

2 Surplus Maximization

How should we decide whether a regulation is good, or bad? The first step is to choose some valuation rule. Suppose, for example, that we are trying to decide whether a rule requiring the arsenic level in drinking water to be less than 23 parts per billion is a good rule or not. A strident environmentalist might say that the more stringent the rule, the better— that a level of 23 is better than 30, but 4 would be an even better rule. Someone else might say that cost should be considered too, and that reducing the level to 4 parts per billion would cost more than the entire budget of the city government, requiring taxes to double.

The standard valuation rule used by economists is what I shall call “surplus maximization,” but which is also variously called “cost-benefit analysis”, “efficiency”, and “wealth maximization”. The idea is simple. Add up how much each person who likes the regulation would pay to have it (his “willingness to pay”, and subtract out how much each person who dislikes the regulation would require to be paid to accept it (his “willingness to reject”). If the resulting number is positive, adopt the regulation.

A concrete example will illustrate. Suppose Anderson and Brown want a stricter arsenic regulation and would pay up to \$30 and \$70 to get it, whereas Corman and Daniels don’t want it, and would require payments of at least \$20 and \$10 to feel that the new regulation had not hurt them. Since supporters would pay up to \$100, and opponents would accept as little as \$60, adopting the regulation does maximize surplus.

This is not at all controversial if the payments actually take place. If we adopt the regulation and make Anderson and Brown each pay \$25 to Corman and Daniels, everybody is happy that the deal went through. We call this a Pareto improvement”, after the economist Pareto who came up with the criterion that if everybody is better off, a policy is good.

What is more controversial is what whether we should say that the regulation is desirable even if the new policy does not include payments to Corman and Daniels. The regulation still maximizes surplus, but Corman and Daniels would oppose it. The standard position of economists is that the regulation is nonetheless desirable, because the winners win more than the losers lose. Absent other considerations, this surely makes sense. The alternative is to make Anderson and Brown unhappy— and more unhappy than Corman and Daniels would have been— by rejecting the regulation.

A Single Market Transaction

Let us now apply the idea of surplus maximization to a market transaction. Smith approaches Jones and asks if Jones will sell a bottle of whisky for \$10. Jones agrees, and the

whisky changes hands. Is the transfer of the bottle from Jones to Smith a good thing? Surplus maximization says that it is. Since Smith offered a price of \$10, we know that his willingness to pay was at least that high, and probably higher— maybe \$15. Since Jones accepted the price of \$10, we know that his value for the bottle was no more than that— maybe it was \$8. Using the figures of \$15 and \$8, the net benefit from Jones giving the bottle to Smith is \$7, an increase in value. It may be that the true gain in value is more than \$7 or less, but since both parties agreed to the trade, we know that Jones’s value must be less than Smith’s.

I will be calling this valuation rule “surplus maximization”, but we can use this example to illustrate why the other terms are used for it. Suppose the values of \$15 and \$8 are correct, and moving the bottle from Jones to Smith at a price of \$10 has a benefit of \$5 for Smith and \$2 for Jones. The effect on their satisfaction is then the same as if the trade were blocked, but by a miracle \$5 suddenly appeared in Smith’s pocket, and \$2 in Jones’s.¹ This is why some people call the idea “wealth maximization”: a surplus-maximizing trade increases the dollar amount at which people value their possessions. It is also why some people call this an “efficient allocation of goods”: by moving the bottle from the hands of Jones to those of Smith, it is as if the economy had found a technology that increase the amount of output by \$7.

It may seem obvious that it is good for Jones to sell the bottle of whisky to Smith. But surplus maximization has an implication which is more troubling. Suppose Smith simply stole the bottle from Jones. Surplus maximization says that this, too, is a good thing. The effect on total surplus is exactly the same as the sale at the price of \$10. The sale benefited Smith by \$5 and Jones by \$2, a total of \$7. The theft benefited Smith by \$15 and hurt Jones by \$8, which also makes for a total benefit of \$7. Surplus maximization treats these the same. All that matters is that the bottle has moved from someone who values it less to someone who values it more.

This is an example of the power of the idea of surplus maximization, but also of its moral neutrality. Economists do not take a stand on morality. To do so would not be controversial in the case of theft, perhaps, but it would be in the case of most government policies. Is it moral to tax rich people at a much higher rate than poor people? Is it moral to make consumers pay more for sugar to benefit sugar producers? Is it moral to forbid racial discrimination— or to require it, as in the case of affirmative action or old-fashioned segregation? Those are important questions, but the economist evades them. This evasion, moreover, is entirely reasonable. It is hard enough to determine whether a policy maximizes surplus or not. We therefore separate that question from the question of whether a policy is moral, which usually requires much different reasoning to answer.

In the case of theft, however, surplus maximization does provide a reason for why the action is bad. In the particular case of Smith and Jones, surplus maximization says that the theft is good. It is good, however, only because we started with a story in which Smith was willing to pay

¹I am ignoring wealth effects here; actually, adding \$5 to Smith’s wealth might have a minutely different effect than giving him \$5 in satisfaction from a bottle of whisky.

at least \$10 and Jones was willing to accept as little as \$10. This information is crucial to whether moving the bottle from Jones to Smith maximizes surplus. If, instead, we announced that we were going to give the bottle to whoever valued it more, and asked each person for their values, imagine what would happen. Smith might say he was willing to pay his entire wealth for it, and Jones might say that he, too, valued that bottle more than anything else in the world. They would lie.

In actual practice, the idea of surplus maximization crucially requires us to know the value of each person affected. The wonderful thing about market transactions is that market prices force people to reveal something about their true values. Smith puts his money where his mouth is when he offers \$10. Jones reveals something about his value too, when he accepts \$10. We do not learn the exact values, but we know there is a net gain in social surplus.

After a theft, we do not know that surplus has increased. We do learn something— that Smith was willing to go to the trouble of stealing the bottle, and Jones was not willing to go to the trouble of guarding it effectively— but that is not enough to guarantee surplus maximization. And in any case, Smith’s stressful thieving and Jones’s worry and precautions are a social waste. If Smith would have paid \$4 to have gotten the bottle legally, and Jones invested \$2 of his time in trying to protect it, those \$6 are a social waste, because transferring the bottle by open purchase would have avoided them. Thus, theft, while it may be good in particular cases, is bad generally, and it maximizes surplus to make it illegal.

I said that the idea of surplus maximization was powerful. One illustration is that we have just used it to derive a reason why theft is bad, rather than having to accept the evil of theft as an independent moral rule. If we command, “Thou shalt surplus maximize”, we can dispense with “Thou shalt not steal”, “Thou shalt not kill”, and “Thou shalt not commit adultery,” which become mere corollaries. Indeed, this is something like the rule of Jesus Christ (and Rabbi Hillel): “Thou shalt treat thy neighbor as thyself,” a rule not unlike surplus maximization in its unbiased treatment of everyone in society.

Another illustration of the power of surplus maximization is in how little information it requires. Return to Smith and Jones, and suppose that Smith already owns 900 bottles of whisky, whereas Jones only owns the one bottle for which Smith offers \$10. How does that affect our opinion as to whether the sale of the bottle is good or bad?

Presented with these facts, some people might say that the sale was bad. Smith, after all, has plenty of whisky already. He does not need another bottle. Jones will be left whiskyless, and we ought not to allow that to happen to him, poor fellow. Under this view, though, we really need more information. If, in addition, we discovered that Jones had 2, 000 bottles of gin, our decision might reverse; Jones has plenty of liquor. If we then discovered that Smith was a millionaire, our decision might reverse again; Smith can buy as many bottles as he wants, from other people. If we then discovered that Jones only had one bottle of whisky because he had drunk the rest of his stock and was too lazy to go and get more, our decision might reverse yet again; he is a lazy pleasure- seeker, and does not deserve his remaining bottle.

Surplus maximization does not require any of that information. All it needs is that Smith is willing to pay more than Jones for the bottle. It does not matter if Smith has more bottles already– he values that last bottle more than Jones. It does not matter if Jones has lots of gin– he does put a positive surplus on that whisky, and it should not be taken from him . It does not matter if Smith is rich– Jones still benefits from the transaction, and is happy to give up his lone bottle for mere cash. It does not matter if Jones is lazy– if we want to punish him for laziness, do it with jail or a fine, not by preventing the innocent Smith from getting a bottle of whisky.

2.2 An Entire Market

The first step to understanding why free markets maximize surplus is to understand why the transaction between Smith and Jones maximizes surplus. The second step is to look at how the market price is chosen. Let us now look at the entire market for whisky, consisting of 300,000 potential buyers and 5,000 potential sellers, each of whom might sell 100 bottles. Buyers vary in their willingness to pay from those who would pay at most \$.01 per bottle to those who would pay as much as \$30. Sellers vary in their minimum acceptable price from as low as \$4 to as high as \$19. The supply and demand diagrams in Figure 1 show this more precisely, as do the supply and demand equations, $P = 30 - 0.1Q^d$ and $P = 4 + 0.03Q^s$ (with quantities of bottles measured in thousands, as in the figure). The demand curve shows that there are 100,000 buyers willing to pay at least \$20 and 200,000 willing to pay at least \$10. (Let us assume that each buyer only wants one bottle.) The supply curve shows that there are 1,000 sellers willing to take as little as \$7 and 2,000 willing to take as little as \$10. Since each seller sells 100 bottles, this means that at the price of \$10, the quantity supplied equals the quantity demanded. This is the “equilibrium price”: the price generated by market forces.

Figure 1: Consumer and Producer Surplus

Why do market forces generate a price of \$10? Think about what would happen if the price were higher– say, \$20 per bottle. Sellers would be delighted to sell everything they had, but only 100,000 buyers would be willing to buy, so some sellers would end up unable to sell. Those buyers would shave the price to \$19.99, causing buyers to switch to them. That would leave other sellers customerless, and those sellers would shave the price to \$19.98. The price of \$20 per bottle is thus unstable, and the same reasoning shows that any price above \$10.00 is unstable. At the price of \$10.00, however, each seller who is willing to sell find a buyer, and there is no incentive to cut price further. A similar argument shows that any price below \$10.00 is unstable. At any price below \$10.00, buyers are more eager to buy than sellers are to sell, and some buyers would be unable to find a seller willing to sell. Those buyers would bid up the price, leaving other buyers stranded, until the price was bid up to \$10.00.

This argument shows that the free market equilibrium price is stable, but it says nothing about whether it is optimal. The price does turn out to be optimal, but to show that we need to

think about the costs and benefits to buyers and sellers. First, let us calculate them for the equilibrium price of \$10.00 and quantity of 200,000 bottles, bought and sold by the buyers and sellers at the left of Figure 1, the buyers with the highest benefits and the sellers with the lowest values. You will see immediately that whatever quantity is optimal, it ought to be bought by those high-valuing buyers and sold by those low-valuing sellers. This is the result in the free market, but it is worth mentioning because it might not be the result under government regulation, which might, for example, result in the bottles going to the buyers who are most morally deserving or the best connected politically.

Think about the benefit to the sellers who are selling those 200,000 bottles (some sellers are inactive, selling nothing, so their benefit is zero). The gross benefit is their sales revenue, which is simply (200, 000 bottles) (\$10/bottle) = \$2,000,000. This is not their net benefit, because the sellers did value those bottles, even if their values were not as high as the buyers. Their values ranged from \$4/bottle to \$10/ bottle, the height of the supply curve. Combined, their values are the area labelled “Seller Cost” in Figure 1, since for typical sellers the value they place on what they are selling is their acquisition or production cost (though the cost could be a different “opportunity cost”— that they cannot drink the whisky themselves!).

We can numerically calculate the size of the seller cost. Geometrically, it is the area of the rectangle \$4/bottle high and 200,000 bottles wide (which is \$800,000) plus the area of the triangle with a height of (\$10/bottle - \$4/bottle) and a width of 200,000 bottles, which is $(1/2)(\$6/\text{bottle})(200,000 \text{ bottles}) = \$600, 000$. That sums to a seller cost of \$1,400,000.

Since the sellers’ net benefit is their gross benefit (revenue) of \$2,000,000 minus their lost value (cost) of \$1,400,000, their net benefit is \$600,000. This is the area labelled “producer surplus” in Figure 1, which is the standard name economists use for net seller benefit.²

Now think about the buyers. The gross benefit the active buyers get from the 200,000 bottles is the sum of the values for each of the buyers. Some buyers have a value of \$30, some of \$29, some of \$28, and so forth down to the last active buyer, whose value is only \$10. (There exist other buyers who are inactive, but they will not be getting any benefit, so we can ignore them.) The sum of the values is the area under the demand curve up to 200,000 bottles. This equals the area of the rectangle \$10/bottle high and 200,000 bottles wide, which is \$2,000,000, plus the area of the triangle above it with height (\$30/ bottle - \$10/bottle) and width 200,000 bottles, which is $(1/2) (\$20/ \text{bottle}) (200,000 \text{ bottles})= \$2,000,000$. Adding up the two areas (which just coincidentally have the same size) gives us the gross benefit of the buyers, their aggregate value for the bottles, which is \$4,000,000.

The net value for the buyers is less than the gross value, because they have to pay the sellers \$10/bottle. This is a payment of \$2,000,000 for all 200,000 bottles, so the net value is \$4,000,000 -

²Readers who have studied economics have probably learned the concept of producer surplus a different way, calculating the area of the producer surplus directly rather than as the revenue rectangle minus the cost trapezoid. I use the slower method because soon we will look at atypical cases where the producer surplus is not a simple triangle.

\$2,000,000= \$2,000,000. In Figure 1, this is the area labelled “consumer surplus”.

Adding together the producer surplus and the consumer surplus gives us the total surplus, the total surplus created by the existence of this market. When the quantity is 200,000 bottles and the price is \$10/ bottle, the total surplus is thus \$600,000 + \$2,000,000=\$2, 600,000.

Having calculated the surplus created by the free market, we must now show that this is more than any government regulation could create. The first step is to understand a remarkable fact: that it is the quantity of 200,000 bottles that determines the total surplus, not the price of \$10/bottle. To see this, suppose that we keep the quantity at 200,000, being bought by the same buyers and sold by the same sellers as in the free market, but raise the price to \$20/bottle. This must be backed up by government force, as a two-part regulation. The first part is those 200,000 consumers must be forced to buy on pain of prison, something not all of them would do at such a high price. The second part is that sellers must be forbidden to reduce their price, on pain of prison, since otherwise they would shade their price to try to acquire more customers.

Having imposed the regulation, let us calculate the surpluses again. This will be easier than before, because the the gross buyer benefit and the seller cost have not changed from their free market levels! Since the same people are buying and selling the same 200,000 bottles, the buyers still value the bottles at \$4,000,000 and the sellers still value them at \$1,400,000. All that has changed is that the buyers now pay a much higher price– a total amount of (\$20/bottle) (200,000 bottles) = \$4,000,000– and the sellers get that higher revenue. Thus, now the consumer surplus is (\$4,000,000- \$4,000,000) = \$0, and the producer surplus is (\$4,000,000- \$1,400,000)= \$2, 600,000. The total surplus is unchanged from its free market level of \$2,600,000; all that has happened is that now the sellers get all of it and the buyers get none of it.

Whatever price is chosen under this two-part regulation, the total surplus will stay the same. When the quantity is fixed, the price is just a transfer from buyer to seller. The total surplus is

$$(Gross\ Buyer\ Benefit - Price * Quantity) + (Price * Quantity - Seller\ Cost).$$

The $Price * Quantity$ terms cancel each other, so the total surplus is (Gross Buyer Benefit - Seller Cost), which does not depend on the price.

The total surplus does depend on the quantity, however, which we were keeping fixed at 200,000 bottles. To see why this quantity maximizes surplus, consider increasing it. We will need the most reluctant seller to become active, one whose value for a bottle is \$10.00 and has already sold some but not all of his 100 bottles, since sellers with lower values are already selling. We will need a new buyer to become active too, one whose value is less than \$10.00, since buyers with higher values are already buying. Even if the new seller’s value is \$10.00 and the new buyer’s value is \$9.99, the highest possible, this new exchange reduces total surplus by \$.01 rather than increasing it. Increasing sales further would reduce total surplus by even more.

How about reducing the quantity below 200,000? If a buyer with a value of \$10.01 or greater

stops buying from the most reluctant seller, with his value of \$10.00, then total surplus drops by at least \$.01. Reducing the quantity below the free market equilibrium level loses some of the gains from trade, from active sellers having lower values for the whisky than active buyers.

Thus, we see that the workings of the free market maximize surplus. First, the free market arrives at the equilibrium price, without any government intervention necessary. Then, it turns out that the equilibrium price elicits an equilibrium quantity which maximizes the sum of producer and consumer surplus.

The fact that this is a two-step process is why the price matters more than it may have seemed in my discussion earlier. Earlier, I showed that if the government required the quantity to be at 200,000, with the same buyers and sellers as in the free market, then the government could require any level it wanted to for the price without altering total surplus. Thus, it seemed that this two-part regulation maximized surplus just as well as the free market could, even if it could not do any better. But notice what I took for granted: that the government had detailed information at its disposal and it could enforce its regulation costlessly.

How would the government know that the optimal quantity is 200,000? Supply and demand curves are not written down in books that the government can consult. Economists measure them with intricate statistical procedures, using market-generated data. Simply asking people how much they would pay, partly because they do not think hard enough about it if they do not really have to pay, and partly because they might not tell the truth. What would happen if, for example, the government said it was going to set the price of a bottle of whisky to one penny per bottle, and asked buyers to step forward if they were one of the 200,000 buyers with the highest values? Moral scruples aside, all 300,000 of the consumers would step forward, since all of them would like the chance to buy whisky at that price. If the government furthermore asked which 2,000 sellers had the lowest costs, no seller would step forward—none of them would want to be forced to sell at such a low price.

The free market, on the other hand, needs very little information. The process described earlier moves the price to \$10.00 without any need for people to tell the truth. The reason is that in an actual market, buyers and sellers put their money where their mouth is. A buyer has no reason to claim that he would only pay \$5.00 for a bottle if his true value is \$13.00 and the result of his claim is that he loses the chance to buy at \$10.00. A seller has no reason to claim that his minimum acceptable price is \$11.00 if it really is \$9.00 and his lie will lose him the sale. This economizing on information is a huge advantage of the marketplace, even more important than the fact that it does not need to pay police to enforce its prices and quantities.

The logic of this is similar to the logic of why the single transaction between Smith and Jones maximized surplus. Indeed, we can think of Smith and Jones as being just two of the thousands of participants in the market. What is different in the case of the market is that anonymous market forces determine that the price will be \$10, whereas in the single transaction example I said that Smith offered Jones a price of \$10 without explaining where that price came from. But

in both the single transaction and the market, the essential idea is that if both buyer and seller voluntarily agree to a transaction, it (a) benefits both of them, and (b) increases total surplus.

Application to Public Goods

Since we will be applying the idea of surplus maximization to government policies, it is worth emphasizing that in two-person interactions the good they exchange does not have to be physical or something that can benefit or harm just one person at a time. Consider the following examples:

Temperature in the Office. A and B are partners who work in the same office, which has a single thermostat. A's preferred temperature is 72 degrees, but B prefers 70 degrees. A would pay 30 dollars per year for his preferred temperature and B would pay 40 for his.

Keeping the temperature at 72 degrees would fail to maximize surplus. Both partners would be better off if B paid 35 dollars to A and they lowered the temperature to 70. The good here is "control over the thermostat". Who owns the good initially is unclear, but does not matter to the surplus-maximizing outcome. If A owns the good initially, B buys it from him. If B owns the good initially, A won't offer him a high enough price, so B will continue to control the temperature. Most likely, neither owns it initially, and they will negotiate. It is hard to say what would happen unless we know the result of failed negotiation, but we would expect B to pay A something under the difference in their values— 5 dollars, perhaps, for being allowed to control the thermostat.

Office Necktie. A and B are partners who work in the same office. A strongly desires that they both wear jackets and ties, and would pay B 5,000 dollars to dress up. B dislikes ties, and would have to be paid 4,000 dollars to be willing to wear one.

Here, the good is "whether B wears a jacket and tie at work". B owns that good, but he would sell it to A and both would be better off as a result. We could imagine A owning the right initially instead— as he would if he were the boss and B were an employee who signed an employment contract that requires neckties at work. In that case, B would offer A up to 4,000 dollars for the right not to wear a tie, but A would turn him down, and that, too, would be efficient.

Church attendance. A and B are partners who work in the same office. A is an evangelical Christian who cares about B's soul and thinks B would be better off if he went to church and were exposed to good preaching. A would pay B up to 2,000 dollars to attend church regularly. B would rather watch TV, but would be willing to attend church if he were paid at least 1000 dollars.

Here, the good is "B's church attendance policy". The surplus- maximizing outcome is for him to attend church, and since he initially owns the good, A must pay him to do so.

These examples are very simple, but I hope they impress up on the reader the idea that surplus maximization is about more than the allocation of physical goods. Anything that people

care about is a good for the purposes of surplus maximization, and the measure of how much they care must always be quantified. Note, too, that the concept is independent of rights. We may think that someone has a right to dress however he likes, but if we force him to do so, he won't be able to accept bribes to wear a necktie, and he will be worse off in his own view. Nor would the infidel B be better off if he were forced not to attend religious services unless he sincerely believed in God. He would prefer to give up the right to spend his Sundays as he pleases in exchange for money, and he would not be thankful for government intervention.

The term "wealth maximization" is easy to misunderstand, because it seems to imply that the wealth-maximizing economy is one that maximizes the dollar value of material goods. That is not so, as the preceding examples tell us.

In the thermostat story, having the thermostat at the optimal temperature raises social wealth by 10 dollars, in the sense that if they move down the temperature from 72 to 70, B feels as much better off as if he had just been given 40 dollars extra to spend at the old temperature, and A feels only 30 dollars worse off, if B did not have to pay for changing the temperature, for a ten dollar gain in social wealth.

In the necktie story, if B offers A 4000 dollars ...

In the churchgoing story, ...

Peevyhouse v. Garland Coal & Mining Company (1962) (1962 OK 267 382 P.2d 109) is a case which illustrates how confusion over the notion of surplus maximization can arise. Mining activity harmed a farm, and the question was what damages the mining company should pay. The court said that the land was of little value, and so the damages should be small. The farmer and his wife, however, had refused to accept the market value of their property as compensation, and to the economist that is a sure sign that the land was of higher value. Looking at the market price was wrong, because the very fact that they had not sold the land showed that their value for it was higher.

I emphasize how the idea of surplus maximization includes nonmonetary values not because this is controversial in economics in theory, but because we often neglect the nonmonetary aspects, and later we will come to examples where some economist readers will balk.

GDP, Wealth, and Surplus

It is important to recognize that surplus maximization is not the same as the maximization of gross domestic product. First, GDP is a flow, not a stock. Rather than measuring wealth's *level* at one instant in time, GDP tries to measure the *creation* of goods and services over the course of a year. More important, GDP is a practical concept that tries to avoid subjectivity in measurement of people's preferences. For one thing, it leaves out most production that is not traded in markets. If I grow corn for myself to eat, that does not enter the marketplace and so

does not enter GDP. If I do my own housework, that does not enter GDP, but if I hire somebody else, it does. The biggest exception is that even if I own the houses in which I live, GDP includes an estimate of the amount of rent I would have paid for that house. But it does not include the value of use of other durable goods such as cars or household appliances. For another thing, GDP is not a measure of surplus, but of output. Suppose I value my leisure at \$40,000 per year. If I do take leisure, I contribute nothing towards GDP. If I take a job for \$41,000 per year, I am better off, by \$1,000. But GDP rises by the full amount of \$41,000.

Thus, GDP is more like a measure of monetized productive activity than a measure of surplus creation.

GDP does share with surplus creation the idea of using subjective value, not objective value, even if, paradoxically, it tries to find an objective measure of subjective value. The paradox is resolved when we realize that GDP is trying to find out how much people value what is produced, and needs an objective measure of how much they value it. That objective measure is the market price. If people are at the margin indifferent about what they spend their money on, they will be earning zero consumer surplus from the last unit of each thing they buy. Thus, we can look at the market price and use that to measure subjective value.

That is how GDP can deal with the problem of, for example, the value of a professor's teaching. It does not try to measure it directly; it looks at how much the university is willing to pay the professor and does not question the university's taste or prudence. Similarly, it does not ask whether people are right to eat as much candy as they do; GDP just adds up the amount paid for all the candy eaten and takes that as the value.

Criticisms of Surplus Maximization

What is the point of the idea of surplus maximization? If a person is selfish, he is going to make the choice that maximizes his own surplus anyway. Even if he is high-principled— maybe especially if he is high- principled— he is going to follow his own principles rather than what maximizes surplus. Ronald Dworkin makes this kind of criticism in his 1986 *Law's Empire*, pp. 276-295.

Cost-benefit analysis may say that a farm should be replaced by a highway for the greater good even if the farmer does not want to take a generous offer of twice the market price, but the farmer will still not be persuaded, nor would somebody who thinks government takings are immoral even with compensation. Cost-benefit analysis may say that if *A* values *B*'s car more than *B* does, it is fine if *A* steals the car, but no policymaker is going to allow the theft because of that.

To this, there are a number of answers. First, though the answer is unsatisfactory to Dworkin, surplus maximization must be looked at as a general policy rather than as isolated interventions. As we have already talked about, theft may be surplus-increasing in a particular

case, but it discourages investment and induces wasteful effort to steal and to protect. We might add here that to reduce the amount of theft, social norms against it are useful, which will make policymakers hostile to allowing it even in special cases, and the gain from having them generally opposed to theft exceeds the loss from their opposing theft in the special cases when it would help.

The criticism above is deeper, though. Its real bite is in the claim that wealth maximization matches poorly with people's common moral beliefs, so they will not follow it, making it of dubious value as either a positive or a normative principle. We must deal with this.

First, surplus maximization has moral significance even if it is not decisive. To the extent that a person cares about the welfare of others, the effect on total surplus is relevant to his political and moral judgements. Even if I believe I have every right to be first in line at a store counter, and count my own convenience as more important than anyone else's, if I see that somebody else is in a desperate hurry, I will take that into consideration. The calculation of surpluses is worth doing as part of a wide range of moral philosophies, even if most of them do not stop with that calculation.

It is worth noting that people's apparent willingness to use overall cost-benefit analysis, rather than just look at their personal costs and benefits, is a problem in contingent valuation surveys. These surveys ask people to estimate the costs or benefits to themselves of public projects— for example, actions to prevent migratory waterfowl from drowning in oil-and-gas company holding tanks. In a study of that example, Schkade & Payne (1994) asked people what they were thinking when they answered the surveys. They found, unsurprisingly, that people had not thought about their valuations before they were asked. They also found that many people thought about the costs and benefit to other people besides themselves, even though that was not what they were being asked about. They did a rough cost-benefit, surplus-maximization analysis themselves, trying to look at the social value of the project rather than the selfish value to themselves. Diamond & Hausman (1994) note that this is a common problem for contingent valuation surveys.

Second, surplus maximization is neutral. It does not require us to decide what preferences are better and worse or which people are more deserving or less deserving. It can serve as a starting point for analysis for people of differing moral beliefs; they can “agree to disagree” and come back later to people's beliefs and willingnesses to pay that they consider illegitimate or misguided.

This is hard to do, but it is what advisors do all the time. Avery Katz (1996) notes that this is what private lawyers do all the time in advising their clients. They do not ask whether what the client wants is what they themselves would want; instead, they try to explain to the clients the advantages and disadvantages— from the client's point of view— of different courses of action. He says that legal academic writing would, if viewing the public or the government as clients, “require the sorts of tasks that private lawyers undertake in regard to their clients; in particular, it would mean consulting the public and paying attention to its actual views” (Katz, 1996, p. 2266).

Third, surplus maximization is the rule we would choose for ourselves if we had to choose a rule without knowing our later moral beliefs. It treats everyone's beliefs and happiness as equally valid. One may question why this sort of contract approach to society matters now that we do know our current positions, but somehow the "veil of ignorance" approach does appeal.

A different angle on this idea is that while deontological moral theories may be best for individuals, consequentialist ones are more suitable for governments.³ Perhaps the role of individuals is to follow moral rules— those commanded by God, or best for personal character, or whatever— but governments, without personal characters or selfishness, are different. The individuals that make up government should act when advancing their own interests in accordance with moral rules that may have very ill consequences for themselves and everyone else, but when acting on behalf of the commonwealth they should do what is best for the commonwealth and for the people they represent, not to satisfy their personal moral beliefs.

Fourth, surplus maximization can serve for positive predictions. It may not match our moral intuitions perfectly, but it is not that far off either, and from a relatively simple principle we can derive an enormous number of social rules. The simple principle is one way of implementing the Golden Rule: Do unto others as you would have them do unto you.

Related to this is that if normatively the social planner does *not* approve of surplus maximization, he should not rely on democratic decisionmaking. Most people in the United States presently believe that marijuana use is bad and are willing to trade off votes on other issues to keep marijuana illegal. Most people in the various states of the United States (for there was never a federal law) used to believe, and perhaps presently believe, that homosexuality is bad and were willing to trade off their votes on other issues to keep it illegal. The aggregate will of the voters can only be thwarted by removing the issue from their jurisdiction. That has not been done with marijuana, but that is what the Supreme Courts *Roemer v. Evans* decision (517 U.S. 620 (1996)) did with homosexuality. The court ruled that regardless of what voters wished— even if they were unanimous, if we follow the logic of the decision— the state could not prevent cities from granting special civil rights protections to homosexuals. This is an example of using a constitutional institution to constrain surplus maximization. Zamir & Medina (1996) explore this possibility

Fifth, surplus maximization can help reveal hidden moral preferences. When someone says that the war against drugs is a mistake because it has high costs and little benefit, for example, he is probably saying that he thinks that drugs aren't that bad and there is little benefit from reducing their use. Most voters seem to disagree, and what our observer is really saying is that they are wrong: that drug use is not immoral, or the effect on people's lives is not so important.

Sixth, some have argued that surplus maximization actually is a good explainer of everyday moral intuitions. Kaplow & Shavell (p. 428) do a good job of that; see also David Friedman in *Law's Order*, who shows how much of law it can explain. Kaplow and Shavell (p.428) use the example of torture for the satisfaction of one person's sadistic preferences. Is it bad in all

³See Zamir & Medina (2006) p. 25 and Goodin (1995) pp. 60-77.

circumstances? Suppose that the torturer becomes fabulously productive after he tortures someone. He becomes another Shakespeare or Einstein— or both combined— and the victim recovers quickly with no memory of the pain. Suppose, too, that the victim admits that the social benefit from the torture is far greater than his own loss, and that he objects only because he is selfish and is not getting that social benefit. Should we really object to this torture any more than we do to, say, tax or trade policies that help society as a whole while hurting particular individuals? As Kaplow & Shavell say, we should recognize that our moral intuitions about torture are not based on this kind of hypothetical, but at most on stories from history or the news. Perhaps, since we are human, we even tend to pass by or forget stories which fit uncomfortably with our moral beliefs. A person does not have to know all the reasons for his opinions. Unless he think he will change them, or will need to elaborate them, the reasons are historical debris, taking up valuable mental space. As Nietzsche somewhere said, “It is hard enough to remember my opinions, without also remembering my reasons for them!”

Paternalism, Mistakes, and Surplus Maximization

Much of the problem we will see in this book about surplus maximization will arise because we want to apply paternalism. There is a fine line between personal preferences and personal beliefs. Suppose many people think that God will be pleased by spending \$50 billion on a large church. Others think that God does not exist, and so the money will be wasted. How does one perform cost-benefit in this case?

Existence value in cost-benefit analysis is a big problem. What is the value of saving an acre of first-growth trees?

Some people like to buy “fair-trade” coffee. In that case, they pay extra, but some might say they are being fooled— that this is not really helping the coffee farmers.

Another criticism of surplus maximization is that people do not have coherent preferences, and even if they do they make mistakes when try to maximize their utility. Sugden (2005) discusses this well. We will return to that later. It is less deep, though no less important, a criticism than Dworkin’s “Who cares?” point, because it can be addressed by asking what decision would maximize surplus if we could somehow find out people’s underlying preferences and help them to avoid psychological mistakes.

The Virtues of Competition

The unfettered market works well for production as well as for distribution. Sellers must choose how to produce a product as well as how much to produce. In this they will be guided by the desire to maximize profit, their surplus, in their choice of different inputs such as capital and labor. If a more labor-intensive method costs \$10 per bottle and a more capital-intensive method costs \$12, the seller will choose the low-cost method without any need for direction from the buyer or the government. If for an initial investment in research of \$10,000 the seller has a 50%

chance of reducing his costs by \$50,000, he will make the investment without any need for outsiders to tell him of the importance of innovation to the economy.

Conclusions

Thus, in ordinary economic transactions, the free market maximizes surplus without the need for government intervention. That is why economists are generally against government regulation. It is an important idea to understand, and much of an introductory economics class is devoted to trying to explain it.

This chapter has set up an argument against government regulation. The question is what counterarguments are available. Economists do recognize that there are situations in which some premise underlying this reasoning fails, so that government regulation could help. These are called situations of “market failure”, and will be our next topic.