

Chapter 6

A Crash Course in Economics

We now move from the natural environment to the economic environment — a transition that requires an understanding of some key economic principles. The first set of questions that we want to answer are when society can reasonably rely on individual self-interest and when not; and how society can and should bend individual self-interest towards the common good. The second set of questions are about marginal thinking and tradeoffs (including issues such as best delay). The third set of questions regard innovation, research, and development. It includes how costs and economic decisions can change with scale and time.

1 Human Self-Interest and Free Riding

The basic premise of economic science is that individuals tend to act in their own self-interests. This premise is, of course, never perfectly true. There are situations in which we are more likely to make voluntary sacrifices, such as when we are close to the recipient and when it costs us less. Most of us will gladly buy dinner for our families or friends.

But beyond sacrificing modest amounts for these chosen few, the evidence indicates that most of us humans put our own interests far above those of others. For example, children in poor countries could be pulled out of privation for the price of one cappuccino a day — and yet the donations from rich nations remain by-and-large pitifully small. (And it is not a valid objection that some of it is wasted in the transfer. The remaining portion would still go far.)

It would be nice if we could assume that people would put the climate problems of the *other 8 billion* high up on their priority lists. Unfortunately, this is not a realistic way to approach the problem. The evidence speaks against it.

It follows that for a voluntary-participation climate-change plan to work, it must be in the self-interest of most people to participate. Each one of us 8 billion people won't make much difference. The world outcome will ultimately remain the same no matter what our friends and we do. (Whatever other people do, most will continue to do so, regardless of what we as individuals will do.)

For example, if avoiding polluting costs you \$1 and this pollution imposes costs on the world's 8 billion people of \$0.0000000125 each, the total world pollution cost would sum up to \$100. For the social good of the world, you should not pollute. Yet for your own good, why spend your \$1 for a return of \$0.0000000125? If you care about 100,000 other people as much as you care about yourself, or if your example somehow convinces 100,000 others also not to pollute, your return increases to \$0.00125. Let us ask again: Do you really want to spend your \$1 for a return of \$0.00125? Your answer matters little, though. What do you think the majority of people will do?

Economists call this free-riding — a situation in which you can consume a common good but without having to pay for it. If you are the free rider, it may be a great deal for you, but it is not a great deal to those who have to pay for you. If they could, they would want to exclude you (to make you choose to pay!) — but this is precisely the problem when exclusion is not possible (as in many cases). There are many examples in which groups cannot start profitable businesses, because they realize that they cannot prevent subsequent shirking by team free-riders. (The most famous example may well be the [prisoner's dilemma](#).)



Economists believe that it is a primary role of governments to ensure that their populations do not become too selfish. Societies cannot function if there is only individual self-interest. Thus, they need governments that can [coerce](#) their citizens.

For example, every sensible person would agree that any civilized society needs public services (streets, schools, police, courts, military, social security, welfare, hospitals). Only failed states, such as Sudan and Libya, are so libertarian that they do not provide these services. (And only a few predators enjoy living in them.)

These services need to be paid for by taxes. And any tax system requires a collection agency and laws to punish tax evasion. Otherwise, who would pay taxes? If paying was a matter of individual choice, everyone — both poor and rich alike and even most idealists — would find good reasons why only others should be paying the taxes. Furthermore, economic evidence suggests that altruism is not contagious at scale. In small communities, there are often individuals with enough concern for the common good to volunteer paying up. Others may pay more when they see these peers doing it (and especially when other villagers notice whether or not they are participating). However, voluntary contagious altruism almost never works in [large societies](#). How many other Americans would be more likely to pay more taxes voluntarily just

because *you* decided to set a good example and pay more taxes voluntarily? Not a meaningfully large fraction.

Nevertheless, although we need forced taxation and large governments, they are less than ideal: it costs money to collect taxes; it wastes money, because tax-rich governments can be twisted by lobbies, voters, demagogues, and disinformation; and it reduces work incentives, because individuals prefer to work for themselves instead of for the government. But taxes are an unavoidable and necessary social arrangement.

The best level of taxation is determined by what economists call a tradeoff — sacrificing one thing to obtain another. In our tax example, true libertarianism makes as little sense as true totalitarianism. Governments should not be too big or too small. (They are also never entirely benevolent and rarely entirely malevolent.) The right amount of government taxing and spending is a difficult and constant balancing act. If it weren't, it would have been a solved problem long ago and concerned citizens would not have been discussing the best scope and scale of governments at least [since Plato's times](#).

The best solution with respect to the tax example is never “no taxes” or “all taxes.” The best solution is somewhere in the middle. It is an eternal question where in the middle it should be relative to where we are at the moment: have we gone too far or not far enough?¹

Unfortunately, institutions and communities are also almost always less efficient the larger they are. This also applies to governments. A village can usually administer local public spending more frugally and efficiently than Washington D.C. can administer national spending or the United Nations can administer international spending. This is why small collective problems are easier to solve than big ones. And climate change is the biggest collective problem of them all.

¹Despite self-denial, the United States is already more “socialist” than many countries in Europe, except at the tiniest 0.01% sliver of its top wealth distribution. [Social Security](#) and [Medicare](#), two hugely popular programs, are the largest socialist redistribution programs in the world.

2 The Tragedy of the Commons

What should government take care of, and what should individuals take care of?

In some cases, the answer is obvious. What would happen if traffic laws were voluntary? Considerate drivers would stop at red lights, but less considerate drivers would just take advantage of them and cross first. Yet at some point, even the most inconsiderate drivers become worse off — once there are too many other inconsiderate drivers on the road. By setting reasonable traffic rules and fairly enforcing them, governments can greatly increase economic well-being for all. Good traffic rules speed up the flow of traffic and reduce accidents. No sane economist thinks it would be a good idea to leave it to drivers to do whatever they please. In this context, “free driving choice” is not in the collective interest.

In other cases, the answer is not so obvious. The best choice is often for government not to intervene. This is so counterintuitive that social philosophers before Adam Smith did not understand it (to the detriment of their states), and many intellectuals still bristle when they are told to keep their hands off because they cannot design a mechanism to allocate resources better than letting individuals make their own decisions. Free-market capitalism can coordinate prices and activities in a way that is not just in the individual (“private”) but also in the collective (“social”) best interest.

Suppose, for example, that you are considering taking a business trip to New York. The cost of a round-trip ticket is \$1,500. All you have to decide is whether the benefits of your trip, net of your other costs, exceed \$1,500. If they do, your private welfare increases when you take the trip. More importantly, if no one else is harmed, total social welfare also increases. If someone else benefits from your trip, e.g., the airline and its employees, all the better! Social welfare increases even more.

However, your purchase could lower the social benefit if all costs of your air travel are not reflected in the price of your ticket. The airline is not likely to be the source of such a problem. Presumably, it would not sell you a ticket for \$1,500 unless it was in its own interest, too, covering the cost of the fuel, personnel, wear and tear, and a fair return for the airline’s investors.

Yet there is a problem with the free-market solution. Airplanes also emit pollution, and neither you nor the airline pay for the damages the pollution

causes. These are called “external costs” — for short, “*externalities*.” In this case, they are negative externalities — because they “spill over” negatively onto others beyond the parties to the transaction. Because your ticket cost does not reflect these externalities, your trip may be beneficial to you and the airline, but not necessarily to society. For example, if your emissions cost others \$500 in health costs, then you should only make the trip in the social interest if the benefits exceed \$2,000.

Figure 1. Los Angeles Before the E.P.A.



Source: [Sierra Club](#).

Even the most ardent free-market economists agree that government intervention may be needed when large externalities are involved, because private transactions fail to take into account all effects on third parties. On the other hand, the opposite alternative is also not desirable. Governments should not be in the business of making airline booking decisions for individuals. They are not as good at knowing what you need as you are.

So how can governments arrange it so that the participating parties will also take the social costs into account? To do so, the government must find a way to include the \$500 external cost in the price of the ticket, so that you and the airline make the best decision in the social interest. Most of this chapter is about explaining, in detail, how and why this is best done with a tax on jet fuel emissions.

If you are still not convinced that we need a government to regulate emission externalities, just learn from history. Prior to the formation of the [Environmental Protection Agency](#) in 1970, there were cities in the United States that looked just like [Beijing](#), Lagos, or New Delhi today. Figure 1 shows off Los Angeles. Smog was often so thick that visibility was limited to one mile. Oil slicks on rivers were widespread. Arsenic, cyanide, and mercury poisoning of groundwater were common. Lead in paints and gasoline poisoned our children. Asbestos killed thousands. Thank God those worst of times are over — and this is due to the intervention of government curbs on self-interested individuals.

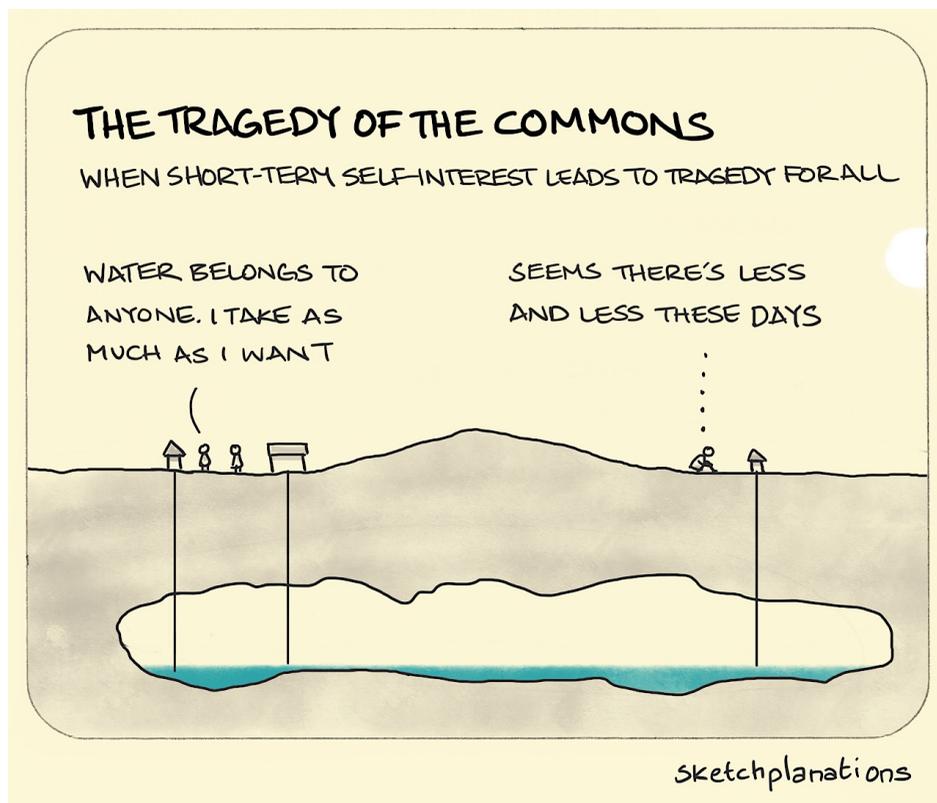
Economists also call such situations “[The Tragedy of the Commons](#),” based on an academic study by Garret Hardin. He looked at examples in nineteenth century England, in which farmers let their sheep overgraze the “commons” (a publicly-shared area). It was so bad that many commons became completely run down. Because each farmer had acted in his self-interest, they collectively lost all benefits of the shared resource in the end. In our case, the analogy is that pollution emitters have been overgrazing the common good of clean air.

Conflicts over shared resources are difficult to resolve.² Even state governments often fail to handle state instances well (and there are good reasons why, which we will explain soon) — much less do so when the problems cross state borders and require coordination among multiple states or even nations.

For example, [water in Southern California](#) (where we live) is a scarce resource. Los Angeles could not exist without managed water access. Yet until very recently, many Californians considered it their civil rights to [drill](#) water wells on their lands. Not surprisingly, some aquifers ran dry and the rest may do so soon. There have been many [water wars](#) fought in the Western U.S. states over the last two centuries. These states are now struggling hard to find a collaborative solution that makes everyone better off than the free-for-all. It is not certain that they will succeed.

Yet it will always be a valid question as to where we are versus where we should be: has government intervention gone too far or not far enough? We study economics to answer thorny questions like this.

²Elinor Ostrom won a Nobel Prize for studying cases where individuals and communities did manage to avoid the tragedy. However, these were *local* situations, not *global* ones.



3 When Should Governments Intervene?

So when should society leave it to free parties to make decisions, and when and how should it involve government coordination and coercion? The answer is easy if an activity is infinitely bad (ban it!) or completely harmless and without externalities (allow it!). But as with all things economic, the tradeoff is interesting and nuanced when it is in the intermediate gray zone.

Greenhouse gases (GHGs) fall into this gray zone. They are not infinitely bad, and they are not totally harmless. We must evaluate how bad the externalities associated with GHGs are. But how? And should it be different if we know that any government we might appoint to work on the problem will be inefficient or even corrupt? And what if some parties lose (like coal miners)? And what if GHGs could later be sequestered again from the atmosphere for a price? Read on.

A Numerical Game-Theory Example

We are now going to try to make a real economist of you with a small dose of [game theory](#). Game theory studies situations in which the payoffs of “players” depend not just on their own actions, but also on the actions of others. If you’ve ever played poker, you already understand this dynamic.

The concepts are so important that they are worth illustrating with a numerical example. However, if you are truly allergic to numbers of any kind, it is possible to skip the numerical example and still understand most, but not all, of what we will discuss later.

Our kinds of “games” always include three elements: (1) players, (2) choices (strategies), and (3) payoffs. The players in our example are three companies, X, Y, and Z. (You can also think of them as three sectors of the economy. The payoffs can also go to both owners and employees.) Our choices allow each company to pick one of two different technologies, \mathcal{C} or \mathcal{D} . Our payoffs need to be shown in a two-dimensional table, because they depend on all three players (X, Y, Z, in columns) and both choices (\mathcal{C} and \mathcal{D} , in rows):

	<u>Players: Company</u>			<u>Total Social Welfare</u>
	X	Y	Z	
w/ Tech \mathcal{C}	\$300	\$400	\$500	\$1,321
w/ Tech \mathcal{D}	\$290	\$421	\$600	

The boldface in each column highlights which of the two technologies yields bigger payoffs for each player. The table shows that player X prefers \mathcal{C} , and that X and Y prefer \mathcal{D} . If each firm acts in its own self-interest, the total social welfare (the sum of the three firms’ payoffs) is maximized at \$1,321 (billion, of course). There is no better collective choice, even if firms could be forced to do whatever society wants them to do. In this game instance, firms can be trusted to make the best social choices.

Now we introduce the next step: pollution. Let’s say that \mathcal{C} is a “clean” technology, and \mathcal{D} is a “dirty” technology. \mathcal{D} leaks paint all over the machine and requires regular cleanup, which costs \$60. Taking pollution cleanup costs into account, here is how each firm will select its strategy:

	Players: Company			Total Social Welfare
	X	Y	Z	
w/ Tech \mathcal{C}	\$300	\$400	\$500	\$1,240
w/ Tech \mathcal{D} – pollution	\$230	\$361	\$540	

Now that \mathcal{D} is polluting, it is less profitable and so there is only firm Z still using \mathcal{D} . Once Y has to take care of its own cleanup costs of \$60, it prefers the clean \mathcal{C} over the dirty \mathcal{D} . The social welfare is lower at \$1,240 because we have changed the technology. Now \mathcal{D} requires cleaning, which was not the case in the previous \$1,321 example.

Most importantly, the fact that the technology is dirty makes no difference in one respect: We still do not need a government. X and Y are choosing the clean technology; Z is not. Government could not do better by telling any company to switch.

A Dirty Externality

Next, we introduce external effects into the example. To do so, we change the manner in which \mathcal{D} is dirty: rather than leaking into the company's own plant, the paint releases its polluting vapor beyond the firm. We assume that the total pollution of using \mathcal{D} is still the same \$60, but the victims change: each of the three companies suffers \$20 in required cleanup costs when any one of them chooses \mathcal{D} . Of the \$60 in pollution, the polluter bears \$20 itself, and the externality imposed on the other two is \$20 each for a total of \$40.

Let us first confirm that if we have a Utopian dictatorial government that forces X and Y to use the clean technology and Z to use the dirty technology, social welfare is still \$1,240. With one polluter, each firm will suffer \$20 in pollution, caused by Z but imposed on each of the three.

	Players: Company			Total Welfare
	X	Y	Z	
w/ Tech \mathcal{C} – pollution	\$280	\$380	<u>Verboten</u>	\$1,240
w/ Tech \mathcal{D} – pollution	<u>Verboten</u>	<u>Verboten</u>	\$580	

It may be unfair to X and Y to suffer the consequences of Z's pollution, but our economic concern is not fairness. Our concern is to achieve the

highest social welfare achievable. If you want to worry about fairness across companies, just realize that nothing prevents the government from taxing any company (such as polluter Z) any amount it wishes and giving the tax revenues to others (such as X and Y).

The Free-Market Solution

We are now ready to explain the problem that arises under the “everyone for themselves” free-market solution. If left to their own devices, what would each company do? For now, assume that each company takes the behavior of others as given, and just focuses on its own choice and welfare. In this case, here is the solution:

Free-Choice Solution: X uses \mathcal{C} ; Y and Z use \mathcal{D} .
(Each company suffers \$40 in pollution [bold “in equilibrium” box].)

	Players: Company			Total Social Welfare
	X	Y	Z	
Proposed Solution	\mathcal{C}	\mathcal{D}	\mathcal{D}	
w/ Tech \mathcal{C} – pollution	\$260	\$380	\$480	1,201
w/ Tech \mathcal{D} – pollution	\$230	\$381	\$560	

Game theory solutions are like puzzles or brain-teasers. Let’s [trust but verify](#) that the solution is correct:

- Cmpny X: Because Y and Z use \mathcal{D} , X will suffer pollution of \$40, over which X itself has no influence. If X chooses \mathcal{C} , X will earn $\$300 - \$40 = \$260$. (The first term, here of \$300, is from the first table on page 9.) If X chooses \mathcal{D} , X will suffer the \$40 pollution from the others, plus an extra share of \$20 from its own pollution. This would give it only $\$290 - \$40 - \$20 = \230 . Thus, we have confirmed that X prefers \mathcal{C} over \mathcal{D} .
- Cmpny Y: Because Z uses \mathcal{D} , Y will likewise suffer pollution of \$20, over which Y itself has no influence. However, Y has control over its own pollution. If Y chooses \mathcal{C} (and thus does not pollute), Y will earn $\$400 - \$20 = \$380$. If Y chooses \mathcal{D} , Y also suffers its own pollution and will earn $\$421 - \$40 = \$381$. Thus, we have confirmed that Y chooses \mathcal{D} .

Company Z: Because Y uses \mathcal{D} , Z will suffer pollution of \$20, over which Z itself has no influence. However, Z has control over its own pollution. If Z chooses \mathcal{C} , Z will earn $\$500 - \$20 = \$480$. If Z chooses \mathcal{D} , Z will earn $\$600 - \$40 = \$540$. Thus, we have confirmed that Z chooses \mathcal{D} .

Summing up all the boldfaced best choices gives the result that social welfare is now only \$1,201. Although the technology still pollutes by the same \$60 as it did in the previous \$1,240 example (in which there were no externalities), the social welfare in the free-market solution is lower when there are these externalities than when there is coercion. This is because Y is not doing the right thing from a social perspective. Y ignores that its pollution has a \$40 negative effect on two other parties.

This is the classic *public goods* problem — the tragedy of the commons in numbers. The reason why we wanted to explain it with a numeric example is that we can now use it to demonstrate what government can do to improve the social welfare.

Government Dictates

► Green Dictatorship

Our Utopian government is not a realistic possibility. So we need to look at other solutions. Our first is a “dictatorship of the environmentalist”: Outlaw all pollution!

	<u>Players: Company</u>			<u>Total Social Welfare</u>
	X	Y	Z	
w/ Tech \mathcal{C}	\$300	\$400	\$500	\$1,200
w/ Tech \mathcal{D} – pollution	Verboten	

Although there is now no pollution and thus no externalities, from a social welfare perspective, this turns out to be a bad solution, too. Government should let Z pollute. (In our example, the government can later clean up the pollution for \$60.) Z’s gain from \mathcal{D} is higher than the total externalities it imposes on X and Y. The problem with this dictatorial green solution is that it prohibits a socially valuable but polluting activity. (Yes, pollution can be in the social interest!)

Most modern activities create pollution. We would not want to outlaw them all. The key is intelligently trading off the benefits of the activities versus the costs of their pollution.

► **Smart Dictatorship and Command-and-Control**

As noted on page 10, the Utopian dictatorial solution, in which a [benevolent government](#) tells X and Y to use \mathcal{C} and tells Z to use \mathcal{D} , is better. In fact, it yielded the best possible outcome, \$1,240. This was the theory behind old-school Soviet-style planned command economies. If the commands are correct, the social outcome is best.

The problem with these command economies is not necessarily that the planners are evil (although they often are), but that even the best planners usually don't have all the necessary information, and their clever subjects will do all they can to exploit any system to their advantage. How would our dictator know whether to tell X and Y not to pollute, and allow Z to pollute? Y will make all sorts of claims why its payoffs with pollution are not just a little but a lot higher. What if the planner's cost and benefit estimates turned out to be wrong?

The economist [James Gwartney](#) tells a wonderful [anecdote](#) of how economic planners in Moscow tried to incentivize managers and employees in the Soviet Union to produce window sheet glass. The planners instituted a reward system based on the tons of glass produced. What could possibly go wrong? The glass plant managers realized that they could earn the rewards by producing sheet glass so thick that it was barely transparent. When the planners realized that they had set the incentives wrong, they changed the rules to reward square meters of glass produced. The results were similarly predictable. Under the new rules, the plants produced glass so thin that it broke all the time.

The Soviet planning system was an extreme example of directed production that was not based on the prices that free-market consumer demand would set. But even capitalist governments must institute incentives to accomplish social tasks. And they all make similar mistakes. For example, under French colonial rule in Hanoi, a government [rodent extermination program](#) paid bounties for each rat tail handed in. Instead of eradicating rats, many hunters caught rats and amputated their tails (so that the rats would continue to breed), while others started farming rats for profit.

Allow us a cautionary tale from the domain of regulating global emissions. In the [Montreal Protocol](#) to protect the ozone layer, firms were awarded valuable credits for destroying the potent greenhouse on hydrofluorocarbon gas HFC-23. However, production of [HFC-22](#) as feedstock ingredient for other products was allowed. The production of HFC-22 creates HFC-23 as a by-product, the subsequent destruction of which would now yield extra credits. Not surprisingly, [HFC-22 production has been rising](#).

sidenote

One possibility for dealing with the public goods problem is for rich countries to pay poor countries for behaving in a more climate-friendly manner. For instance, rich countries could pay Brazil, Congo, and Indonesia to retain their rain forests. In effect these countries could probably be convinced with very modest subsidies to harvest the wood and replant it rather than burning it to clear agricultural land. (We will discuss international treaties in Chapter 7 in more detail.)

Unfortunately, if one pays for reducing harmful actions, the incentives to commit those actions in the first place in order to obtain the benefits for the reduction also increase. If rich countries were to pay Brazil and Congo to stop burning down rain forests, would it be a reward for their current practices of aggressive burning? Furthermore, how many American voters would be willing to pay taxes so that the money can be transferred to Congolese warlords? (And what would they really do with the money if they got it?)

But don't laugh too hard at the follies of Soviet planners, French colonialists, and climate treaties. Firms are essentially small economies, in which management tries to coordinate economic activity without resorting to free-market price-based contracting. (Otherwise, why have a firm at all? Individuals could just transact in free markets.) The CEO's role is essentially that of a planner and even dictator — and many firms have failed for the same reasons of poor information, mistaken incentives, and clever employees as government programs have failed.

Information and incentive problems are at the heart of much modern economics. And these same problems also lie at the heart of solving the climate-change problem.

A Managed Market With a Pollution Tax

Returning to the numerical example, we still need a better solution in which firms do not pollute when it is not in the social interest but a dictatorial planner does not have to tell firms how to operate, either. After all, firms should know best about the relative costs and benefits of their (C and D) production technology choices. We now explain how to do this. Government can set a tax equal to the externality (the “social harm”), and *then* let firms decide for themselves.

Recall that the externality we discussed above was \$40. What would happen if government taxed the use of the dirty technology by exactly this amount? Again, assume that each company takes the behavior of others as given and focuses on its own behavior and welfare. With a \$40 tax, the firms will choose the following:

Solution With-\$40-Tax on D: X and Y use C; Z uses D.
(Each company suffers \$20 in pollution [in bold “in equilibrium” box].)

	Players: Company			Total Welfare
	X	Y	Z	
w/ Tech C – pollution	\$280	\$380	\$500	\$1,240
w/ Tech D – pollution – \$40 tax	\$210	\$341	\$540	

The \$1,240 is the sum-total of \$1,200 to the firms plus \$40 in tax revenues, which belong to society collectively. Again, we need to check that we have solved the puzzle correctly, i.e., confirm that this solution is in every company’s self-interest:

- Cmpny X: Because Z uses D, X will suffer Z’s pollution of \$20 over which X itself has no influence. If X chooses C, X will earn $\$300 - \$20 = \$280$. (The first term, here of \$300, is from the first table on page 9.) If X chooses D, X will still suffer Z’s pollution of \$20, plus another \$20 of its own pollution, plus the tax of \$40, for a private benefit of $\$290 - \$20 - \$20 - \$40 = \$210$. Thus, we have confirmed that X chooses C.
- Cmpny Y: Because Z uses D, Y will suffer Z’s pollution of \$20 over which Y itself has no influence. If Y chooses C, Y will earn $\$400 - \$20 = 380$. If Y chooses D, Y will still suffer Z’s pollution of \$20, plus another \$20 of its own

pollution, plus the tax of \$40. This gives Y $\$421 - \$20 - \$20 - \$40 = \$341$. Thus, we have confirmed that Y chooses \mathcal{C} .

Company Z: Because X and Y use \mathcal{C} , Z only suffers pollution over which it has no influence. If Z chooses \mathcal{C} , Z will earn \$500. If Z chooses \mathcal{D} , Z will earn $\$600 - \$20 - \$40 = \540 . Thus, we have confirmed that Z chooses \mathcal{D} .

The social welfare is the same as that chosen by the Utopian benevolent dictator — except that all that this limited government planner had to measure was the social externality and set a tax equal to it. Government could then leave it to the companies to measure their relative benefits from the two technologies and make their own decisions. The pollution tax provided the proper incentives.

A subsidy to clean technology instead of taxes, or a combination of clean technology subsidies and dirty technology taxes, could also work. Again, the point of these policies is not for the government to collect or distribute funds; it is to alter the incentives of firms so that it is in their self-interests to do the right thing.

It is important to recognize that there is a distinction between pollution costs and taxes. Pollution is a negative, both from a private and a social perspective. Taxes are a negative only from a private perspective, not from a social perspective. The taxes paid to the government do not evaporate. Instead, they can be used for the common benefit — such as removing the paint vapor from the air or paying X and Y to remove the \$20 in vapor damage caused by Z.

sidenote

There is a close alternative to a pollution tax, called [cap-and-trade](#). Emitters need permits ([carbon credits](#)) to pollute, but they can trade these permits. We will not go into further details on the advantages vs. disadvantages. There are many online explainers, such as one at [Brookings](#) or at the [World Resources Institute](#).

4 Practical Problems of Pollution Taxes

A government tax on pollution equal to the externality is a great solution, because it puts each party in charge of what it knows best. Firms know their businesses, but they are often not aware of how much of an externality they impose on others. A firm should be in the business of its business. We cannot expect it to research the social good. Government should be in the business of finding out how its members work together and making them collectively better off. Thus, judging social harm, which affects non-contracting parties, is what governments should be in the business of.

Good economics suggests that governments should tax more when the pollution problem is worse. The challenge is finding a good balance. If you want certainty, stick to death and taxes — usually no one knows exactly. In the case of pollution externalities, even the best level of taxes are often uncertain. And the implementation problems go beyond just the uncertainties of estimation. The theory is easy, but the practice is hard. In the real world:

- What if it is not easy to judge harm? What if it costs \$100 to investigate whether there is any externality at all?
- What if the government is inefficient? For example, what if administering a \$40 tax costs \$10 or \$100? What if the government does not use the funds to clean up the pollution as it is supposed to do, and this allows the pollution to cause a lot more harm?
- What if the government is corruptible and untrustworthy? What if one of the three parties in our example has so much influence that it practically “owns” the government? In the non-legal sense, all governments are corrupt to some extent. We cannot trust governments to always do the right thing — yet we have no better alternative.
- What if voters clamor for subsidies and lower energy prices? This is not far-fetched. When energy prices experienced a small spike in 2021, even the voters in [Norway](#) suddenly wanted relief from high energy taxes.

Economics generally suggests that the worse governments are, the less they should be involved; and, in extremes, it is sometimes better simply to accept a bad free market-solution than to replace it with a worse government “solution.”

But what if the pollution is seriously harmful? What if the pollutant is not just paint but cyanide? Would you want to allow the [Cyanide Manufacturers](#)

Trade Association of America to provide the information to judge how harmful cyanide is to the people in the neighborhood just outside of their factories? Would you leave it to their profit margins and industry wages to decide how much cyanide they deem safe to emit? Would you want to allow them to fund the election campaigns of politicians? Would you want them to fund those scientists who (truly or ingeniously) happen to believe that cyanide is not that bad? Would you want to have an EPA administrator who was a lawyer lobbying for the cyanide manufacturers in a previous life...or will be again in a future life?



Your Majesty, according to our study, the shoe was lost for want of a nail, the horse was lost for want of a shoe, and the rider was lost for want of a horse, but the kingdom we lost because of overregulation.

Industry do as they please without any control in a free market could be even more disastrous. Ideally, administrators should be free of conflicts and well-informed — but this is simply not how the real world works.

In real life, both Congress and the administration (including the EPA) are often partly captured by industry and by single-issue or single-industry voter interests. The average top-level U.S. government agency administrator worked for a political or industry think-tank and will do so again after leaving the (low-paid) job. However, the alternative of letting every in-

Winners and Losers

But even with a perfectly informed and benevolent government, there is still another big problem. Higher social welfare does not mean that each and every party will be better off. There will be winners and losers. In our example, Z will be worse off, because it will have to pay taxes (which all parties, collectively and jointly, will get to keep and use productively). Y will be miffed, too. Y did not have to pay the $\$$ tax but could have made more profits without it. Similarly, we did not concern ourselves with how to compensate victims of pollution. Should the tax receipts flow to the pollution victims? Seems fair, but what if more powerful lobbies have better lawyers and are experienced in getting government grants?

The problem with dictatorships is that they rarely have the social interest at heart. Benevolent dictatorships are rare. The problem with democracies is that even a government that wants to have the social interest at heart is often ill-suited to solving the social pollution problems. We may chuckle that “poor” Z is now rightly forced to pay taxes — until Z starts a public relations campaign about how unfairly it has been treated, sues based on a claim that the EPA has overstepped its bounds, gets its employees to vote against the current administration, and funds the opposition presidential candidate. If Z employs more voters than X and Y, it could even win. Z may not even need a majority of voters; it may be enough if it has a majority in pivotal swing states.

In the United States, coal miners in Pennsylvania and West Virginia may be few, but there may be enough of them to tip the balance in favor of one or the other presidential or senatorial candidate. Perhaps the most prominent nonsensical political subsidy in the United States for decades is for [ethanol](#) (alcohol as fuel). It exists only because corn-growing farmers in Iowa are the earliest primary voters. Ethanol is not a green fuel, because it costs more diesel fuel to farm ethanol than it creates.

Elasticity of Escape

But even these dire problems in choosing among lesser evils still paint too rosy a picture of real life. The most serious problem of taxes (other than the fact that they may be set wrong) is that payers can seek to escape them. If government taxes technology \mathcal{Z} , what prevents Z from setting up shop in a neighboring country, right across the border (which keeps the pollution externalities we suffer the same), which happens to be a country that is looking for more employers and has thus not imposed any tax? (Subsidies for \mathcal{C} don't solve the problem, either, because they will attract new "pretend-polluters" that will emerge primarily to collect subsidies on \mathcal{C} .) As far as CO_2 is concerned, it does not matter whether it is emitted in the US or in China.

And, for an extreme example of this problem, what if X, Y, and Z were not firms but countries? Would Z ever impose a pollution tax on itself? Once you answer "no," then you understand the single most important aspect of our global CO_2 tax problem. The fox will not protect the henhouse. When it is strongly against their self-interests, countries will not volunteer.

joke

[\[Isn't instituting a pollution tax in just one region\]...like making a peeing section in a swimming pool?](#) — George Carlin, comedian.



Economists have earned a deserved bad rap for being more pro-capitalist and less interventionist than other social scientists. In defense of our discipline vis-a-vis the social science disciplines of our colleagues, most of us economists are not so much enamoured with the free market as we are more skeptical about real-world governments. Just because the free market has a problem does not mean that the government, especially over the long run, will solve it better.

5 Margins, Costs, and Scale

Understanding pollution taxes and the limits of both the free market and of governments was the most important point of our chapter's lesson in economics. But there are just a few more important concepts that play central roles in the economics of climate change. This section discusses margins (perhaps the central concept of economics for most students) and the economics of costs.

Marginal Thinking on Global Warming

As we explained in the previous chapter, climate change is real and it is big. The global change of 2-3°C will cause large changes in climate. Worse, this large a change could push Earth over tipping points and into feedback loops that scientists do not fully understand.

However, when climate activists advocate for reductions in CO₂, e.g., at UN COP conferences, many make a basic mistake. They compare apples with oranges. They fail to understand that they are not fighting against the terrible consequences of 2-3°C warming. They are fighting “only” for reducing global warming by 0.3-0.5°C. (It is realistically not possible to fight against 2-3°C warming any more than it is possible to fight against the sun going down at night.) It is only the *marginal* benefit of *lesser* harm that world leaders have to weigh against the *marginal* cost of GHG reduction measures.

Scientists are of course aware of this. In the next chapter, we explain that the *global marginal benefits* of reducing warming by “only” 0.3-0.5°C relative to the *global marginal cost* of curbing CO₂ are still large enough to warrant a global carbon tax of about \$50/tCO₂. The result would be fewer GHG emissions and an increase in global social welfare of about 0.1-0.2% — a large sum, on the order of about \$100 billion per year. However, although it would be collectively suboptimal for humanity to forego this 0.1-0.2% gain from a reduction of 0.3-0.5°C via CO₂ reductions, it would *not* be world-ending. In fact, in aggregate economic terms, it would make only a tiny difference.

Diminishing Returns to Scale

The [majority of dwellings](#) in the United States house more cockroaches than people. To exterminate 90% of the roaches in a home may cost only \$100 in spray cans. However, it could easily cost \$1,000 to get rid of 99% and \$100,000 to get rid of all 100%. (Roaches are good at hiding.) To get rid of all of them, the average cost is \$1,000 per roach-percent. But after the first 99% have been wiped out, the last 1% cost an incremental \$99,000. Economists call this \$99,000 the [marginal cost](#) of the last 1%. (Most residents follow the sensible route and wipe out only the first 95% year after year.)

Here is another example. Let's say you are the typical American driving 15,000 miles per year. If we asked you to drive 100 fewer miles per year (e.g., to reduce pollution), we may not even have to pay you. You could easily find a few unnecessary trips that you could omit. If we asked you to drive 1,000 fewer miles per year, it could become uncomfortable. You might have to curtail some fun trips that you would otherwise have taken. If we asked you to curtail 10,000 miles, it would seriously affect your life and work. We would probably have to pay you a lot of money to get you to agree. Because the 10,000 miles is the sum of 100 trips of 100 miles each, you can think of the first 100 mile reduction as costing you almost nothing, the tenth reduction costing you more, and the one-hundredth reduction costing you a lot. Every additional cut is progressively more painful.

The marginal cost of the first trip reduction is lower than the average cost, which is lower than the marginal cost of the last trip reduction. You can think of the average of all earlier marginal cuts — from the first zero-cost reduction to the last most expensive marginal-cost reduction — as the average cost. Non-economists often confuse the two, and the most important lesson for new economics students is to get them thinking in terms of margins for decision-making, not averages.

When marginal thinking is applied to purchases, economists call the increasing cost phenomenon the [Law of Diminishing Returns](#) (LDR). You probably enjoy the first hot dog, but when forced to eat the [70th hot dog](#), you would probably no longer enjoy it as much as the first.

The LDR is of pervasive importance in economics. From our perspective, the most important consequence is that it is likely much cheaper to clean up 2% of each of 50 U.S. states (when each dollar of cleanup still goes far) than to clean up 100% in one state (where the last 1% of pollution cleanup

could cost billions of dollars). Similarly, it is likely to be relatively cheap and easy to cut the first 1-10% of global fossil fuel use. It could even turn out to be profitable — we will discuss the options in the third part of our book. Entrepreneurs could plant trees and sell the lumber. But eventually it would become harder and harder to find good land for profitable tree farming. Similarly, governments could purchase the dirtiest least profitable coal plants very cheaply in order to retire them. But eventually they would run out of bad cheap coal plants. At some point, it would become more expensive to reduce CO₂.

The conclusion? It will be increasingly more expensive, per ton of CO₂, for the world to cut, say, 90% rather than 10-50% of global CO₂. And it will likely be almost impossibly expensive to cut 100%, at least for many decades.³ Humanity should not be dogmatic. An ambitious goal to reduce as much CO₂ as is economically feasible is good. A stupid goal to eliminate every last bit of CO₂ is not.

The LDR also implies that the world should immediately take care of the “lowest-hanging fruit” (CO₂-emitting activities that can be cleaned up), because they are very cheap to take care of. And this is exactly our point—let’s *immediately* take care of the first 10%. *Move the needle*. We know it’s economically worthwhile. From a collective perspective, the world is recklessly delinquent in not moving this needle now.

Sunk Costs

The marginal concept has another startling consequence. After any generating plant has been built, its marginal cost of electricity production is far lower than the average cost (which does include plan construction cost). The extreme cases are nuclear power plants. It costs \$5-10 billion to build one; but, once built, it can deliver power at a per-MWh cost lower than other technologies. (Uranium is dirt-cheap.) Consequently, closing already constructed nuclear power plants prematurely is wasteful. (We will explain nuclear technology and its alternatives in greater detail in Chapter 10.)

The average cost of electricity from many large fossil-fuel plants is also much higher than the marginal cost, although fossil fuels are more expensive

³For example, polar researchers should probably never build a clean power plant for their Antarctic winter research station. Diesel and diesel generators are just too efficient and cost-effective to be replaced by alternatives.

fuel than uranium. Coal plants have lifetimes of about 30-50 years. Once built, it becomes inefficient not to use them. Thus, the game is largely rigged at construction time.

It is ominous that China is building new coal plants this decade at a record pace. It is an easy prescription that the world should do everything in its power to convince China not to build them. If the world could only get China to replace its aging coal plants not with newer coal plants but with nuclear plants, it would make a huge difference in global CO₂ emissions. Unfortunately for the world, it may already be too late to convince China. China already suffered acute [power outages](#) in 2021. These outages make it even less inclined to put its current coal plans on hold. It may be too late for China, but it may not yet be too late for India, Africa, and other regions. They need to get to the point where it is in their economic interests not to build coal but clean electricity plants!

Learning Curves and Returns to Scale

Over small quantities and in the short run, most — but not all — activities follow the LDR, with its decreasing marginal benefit and increasing marginal cost. Over large quantities and in the long run, this is not so clear. Learning and mass production can often reduce marginal per-unit costs.

For example, when [Thomas Thwaites](#) decided in 2021 to find out how much it would take to build just one toaster, it cost him \$1,792. If you add his time, it would probably be ten times this. Engineers call the first pilot product [FOAK \(first of a kind\)](#). If Thwaites had to build a few more, the average cost per toaster would be much lower — in engineering lingo, these toasters would be [NOAK](#) (next of a kind). The cost decline would reflect his [learning curve](#).

Furthermore, if Thwaites built 1 million toasters and appropriately farmed out tasks, he could probably do so for as little as [\\$20/unit](#) because of [economies of scale](#). For example, mass producers do not hand-make single pieces of plastic. Instead, they make casts (that costs a few thousand dollars one time) and then pour thousands of plastic pieces at almost zero marginal cost.

One central question when it comes to re-engineering humanity's energy and pollution systems is: "When and how will per-unit-costs increase or decrease?"

It is almost surely the case that some clean technologies will become cheaper with more experience and mass production and thus have increasing returns to scale. Other technologies will become more expensive with declining resources and thus have *decreasing returns to scale*. We don't yet know which are which. Will electric energy storage become more expensive as we need to mine increasing amounts of lithium, cobalt, etc., to build more batteries; or will we learn how to build better cheaper energy storage solutions from other ingredients altogether? (History suggests the latter.) Will CO₂ sequestration become more expensive as we run out of space to plant more trees or will we find cheaper ways, e.g., by breeding trees that are more efficient in converting CO₂ into wood?

The Economics of Delay

Virtually all instances of increasing returns to scale due to mass production share one feature: They require a high upfront fixed cost. The construction of large-scale factories is expensive and not undertaken lightly. In many cases, it is more efficient not first to build factories but to invest more in research and development to bring down the factory cost as much as possible, and only then to spend the much larger amounts necessary for deployment of better plants in large quantities.

Depending on the specific technologies, humanity today may not yet be well-served to pursue zero-CO₂ solutions. This does not mean it could not be well-served investing in researching how to build zero-CO₂ solutions in the near future. Research and development is often orders of magnitudes cheaper than deployment. By saving on deployment now, scientists might be able to spend ten times as much on research and development. Most scientists and engineers believe that investing in clean-tech R&D should be highest priority for humanity today. Have we already mentioned that we are strongly in favor of more clean-tech R&D?

When should entrepreneurs deploy a new technology? The right time is not obvious. If a technology costs decline by 10-15% every year (as battery costs have declined over the last 10 years), then it will cost only half as much to deploy the batteries five years later. Ergo, is it better to buy a battery that saves emissions of 1 MtCO₂ (over the next 30 years) today, or a battery that saves 2 MtCO₂ but beginning in 5 years? Or half-half? Applied economists build models to inform such decisions. Generally, when costs are coming

down fast, it makes more sense to delay. However, delay is also dangerous. What if progress will be slower than anticipated? (For example, controlled [nuclear fusion](#) has been 30 years in the future seemingly forever.) And what if delay leads to procrastination?

6 The Economics of Innovation

Our final topic combines the previous two. It is about innovation, and how the world should foster and pay for it.

► The Research Externality

Whereas pollution is a classic negative externality, research is a classic positive externality. Few for-profit companies conduct [basic fundamental research](#) (not related to specific products), because the benefits would rarely accrue primarily to themselves even if the research succeeds. Instead, their fundamental-research insights will likely also benefit their competitors, other firms, and eventually all of us. We are in effect free-riders. Ergo, companies do not invest the much larger fundamental-research sums that would be in the social interest.

For argument's sake, let's say there is a good chance (say 10%) that chamomile tea would create free energy, and it would only cost \$10 billion to do the basic research to find out. If \$10 billion would possibly solve the world's energy problems, it would be a bargain for the world! Yet, no private company would spend the money. If the research succeeded, every company could just brew chamomile energy. Thus, each company would prefer to follow rather than lead.

It is because of positive externalities like this that many countries fund basic research at their universities rather than leave it to for-profit companies. Although funding for the [National Science Foundation](#) and some other agencies helps reduce the externality problem in the United States, it does not solve the problem fully. Beside the fact that science funding levels are generally quite low, why should the United States fully support basic research that will end up also helping China and Europe (and vice-versa)? Most silicon technology (including computers, displays, and solar cells) was invented in the United States, but nowadays are produced overwhelmingly in Taiwan,

Korea, and Japan. [Vanadium flow batteries were invented with U.S. taxpayer money but are now manufactured in China](#). The Far East has been free-riding on American inventions for much of the 20th century — just as America was free-riding on the [European inventions](#) of the 18th and 19th century.

One solution is [patent](#) protection. Yet global patents are remarkably expensive to obtain and enforce. (Except for medical inventions, it is unclear if they have benefitted anyone except the legal industry.)

► The Sharing Externality

This leads us to the second related problem. Let's presume that chamomile energy and patent protection were perfect and some company had invested its \$10 billion, solved the problem, and now owned the perfect chamomile energy patent. At this juncture, our firm would not want to share its technology (which would reduce the global pollution externality). Instead, the company would want to charge a few trillion dollars as a reward for its inventors and investors. After all, they had staked their fortunes and sometimes their livelihoods on exploring chamomile when no one else was chipping in. Why give it away? However, this also means that chamomile won't be used as widely as it should be to solve the planet's emission problems.

Bill Nordhaus calls the issues of insufficient research and insufficient sharing the *double externality* of clean energy innovation. It is a central problem in fighting climate change.

► Perspective

Lack of sharing of proprietary technology, often developed at great expense, and fear of expropriation has and will continue to hobble many otherwise planet-optimal solutions. China has a reputation for ruthless [expropriation and reverse engineering of foreign technology](#) in joint ventures in the past. Such behavior now keeps fearful Western companies from offering to replace China's planned coal plants with cheaper and better nuclear-power technology. This situation hurts not only the Chinese, but all of us when it comes to CO₂ emissions. But who can blame the nuclear engineering companies? Why should they commit economic suicide?

Like many others, we believe that accelerating innovation in clean technology ranks at the top of the best ways to stem global climate change. We

are excited about Bill Gates' *Breakthrough Energy Fund* (BEF), which invests with a part-philanthropic mandate to take the social benefits of inventions into account. (If we had enough money, we would have invested, too.) However, BEF also has a financial mandate. Its hardest decisions may yet be to come: If BEF succeeds in developing breakthrough inventions, will it share the know-how with parties that cannot afford to pay — or simply will not want to pay? Will BEF share its inventions with Chinese companies, which will then likely copy the technology and sell it cheaply in mass all over the world? It would be good for the planet but not for BEF.

In general, economists know of no universal solutions that solve the double externalities. The socially best partial solutions so far have typically involved individual countries funding universities to conduct the basic research with their public money. This is obviously less than ideal.

In real life, this approach has run into many difficulties. Universities and countries focus on their own self-interests, too. Universities do not even want to allow other universities to exploit their research breakthroughs. They want to patent their inventions and earn licensing fees, often by partnering with just one company, which can then earn monopolistic profits by restricting access.

Finding good economic solutions is a constant struggle. If the problem were easy to solve, we would not have needed to talk about it for so long.

► A Pollution Tax or Clean-Technology Innovation Subsidies?

Although our book advocates increasing subsidies for research and development in green technology, it would be incorrect to view our endorsement as a substitute for a pollution tax. Pollution taxes help redirect innovation to where it is more useful. By reducing the pollution (and thus the incurred tax), pollution taxes increase the incentives for clean-energy research.

7 Moving the Needle Now

We chose the subtitle of our book, *moving the needle*, to signal that the *marginal* costs of reducing emissions and taking action right now are already more than low enough to act on many fronts. Our book's goal is to help nudge the needle in the right direction. We hope that the world will invest in more clean research, learn as it goes, and as it learns from research, decide where it should then push next.

We understand and sympathize with the concerns of and the disagreements among scientists about how far humanity will eventually have to go over the rest of the century. But this long-term debate is not as important as picking the low-hanging fruit over the next decade.

A journey of a thousand miles begins with a single step. The world needs to start the journey sooner rather than later, and the first steps are exceptionally easy because they are cheap enough so as to be in the interest even of individual countries. We will offer some specific recommendations in Part III. But first, we must explain how integrated models of economics and climate change work.

Further Readings

BOOKS

- [Cowen, Tyler](#), 1991, [Public Goods and Market Failures](#), Routledge.
- [Dixit, Avinash K.](#) and [Barry J. Nalebuff](#), [Thinking Strategically: The Competitive Edge in Business, Politics, and Everyday Life](#), Norton.
- [Ostrom, Elinor](#), 1990, [Governing the Commons: The Evolution of Institutions for Collective Action](#), Cambridge University Press.
- [Ridley, Matthew](#), 2020, [How Innovation Works](#), Harper.
- [Sowell, Thomas](#), 2015, [Basic Economics](#), Basic Books, New York, NY. An intuitive book on the basic principles of economics.

REPORTS AND ACADEMIC ARTICLES

- [Nordhaus, William D.](#), 2019, [Climate change: The ultimate challenge for economics](#), American Economic Review, 109 (6): 1991-2014. Includes an explanation of why climate change is the “mother of all externalities.”
- Hsiang, Solomon and Robert E. Kopp, 2018, [An Economist’s Guide to Climate Change Science](#), Journal of Economic Perspectives.