

Perfect Markets and the “World of Truth”

You might not expect Jim Carrey films and economics to have much in common, but in fact there is much we can learn from the rubber-faced comedian. Consider the film, *Liar, Liar*, which tells the story of Fletcher Reede. As a result of his son’s birthday wish, Fletcher Reede finds that he is compelled to tell the truth for twenty-four hours. This is problematic for Fletcher because he is a lawyer—or a *liar*, as his son understands it—and hilarity predictably ensues as a horrified Fletcher incriminates himself by helplessly blurting out truthful answers to every question he is asked. They don’t make as much of a feel-good movie, but free markets are just like Fletcher Reede’s son—they force you to tell the truth. Yet while the results were humiliating for Jim Carrey’s character, we will discover that a world of truth leads to a perfectly efficient economy, one in which it is impossible to make someone better off without making someone else worse off.

In this chapter we’ll see what truth means in economic terms, how it leads to efficiency, and why efficiency is good. We’ll also explore efficiency’s shortcomings: how efficiency isn’t always fair, and why we have taxes. As we’ll see, taxes are like lies: they interfere with the world of truth. But I’ll reveal one way in which taxes can be implemented, which is both fair *and* efficient. This could be good news for seniors struggling to pay their winter heating bills, but bad news for Tiger Woods.

Imagine if you will that Fletcher's son gets his birthday wish, not just for his smooth-talking dad but for the whole world. So, let's buy a cappuccino in the world of truth. Before frothing up the half-and-half for you, the barista looks you up and down and asks:

"What's the most you're willing to pay for this coffee?"

You'd like to lie and pretend that you don't really want it, but the truth just slips out:

"I'm in caffeine withdrawal. Fifteen bucks."

With a smirk, the barista prepares to ring up the extortionate sum, but you have a few questions of your own:

"How much did those coffee beans cost?"

"How much did you pay for the plastic lid and the cup?"

"How much does it cost to raise a cow, and how much milk can you get from one?"

"How much did the electricity cost for the refrigeration, heating, and light in here?"

Now it is the barista's turn to have a Fletcher Reede moment. No matter how she tries to evade the questions or froth up the cost of the cappuccino, she cannot tell a lie. It turns out that the cappuccino costs not fifteen dollars, but less than one. The barista tries to haggle, but you have one more killer question:

"Are any other places within thirty yards selling coffee like this?"

"Yes . . ." she moans, her head thudding to the counter in a gesture of abject defeat.

You walk out of the shop with the coffee safely in your possession for the bargain price of ninety-two cents.

Prices are optional, which means they reveal information

There’s a basic truth incorporated into any system of prices. That truth comes from the fact that stores and consumers do not have to buy or sell at a given price—they can always opt out. If you’d been willing to pay only fifty cents for the coffee, nobody could have forced you to raise your offer or forced the barista to drop the price. The sale simply would not have occurred.

Of course, you sometimes hear people complaining that if they want something—say, an apartment on Central Park West—then they have to pay the exorbitant asking price. That’s true, but although prices sometimes seem unfairly high, you hardly ever *have* to pay them. You could always use your money to buy an apartment in Harlem or a house in Newark or a million cups of coffee instead.

In a free market, people don’t buy things that are worth less to them than the asking price. And people don’t sell things that are worth more to them than the asking price (or if they do, it’s never for long; firms that routinely sell cups of coffee for half of what they cost to produce will go out of business pretty quickly). The reason is simple: nobody is forcing them to, which means that most transactions that happen in a free market improve efficiency, because they make both parties better off—or at least not worse off—and don’t harm anyone else.

Now you can begin to see why I say that prices “tell the truth” and reveal information. In a free market, all the buyers of coffee would prefer to have coffee than the money the coffee cost, which is shorthand for saying they prefer coffee to whatever else they might have spent ninety-two cents on. That is, the value of the product to the customer is equal to or higher than the price; and the cost to the producer equal to or lower than the price. Painfully obvious, perhaps, but the implications turn out to be dramatic.

It may seem trivial to say that in a free market we know customers value coffee more than the money they pay for it. Yet it’s

not quite as trivial as it looks. For a start, this “trivial” piece of information is already more than we can say about anything that is paid for outside the market—for example, Washington DC’s hugely controversial new baseball stadium. The Montreal Expos baseball team agreed to move to DC on the condition that the DC government subsidize the cost of a new stadium. Some say the subsidy will be \$70 million, others that it will be far higher. Maybe this is a good idea, and maybe not. It’s not clear how we decide whether this is a good way of spending taxpayers’ money.

When decisions are made inside a market system there’s no such controversy. If I decide to pay \$70 for a ticket to see a baseball game, nobody questions whether it’s worth it; I made my choice, so obviously I thought so. This free choice produces information about my priorities and preferences, and when millions of us make choices, market prices aggregate the priorities and preferences of us all.

Perfect markets: *The truth, the whole truth, and nothing but the truth*

So the trivial piece of information that in a free market customers value cappuccinos more than the money they pay for them is not so trivial after all. But we needn’t stop there.

Imagine now that the coffee market is not only free but extremely competitive, that entrepreneurs are always starting new firms with fresh ideas and entering the market in an attempt to undercut the incumbent companies. (Profits in a competitive industry are high enough only to pay workers and persuade entrepreneurs that their money isn’t better off in a savings account—no higher.) The competition will force the price of coffee down to the “marginal cost”—the cost the coffee bar incurs when making one more cappuccino, which we may remember is just under a dollar. In a perfectly competitive market, the price of coffee would equal the marginal cost of coffee. If the price were lower, firms would go out of business until it rose. If the price were higher, new firms would enter or old firms would expand their output

until it fell. Suddenly, the price is not conveying a vague fact (“this coffee is worth ninety-two cents, or more, to the buyer, and it cost the coffee bar ninety-two cents, or less”) but a precise truth (“this coffee cost the coffee bar exactly ninety-two cents”).

What if other industries were also perfectly competitive? That would mean that for every product, the price equaled the marginal cost. Every product would be linked to every other product through an ultracomplex network of prices, so when something changes somewhere in the economy (there’s a frost in Brazil, or a craze for iPods in the US) everything else would change—maybe imperceptibly, maybe a lot—to adjust. A frost in Brazil, for example, would damage the coffee crop and reduce the worldwide supply of coffee; this would increase the price coffee roasters have to pay to a level that discourages enough coffee drinking to offset the shortfall. Demand for alternative products, like tea, would rise a little, encouraging higher tea prices and extra supply of tea. Demand for complementary products like coffee creamer would fall a little. In Kenya, coffee farmers would enjoy bumper profits and would invest the money in improvements like aluminum roofing for their houses; the price of aluminum would rise and so some farmers would decide to wait before buying. That means demand for bank accounts and safety deposit boxes would rise, although for unfortunate farmers in Brazil with their failed crops, the opposite may be happening. The free-market supercomputer processes the truth about demands and about costs, and gives people the incentive to respond in astonishingly intricate ways.

That may seem like a ridiculous hypothetical scenario. But economists can measure and have measured some of these effects: when frosts hit Brazil, world coffee prices do indeed rise, Kenyan farmers do buy aluminum roofing, the price of roofing does rise, and the farmers do, in fact, time their investment so that they don’t pay too much. Even if markets are not perfect, they can convey tremendously complex information.

Governments—or any organizations—find it hard to respond to such complex information. In Tanzania, coffee is not produced in a free market, and the government, rather than the farmers,

receives any windfalls from high coffee prices. Historically, the government has failed to spend the money sensibly, blowing too much on unsustainable salary rises for civil servants, and failing to realize that the price spike was temporary.

To appreciate why markets do such a good job of processing complex information, first think about the customer. We know that he won't buy a cappuccino unless he values it more than anything else he could buy with the same money. But what else could he buy with the same money? In our world of truth, he could buy *anything* that costs the same as, or less than, a cappuccino. If he chooses the coffee he's saying that of all the things in the world that cost the same as coffee, he would like coffee to be made.

Elsewhere, of course, there are other people spending their money not on coffee but on movie tickets, bus fares, or underwear; and there are others choosing not to spend their money at all and to put it in the bank instead. All of these competing demands pull producers to respond. If people want computers, then manufacturers will build factories, hire workers, and buy plastics and metals, which will be diverted from other uses to go into computers. If people want coffee instead of underwear, then more land will be devoted to coffee and less to other uses, like parks or housing or tobacco farming. Lingerie shops will be replaced by coffee shops. Of course, start-up companies will borrow money from banks, and interest rates will rise or fall, depending on the balance between the number of people wanting to save and the number of people wanting to borrow. Interest rates are just another price: the price of spending today instead of next year. (You might have thought that interest rates were set by central bankers like Alan Greenspan at the Federal Reserve or Mervyn King at the Bank of England. Actually, Greenspan and King chair committees that set “nominal” interest rates. True interest rates are interest rates after inflation—set by the market in response to the central bankers.)

The changes don't stop there. The ripples in the price system continue outward. They whip through some parts of the economy

at tremendous speed and cause slow but powerful seismic shifts in others, like education or technology. For example, if there aren't enough trained workers to produce computers, manufacturers like Dell and Compaq will have to train them, or raise wages to poach them from other manufacturers like Apple and Gateway. As the wages for skilled workers rise, people will see that it's worth taking time off and paying to go to college. Manufacturers' interest in producing cheaper or better computers will give a boost to research labs and engineering schools. Higher demand for plastics will raise the price of the raw material—crude oil—which will in turn encourage those who use oil for energy to switch to cheaper substitute fuels or to invest in energy-saving technology. And so it continues. Some of these effects will be tiny. Others will be enormous. Some will have an instant effect. Others will not be realized for decades. But in the world of truth—the world of perfect markets—all of them will have an impact.

What is the result of a set of perfectly competitive markets interconnected like this?

Companies are making things the right way. Any company that wastes resources, over-produces, or uses the wrong technology, will go out of business. Every product is produced in the most efficient way.

Companies are making the right things. The price of a product equals the cost to make it. The price also reflects the terms at which customers can trade off one priority against another. (Two cups of coffee cost the same as one Danish; which would you prefer?) The price is a direct line of communication from what products cost to what customers prefer, and back again.

Things are being made in the right proportions. If too much coffee were being produced, manufacturers would cut prices; and if too little, prices would rise. Either way, the situation would correct itself. In the competitive market, price equals cost; there is no incentive for anyone to produce less (giving up profitable

sales) or to produce more (creating products that cost more than anyone is willing to pay). The competitive rule—price equals cost equals value to the consumer—keeps things efficient.

Things are going to the “right” people. The only people who buy products are the people who are willing to pay the appropriate price. Let’s say I confiscate a cappuccino from Axel and give it to Bob. In the world of truth, this is wasteful. Axel was willing to pay for coffee, and Bob was not, which means Axel values coffee more than Bob, and my confiscation is inefficient. Notice that here I am equating “right” with “efficient,” an assumption we’ll examine and challenge shortly.

So: if the right things are being made right in the right quantities and going to the people who value them most, there is no room for any gains in efficiency. To put it another way, *you can’t get more efficient than a perfectly competitive market*. And it all follows perfectly naturally from the truth contained in the price system: prices are true representations of cost to firms, and also true representations of value to customers.

Life without markets

Because Western society relies heavily on free markets, we find it difficult to imagine what it would be like if we didn’t, or to take a step back and see quite how profound the effect of the market is. Yet any modern democracy provides goods outside the market system, and looking at the way such goods are provided gives us a hint of the strengths and weaknesses of markets. Think of your friendly local police force, which is paid for by a nonmarket system of taxation. The nonmarket system has some advantages—for one thing, when you dial 911 nobody asks for your credit card details. The government is supposed to afford the same level of protection to the rich and poor, although it does not always seem that way.

But the nonmarket system also has some disadvantages. For instance, if a police officer is rude or incompetent, you don’t have

the option to shop for a different police force. If you think that the level of police protection you receive is excessive, it's not up to you to cut back a bit. Neither can you spend more if you decide that you'd like extra service. No, you have to lobby your local politicians and hope they consider your demands.

Government-provided schooling is another example of a nonmarket service that many of us use. In both Britain and the United States, most people send their children to government-funded schools. But those schools are different from each other—different atmospheres, different academic emphases. Most importantly, some are good schools, and some are not. The market solution for schools is similar to the market solution for food: the best food goes to the people who are willing—which also implies able—to pay most for it. But within the government sector there are no prices. What happens instead? Parents line up, haggle, and protest. They move to districts with better schools. In Britain, government-run religious schools often have the best academic records, so atheists take their children to church every Sunday in order to get good references from priests and get their children into these schools.

As with the police, the nonmarket system has the cozy advantage of concealing the fact that the poor don't get the same quality of education that the rich do. But again, the nonmarket system suffers from a serious problem: the truth about values, costs, and benefits has disappeared. It is impossible to tell which parents enroll their children in church schools for religious reasons and which parents are just looking for better results. It is also impossible to know how much parents would be willing to pay for more teachers and better materials. In a market system the truth would emerge about how much it costs to provide good schools, and who would be willing to pay for them. The nonmarket system struggles with these basic questions.

It seems that there is a willingness to pay for good schools, and we see it emerge because house prices are higher in the areas of schools with the best reputation. The nonmarket system, which gives preference to local children, channels the

money that parents are willing to pay for a good school into the hands of property owners near existing good schools. This hardly seems sensible. A market system would simply direct the money to pay for more good schools.

The signaling function of prices

Prices perform two functions, not just one. In a market system, prices provide a way of deciding who gets to enjoy a limited supply of schools: whoever pays most gets to send their children to the best schools, an uncomfortable state of affairs, which the government-school system is designed to prevent. But prices also give the signal to build more schools, hire more teachers or raise their wages if they're in short supply, and buy better materials. In the longer term, a price system will transform a high willingness to pay for good schools into a lot of good schools, just as surely as it will transform a high demand for coffee into a lot of cappuccino.

Don't politicians know that we value good schools already? Should they be making government money available? The difficulty is that politicians hear that we want good schools, but they also hear that we want more police on the streets, a better health service, lots of big roads, excellent welfare benefits, low taxes, and a double-shot caramel Venti latte. It's easy for us to demand all of these things, but prices, by forcing us to put money where our mouths are, uncover the truth. Taxes have their advantages, but many don't contribute to truth because we cannot choose whether or not to pay them, depending on whether each penny is spent according to our wishes. Because prices are optional, they reveal information.

None of this amounts to a knockdown argument against providing a police service or a school system with a nonmarket process. Nonmarket systems have their advantages, but they also lose something important: information, information about wants, needs, and desires, and about inconveniences and costs. Sometimes the loss of information is worthwhile because it is offset by gains in equality or stability. But sometimes the loss of informa-

tion can leave an economy, and a society, floundering in waste and confusion. We think that the value we get from schools and police are more than what they cost us in taxes, but we don’t know for sure. Not so with the cappuccino.

**Efficiency versus fairness:
Can we handle the truth?**

A perfectly competitive market is like a giant supercomputer network. With amazing processing power and sensors in every part of the economy—reaching even inside our brains to read our desires—the market is constantly reoptimizing production and allocating the results perfectly. Remember that when economists say the economy is inefficient, they mean that there’s a way to make somebody better off without harming anybody else. While the perfectly competitive market is perfectly efficient, efficiency is not enough to ensure a fair society, or even a society in which we would want to live. After all, it is efficient if Bill Gates has all the money and everybody else starves to death . . . because there is no way to make anybody better off without making Bill Gates worse off. We need something more than efficiency.

So it’s hardly surprising we sometimes prefer the cozy white lies: it is expensive, for example, to heat the house of an elderly lady in Minnesota, but we may prefer to subsidize the fuel, not wanting her to face the truth of that expense.

Even more than subsidies, taxes are a common cause of inefficiency: the government taxes market transactions and spends the money on, we hope, good things like police forces and schools. Why are taxes inefficient? Because they destroy the information carried by prices in perfectly competitive, efficient markets: price no longer equals cost, so cost no longer equals value. For example, a sales tax of 10 percent creates a “lie” in the following circumstances:

Cost of cappuccino: ninety cents

- Price of cappuccino in perfectly competitive market: ninety cents

- Price of cappuccino after tax: ninety-nine cents
- Willingness to pay for cappuccino: ninety-five cents
- Cappuccino sold: **none**
- Tax raised: **zero**

There was a sale that could have created five cents of efficiency gains (cappuccino cost ninety cents but was valued at ninety-five cents) but which never happened because of the tax. What's worse, the tax wasn't even paid. If the government were able to waive the tax in such circumstances, they would be no worse off, but the coffee buyer would be better off: a clear efficiency gain.

It's hard for tax officials to know when to charge the tax (situations where taxes will not change buyers' behavior) and when to waive the tax (because potential buyers would have avoided it anyway, by not buying coffee). But they try to do so using the kind of price-targeting strategies outlined in chapter 2. Taxes are often higher when price-sensitivity is low. For example, the government charges high taxes on gasoline and cigarettes, not for environmental and health reasons but because people who buy these products need to drive and are addicted to smoking; they won't change their behavior much even in the face of large taxes.

We are faced with a dilemma. We want to avoid inefficiency, because that would leave us passing up an opportunity to make somebody better off at no cost to anyone else. But taxes cause inefficiency, and most of us think we need taxes to redistribute income (to a greater or lesser extent) from the rich to the poor. We seem to be facing two contradictory imperatives: avoid the needless waste that is "inefficiency," but make sure that wealth is at least somewhat evenly spread. What we need is a way to make our economies both efficient and fair.

Can we enlist markets to help with fairness?

Is it true that we have to choose between the efficiency of perfect markets and the fairness of benevolent government intervention? This seemed to be the conclusion of governments throughout

the free world after the experience of the Great Depression and World War II. President Roosevelt’s “New Deal” programs of the 1930s expanded the role of the United States government, in response to the Great Depression. In Britain, Clement Atlee’s postwar government took control of much of the health, steel, air travel, petroleum, rail travel, and telephone industries. Government-owned businesses took over partly because in the deprived, exhausted yet hopeful years after the war, economists had some confidence in the experts who had masterminded the war effort and thought they might not do a bad job of organizing the economy efficiently. Few people foresaw the later collapse of government-run economies, whether vast like the Soviet Union and China, or small like Tanzania or North Korea. But even if they had believed that private markets were more efficient, this was neither here nor there in the 1940s: the postwar Labour government in Britain would have been content to live with some inefficiency if it meant a fairer society.

But the old dilemma between efficiency and fairness was about to be shattered by a young New Yorker called Kenneth Arrow, who knew all about unfairness after watching helplessly as a teenager while his father lost his successful business and all his savings in the Great Depression. The desire for social justice stayed with Arrow, but intellectually he couldn’t just ignore the question of efficiency. The young economist set his logical mind to wrestling with the tension between the unerring efficiency of the free market and the imperative that some kind of fairness should prevail. His solution was brilliant, twisting the traditional thinking about competitive markets and efficiency on its head. He proved that not only are all perfect markets efficient, *all efficient outcomes can be achieved using a competitive market, by adjusting the starting position.* Arrow went on to win every plaudit available to an economist, and he remains the youngest man ever to win the Nobel Prize for Economics. But why was his insight so important?

I call it the “head start theorem.” Instead of focusing on the enormous complexity of a real economy, think of a very simple one-dimensional human challenge: the 100-meter sprint. The fastest

sprinter will win the race. If you wanted all the sprinters to cross the line together, you could just change the rules of the race, ordering the fast runners to slow down and everyone to hold hands as they crossed the line. A waste of talent. Or you could move some starting blocks forward and some starting blocks back, so that although each sprinter was running as fast as he could, obeying the usual rules and objectives of sprinting, the fastest had to cover enough extra ground that he would end up breaking the tape neck-and-neck with the slowest.

Arrow demonstrated that the same approach could work when trying to balance the excesses of competitive markets: instead of interfering with the markets themselves, the trick is to adjust the starting blocks by making lump-sum payments and levying one-time taxes.

An example of a lump-sum tax would be the government taxing everybody eight hundred dollars; or alternatively, taxing everyone over the age of sixty-five eight hundred dollars; or alternatively, taxing everybody whose surname on the birth certificate starts with *H* eight hundred dollars. The point is that unlike an income tax or a sales tax on coffee, a lump-sum tax doesn't affect anybody's behavior, because there is nothing you can do to avoid it. So unlike sales tax, it doesn't lead to an efficiency loss. Similarly, an example of lump-sum redistribution would be to *give* eight hundred dollars to everybody whose name starts with *H*, a policy for which I would be happy to vote.

In the 100-meter sprint, lump-sum taxation is like moving the starting blocks back a few paces. Income tax and sales tax are like asking the best runners to run backwards. Both would have the effect of ensuring a more equal finish, but moving the starting blocks around doesn't slow anybody down.

In the context of a sprint, it's fairly obvious that one of the ways to get a close result is to give the slower runners a head start. In the context of an economy, with literally billions of different goods, desires, raw materials, and talents, the head start theorem is a much bolder claim. But it's true: you can allow the competitive economy to use every skill and every raw material,

take advantage of every opportunity to trade, cooperate, educate, or invest . . . but still get a fair outcome by moving around the starting blocks and letting perfect markets do the rest.

The implication is that in a world of perfect markets, the only thing needed to ensure both fairness and efficiency is a “head start” strategy: a program of appropriate lump-sum taxes and subsidies that puts everyone on equal footing. The perfect markets then find every possible opportunity to make everybody better off from their revised starting points. The question is, can this be done in practice?

Impractical examples

Let’s take an example. American political philosopher Robert Nozick deployed a famous argument against taking a view of “justice as fairness.” In other words, he disputed the notion that one particular allocation of wealth could be deemed the “best” or “fair” allocation. Nozick’s argument invokes Wilt Chamberlain, a basketball star famous in the 1960s and ’70s, when Nozick was writing. Chamberlain’s talents made him wealthy; Nozick felt this was “just” because Chamberlain’s wealth was the outcome of legitimate decisions by fans happy to pay to see him play. The situation may have been “just” in Nozick’s sense of the word, but can any situation that leads to a highly unequal distribution of cash be considered “fair”?

Perhaps taxing Chamberlain’s income heavily would make the situation fairer, but Nozick warns that if Chamberlain did not really enjoy playing basketball *and* he was loaded down with heavy taxes, he might stop playing altogether. So although this situation might seem more “fair,” there would be neither the tax revenue, nor the basketball game: the problem of the cappuccino sales tax all over again. So how is it reasonable to call a distribution of income “fair” when everybody concerned, both fans and player, would prefer the “unfair” outcome?

Thanks to Kenneth Arrow, we now know that, when faced with a modern-day sports star like Tiger Woods, the solution is

to levy a one-time lump-sum tax of several million dollars on him. He would still have the incentive to earn money by playing golf, since he could not avoid the tax by playing less, as he would have to do in order to avoid a heavy income tax. He would no doubt earn enough to pay off the tax bill and still retain enough to buy a family car and a nice house somewhere unassuming. In this scenario, there is no waste or inefficiency, but the result is “fair” in that it produces a much more even allocation of wealth.

The only trouble with this plan is that it’s wildly impractical. The problem is not that it’s impossible to have taxes that only apply to one individual: President Franklin Roosevelt introduced an income tax rate of 79 percent, but the threshold was so high that the tax was paid by only John D. Rockefeller. Rather, the problem is more that a true lump-sum tax isn’t supposed to change behavior at all. Ideally it would have been decided before Tiger Woods was born, because if he could have predicted that he would be liable for a tax as a result of his success he might have chosen a different profession.

This is, of course, quite impossible. But we shouldn’t abandon the head start theorem quite yet. While we can’t always use lump-sum taxation and redistribution, we can sometimes: and when we can, it’s worth considering because it preserves the efficiency and the truth of the competitive market while adding a welcome dose of fairness.

A practical example

A more practical application of the head start theorem could be used to prevent elderly people from getting cold in winter, without damaging the environment. In a typical winter in Britain twenty-five thousand seniors die as a result of inadequate heating. To address this concern, domestic fuel is subject to lower taxes than many other things. But that’s a slightly odd way to deal with the problem—an equivalent to the “running backwards” solution. If governments need to raise tax revenue—and all of

them do, it seems—then a first approximation of an efficient strategy would be to have the same sales tax on everything, because that wouldn’t distort people’s buying decisions too much. A more refined view would recall the “price-targeting” of chapter 2. Because customers cannot easily cut down on fuel consumption, they are not very sensitive to the price of domestic fuel, hence the government should levy a bit more tax on fuel and a bit less on other goods: customers would not change their behavior much and so the inefficiency would be small. An even more sophisticated view (perhaps acquired from a peek ahead at chapter 4) would note that domestic fuel is a nonrenewable resource and using it causes pollution, so the case for higher tax on domestic fuel becomes even stronger.

The case for lower taxes on domestic fuel and higher taxes on other goods is hard to understand, until we start to worry about the elderly shivering in front of a lifeless gas or oil furnace that they cannot afford to switch on. Is this just one of those hard choices that governments sometimes have to make? Not necessarily. Instead of levying the wrong rate of tax on everyone else, better to choose a more sensible rate but give the elderly a head start—because of their poverty and because, being frail, they have an additional need for heating. The simple policy remedy is to raise fuel tax but give extra money to the elderly, money that they could use to switch that furnace on and stay warm.

We know from the head start theorem that given the money, pensioners will find their way to the efficient outcome—which, incidentally, may not involve more fuel being burned. Not every pensioner feels cold, and those who do may find better solutions to the problem. Some may use the money to move to Florida. Some may insulate their homes. Those who did not feel the cold in the first place can spend the money on other things. Nobody will burn extra fuel unless they need to, and if they need to they’ll have the money to meet that need.

The lesson of the head start theorem is that when a problem arises, it’s worth asking whether the problem can be addressed

by rearranging the starting blocks rather than interfering with the race. This strategy isn't always practical, but because free markets are efficient, it's worth trying to harness that efficiency to meet other goals.

Throughout this chapter, we've been on a flight of fantasy no more plausible than the story of Fletcher Reed. The "world of truth" is a world where markets are complete, free, and competitive. In reality we're about as likely to achieve a world with complete, free, and competitive markets as hotshot lawyers are to start telling the truth to everyone.

You might therefore be asking yourself why you've read a chapter, even a brief one, about some bizarre economists' fantasy. The answer is that the fantasy helps us understand why economic problems arise and also helps us to move in the right direction. We know that a world of perfect markets combined with the head start approach is as good as we're going to get. When real world economies malfunction, we know to look for the market failures—and to do our best to patch them up.

We've already explored one of those failures: some companies have scarcity power and can set prices that are far above their true cost, which is where they would be in a competitive market. This is why economists believe there's an important difference between being in favor of markets and being in favor of business, especially particular businesses. A politician who is in favor of markets believes in the importance of competition and wants to prevent businesses from getting too much scarcity power. A politician who's too influenced by corporate lobbyists will do exactly the reverse.

Whether abetted by politicians or otherwise, companies with scarcity power are one market failure. There are two others. In the next two chapters, we'll encounter them, leaving the curious world of truth behind us and facing up to the real world once again.

Crosstown Traffic

We've just learned that in the world of perfect markets, everything is for the best. We know that perfect markets are completely efficient, delivering outcomes that are flawless in every respect except distribution. We also know from the head start theorem that we can fix any complaints about distribution in advance. Presto, every problem solved, or at least every problem concerning the allocation of goods and services.

That's nice to hear, but then why did I spend two hours stuck in traffic on the way to work this morning? The bumper-to-bumper traffic was a stupid waste. All of us could have been riding buses, or carpooling, and we would have reached our destinations in downtown DC in fifteen minutes. Where is the perfect market there? The obvious answer is that, of course, there is no market perfect or otherwise for driving around on the streets. What may be less obvious is that there could be.

Economies that work smoothly because they are full of perfect markets are neither interesting nor realistic. But because perfect markets provide such a clear benchmark, economists find it much easier to start from them and work out what is going wrong, rather than start from scratch and work out what is going right. And this method of thinking about the world will lead us to the cure for crosstown traffic.

What's wrong with my world

I am a happy man, but there are things in my life that infuriate me and that I wish could be different. I wish I didn't have to upgrade my computer software every couple of years at great expense. I wish that I could rely on my doctors to give me appropriate medical treatment when I am ill. I wish Washington's streets were not clogged with traffic and filled with pollution.

These three personal, if common, grumbles correspond to the three key ways in which markets fail to live up to chapter 3's lofty ideals of perfection. Markets fail to work well in the face of scarcity power, as we saw in chapter 2. That is one of the problems with buying computer software—the market is dominated by a single company, Microsoft, which has tremendous power to set high prices. Markets also fail to work well if some decision makers lack information. When I go to my doctor I have no idea if he is giving me good treatment, while he has no need to take into account the cost of the treatment, and my insurance company has every incentive to refuse to pay, without knowing the true situation. (We'll deal with health care in chapter 5.) Finally, markets fail to work well if some people make decisions that affect bystanders: when a driver buys gasoline from a gas station, that is all very well for the driver and the gas station but not for the bystanders, including other drivers, who have to breathe the resulting carbon monoxide.

These three big problems are called “market failures”: scarcity power, which we discussed in chapters 1 and 2; missing information, which we will discuss in chapter 5; and the subject of this chapter, decisions that have side effects on bystanders. Economists call the side effect an “externality” because it lies outside the original decision, for instance, the decision to buy gasoline. Whether because of scarcity power, incomplete information, or an externality, when the economy fails to live up to the idealized “world of truth,” trouble is in store.

How drivers affect bystanders

Washington DC, London, Tokyo, Atlanta, Los Angeles, and Bangkok, and indeed any of the world's great cities, are full of cars, buses, and trucks. Those vehicles seriously damage the happiness of innocent bystanders. They cause severe air pollution. Admittedly, London's current air pollution is not as severe as the "Great Stink" of the 1850s, in which tens of thousands died of cholera. But still, air pollution from traffic is not trivial: many thousands of people die because other people want to drive. Around seven thousand people a year die prematurely because of traffic pollution in Britain, a little more than one in ten thousand. In the United States, the Environmental Protection Agency estimates that fifteen thousand people die prematurely because of the particulate matter produced from sources such as diesel engines. Within urban areas like London, the cost of delays from congestion are even worse, if you consider the number of hours spent sitting in traffic as being in any way a significant loss of productive or enjoyable life. Then there is the noise, the accidents and the "barrier effect," which discourages people, and particularly children, from walking to school, the local stores, or even to meet their neighbors across the street.

People are not fools: it's almost certainly true that anyone taking a trip in a car is benefiting from driving. But they are doing so at the expense of everyone else around them—the other drivers stuck in traffic, the parents who dare not let their children walk to school, the pedestrians who risk their lives dashing across the street because they are tired of waiting for the light to change, the office workers who even in the sweltering summer cannot open their windows because of the roar of the traffic.

Because each driver who gets into his car is creating misery for other people, the free market cannot deliver a solution to the problem of traffic. The external effects of congestion and pollution are important departures from the "world of truth." In the "world of truth" every act of selfish behavior is turned to the common good. I selfishly buy underwear because I want it, but in

doing so channel resources into the hands of underwear manufacturers, and do nobody any harm. Textile workers in China, where the underwear is made, selfishly look for the best job, while manufacturers selfishly look for the most capable employees. All of this works to everyone's benefit: goods are manufactured only if people want them, and they are manufactured only by the most appropriate people to do the job. Self-centered motives are put to work for everybody.

Drivers are in a different situation. They do not offer compensation for the cost they inflict on other people. When I buy underwear, the money I spend is compensation for all of the costs incurred in making it and selling it to me. When I take the car for a drive then I do not even need to think about the costs incurred by the rest of society as I avail myself of the free roads.

Different kinds of prices: Marginal and average

It is not quite fair to say that drivers can use the roads for nothing. In the United Kingdom, it's not legal to drive a car, or even to park it on public streets, unless you have paid a sizable annual tax called "Vehicle Excise Duty." Many states in the United States have a similar tax. Gas and diesel fuel are also taxed heavily enough to cause great resentment. In the autumn of 2000, for example, a series of protests against high fuel prices prevented fuel reaching the country's gas stations, and brought Britain to a standstill. In Britain, drivers pay £20 billion in taxes on cars and fuel every year; in America, the figure is around \$100 billion. To ask "have they paid enough?" is to ask the wrong question. The right question is, "are they paying for the right things?" The answer is no.

There are two different concepts of price at play here, and the distinction matters. The average price that a driver pays for a journey across a city is quite high if the driver is paying an annual license fee. But the price that the driver pays for one extra trip across the city is low: the trip doesn't burn much fuel and drivers are not charged for extra trips. Once you've paid for the right to

take the car on the street in the first place, you don't get a discount for low mileage: you might as well drive and drive, because it won't put a penny on your tax bill. That is the difference between the average price and the marginal price, which is the price for one extra trip.

To understand why the difference matters, let's turn to alcohol. When I was in college, clubs and societies used to have big parties where some people didn't drink at all and, less surprisingly, most people drank far too much. This was because there were two types of ticket. "Alcoholic" tickets allowed unlimited boozing after payment of an up-front fee of, say, ten pounds (at that time, about fifteen dollars). The other type of ticket was a lot cheaper, and you had to drink rancid orange juice instead and stand in a corner while the drinkers got more and more obnoxious. Turning up and having a couple of beers was a pretty expensive proposition, so most people preferred either to maximize the value of the unlimited drinking opportunity, or opt out of drinking alcohol completely. Of course, the result was chaos, although some people felt it made for pretty good parties.

Since the university felt that the drunkenness represented a problem, they considered dealing with it at the next party by raising the up-front fee to, say, twenty pounds (about thirty dollars). But the likelihood would be that while a few people would switch to being disgruntled orange-juice drinkers or give up on the society altogether, most of the drinkers would decide there wasn't much point in a party without drinking. Grumbling, they would empty the contents of their wallets. Later in the evening, many of them would empty the contents of their stomachs.

The university misunderstood the problem. They understood that people were drinking too much and correctly thought that the solution probably involved raising the price of drinking. The problem is that there are different ways of describing the price of drinking. There's the price of being a drinker: ten pounds. There's the average price of a drink: for the typical student who has twenty

drinks, this is fifty pence. Then there's the *marginal* price of a drink, which is zero. Once you've paid the up-front fee, you might as well keep drinking.

Question: if you were running the university, would you deal with the problem by: (a) raising the up-front fee for drinking?, (b) buying better orange juice?, or (c) scrapping the up-front fee and charging people for what they drank?

Better orange juice might be nice, but the Undercover Economist would humbly suggest that the solution to the underlying problem is *c*.

Now, back to traffic congestion. If you were advising the secretary of transportation you might suggest an analogy with student parties. Currently, potential drivers have two options: they can cough up a large up-front fee and drive as much as they like; or they can not drive at all. This second option, the "orange-juice" option, requires them to bike, use public transportation, or walk—although as with the student party, the more people who choose the first option, the less attractive the second option becomes.

You might even propose some policy options: (a) raise the up-front fee for driving; (b) supply better "orange juice" (more buses, better trains, cycle routes, pedestrian crossings); or (c) scrap the up-front fee and charge people for the trips they drive.

All of these options could be expected to reduce traffic congestion to some extent, perhaps to an important extent. But it is option (c) that attempts to deal with the cause of the problem. Drivers do not live in the "world of truth"; that is they do not pay the true cost of their actions, including the "externalities" or side effects that affect bystanders. Option (c) tries to make them pay that cost; we might call it an "externality charge."

Currently, every potential driver is being offered the same kind of deal as partygoing students: put up a wad of cash in exchange for an unlimited binge, or pay nothing and receive nothing in return. There are no half measures.

Student parties were not livened up excessively by the fact that drinks worked out at fifty pence (less than one dollar) on average: they were livened up excessively by the fact that the next drink was always free. Similarly, congestion is not caused by the fact that the tax on a car trip is fifty cents on average: it is caused by the fact that the next trip is always free.

We must not get obsessed with the question of how much drivers pay on average. Certainly, how much tax any type of person pays on average is an important question of distribution. While distribution is important, however, it doesn't have a big impact on whether our streets are clogged up and our cities are polluted.

What matters much more for congestion is the price drivers pay at the margin; or, to put it another way, the price drivers pay to make one extra trip. Cars don't cause much pollution or congestion, after all: car *trips* are the problem. Universities would encourage appropriate levels of drinking by charging students per drink. Similarly, the Department of Transportation would encourage appropriate levels of driving by making drivers pay for each trip.

Pricing should reflect the damage

I've been oversimplifying, as usual. In most European countries, drivers do pay a tax per mile in the form of a high tax on fuel. But the tax on fuel doesn't closely match the costs that drivers inflict on each other and on nondrivers. People in rural areas pay the taxes (typically they spend between a quarter and a third as much again on gasoline as those in urban areas), but it is the commuters in the London, New York, or Paris rush hours who are causing the most serious congestion, severe air pollution, and noise. The same trips made in the small hours of the morning do not cause congestion, although pollution and noise are still a problem. Make a similar trip between two houses in Alaska and you do not cause congestion. The noise is likely to be heard only by the occasional stray caribou. The damage caused by pollutants is much reduced, because many of them will disperse harmlessly. If

the idea of a charge on driving is that each driver faces the costs of his actions, the rush-hour New York driver should pay more because he is causing more harm to others. Whatever level of externality charge turns out to be appropriate, if it is to reflect the truth, it should vary according to time and place.

The idea of an externality charge is not to discourage everyone from doing anything that might inconvenience anyone else; it is to get them to take into account the inconvenience they cause to others. To take an extreme example: if I go walking in Virginia's Blue Ridge mountains, it is nice to be able to take in the natural beauty of the place in relative solitude, and so it's mildly annoying to find the trails cluttered with other people. They may be inconveniencing me, but it would not be efficient to forbid their trip because it gives them so much pleasure and me so little trouble.

Externality charging needs to strike the right balance between pleasure and trouble; it must reflect the cost of the externality . . . but no more. We should aim to make ours a world where people feel free to do things they enjoy, even if others are mildly inconvenienced, but also one where we all refrain from harming other people if the effort involved to avoid harming them is small. We discovered in chapter 3 that perfect markets deliver this world, at least within the sphere that markets operate. Perfect markets cannot make us smile at passersby or love our families, but they can make sure that we get a cappuccino if and only if we are willing to pay more than the true cost—which includes the cost in time and trouble of the baristas, the bean pickers, the entrepreneurs, the machine manufacturers, and the rest. In other words, perfect markets allow us to feel free to do things that we enjoy only if our enjoyment outweighs the trouble caused to make it all possible.

This is why economists are fairly relaxed when markets seem to be working well. But we are also vigilant for the many market failures. So how do we make sure that when deciding whether to drive across town, I can be sure that the benefit to me outweighs the cost to everybody else? There's no need to worry about costs and benefits that are part of an efficient market transaction. So, if

oil refining and gasoline retailing are perfect markets (contrary to popular belief, they are not far off), then the trouble it took to refine and distribute the gasoline is fully represented in the price. I will not buy gasoline unless the benefit I get from it is greater than the trouble it took to refine and distribute.

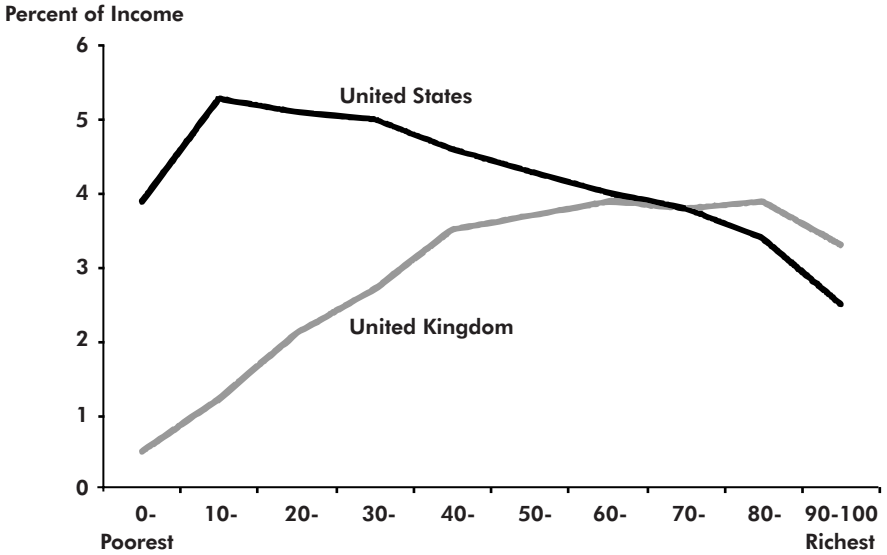
Instead, we should worry about costs and benefits missing from the market transaction. The pollution from the gasoline causes local poisoning and global warming, and the majority of the pollution damage when I burn a tank of gasoline is not caused to me or to the oil company. The trick is to mimic perfect markets by getting drivers to pay all of the costs of their actions: since they have already paid the market costs to the oil company, they also need to pay, on top of that, the externality costs. These externality costs are the costs inflicted on others but not borne by the driver or the oil company.

We now have all the elements in place to design an externality charge. We know that there may be costs and benefits that spill over from an individual choice or a market exchange, and if so, this will be inefficient (translation: we could do better, making at least one person better off and nobody worse off). We also know that if we want to change behavior to correct the inefficiency, we need to address prices at the margin, not average prices. Third, we do not need to worry about costs, which have been incorporated into a well-functioning market transaction, only externality costs, which have been left out. Fourth, our marginal pricing should reflect those externality costs accurately. It's not enough simply to ban any behavior we don't like. Instead, we should be focusing on cases where the active person gains small benefits but causes large costs for others.

Two objections to externality charges

A charge for externalities is effectively a government tax, and all government taxes are controversial. Externality charges are often attacked from two opposing ranges of moral high ground.

How much do people spend on fuel?



Source: Smith 1992. Chermick and Reschovsky 1997.

On one flank comes the objection that the externality charge is an unfair tax aimed at a disadvantaged group. Consider the idea of charging drivers to drive at congested times. To such proposals (and they have been widespread) the pro-car lobby argues that drivers pay enough, and it is not fair to price poor drivers off the road. On the other flank come those who strongly object to the activity that is to be taxed on the grounds that after the externality charge has been imposed, the rich will still be able to do whatever it was that was objectionable. In the case of traffic, the anticar lobby claims that it is outrageous that rich drivers can afford to drive around as much as they like, given the environmental damage caused by cars.

Are externality charges unfairly redistributive? They are not aimed at poor people but at voluntary activities: if you decide to stop causing trouble for others, you don't have to pay an externality charge. It is true that the rich can afford to drive more than the poor, but it is just as true that the rich can afford to eat more than the poor. This is unfair too, but if you accept the work-

ings of the price system for typical goods like food, why not road space or clean air? We recognize that food, clothes, and houses cannot be free or we would quickly run out of them. It is because roads are free that we have run out of spare road space.

Furthermore, since the rich do more of most things, externality charges often redistribute money in a desirable way. In the case of congestion charging, the truth is striking: in the United Kingdom, poor people do not drive—they bicycle, walk, or take the bus. The poorest tenth of the population spends almost seven times less on fuel than the richest tenth, as a percentage of their much smaller income. The total spending on fuel by the richest 10 percent is at least thirty times more than by the poorest 10 percent. The conclusion is that congestion charging not only improves efficiency, it also redistributes money by raising more tax from the rich.

That's nice for the defenders of congestion charging in Britain, but useless in the United States, where the poor still drive a lot and so pay larger amounts of tax as a percentage of their incomes. But this needn't be an impossible objection, because externality charges can be designed not to redistribute very much. In the case of roads, the government could scrap the vehicle excise duty, which is a large up-front tax, while starting to levy congestion charges on each trip. This would capture the efficiency benefits of a congestion charge without having major effects on distribution. It is possible to neutralize much of the redistribution caused by the externality charge, while keeping its efficiency-boosting effects. This is a variant of the lump-sum tax on Tiger Woods proposed in chapter 3: we can use lump-sum taxes to redistribute without destroying efficiency.

Having met the attack from the redistributive flank, the economist must face the other way and deal with the enthusiastic charge from the moral high ground of environmentalism. Not every environmentalist opposes pollution and congestion charges, but some do. The reason is that they feel that pollution should simply be illegal, rather than illegal for the poor and affordable for

the rich. Why should the rich be allowed to pollute? More generally, some pressure groups protest against externality charges on the grounds that they allow people to pay and thus continue doing whatever objectionable thing it was that they were doing.

A partial response is to say that even the rich do not pollute for fun. It is true that the rich are more likely to be able to pay a congestion charge, but they will not ignore it. Perhaps they will be careful to make one trip to the store rather than two, or even walk to the local shop rather than drive to somewhere farther away. Externality charges make other alternatives look more attractive, both to rich and to poor.

More fundamentally, we must not confuse the strictness of the externality regulation with the method of the regulation. A congestion charge can be set at one dollar a day, or ten dollars a day, or a thousand dollars a day. What we know is that whatever society decides about the seriousness of the externality, externality charges are the most efficient way to deal with it. Well-designed congestion charges, for instance, are the most efficient way to achieve any given reduction of road use. How much of a reduction is desirable is an open question, but congestion charging can deliver no matter what the answer is.

There are alternatives to congestion charging which stop short of the extreme of an outright ban on driving. The trouble is, they don't work as well. For instance, the government could give everybody vouchers, which allow them to drive up to twenty miles per week. The immediate result of such a plan is that some people, mostly poor, would want to sell their vouchers to others, mostly rich: the poor would rather have the money, the rich would rather have the right to drive. If the government allows vouchers to be traded, they have simply imposed a congestion charge by another means and probably a slightly less efficient one, given the hassle of trading. (The charge is whatever the market value of the vouchers turns out to be.) And if the government bans voucher trading, the plan is clearly inefficient, because people who wish to trade are being prevented from doing so.

Other alternatives such as high parking charges are probably less efficient still, although it is harder to prove this in a paragraph. For instance, high parking charges do discourage some drivers; but the relationship between driving and parking is rather indirect. Some drivers spend extra time on the streets, looking for free spaces. If the government wishes to discourage driving by causing drivers some expense, better to charge them directly and spend the revenue on something useful.

Some interest groups will always complain that externality charges are not tough enough, while others squeal that they are draconian. The economists' defense is that however tough we agree to be, an externality charge is the most efficient way to be tough. For any other policy, the economist can propose an alternative, using an externality charge, which would make some people better off and nobody worse off.

How much is your life worth?

The previous section makes it very clear that the level of any externality charge is bound to be a matter for controversy. For the Undercover Economist, seeking to re-create the "world of truth," the ideal externality charge addresses all the real external costs and *only* the real external costs.

It's worth thinking about what the ideal system would look like for the externalities surrounding driving. Any driver making a trip, which emitted pollutants that damaged the local area, would be charged for the pollution if he drove in a densely populated area. There would be a different charge, which would be imposed for every trip, for emitting carbon dioxide, because it contributes to climate change no matter where on the surface of the planet it is emitted. In each case, the price of the trip would also depend on how clean the vehicle emissions were. Drivers would face additional charges for trips in congested areas at congested times. Older buses, which emit the worst pollutants, would be heavily taxed, and this is likely to lead to an upgrading of their engines. Heavy vehicles would incur a charge for trips on fragile

roads and bridges. SUVs would be taxed because they are more likely to kill other road users in an accident.

Does this mean we should bring back the “luxury tax” on expensive vehicles? Not at all. Such a tax was almost certainly environmentally counterproductive, because it encouraged people to hang on to old, more polluting vehicles: a cheap old car will generally be much dirtier than a fancy modern car. SUVs may be more heavily taxed because they are fuel-inefficient and because their weight and height poses a danger to other vehicles—but not because they are expensive. The aim is to encourage people to drive smaller, lighter, more efficient vehicles, not to encourage people to drive cheaper vehicles.

That sounds complicated. Could it possibly work? It’s easy to imagine each car having a little computer linked to a global positioning system to track congestion; the computer would also monitor the vehicle’s exhaust. A display on the dashboard would flash up the rate at which the charge was being incurred, perhaps with helpful tips. “Tim, your trip is currently costing you nine cents per minute. Did you know that you could halve this if you got your engine tuned?”

The technology will come; much of it is already available. But there is another difficulty: working out what the costs of the externalities really are. The computer can measure congestion and pollution, but what is the cost of wasting other people’s time in a traffic jam? What is the cost of poisoning people with particulates or benzene? Many other externalities involve real external costs and benefits that are very hard to measure: time, health, peace, even death.

For the sake of clarity, it may help to focus on the specific example of pricing the externalities caused by driving. It is hard enough to measure physical facts: how much road damage does one more vehicle do? How much noise does one more vehicle cause? How many accidents? How much delay to other vehicles? How much pollution? What ill-health does that pollution cause? But it is harder still to measure psychological consequences. How

much do people care about various annoyances: foul air, noise, delay, and stress—even illness and death? Not to mention that every individual values these things differently.

It is very tempting to give up in the face of these problems. Surely it is not possible to set a value on noise or delay, and certainly it seems impossible to set a value on human life. But we're kidding ourselves if we think we can opt out of these decisions. Every policy the government adopts, and every individual choice you make, implies that a valuation has been made, even if nobody has been honest enough to own up to it or even admit it to themselves.

Individually, we constantly make decisions that put a value on our own environment, our own time, and even our own lives. If you pay more to avoid a noisy area when you rent an apartment or a hotel room, then you have implicitly put a value on peace and quiet. If you decide to wait for the bus rather than flagging down a cab, you are implicitly putting a value on your time. If you decide you can't be bothered buying a smoke alarm, you have traded off saved time and expense against an increased chance that you will die. However, when you make any of these decisions, you probably don't come clean to anyone, even yourself, about the price you've put on quiet, time, or life.

Governments, too, make decisions that imply that they have worked out how much our lives are worth. Should the government install extra street signs and markings, or spend more money on speed cameras, or improve health care, or fund cancer research—or indeed not do any of these things but cut taxes, improve the quality of universities, or spruce up national parks? Such decisions have to be made; when they are made, embedded within them are assumptions about subjective values, including the value of human life. Estimating externality charges is simply more awkward because, if done properly, it requires that those assumptions be justified and made explicit. Leaving them implicit and unjustified leaves us at best at the mercy of the random drift of political processes, and at worst acquiescing to the self-serving demands of interest groups.

One of the best ways of estimating these subjective values is to look at what people actually do. Economists have a theory of “revealed preference,” which is that people’s preferences are revealed by the choices that they make as consumers. You bought apples when you could have afforded pears: therefore you preferred apples to pears. For an economist, preference is not just deduced but actually defined by such choices. It is a short step to conclude that people are also rational consumers when it comes to less tangible factors, even when it comes to their health and safety. If you are not willing to pay five dollars to catch a cab and save twenty minutes, then the Undercover Economist concludes that you would rather spend the five dollars on something else. This is not a very dramatic conclusion, but some people find it controversial. He also concludes that peace and quiet is worth an extra fifteen dollars a week to you, on the basis of your rental decision; and observing that you have no smoke alarm, he presumes that you are not willing to spend an hour and pay twenty dollars to reduce your chance of death by one in a million.

Two important sources of information about people’s preferences are house prices and wages. House prices contain embedded information about the value people place on all kinds of amenities: shops, greenery, low crime, quiet, the sun through the window in the morning, and so on. Some of these can be quite accurately measured: for instance, the price of two identical houses facing each other across the same street will probably reveal how much people prefer a house that faces the sun. Meanwhile wages can reveal information if there is a salary differential for jobs with very similar skill requirements but different levels of danger.

There are flaws in this method: in particular, what if the peace and quiet comes hand-in-hand with a cul-de-sac, which is safe for your children to play on, and insulation, which will save your heating bills? How much of the fifteen dollars a week is really a payment for quiet surroundings? What if the well-paid but dangerous job on the oil rig also requires you to not drink for six weeks at a time and spend all your spare time indoors? Perhaps

the payment is nothing to do with danger, and everything to do with inconvenience. It will always be hard to disentangle these different factors, and it is impossible to know how well you have succeeded. But with enough information, economists can have what they think is a decent try.

A second problem is that when you bought your smoke alarm, perhaps you thought that it would reduce your chance of death by only one in fifty million, not one in a million. So before we leap to conclusions about how much you value your own life, we really need to find out how likely you thought the smoke alarm was to save it, and recognize that you may quite rationally not bother to invest too much time in finding out.

Controversial and imperfect as these methods are, they reflect an important presumption of mainstream economics: nobody has your best interests at heart quite as much as you do yourself.

Two different gaps in our knowledge

The use of externality pricing does rely on shaky information about how much it is really worth to us to reduce externalities such as noise, accidents, pollution, and congestion. But this is not the only gap in our knowledge; we also do not know the cheapest way of reducing noise, accidents, pollution, and congestion. It is with this second gap that externality pricing comes into its own.

Externality pricing is no worse than any other policy when it comes to facing up to the first kind of shaky information. We now know that any policy—of regulation, pricing, command and control, tax, or “laissez-faire”—contains implicit or explicit assumptions about the scientific evidence on externalities like pollution and congestion, and the subjective preferences of people about their time, convenience, and health. No policy can be more successful than the accuracy of its assumptions.

The real advantage of externality pricing is that it circumvents the second gap in our knowledge. Nobody knows the cheapest way of solving our traffic problems—yet. But externality pricing

brings pollution, congestion, and the rest inside the world of truth, which markets create for us. As long as individuals have to face the truth, or at least our best estimate of the costs of their actions, they will find a way to reduce those costs. The longer they have to respond, the more surprising and innovative the responses can be, as we are about to see.

The New Orleans effect

A visit to New Orleans tells us how profoundly people can react to price signals. New Orleans displays a unique architectural style—the “Camelback” house—based on avoiding tax. In the late nineteenth century, houses were taxed based on the number of stories at the front, so the Camelback design had one story at the front and more at the back. They are charming, but if it was a practical design for a house it would have caught on elsewhere. There is a similar story in Britain, which is full of dingy houses in response to the policy, in force from 1696 to 1851, of taxing people based on the number of windows their homes had.

Advocates of congestion charging believe that it must be easier to persuade people to find a way to make fewer trips by car than it is to persuade people to build their houses in an architecturally innovative but wasteful style. Their expectation is that things wouldn’t change much for a few weeks, but over the months and years, we would be living in a society where we could all get around safely and quickly.

Congestion charging can change the small decisions we make every week about whether to drive to a supermarket, or catch the bus, or walk to a local store, or buy food on the Internet. But it will also weigh in the balance with the big decisions. Each year, one in three people change jobs and one in seven people move; every time that happens, there is a clear opportunity to reconsider travel choices in the light of congestion charging.

There’s also a domino effect here, as changes in behavior reinforce each other. If more people begin riding buses, there will be more room on the streets and buses will move more quickly . . .

and can cost-effectively run more frequently. If more people join car pools, each person will find potential pool members more quickly and with more similar trips. If more people try to save the congestion charge by working at home a couple of days a week or commuting at a different time of day, more companies will find ways of accommodating them. People may try to live closer to their jobs; or companies may move to more rural areas to allow staff to commute without paying a high congestion charge.

We simply do not know. The attractive thing about externality pricing is that it attacks the problem but makes no assumption about the solution. The congestion charge gives drivers a signal: by bringing your car into town in rush hour, you are imposing a cost on everybody else. The drivers then have a choice: pay compensation, or find a way to avoid imposing the cost. There are many, many ways to avoid that cost, and markets can produce the ingenuity needed to uncover them. When no externalities are present, markets automatically take account of costs and provide incentives for producers to reduce them. When externalities are present, those costs are invisible to the market, but systems such as externality charging provide the missing signal that the cost exists.

When London introduced a congestion-charging zone in early 2003 (charging £5 or about \$9 per day to drive into the city center) people responded far more quickly than many critics had expected. After a year, car rides fell by nearly a third. Trips that were exempt from the charge became more popular: there were 15 percent more bus rides, 20 percent more motorcycle rides, and 30 percent more trips by bicycle. Drivers who no longer enter the charging zone have chosen a variety of responses: one quarter drive around it, 55 percent have switched to public transportation, and 20 percent use alternatives like bicycles, car pools, or working from home on some days. While the number of trips by car fell, the total delays caused by congestion fell by much more, which suggests that the congestion charge allowed the

streets to be much more efficiently used. And as people have more and more time to adjust to the congestion charge, the cost of dealing efficiently with this externality will fall further.

Battling pollution on the cheap

In the 1990s the Environmental Protection Agency (EPA) in the United States discovered how cost-effectively an externality charge could fight pollution when it decided to attack acid rain. The EPA wanted to reduce sulfur pollution from power stations. It seemed likely that some reduction would be efficient, but reducing pollution has costs as well as benefits. So the regulators were unsure by how much they should demand that pollution be reduced.

The trouble is that polluters will lie to regulators about what the cost of abatement really is. After all, even breathing emits a pollutant, carbon dioxide. But regulators could hardly demand that we all stop breathing to prevent pollution. So which pollution should be reduced? And how? By switching to different methods of power generation? Or reducing power consumption? Or something else? Ask the polluters and they will all tell you that reducing their pollution is like stopping breathing—it would be very expensive to stop, and so somebody else should make the changes.

But it's not really hard to find out the truth. Regulators can find out how much it costs to reduce pollution by telling people either to change their ways or pay a charge. Watch which decision they make. Judge them by their actions.

The EPA tried this in the case of sulfur emissions. They set up an auction for the right to emit sulfur dioxide, which causes acid rain. Polluters were given a quota of emission permits and could either buy more permits in the auction or reduce their emissions by shutting down, installing sulfur scrubbers, or buying cleaner coal. When the EPA simply tried to tell them to install sulfur scrubbers, the power generators argued that it would be very expensive to do so, and they lobbied hard to stop the mandatory

regulation. Even the EPA estimated that the cost of reducing sulfur dioxide emissions by one ton would probably be in the range of \$250 to \$700 and might be as high as \$1,500. But when the EPA conducted the auction in 1993, very few polluters made high bids. The companies had been exaggerating their costs. By 1996 permit prices had fallen to \$70 a ton, and even at that price many polluters were buying cleaner coal or installing scrubbers rather than buying permits to continue polluting.

The regulators discovered that getting rid of sulfur dioxide was so cheap that few people were willing to pay much for the right to keep producing it. In the end, the only people willing to pay high prices for permits were student environmental groups buying single permits in an attempt to win fifteen minutes of fame. The clever thing about the auction was not that the sulfur emissions were reduced—that could have been required by law—but that legislators all over the world found out how much sulfur scrubbers *really* cost. It created a basis for further legislation: not making rules in the dark but in full knowledge of the (modest) cost. And it has set an example to the world; for instance, Taiyuan in North-East China is putting a similar plan into place.

Now economists are designing the same kind of auction for carbon dioxide emissions in the hope of reducing the effects of climate change. There is massive controversy about how much emission reduction will cost, but an auction of permits to extract oil, coal, and gas would soon start to tell us. An auction could start gently: in 2007, auctioning permits to extract the same number of tons of carbon as were extracted in 2006. This would require that economic growth take place without any growth in carbon emissions. If many environmentalists are to be believed, the auction wouldn't even sell all the permits, because basic energy efficiency measures cost nothing. We'd soon find out.

Then over the next few years, we would auction fewer and fewer permits. Carbon emissions would probably fall faster than the number of permits, because carbon speculators would be buying the permits and hoarding them. This would cause no problems: the same emissions take place in the end but are delayed. If

it turned out that the permits were expensive, then we would have the information for an informed debate. We could ask if the costs of climate change were worse than the cost of emission reduction. But many economists believe that, like sulfur permits in California, the carbon permits would quickly reveal that decarbonization is cheaper than we expected, and we will wonder why we took so long to start.

Is the environment too important to be a moral issue?

“How did you travel here today?”

“I’m sorry?”

I’m puzzled. Here I am, going to a panel discussion organized by an environmental charity, and a very earnest young member of staff is grilling me before I even get past the door of the lecture hall.

“How did you travel here today? We need to know for our carbon offset program.”

“What’s a carbon offset program?”

“We want all our meetings to be carbon-neutral. We ask everyone who attends to let us know how far they came and on what mode of transportation, and then we work out how much carbon dioxide was emitted and plant trees to offset the emissions.”

The Undercover Economist is about to blow his cover.

“I see. In that case, I came here in an anthracite powered steamer from Australia.”

“Sorry . . . how do you spell anthracite?”

“It’s just a kind of coal—very dirty, lots of sulfur. OW!”

The Undercover Economist’s wife gives him a sharp dig in the ribs.

“Ignore him. We both cycled here.”

“Oh.”

Apart from being a good example of how irritating an Undercover Economist can be, this true story should, I hope, provoke a few questions. Why would an environmental charity organize a carbon neutral meeting? The obvious answer is “so that it can engage in debate without contributing to climate change.” And that is true, but misleading.

The Undercover Economist in me was looking at things from the point of view of efficiency. If planting trees is a good way to deal with climate change, why not forget about the meetings and plant as many as possible? (In which case, everybody should say they came by steamship.) If the awareness-raising debate is the important thing, why not forget about the trees and organize extra debates?

In other words, why be “carbon-neutral” when you can be “carbon-optimal,” especially since the meeting was not benzene-neutral, lead-neutral, particulate-neutral, ozone-neutral, sulfur-neutral, congestion-neutral, noise-neutral, or accident-neutral? Instead of working out whether to improve the environment directly (by planting trees), or indirectly (by promoting discussion), the charity was spending considerable energy keeping itself precisely “neutral”—and not even precisely neutral on all externalities, nor even a modest range of environmental toxins, but preserving its neutrality on a single, high-profile pollutant: carbon dioxide. And it was doing so in a very public way.

A kind view would be that the charity was setting a “good example,” if acting nonsensically can ever be a good example. An unkind view would be that it was indulging in moral posturing.

This line of reasoning may make economists look too smug for their own good, but it is an important illustration of a wider point. The ethical showboating of an environmental charity can be directly connected to the fact that public policies do not make

evident the environmental costs of our actions. If they did, environmentalists could argue their points from an economic standpoint; much of the moral tone would drain out of the environmental debate, but the environment itself would be much more effectively dealt with.

In a world where environmentalism is merely a moral issue, even the environmentalists themselves cannot work out the environmental impact of everyday decisions. Which is worse: disposable diapers (which clog up landfill sites) or washable diapers (where the washing process uses electricity and releases polluting detergents)? Even with the best will in the world, it is hard to know how to make the right choice.

More importantly, the diaper problem, like any other environmental issue large or small, will certainly not be solved by a tiny minority arguing inconclusively over the morally appropriate individual action. While the Green minority lacks the right signals about environmental damage to act appropriately, the majority of people would not inconvenience themselves even if they understood environmental problems. Both information and incentives are necessary, and as we discovered in chapter 3, markets can provide both.

Economists have long been in the forefront of analyzing environmental problems, and this double difficulty is why they advocate externality pricing. Economists care about the environment but dream of a world where it is no longer an issue that invites moral posturing, but is properly integrated into markets and the world of truth, which would provide both the information and the incentives necessary to persuade ordinary people to behave in an environmentally responsible way. In such a world, we would all have clear signals about the costs of our actions within a market price. Plastics might well be taxed, because they do not biodegrade and so fill landfill sites. This would discourage the use of plastic packaging, disposable plastic bags, and plastic diaper liners. People would use only the more expensive plastics if the convenience it provided was worth the extra money—as it probably is in the case of diapers but might well not be in the case of plastic

packaging. Electricity generation that contributed to climate change would be taxed, too, which would raise electricity prices unless we could develop cleaner fuels. Diaper-washers, and everybody else too, would have an incentive to buy more efficient washing machines and cut down on energy use generally.

Instead of fretting about the environmental impact of our decisions, we would be well aware that if we were willing to pay the externality charge on a product such as a diaper, we would be compensating others for the harm done by our actions, and at the same time we would be confident that that harm was less than our own convenience. We might even find that there are easier ways to improve our environment than by messing around with diapers. They're messy enough already.

Being positive

We've spent a lot of space on what economists call "negative externalities"—unpleasant side effects of actions people get away with scot-free.

Once you start to think about the idea of "negative externalities," you quickly realize that there must be "positive externalities," too. These are *pleasant* side effects of things people do, for which they are not rewarded. If Abraham paints the front of his house and sorts out his garden, the whole street looks better as a result, but nobody will offer to pay for his paint or pruning shears. If Belinda opens an attractive sidewalk café, the streets are more pleasant to walk along, but her clients will be willing to pay only for their own pleasure, not for the pleasure of bystanders. And if Craig decides to vaccinate his son for measles, mumps, and rubella, this means other children are less likely to catch the diseases, but the government can only go so far in encouraging Craig to do this.

"Positive externalities" all seem very agreeable, until you realize that Abraham may decide he can't be bothered to paint his house, Belinda, fearing bankruptcy, may not set up her café, and Craig may decide that he is too worried about the possible side

effects of the vaccines to take his son to the doctor. The rest of us would have benefited if they had gone ahead, but they each decided that on balance it wasn't worthwhile. Just as negative externalities will tend to lead to too much pollution or congestion, positive externalities will leave us undervaccinated, with scruffy neighbors, and a dearth of pleasant cafés. And while negative externalities attract all the attention, positive externalities may be even more important: so many of the things that make life worth living are, in fact, subject to positive externalities and are underprovided: freedom from disease, honesty in public life, vibrant neighborhoods, and technological innovation.

Once we realize the importance of positive externalities, the obvious solution is the mirror image of the policies we considered to deal with negative externalities: instead of an externality charge, an externality subsidy. Vaccinations, for example, are often subsidized by governments or by aid agencies; scientific research, too, usually gets a big dose of government funding. But we need to be realistic about how far all this should go, because although externality charges and subsidies seem a great fix for externalities, there may be an unexpected hiccup.

Too much of a good thing?

Solving externalities without the government

When is an externality not an externality? Here's an example. I may complain about my neighbor's tree damaging my wall, but if it really bothers me I can pay him to let me cut it down. If he refuses the offer then I have to conclude that he gets more pleasure out of the tree than I get nuisance, and in fairness it should stay up. Or perhaps I have the right in law to force my neighbor to cut down the tree. But in that case he can pay me not to exercise that right, and I can spend some of the money fixing my wall. If I have the right to decide then I end up richer; if he has the right to decide he ends up richer. But either way the tree stays up if it's worth more to him and it comes down if it's more irritating to me.

Externalities simply aren't externalities if people can easily get together and negotiate. Remember that they were called "externalities" because they stood outside market transactions. But some of the things that we imagine to stand outside the market can easily be brought into it.

Since these pseudoexternalities can in fact be dealt with very well by the private sector, if the government also steps in with an externality charge, we may find ourselves "solving" the externality twice. This would be just as undesirable as not solving it at all: instead of having too much pollution from power generation, we exaggerate the cost and overcompensate by switching off our freezers and streetlights and walking to restaurants every evening, in the dark.

How would you get this kind of "overdose" of a remedy for externalities? When mentioning the case of Abraham painting the front of his house, I glibly said that the neighbors wouldn't compensate him for his paint, but in fact such things have been known to happen. It's particularly common for landlords to buy paint for tenants to deal with a positive externality: if the tenants repaint, they get a more pleasant apartment, but the landlord will also find it easier to rent it out in the future. Taking into account the advantage to both landlord and tenants, the apartment is worth painting, but without cash from the landlord, the tenants may decide it's not worth the bother. By providing paint the landlord shares the cost as he will also enjoy the benefit. It might well be that in this case, what appeared to be an externality has been "internalized" by bargaining over sharing the costs.

But what if the government has been thinking about subsidies for positive externalities? Imagine that the policy wonks wrongly believe that wherever an apartment is being rented, landlords and tenants will renovate too infrequently and let the place fall apart. It isn't worth it either for the tenants or for the landlord to refurbish more often, but if either of them were taking the other's interests into account, they would agree to do the work. Seeing this positive externality, the government starts dishing out externality subsidies of \$500 to tenants who renovate their apartments.

(If you find this implausible, remember that governments certainly grant subsidies for improvements to the energy-efficiency of homes, which produce a positive externality.)

Imagine that the redecorating effort is worth \$300 of cleaner living to the tenants and \$500 of higher future rents to the landlord, but that it costs \$1,000. The government has identified the externality perfectly: \$500 of benefits (higher rents) to the landlord that the tenants do not take into account. But note that since \$300 plus \$500 is less than \$1000, the subsidy is not enough to persuade the tenants to redecorate: nor should it be, because the redecorating is clearly more trouble than it is worth to both landlord and tenants.

Unfortunately, the tenants have every incentive to pocket the \$500 subsidy and ask the landlord to chip in another, say, \$350. Then they are doing very nicely, enjoying a \$1,000 renovation for only \$150—because 85 percent is paid for by the landlord or the government—and which is worth \$300 to them. The landlord is willing to play along, because he pays \$350 for a renovation, which is worth \$500 to him. But the redecoration should never have happened: after all, the government has spent \$500 but managed to make tenants and landlord only \$150 each better off: not a very effective way to give money away.

Why does the problem arise? Because the positive externality has been dealt with twice over, once by a government subsidy and once by a process of bargaining. Either solution alone represents an efficient way for a society to deal with externalities and to come to the right conclusion, which in this case is that the positive externality is not large enough to justify the decoration work. Both solutions together mean that there is too much subsidy of positive externalities. The same thing could happen for negative externalities too. If the government taxes my neighbor's dirty gas-powered lawnmower and I also offer to pay him to get rid of it because I dislike the noise and smell, the combination of the taxes and my own offer may persuade him to ditch the monster, even though the fun and convenience he gets from

the thing outweigh any damage to anybody else and he really should be keeping it.

Yet, many externalities are very real. Outside the cozy (if cracked) walls of our gardens or our unpainted apartments, we still have clogged, choking streets. Traffic congestion is not easily solved by sitting down over a cup of coffee and making a deal. There are too many people involved, enforcing the agreement would be impossible, and there would always be a temptation for many to avoid the costly negotiations but hope to enjoy the benefits for free.

Government-imposed externality charging is far more likely to be appropriate in situations where a negotiation over the externality will not work, as in the case of the noise from low-flying aircraft. The more likely people are to be able to sit down around the table and work something out, the more likely government intervention is to screw things up: first, because governments can be swayed by interest groups and so they do not always act in the public interest; second, because of the “overdose” problem; and last, because people know the truth about their costs and benefits better than any government ever could. Externality pricing will work very well for problems such as congestion and climate change, for which individual negotiation is nearly impossible. For smaller-scale situations we must ask whether the government-imposed cures are worse than the disease.

Epilogue: what is economics really about?

This chapter has proposed ways of dealing with some of the major blights on our society: pollution, congestion, and fights with neighbors. We’ve learned that an externality charge on waste or on driving in congested areas, or a subsidy for research or for vaccinations, is the most efficient way of dealing with many of the problems that the market leaves to one side. Externality charges give people both the information to make the right choices, and the incentives to do so. Such charges do not automatically answer the question of how tough regulation should

be, or what should be regulated, but once our political processes have produced a view about what we want, they provide the most cost-effective way of getting there.

Yet you will often hear so-called experts complaining that taxes on driving or on pollution would be bad for the economy. That sounds worrying. But what is “the economy”? If you spend enough time watching Bloomberg television or reading the *Wall Street Journal* you may come to the mistaken impression that “the economy” is a bunch of rather dull statistics with names like GDP (gross domestic product). GDP measures the total cost of producing everything in the economy in one year—for instance, one extra cappuccino would add \$2.55 to GDP—or a little less if some of the ingredients were imported.

And if you think this is “the economy,” then the experts may be right. A pollution tax might well make a number like GDP smaller. But who cares? Certainly not economists. We know that GDP measures lots of things that are harmful (sales of weapons, shoddy building work with subsequent expensive repairs, expenditures on commuting) and misses lots of things that are important, such as looking after your children or going for a walk in the mountains.

Most economics has very little to do with GDP. Economics is about who gets what and why. Clean air and smooth-flowing traffic are part of the “economy” in this sense. It’s possible that congestion charging would increase GDP because people would get to work more quickly and produce more, and prices in stores would be lower because of more efficient distribution. But it’s perfectly possible that congestion charging would reduce GDP. This does not, in fact, matter in the slightest. We know for certain that it would make us better off in a much more meaningful sense: that we would have many new choices open to us about where we go and what we do. There is much more to life than what gets measured in accounts. Even economists know that.