Market Failures and Efficiency

This topic considers *efficiency* and whether or not market outcomes are efficient. We begin with notion of Pareto Efficiency. *Pareto efficiency* means that it is not possible to make anybody better off without making at least one person worse off. Suppose, as a simple example, that \$100 may be distributed between me and you. Assuming that we all like more money, which of the following distributions are Pareto Efficient?

a. \$50 for me, \$50 for you

b. \$100 for me, \$0 for you

c. \$25 for me, \$25 for you.

$a.$ is Pareto Efficient. The only way that I can increase my share is to reduce yours, or vice-versa. By the same reasoning, b is Pareto efficient as well. There can thus be multiple Pareto efficient allocations and the definition of Pareto efficiency is unrelated to equity. Because c does not exhaust all of the \$100, I can, however, give some of the undistributed \$50 to either me or you, making at least one of us better off without doing the other any harm.

In this example, there are a few ways that an outcome may not be Pareto efficient. Like c, it may not use all of the available resources. Another is that people could be altruistic in that they care about others. I might prefer a to b because I would rather you have the extra \$50 than me.[1]

Because there are typically many Pareto efficient outcomes, an allocation being Pareto efficient does not imply that it is optimal. But an allocation not being Pareto optimal typically does suggest that it is sub-optimal, why wouldn't we want to make someone better off if it does nobody else any harm?

To compare different allocations, we could employ a *social welfare function* which is a way to combine individual's welfare into a measure of societal well-being.

[1] Total bullshit, of course. I will take the cash.

Question: Compare two social welfare functions: First, Welfare=You^2+Me^2. Second, Welfare=$\sqrt{You}+\sqrt{Me}$ Which allocation from $a-c$ is preferred under each welfare function? Which makes more sense to you?

Our simple graph of supply and demand allows us to examine the Pareto efficiency of the market for soda. The demand curve is the marginal utility from each unit. The difference between the demand curve and the price is thus the surplus utility obtained from purchasing that unit. Adding up all such differences, we obtain an area known as *consumer surplus*, this is the benefit to households from being able to participate in the market. This is what households would loe if the market for the good or service ceased to exist.

Graph: Consumer and Producer Surplus:

The supply curve is the cost needed to produce each additional unit. The difference between the price and cost is the surplus obtained by firms. Adding up, we obtain an area known as *producer surplus*. Producer surplus minus fixed costs are firm profits. We can think of producer and consumer surplus together as the gains of the market existing.

There is no way in this model to increase consumer surplus without decreasing producer surplus or vice-versa. The equilibrium is thus Pareto efficient. Anything that benefits consumers harms producer and vice-versa. In this case, no additional governmental policies are needed to ensure

optimality. The Scottish economist Adam Smith called this the "invisible hand" of the free market.

The notion of the invisible hand is formalized by the *first theorem of welfare economics*. It states that our basic model of supply and demand yields a Pareto efficient outcome. We cannot make someone better off without harming someone else.

Market Failures

Important: The first theorem of welfare economics does not imply that markets always deliver the best outcome. Rather, if the assumptions of the model are correct, then equilibrium is Pareto efficient. And remember that Pareto efficient does not mean best.

Market failures refer to cases where the assumptions of our model are violated in a way that leads to Pareto inefficiency. Many important economic policy debates boil down to whether market failures exist, how severe they are and, if a policy maker can try to fix them. We proceed with a non-exhaustive list:

1. Price floors/ceilings. Suppose that a government imposes a price floor where suppliers may not legally charge a higher price. Assuming that this is enforced and that the price floor is lower than the price that would occur absent it (the place where supply and demand intersect), a shortage will result where there is not enough supply to meet demand.

Graph: Price Floor

Notice that the sum of producer surplus and consumer surplus is reduced. The difference is called *deadweight loss* and represents the loss of efficiency.

It is possible that consumer surplus is higher because of the price floor. If consumers are better off, then how can this be an example of Pareto inefficiency? The answer is that removing the price floor, and then imposing a lump-sum tax on producers and transferring it to consumers, could make both sides of the market better off.

The opposite of a price floor is a *price ceiling* where sellers cannot charge above a mandated price. Like a price floor, an effective price ceiling creates deadweight loss. A price ceiling is effective only if it is set below the price that would otherwise emerge (again, where supply and demand intersect). A classic example of a price ceiling is rent control where landlords may not charge above a prescribed rent.

Current policy debate: As of July 2026, The Federal minimum wage is $7.25 and has not changed since 2009. This price floor affects relatively few workers, especially in state and municipalities with a higher minimum wage. In April 2026, some member of Congress proposed a $25 per hour Federal minimum wage. Were it to pass (it has no chance), it would benefit workers who keep their jobs. It would harm firms and workers who are no longer able to find work.

2. Externalities. We have assumed that the private costs and benefits (those experienced by the households purchasing the product and the firms producing it), equal the total costs and benefits. Suppose, instead, that there is an externality where the purchase of the product affects other people in the model, but that the buyer and seller do not care about this.

The classic example of an (negative) externality is pollution. Suppose that a paper mill is free to dump chemicals into public rivers. The private costs are those borne by the firm and determine its production decisions. The total costs, however, include the costs to clean up or tolerate yellow, foamy rivers that smell like shit. Firm will therefore produce too much pollution because they do not have to pay all the costs.

Figure 1: The Androscoggin River in Lewiston, 1973: Not Pareto Efficient

Graph: Supply and Demand

The solution to this type of externality is to either tax firms to align public and private costs or to restrict their production thorough a quota. This was the motivation behind the Clean Water Act which most agree was good policy.

Externalities may also be positive where production creates a non-private benefit. One example is production spillovers. Suppose that one firm's production produces innovation that benefits other producers in the industry. In this case, the total benefit exceeds the private benefit so that firms produce too little. Here, a common solution is a subsidy (the opposite of a tax).

Graph: Supply and Demand

Current policy debate: Climate change is a salient example of a negative externality. If my private jet emits a ton of carbon dioxide, it contributes to climate change and imposes a cost to the rest of society. One way to correct this externality would be to tax me. The problem is quantifying the cot to the rest of society. This cost os known a th e*social cot of carbon* and estimates vary widely. Under President Obama, the U.S. Environmental Protection Agency set the social cost of carbon at $43 per ton. This was used in some regulatory settings but not as part of a broad carbon tax. Under President Biden, the social cost of carbon rose to $190 per ton. President Trump has since lowered it to $4 per ton.

3. Lack of property rights. The model has implicitly (*i.e.* we did not list it) assumed that households can confidently use the goods and services they buy and that suppliers can confidently keep the profits that they make. Suppose instead that there was a high risk that either could be expropriated. In this case, the market might fail because agents would have to account for this risk. For this reason, most economists view protecting property rights as a fundamental purpose of government.

Current policy debate: The following map measures expropriation risk with darker shades indicating a higher risk.[2] Source: High risk countries are often politically unstable or those with a history

[2]Credendo Group

Figure 2: Expropriation Risk

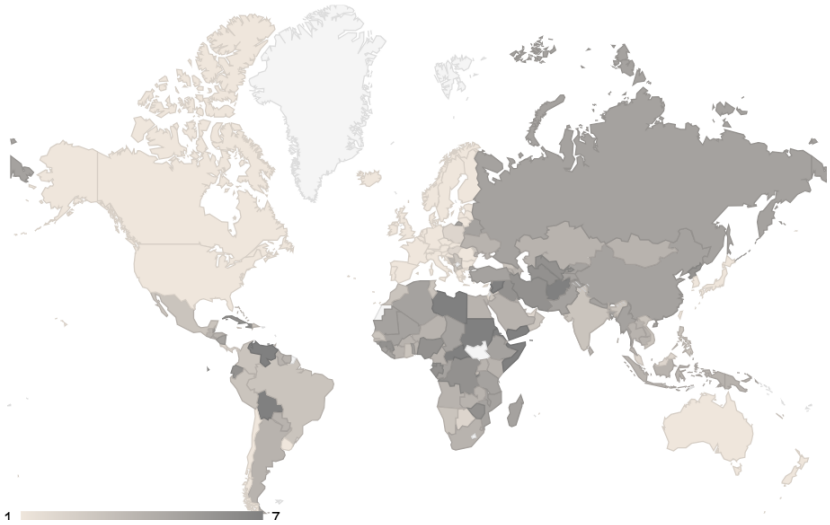

of or ideological inclination towards seizing private property. Note that these countries are often low income. The causation almost surely flows both ways, Wealthier countries are able to afford institutions such as courts that protect property rights. But weak property rates also cause poverty. 4. Public Goods. How much would you pay for national defense if you could choose? Most people would answer zero because anything you contribute is unlikely to make a discernable difference. The problem is that without intervention, everyone would make a similar choice, and there would be too little national defense. This is the basic public goods problem. In a free market, people choose to free ride.

A public good is one that exhibits non-rivalry of consumption: my consumption of it does not hinder your ability to consume it, and which is non-excludable, we cannot easily prevent someone from using it. Many goods meet these criteria imperfectly and there is often a grey area. But most economist consider other examples of public goods to include roads, infrastructure, clean air, etc.

The most common solution to the public goods problem is for a government to provide a public good. Of course, there is often disagreement over the right amount of a public good that should be provided.

5. Distortionary taxes. These are taxes which change peoples' behavior. Go back to the example of soda. Recently, some municipalities have enacted a tax paid by the purchasers of each unit of soda. This tax creates a gap between the amount paid by households and the amount received by suppliers from selling soda.

Graph: Supply and Demand

Instead of taxing households for buying soda, the government could have taxed the firms which are selling it. This would create a gap between the marginal cost paid by suppliers and the amount needed to actually produce the product.

Graph: Supply and Demand

Notice that it doesn't matter whether the tax is paid by the demander or the supplier. The price paid by the household is the same in both cases and the quantity is reduced by the same amount in both cases. Also note that the combined producer and consumer surplus region are reduced. The best known distortionary taxes are income taxes that collects a fraction of a worker's earnings. This tax changes their after tax wage rate and thus affects their behavior.

Current policy debate: Considering #4 and #5 collectively, it become clear tat a government cannot eliminate all market failures. Taxes are needed to provide public goods, although there is major disagreement about the appropriate level of funding for national defense, roads, etc. In addition, most people support the government providing some goods that are not public, such as education. This requires tax revenue. The only alternative to distortionary taxes, however, are lump sum taxes where people pay a fixed amount independent of income. It is not realistic to fund a modern government through such taxes alone. As a result, one of the most significant economic debates is on the role of the government in the economy. The fiscally conservative position is that the government should collect low levels of taxes while providing relatively low levels of public or private goods. Alternate positions call for combination of higher taxes/services.

6. Incomplete Information. We assumed that both suppliers and demander had good information about the market. But if some agents have incomplete or bad information, then inefficiency may result.

Consider the market for used cars. Suppose that suppliers are able to observe the true quality of cars but that potential buyers are not able to distinguish between good cars and crappy ones (lemons). Now suppose that a seller offers a buyer a car for $1000. If they cannot observe the true value of the car, they might extract from the seller's offer that the car is worth less than $1000. If the seller offers $500, they might infer that it is worth less than that. In extreme cases, the market breaks down altogether. This is known as *moral hazard*.

Another example is the market for car insurance. Insurance exists to spread risk. Everyone in the pool pays premiums and these premiums are used to pay claims made by the unlucky few who are in accidents.

Suppose there are 10 drivers who each have a $\frac{1}{10}$ chance of getting into an accident that will result in $1000 in damages. If each pays $100 in premiums then we expect that the $1000 will be able to pay for the damages.

But now suppose that there is one really awful driver who has a $\frac{1}{2}$ chance of getting into an accident. To cover the potential damages, the premiums will have to be higher and it is possible that the other drivers will choose not to pool with the bad driver. The market for insurance might fail. This is known as *adverse selection*. This simplified example motivates why most states mandate automotive insurance.

Current policy debate: Adverse selection is a major problem in health insurance markets where younger and healthier people often opt not to buy insurance. This makes the pool seeking insurance riskier, requiring higher premiums. The Affordable Care Act, passed in 2010, initially mandated that all Americans have health insurance (it provides subsidies for low-income households to buy it). This was an effort to fix adverse selection. In 2017, Congress passed legislation eliminating this mandate.

7. Market Power. We assumed that firms and households are too small to have any influence over the price. But what if there are barriers to entry, things that prevent additional firms or households from entering the market? In this case, either demand or supply could consist of only a small number of agents. If so, then these agent might be able to influence the price.

Agents who have market power exploit it by restricting quantity in order to affect the price in their favor. Market power can be had by either the demand side or supply side. On the supply side, examples include a monopolist who is the only supplier in a market, or a labor union which exploits market power in the labor market. Graphically, we see a higher price and lower quantity.
Graph: Supply and Demand

How exactly market participants exploit market power is left to the next topic.

An Algebraic Example:
We now consider a simple example. Consider the market for vaccines. The private demand for a vaccine is given by:

$$P = 10 - Q \tag{1}$$

Recall that demand is simply marginal utility. This (inverse) demand curve thus represents the benefits to a consumer of getting a vaccine. In this example, it is from reduced risk of they themselves.

Now assume that supply is simple $P = Q$. Solving for the private equilibrium is thus:

$$Q = 10 - Q \tag{2}$$

which is simply $P = 5$ and $Q = 5$.
Graph: Supply and Demand:

We can then calculate consumer and producer surplus. Both are triangle with a base of 5 and a height of 5. These area both equal 12.5 for a combined area of 25.

So far, this example only captures the benefits to the firms making the vaccines and the people receiving them. Now suppose that there is a *positive* externality from vaccines. By getting one,

a consumer makes it less likely that third-parties will become infected. Suppose that this public benefit raises the total marginal utilty for each unit by 10. Left uncorrected, the market does not consider this benefit. But suppose that it did. The total demand curve is then:

$$P = 20 - Q \tag{3}$$

Re-solving:

$$Q = 20 - Q \tag{4}$$

which is simply $P = 10$ and $Q = 10$. This is the efficient equilibrium.
Graph: Supply and Demand:

We can now re-calculate consumer and producer surplus. Both are now triangles with base and height equal to 10. Each area is thus 50 and a combined area of 100. This is 75 more than before. This difference is deadwight loss. It represents the loss to society because the private market does not consider the public benefit of vaccines.

How could a policy maker fix this externality? One way would be to simply pay consumers 10 units per vaccine. This would increase their marginal utility so that their demand curve becomes $P = 20 - Q$. This would make consumers behave a if their care about the 10 unit marginal benefit that their vaccine provides the third-parties.

Suppose that the policy maker instead pays vaccine manufacture 10 units for each vaccine sold. This makes the supply curve:

$$P = -10 + Q \tag{5}$$

Recall that supply is marginal cost. The marginal cost for the first unit is thus -9. This may seen odd, but keep in mind that it include the 10 unit payout from the policymaker. Solving for equilibrium:

$$-10 + Q = 10 - Q \tag{6}$$

which yields $P = 10$ and $Q = 10$. This is again the efficient equilibrium. Notice that it makes no difference whether the policy maker pay the subsidy to the producer or the consumer, the effect is the same. This is a common result.