what a good model does, how should someone interested in public policy think about applying the insights of a model? I like to think of a model as

a really smart adviser who understands only one narrow aspect of a policy issue. You, the policymaker, should listen to your model and the ideas it suggests. But you should not do so slavishly. After all, you know a lot more about the world you are dealing with than your model does. A third question, then, that you might want to ask before applying the insights or policy suggestions that come out of your model is something like this:

3. Is there something about the world that is missing from my model that I believe would materially change the conclusions of the model were it included?

If the answer is no, then it seems the insights of your model are applicable. If the answer is yes, however, then you should proceed to apply the model's insights only in a cautious and somewhat skeptical way.

With all of this in mind, I should perhaps say a word or two about how I chose and constructed the models that appear in this book. My goal is not to bring you to the technical frontier. Rather, it is to introduce what I take to be some of political economy's most important insights in a rigorous, but accessible, way. As such, two principles guide my modeling choices.

First, I always present the absolutely simplest, least technical version of a model that makes all of the points I want to make. The goal is to make substantive insights clear and rigorous. This requires some technicality. But I emphatically do not want you to get lost in the math.

Second, I only present models whose messages I believe are general. I am not interested in presenting simple models whose main point is fragile. To do so would privilege clever formalization over substance. That's not what I'm about. So, for every simple model I present, I'm asking you to take my word that there is a body of scholarship (too technical for our purposes) that has developed the arguments in sufficient depth to convince me that the main idea is both important and robust. At the end of each chapter, I point the interested reader to some of these further resources.

Why Rationality?

Almost all of our models build on the assumption that people are rational. Rationality, here, means several things. In its most stripped down form, rationality simply means that people have coherent preferences over outcomes and they act to pursue those preferences. I discuss this a bit more formally in Appendix A.1.1.

But we will generally assume something stronger than just coherence of individual preferences. In particular, we will assume that people have coherent preferences over outcomes, that they pursue those preferences, and that those preferences are mostly centered on their own personal welfare. That is, we will

assume that people are not primarily motivated by altruism or concern for others, but rather by making themselves better off.

Notice, the assumption of rationality of this kind is already a model. It is a model of human agency. Why should we adopt this particular model? After all, we regularly hear stories of great acts of charity and heroism. Surely a coherent analysis of social life would take into account the fact that human beings are not entirely selfish and avaricious.

I don't disagree. People are, at times, altruistic. (Though the fact that such stories are newsworthy suggests that they may be the exception, rather than the rule.) But, remember, we are building a model. We are trying to simplify the world in a way that is useful for the problems with which we are concerned. And the problems with which we are concerned here are problems of public policy.

Public policy is, at least in one conceptualization, what society does when people, left to their own devices, do not act in one another's interests. Given this, if we are going to think about a model of people to motivate the making of public policy, surely we do not want to assume that people are primarily motivated to look out for one another. If we base our thinking on such an optimistic model of human motivations and proceed to design policies on the basis of that assumption, we may find ourselves disappointed by the actual behavior of our fellow humans. If, instead, we start with the assumption that people are basically selfish, we put ourselves in a position to think about policy interventions that will work even under the worst of circumstances. To be sure, such an assumption may be overly pessimistic, leading us to take precautions against bad behavior in excess of what is necessary. But, in the context of policymaking, this seems like the safer kind of mistake to make, at least as a starting point.

There is another (perhaps more pragmatic) reason for restricting our attention to rational individuals. Doing so massively simplifies many of our analyses. And, remember, the goal of model building is to come up with something that corresponds in some way to the world, but that is simple enough that you can think it through and learn something you didn't see before. For this purpose, the assumption of rationality (in the sense of having coherent preferences, not the assumption of selfishness) is indispensable.

Further Reading

The appendices contain all the game theory you need to understand this book. But if you are looking for an alternative introduction to game theory, there is none better than Martin Osborne's *An Introduction to Game Theory*. If you need a basic calculus refresher, I recommend Daniel Kleppner and Norman Ramsey's *Quick Calculus*.

Many philosophers of science have deep things to say about the role of models. My favorite is Ian Hacking's *Representing and Intervening*. In the chapter, I mentioned Ronald N. Giere's *Scientific Perspectivism*. Perhaps most relevant for the kind of models explored here is Mary S. Morgan's fantastic book *The World in the Model*. The map analogy as a means of explaining the pragmatics of models is certainly not mine. For instance, in her *Essays in the Theory of Economic Growth*, Joan Robinson wrote, "A model which took account of all the variegation of reality would be no more use than a map at the scale of one to one." Paul Krugman's essay "The Fall and Rise of Development Economics," which is the first chapter of his *Development, Geography, and Economic Theory*, has a fantastic discussion of the history of map making, its role in thinking about economic development, and the conceptual importance of simplified models in social science.