Policy Goals

September 14, 2022

We have discussed a number of fundamental economic models. In each of those models, a primary focus is on equilibrium outcomes. To this point we have not discussed whether those outcomes are desirable from society's perspective, merely that the models are in equilibrium so they are not changing. We turn our attention now to studying some properties of these models, particularly efficiency and equity. The equilibrium concept will play an important role because we want to analyze the models when they are in equilibrium.

1 Efficiency

I have mentioned that the most important result across all the models in economics is that actions should be taken until the marginal benefit (in whatever form – marginal utility, marginal revenue, etc.) exceeds the marginal cost. If we evaluate an individual action, an economist would define taking the action as efficient if the marginal benefit is greater than or equal to the marginal cost. But the models we have discussed involve more than taking a single action, so we will evaluate them in different ways.[1] However, keep in mind at the core these definitions of efficiency come down to whether marginal benefits are greater than marginal costs.

1.1 Supply and Demand Model

We begin by returning to the supply and demand model. We will use the same equations from the second set of notes, $P = 5Q_S + 7$ for the inverse supply function and $P = 37 - 5Q_D$ for the inverse demand function. The green dashed lines in Figure 1shows the equilibrium quantity (3) and price ($22) in this market. Earlier we discussed reasons the curves would shift and how equilibrium price and quantity would change with each respective shift. Our focus now is analyzing supply and demand on the basis of efficiency.

From our model of consumer choice we related an individual's demand curve for a product to that individual's marginal utility. From the perspective of the market, a market demand curve could be viewed as a marginal benefit curve for society. Similarly, in our model of profit maximization, we discussed how a firm's marginal cost curve could be considered a supply curve, at least in perfectly competitive markets. From the perspective of the market, a market supply curve could be viewed as a marginal cost curve for society. At equilibrium, the price equilibrates supply and demand, and we can see that the demand curve is always higher than the supply curve until the equilibrium quantity is reached. Thus the equilibrium quantity is the point at which the marginal benefit and marginal cost are equal. Producing beyond the equilibrium quantity would mean that we are now producing where marginal cost exceeds marginal benefit. For instance, the price that consumers are willing to pay for the 4^{th} unit is $17, while the price that producers are willing to sell the 4^{th} unit at is $27.[2] It is clear that, from the perspective of marginal analysis, the 4^{th} unit (or unit 3.5, unit 3.1, unit 3.00001) should not be produced.

1.2 Gains from Trade

Another way of examining supply and demand is through the lens of gains from trade. If one individual values an item at $20 and another individual is willing to sell that item for $8, then there is $12 of surplus

[1] If you look up "economic efficiency" in a dictionary you will get different answers depending upon the particular question being asked.

[2] I obtained those prices by using the inverse supply and inverse demand functions from the example.

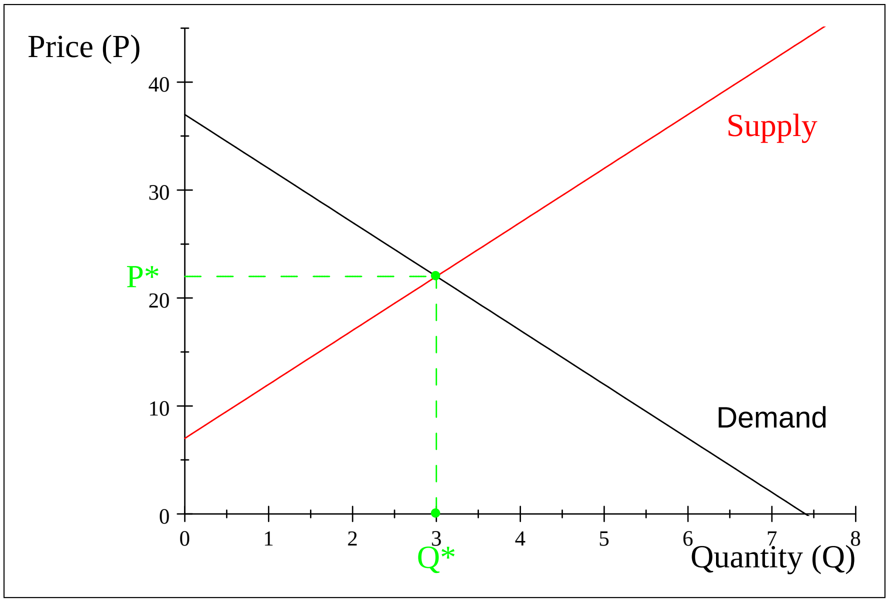

Figure 1: Determination of equilibrium price and quantity using supply and demand.

that can be created if the individuals are able to complete an exchange. If there is surplus then it would be efficient for an exchange to occur. For now we are not concerned with the distribution of surplus, only that there is surplus from a potential exchange. To find the gains from trade we find the difference between the demand and the supply curve, starting from a quantity of zero and until we reach the quantity traded, which is the equilibrium quantity when the market is in equilibrium. The shaded area in Figure 2 represents the gains from trade in the market. The dollar value of the gains from trade in this market is $45.[3]

Another result from the supply and demand model is that the gains from trade are maximized at the equilibrium price and quantity. Recall that the 4^{th} unit would cost an additional $27 to supply, but that consumers would only be willing to pay $17 for it. This trade could be made if we could shift some of the value created from trading the first three units around to the supplier and consumer of the 4^{th} unit, but we would need to redistribute $10 of the gains from trade from the first three units. At the equilibrium price and quantity we had a gains from trade of $45 – if we took $10 of that and redistributed it so that the 4^{th} unit could be traded we would now only have $35 in gains from trade.

We could have the highest valued user pay a higher price than the second highest valued user, and the second highest valued user pay a higher price than the third highest valued user, etc., until we reach the point where the final consumer is purchasing the item at a price equal to the lowest point on the supply curve. That process would allow us to maximize the *number of trades*, but there would be very little, if any, gains from trade from that process as consumers would essentially be paying their value for each unit of the item they consume. In addition to reduced gains from trade, that process would suffer from two other issues. First, we would need to know each consumer's value and each supplier's cost for each unit of the good. Second, we would need to match each consumer with the correct supplier to ensure that we are maximizing the number of trades. In the market setting, knowledge about individual values and costs is not needed nor is ensuring consumers and suppliers are appropriately matched to ensure that value is created from each exchange.

1.2.1 Consumer Surplus

In a voluntary exchange, both the consumer and producer should be receiving some gains from trade from the exchange. There is some misconception that only producers (firms) receive value from exchange in the

[3] For the purpose of creating an example it is helpful to have a numerical value for the gains from trade. The appendix of this set of notes discusses how to calculate the gains from trade using geometry and calculus.

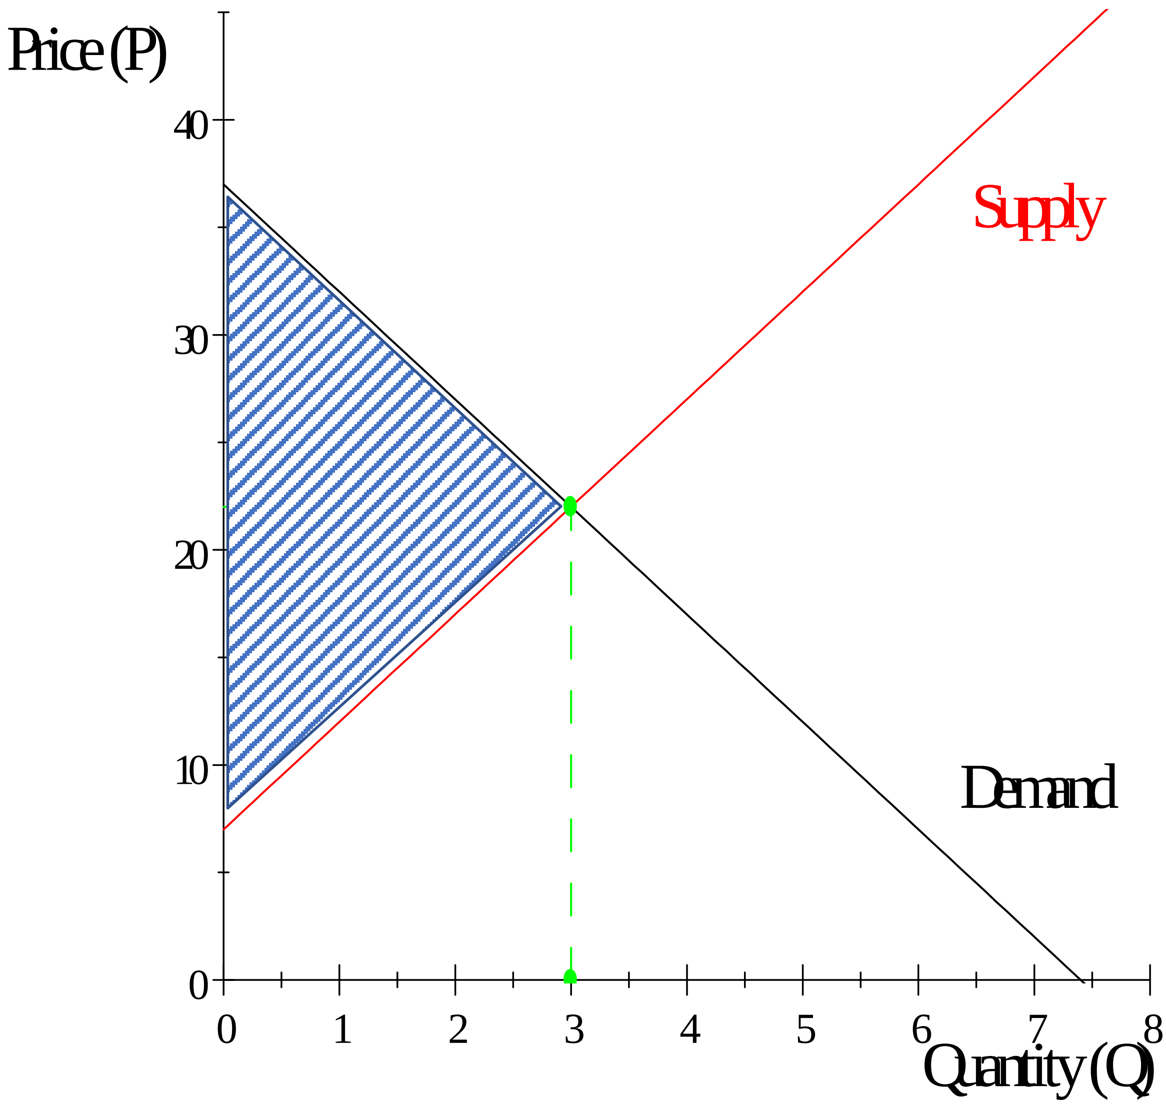

Figure 2: The shaded area illustrates the gains from trade in a market.

market system; that misconception occurs because society is typically conditioned to thinking about firms as wanting to make a profit and that profit is recorded in dollars. However, consumers also receive value from voluntary exchange but as it is not recorded in dollars that value is not as obvious as a firm's profit.

Consumer surplus refers to the value that a consumer receives in a voluntary exchange. As an example, if a consumer were to pay \$6 for a shirt for which the consumer would have paid \$22, that consumer has received \$16 in consumer surplus. The consumer was willing to pay \$22 but only paid \$6, so the consumer has received additional value. That value tends to be hidden because the consumer is not receiving money directly, but rather indirectly in the form of money not being spent on the shirt. In general, any time you think you "got a steal" or "got a great deal" on some purchase, you have received consumer surplus. When considering the entire market, consumer surplus would be defined as the area below the demand curve but above the equilibrium price, assuming that consumers all pay the same price. The shaded area in Figure 3 shows the consumer surplus in the example supply and demand framework we have been using. Numerically, the consumer surplus in this market is \$22.5.

1.2.2 Producer Surplus

Producer surplus is a similar concept to consumer surplus. If a producer was willing to sell a unit for \$9, but sold the unit for \$22, then the producer surplus would be \$13. While this calculation seems (and is) similar to calculating profit, producer surplus and profit are not the same. The primary difference is that producer surplus does not take into consideration fixed costs, whereas profit does. The producer surplus in a market is the area below the equilibrium price but above the supply curve. The shaded area in Figure 4represents the total producer surplus in the market. Numerically, the producer surplus in this market is \$22.5.

The gains from trade can also be defined as the sum of consumer and producer surplus in a market. In this market, the consumer surplus and producer surplus happen to be equal, but that does not have to be the case in all markets. For instance, suppose we held the demand curve in our example the same, and also held the intersection point of the supply and demand curves at \$22 and 3 units. But now rotate the y-intercept of the supply curve up or down so that the slope is flatter or steeper. By conducting that exercise, we can see that the resulting producer surplus will either be smaller or larger than the original producer surplus.

1.3 Profit Maximization Model

In our discussion of profit maximization by firms, we noted that the outcomes from the perfectly competitive market and the monopolist market differed. We begin with a discussion of efficiency concepts in the perfectly competitive market and then examine those concepts under monopoly.

1.3.1 Efficiency in Perfect Competition

In our initial discussion of perfect competition, our focus was on the individual firm. Figure 5 shows the MR, MC, and ATC for a representative firm in the market. Recall that the firm's demand curve is the same as its MR curve, so it is the perfectly horizontal red line, which also happens to be the price. Using our concepts of gains from trade, consumer surplus, and producer surplus, it seems like all the gains from trade in this market are producer surplus as the price and demand curve are the same. However, that analysis would be incorrect because Figure 5 represents an individual firm in the market, not the entire market. The individual firm's demand curve comes from an underlying market supply and demand model that looks like our standard supply and demand model where the equilibrium price is \$1.93 in the market. Thus, from the perspective of gains from trade, for a perfectly competitive market the results in Section 1.2 hold, and the gains from trade in the market are maximized.

Figure 5 also shows that a perfectly competitive firm in long run equilibrium will be producing the good at a price equal to the minimum of its average total cost. If a firm were to price any lower it would be making an economic loss, so a price equal to the minimum of a firm's ATC is the lowest possible sustainable price for this item. Allocative efficiency is defined as the allocation of resources to the production of goods most desired by consumers, at the lowest possible cost. In addition to the gains from trade being maximized in the perfectly competitive market, it also meets the standard of allocative efficiency.

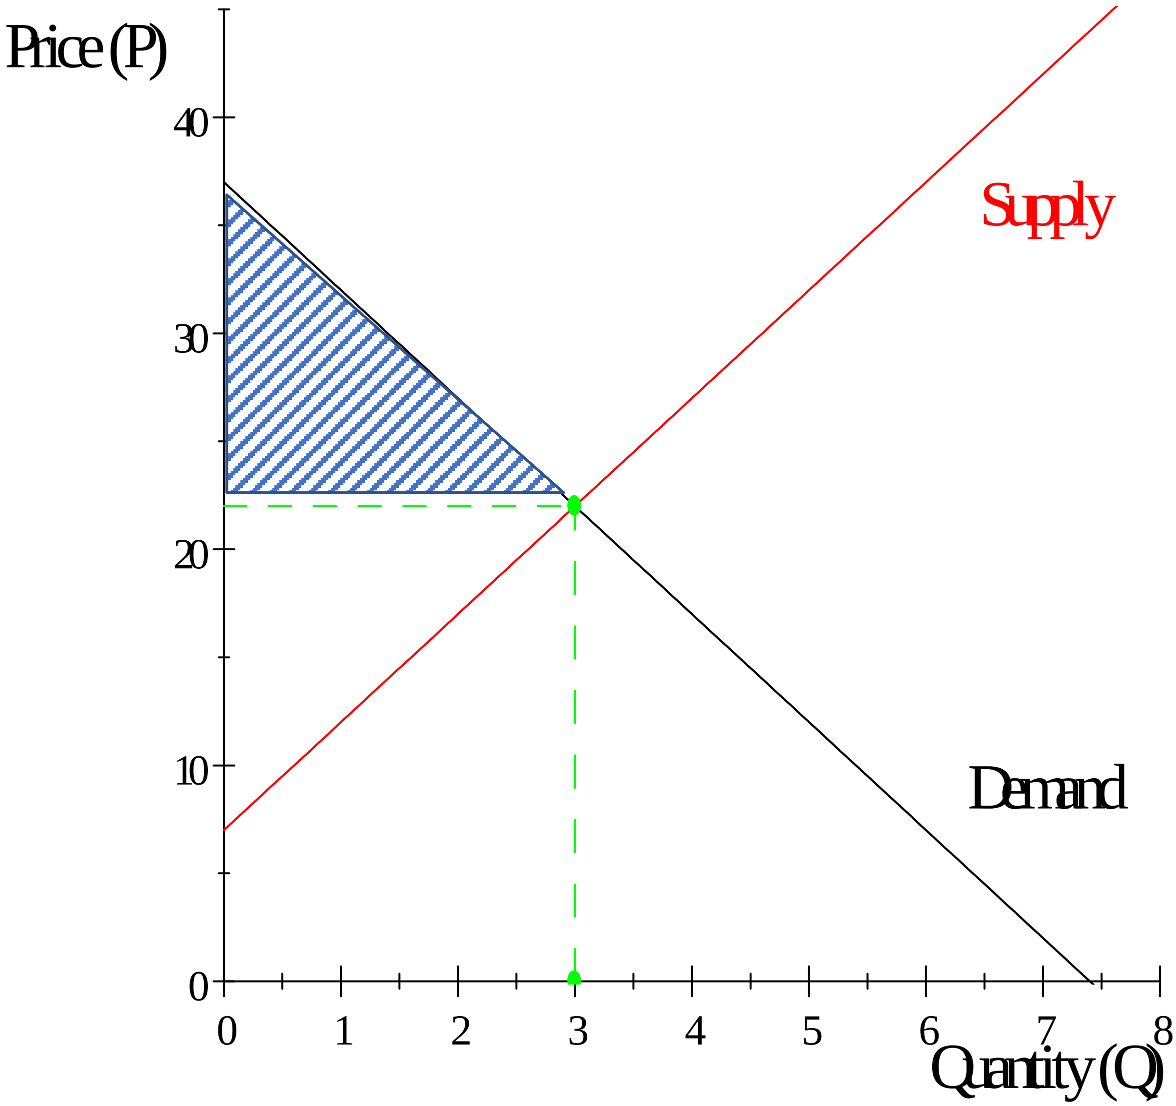

Figure 3: The shaded area represents the total consumer surplus in this market.

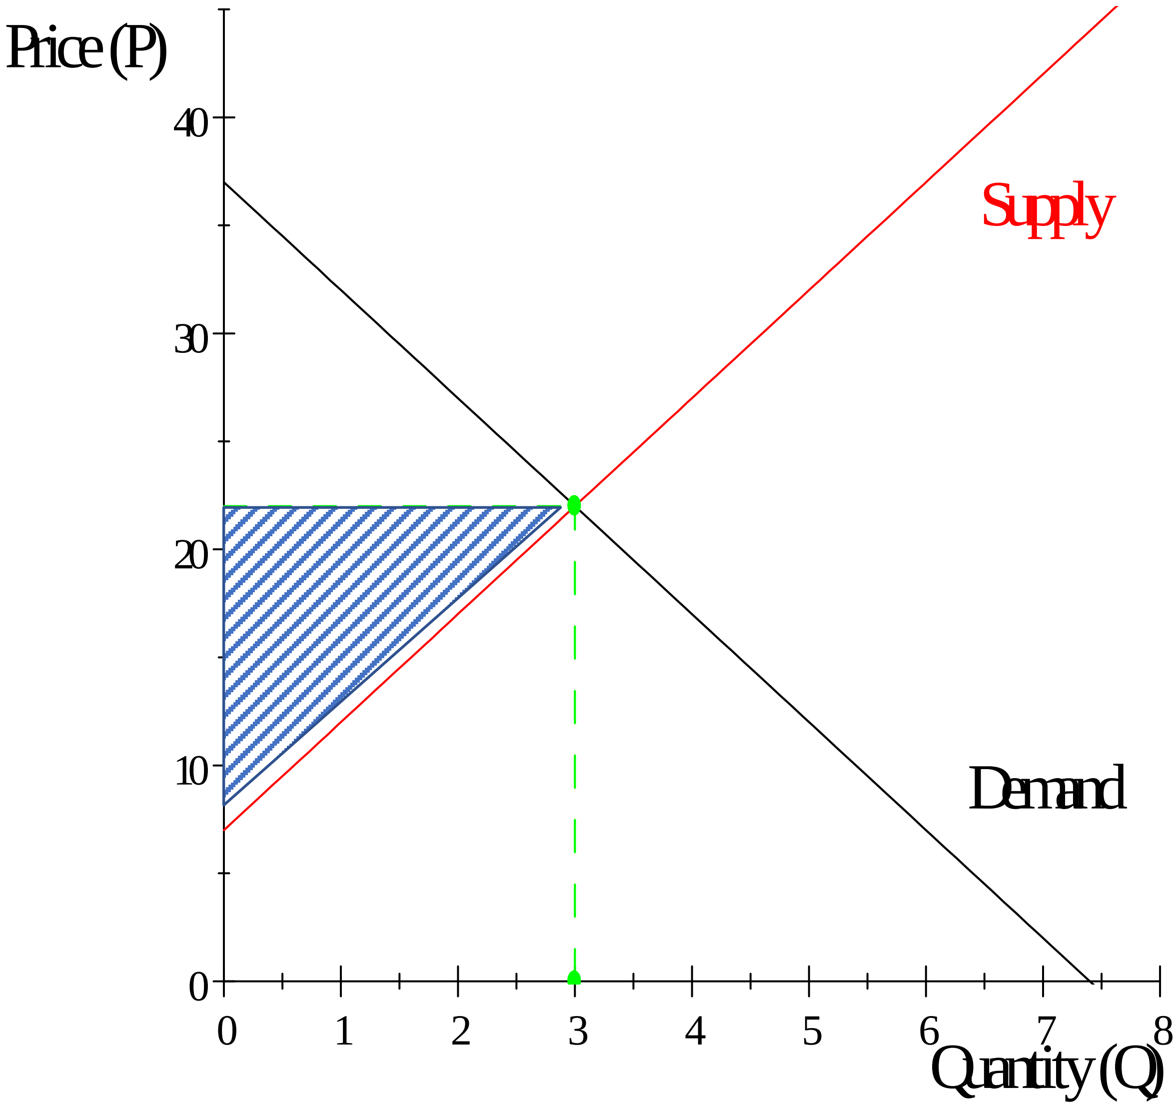

Figure 4: The shaded area represents the total producer surplus in the market.

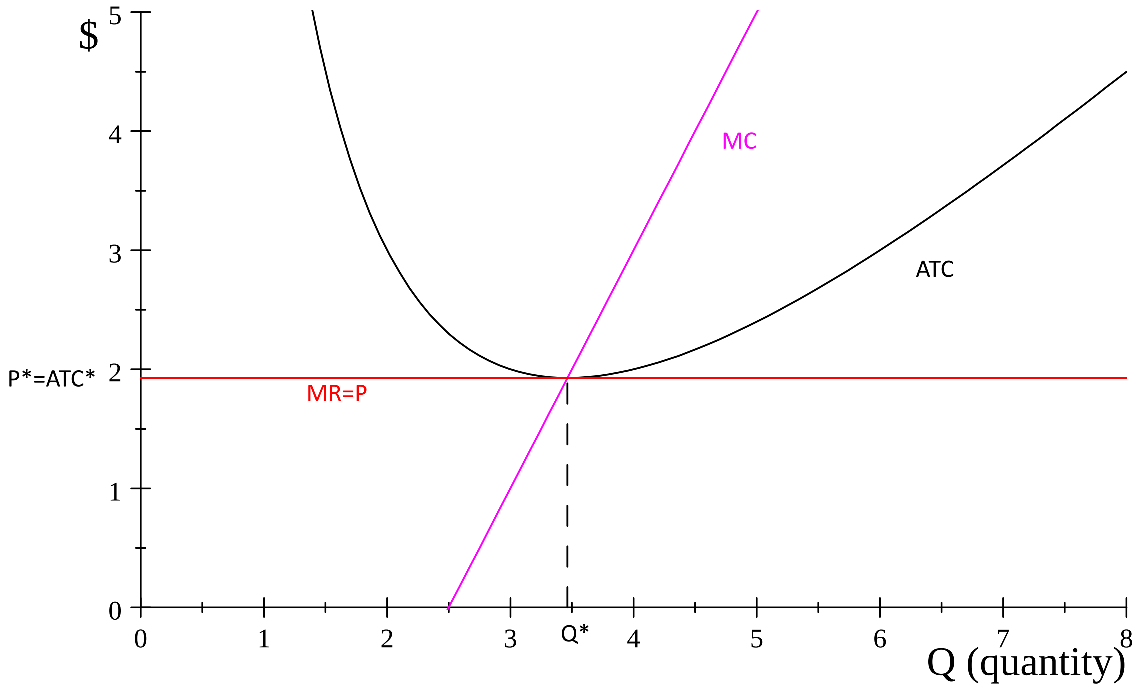

Figure 5: A firm in a perfectly competitive market earning zero economic profit. The ATC is in black, the MC is in purple, and the MR is in red.

1.3.2 Efficiency in Monopoly

Unlike the perfectly competitive market, there is no underlying supply and demand model that generates the market price for a monopoly; the market for the monopolist is also the market for society as a whole. Figure 6 shows the demand (D), MR, MC, and ATC for a monopoly, as well as the profit-maximizing price of \$63 and quantity of 17 units. Recall that the monopolist in Figure 6 is in long-run equilibrium, and that the monopolist is not choosing a price that is equal to the minimum of its ATC. Thus, the monopolist does not meet the standard for allocative efficiency.

We can also examine the monopolist market from the perspective of gains from trade. Recall that, for a monopolist, we can consider the monopolist's MC as a supply curve for society. Figure 7 shows a monopoly market with the consumer surplus shaded in green and the producer surplus shaded in red.[4] The definition for each term is the same as earlier, with consumer surplus being the area below the demand curve but above the price in the market, and producer surplus being the area below the price but above the supply (marginal cost) curve. Even with a monopolist, there are gains from trade.

But are the gains from trade maximized when the monopoly market is in equilibrium? No – from Figure 7 we can see that there are units that could be traded where the marginal benefit exceeds the marginal cost. When the monopolist chose its profit-maximizing quantity, it set $MR = MC$. In this example, we have $MR = 97 - 4Q$ and $MC = 2Q - 5$, which leads to $Q = 17$. The demand (D) in the model is given by $P = 97 - 2Q$, and we then found the price associated with $Q = 17$ from the demand curve, which is $P = 63$. However, trades are still possible until marginal cost exceeds marginal benefit, which means we would need

[4] Figure 7 also illustrates that producer surplus and profit for the monopolist are not the same.

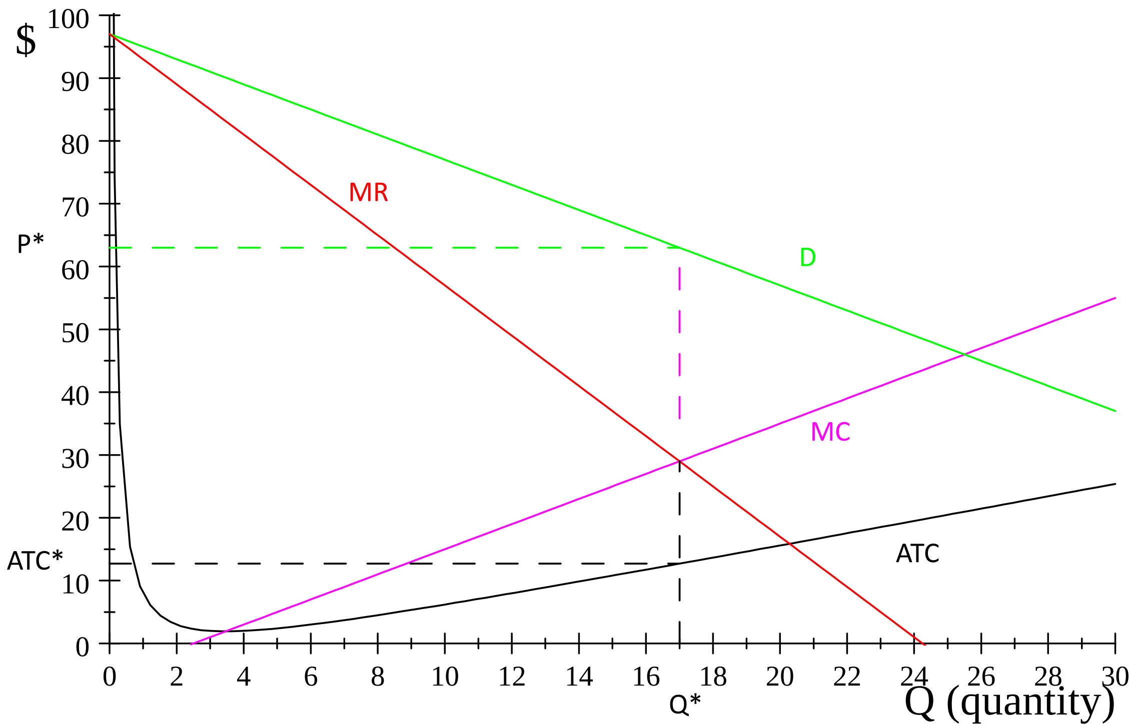

Figure 6: The demand (D), MR, MC, and ATC for a monopoly. The demand is in green, the MR is in red, the MC is in purple, and the ATC is in black.

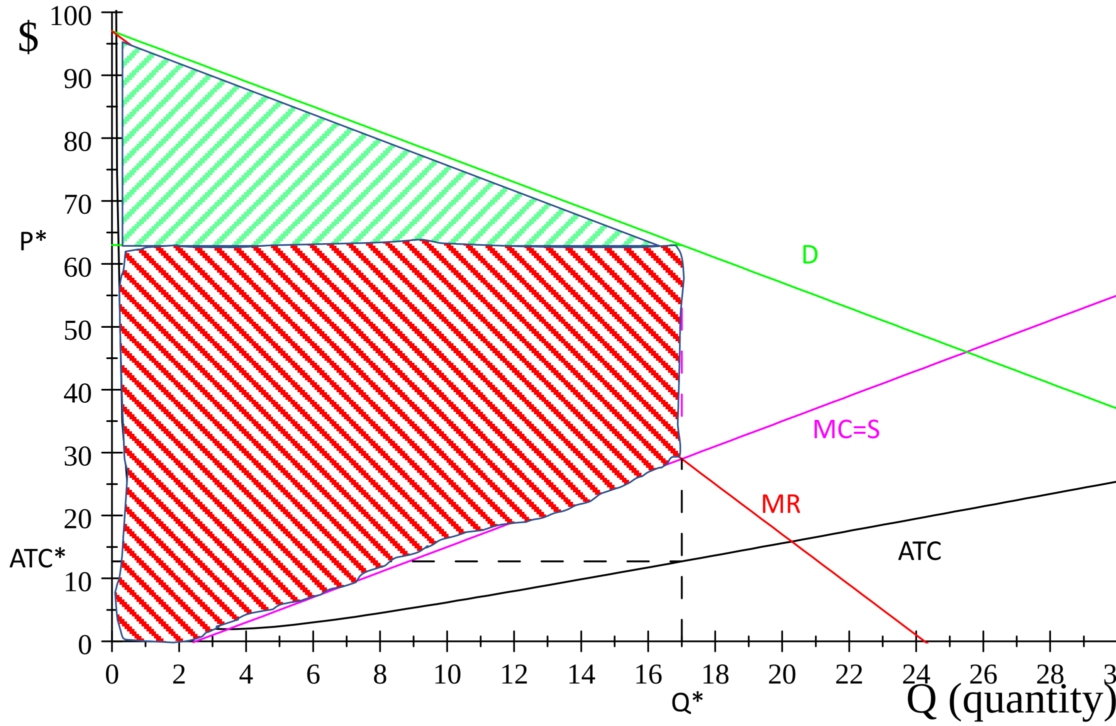

Figure 7: The gains from trade in a monopoly market. Consumer surplus is shaded in green and producer surplus is shaded in red.

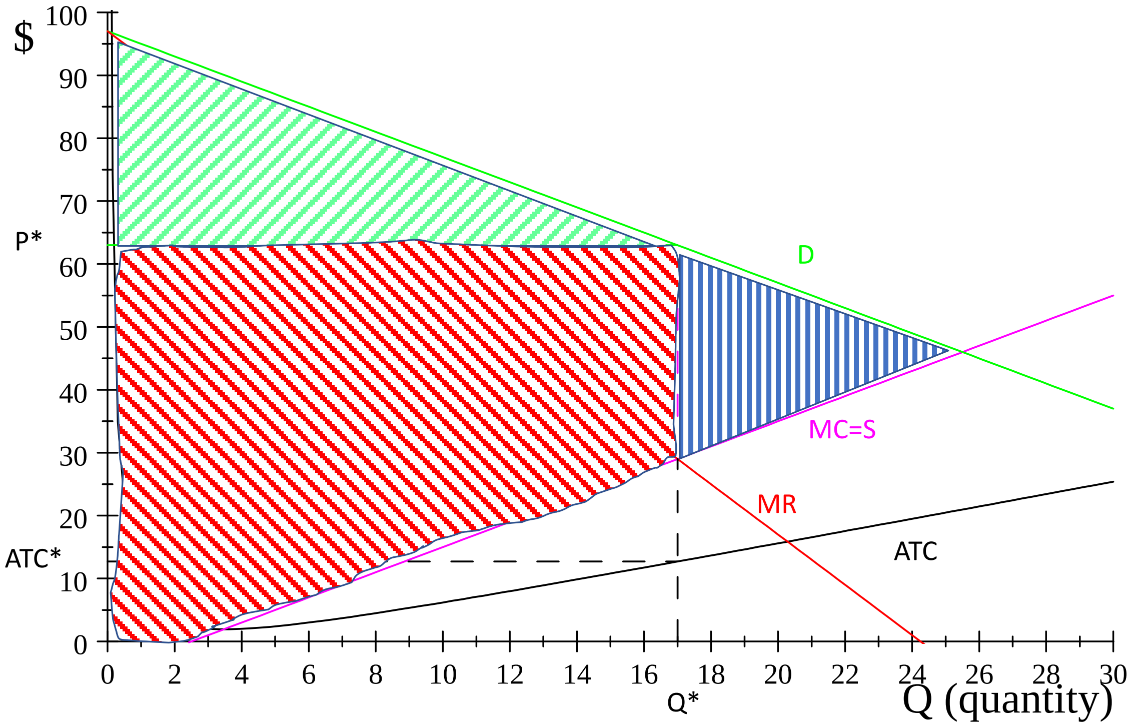

Figure 8: The gains from trade in a monopoly market. Consumer surplus is shaded in green and producer surplus is shaded in red. The deadweight loss is shaded in blue.

to find the quantity where $D = MC$. When we make that calculation, we find:

$$\begin{aligned} D &= MC \\ 97 - 2Q &= 2Q - 5 \\ 102 &= 4Q \\ 25.5 &= Q \end{aligned}$$

Returning to Figure 7, $Q = 25.5$ is consistent with the intersection of demand and the monopolist's MC. So there are 8.5 units that could be traded that would provide positive gains from trade. In the monopoly market, these units are not traded because the monopolist's marginal benefit (which is marginal revenue) is not the same as society's (which is demand).

Figure 8 again shows the gains from trade in the monopoly market, with consumer surplus shaded in green and producer surplus shaded in red, but now also shows the unrealized gains from trade shaded in blue. Economists use the term *deadweight loss* to refer to these unrealized gains from trade, because they are trades that could occur but do not. The concept of deadweight loss is not unique to a monopoly market and we will return to it periodically throughout the course.[5]

1.4 Consumer Choice Model

In our basic model of consumer choice we considered a single consumer who has a fixed income, knows prices of all goods in the market with certainty, and chooses a consumption bundle to maximize utility. At the

[5]The appendix to these notes explains the basic process for calculating the gains from trade in a market numerically. The same techniques can be used to calculate consumer surplus, producer surplus, and deadweight loss, though the range of quantity over which these numbers are calculated varies depending upon the calculation.

optimal consumption bundle, we know that the consumer spent all income and had chosen a consumption bundle on an indifference curve that was tangent to the budget constraint.[6] We want to expand that model to two consumers, with the understanding that the same basic results that occur for two consumers hold for $N > 2$ consumers.

1.4.1 General Equilibrium

When discussing supply and demand it was mentioned that the model was a partial equilibrium model; we are focused on the outcome in that particular market while ignoring any effects changes in the market under study have on other markets. The same is true of our model of individual consumer choice – we are only concerned about the choices that particular individual makes, and not how those decisions affect other consumers or other markets.

General equilibrium is used to provide a broader overview of how changes in one market affect other markets. Even an empirical general equilibrium model is unlikely to capture the effects of a change in one market on all markets,[7] but the idea is to provide a framework of whether this decentralized multiple market structure can lead to "efficient" outcomes. At a fundamental level, general equilibrium is intended to provide some structure to Adam Smith's invisible hand theorem in the *Wealth of Nations*. According to Smith, each individual in the economy, acting in self-interest, will maximize society's welfare through interdependent market actions. If individuals are free to consume and produce as they please, and if what they please provides value, the market will reward them; if not, the market will not reward them and they can choose other pursuits. There is no individual or organization mandating others consume certain goods or produce certain goods, but an invisible hand (market forces) that guides individuals to productive uses of their time.

Exchange Economy The initial structure of our general equilibrium model will be to focus on a pure exchange economy and then we will briefly discuss introducing production into the model. As with the other models in the course, the goal is not to solve these models, but to use them to generate intuition. We have two consumers, Consumer J and Consumer K, and there are two goods, Good 1 and Good 2, in the model. Consumers have preferences over the two goods that we can represent with a utility function, $u_i(q_1, q_2)$ for $i = J, K$. The consumers do not have income given to them, but they each have endowments of the two goods. Let e_1^J and e_2^J represent consumer J's endowment of Goods 1 and 2, respectively, and e_1^K and e_2^K represent consumer K's endowment of Goods 1 and 2, respectively. Thus the total endowment of Good 1 for society is $e_1 = e_1^J + e_1^K$ and the total endowment of Good 2 for society is $e_2 = e_2^J + e_2^K$. This endowment of the two goods is important for two reasons. First, it provides the consumer with an initial consumption bundle. Second, it provides the consumer with an initial income. Prices are not known, as they are in the individual consumer choice model, but the goods have value to the two consumers. Thus, if p_1 represents the price of Good 1 and p_2 represents the price of Good 2, then a consumer's initial income Y is: $Y = p_1 e_1^i + p_2 e_2^i$ where $i = J, K$. This problem has the same basic form, only now the consumers have to determine the rate of exchange for the two goods.

Edgeworth Boxes We can use the concept of an Edgeworth, or Edgeworth-Bowley, box to illustrate an equilibrium in this framework.[8] The basic framework focuses on the two consumer, fixed endowment economy. Begin by creating the Good 1 axis with an amount equal to e_1 and the Good 2 axis with an amount equal to e_2. We will let $e_1 = 30$ and $e_2 = 35$. Then draw the indifference curve for Consumer J that contains Consumer J's initial endowment. We will assume that Consumer J begins with 6 units of Good 1 and 24 units of Good 2. Figure 9 shows Consumer J's indifference curve through Consumer J's initial endowment.

Now do the same for Consumer K, rotate it 180 degrees, and place the origin at the point (e_2, e_1) on Consumer J's graph. It will be easier to understand looking at a picture. Figure 10shows the Edgeworth

[6] We assumed that there was an interior solution (both goods were consumed in non-zero quantities); if the solution was a corner solution (in which one good was not consumed at all), the result about tangency need not hold.

[7] We may be unconcerned with the effect of a change in the market for peanut butter on the market for wood picture frames because those are likely not related markets, but we may be much more concerned about the effect of a change in the lumber market on the market for wood picture frames.

[8] The general framework was developed by Francis Ysidro Edgeworth in 1881, further developed by Vilfredo Pareto in 1906, and popularized by Arthur Bowley in the 1920s.

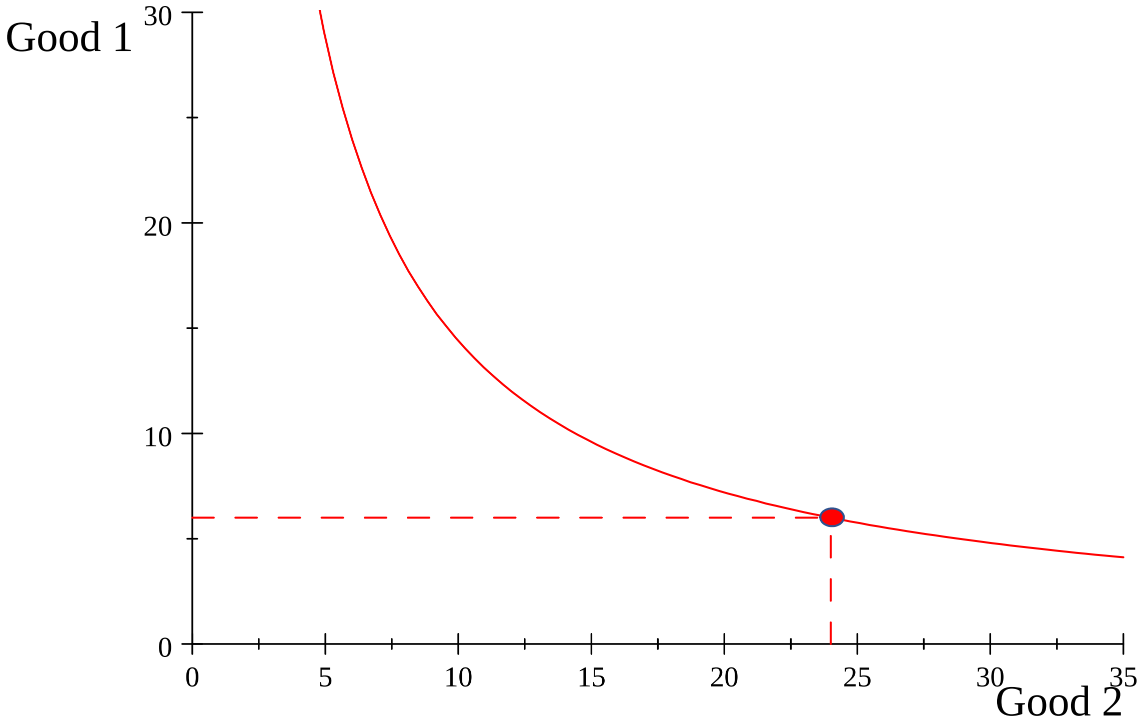

Figure 9: Consumer J's indifference curve through Consumer J's initial endowment.

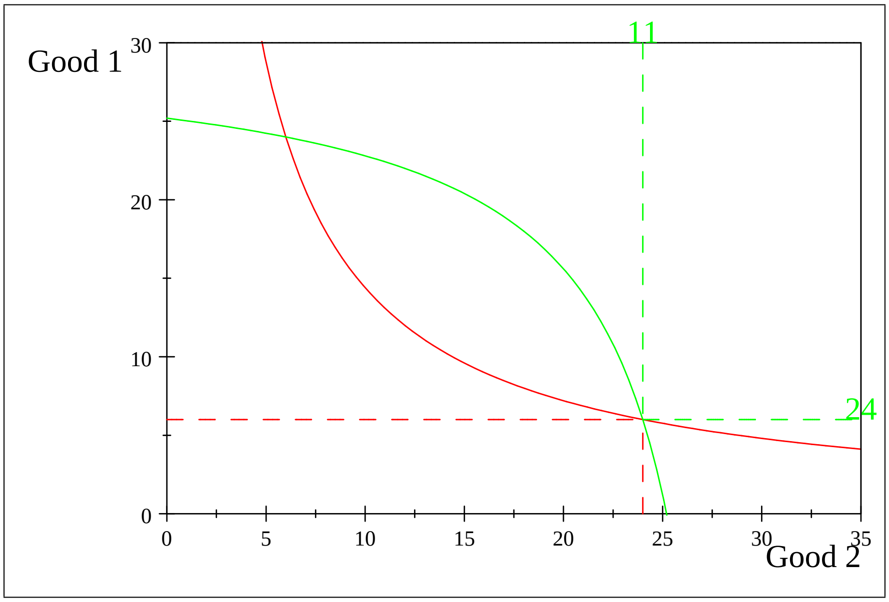

Figure 10: An Edgeworth box with the indifference curves for Consumer J and Consumer K at their initial endowment. Consumer J's indifference curve is in red and Consumer K's indifference curve is in green.

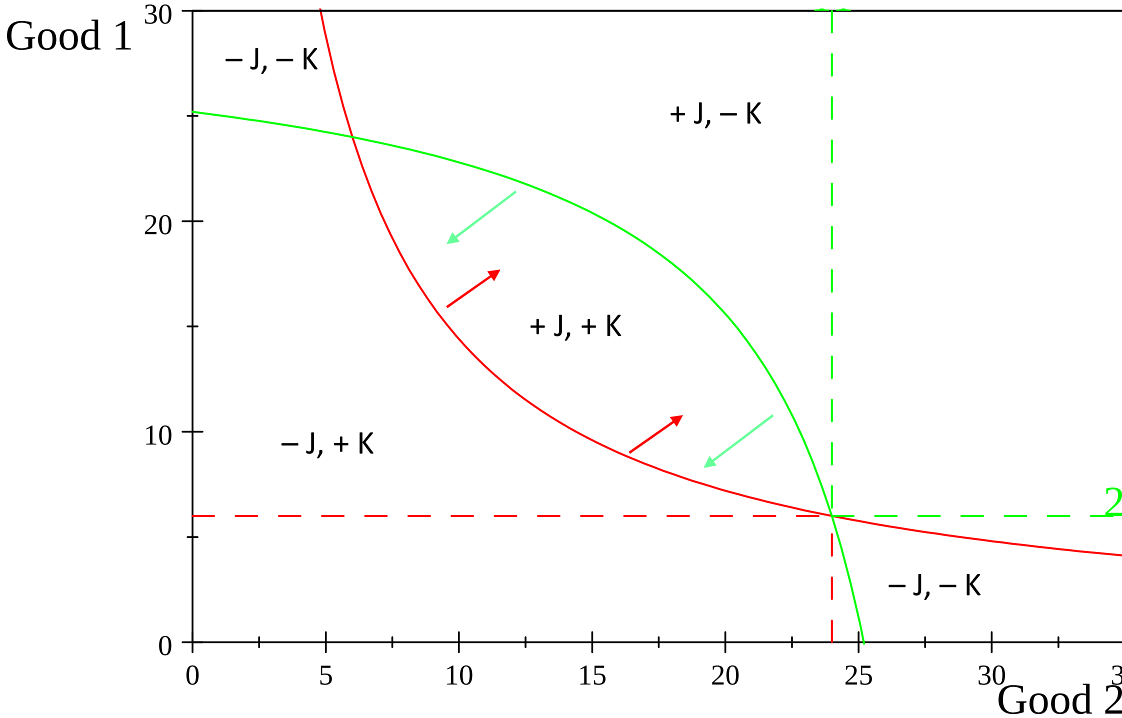

Figure 11: An Edgeworth box with labels representing areas of consumption bundles where both consumers are worse off, one consumer is better off and the other is worse off, and both consumers are better off.

box for Consumer J and Consumer K. One point of intersection of the indifference curves is the initial endowment. Consumer J, with the red indifference curve, has a standard orientation; Consumer K, with the green indifference curve, has the rotated orientation. Consumer J has an initial endowment of 6 of Good 1 and 24 of Good 2; Consumer K's initial endowment, 24 of Good 1 and 11 of Good 2, is labelled with their quantities running in the opposite direction.

Figure 11 shows the same initial starting point as Figure 10 but has a few additional notes. The direction that the indifference curves increase to show higher utility is shown by the red and green arrows, for Consumer J and Consumer K, respectively. There are also five areas labeled. Two areas are labeled $[-J, -K]$. These areas represent consumption bundles that would make both consumers worse off than their initial endowments. There is one area labeled $[-J, +K]$ and another area labeled $[+J, -K]$. These areas represent consumption bundles that make one consumer better off and one consumer worse off than their initial endowments. In the middle of the graph, between the two indifference curves, is an area labeled $[+J, +K]$. This area represents consumption bundles that make both consumers better off as bundles in this area will put both consumers on a higher indifference curve. This area where both consumers are better off is called the lens; if consumers are going to make voluntary exchanges, the outcome will be at some bundle in the lens.

The process of determining prices in this type of market is typically specified as one where an auctioneer calls out prices for the goods, receives supply and demand orders from the consumers, and then checks to see if the markets clear (meaning the quantity demanded for each good equals the quantity the consumers are willing to supply). If all markets clear at those prices, then those are set as the market prices and trades occur.[9] We will not work through the math, nor will we spend time discussing whether that auction process is representative of how markets work, but we do want to focus on the equilibrium outcome in the Edgeworth box. Figure 12shows the resulting equilibrium outcome. Note that the indifference curves for both consumers

[9] The auctioneer is sometimes called a Walrasian auctioneer, and the process is called tâtonnement.

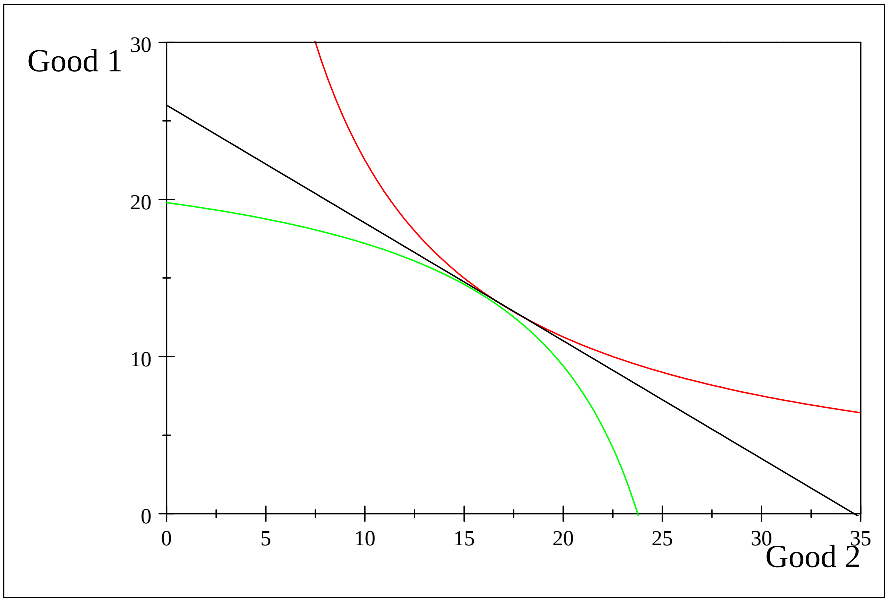

Figure 12: An Edgeworth box at an equilibrium allocation.

are tangent to each other, and the price ratio, given by the downward sloping black line, is also tangent to each indifference curve. At this point, there are no other trades that can be made that will make either consumer better off (in terms of shifting to a higher indifference curve) without making the other consumer worse off (shifting the other consumer to a lower indifference curve). An allocation with the property that no consumer can be made better off without making another consumer worse off is known as a *Pareto optimal* allocation.[10] Also, recall from our study of a single individual making an optimal choice that at the optimal bundle the slope of the budget constraint (which we called the Marginal Rate of Transformation) is equal to the slope of the indifference curve (Marginal Rate of Substitution). The same is true in the Edgeworth box, with the additional result that the Marginal Rate of Substitution for both consumers is equal.

While the focus has been on an exchange economy, we can add production (firms) to the model as well. If all goods are produced by firms in perfectly competitive markets, then we have a very similar result. There is a little more detail as some endowments will be used in the production process and consumers may own shares in some of the production processes that provide income to them, but the end result is similar. There is a set of prices at which each consumer is maximizing consumption given their preferences and initial endowment, each firm is maximizing profits given its production process, and quantity demanded at that set of prices equal quantity supplied. The resulting outcome is called a *competitive equilibrium*.

1.4.2 Fundamental Theorems of Welfare Economics

Now we turn to why the Edgeworth box analysis is useful in establishing some fundamental results of welfare economics.

First Fundamental Theorem of Welfare Economics: Any competitive equilibrium, which is defined as a set of prices at which each consumer is maximizing consumption given their preferences and initial endowment, each firm is maximizing profits given its production process, and quantity demanded at that set of prices equal quantity supplied, is Pareto optimal.

This result is a powerful one. It states that, provided certain assumptions are met (and they are unlikely to be met), if the market is allowed to operate without interference then the economy will reach a point at which no individual can be made better off without making someone else worse off. Importantly, there is no central authority or governing agency that mandates what should and should not be produced; if there is value to be created, individuals will produce the goods.

[10] You may also see this outcome referred to as Pareto efficient.

Now, Pareto optimality makes no mention of whether outcomes are equitable and in these economies there are often multiple equilibria. As a segue into the next section on equity, we have the second fundamental theorem of welfare economics.

Second Fundamental Theorem of Welfare Economics: Any Pareto optimal allocation can be reached through a competitive equilibrium process, provided that there are transfers of the endowments to individuals before the process begins.

This result is also a powerful one. It states that as a society we may not believe the initial distribution of endowments leads an equitable outcome. However, with the correct transfers of endowments, the competitive equilibrium process would lead us to the desired societal Pareto optimal allocation. Note that "transfers" here are a takings – some quantity of an endowment from one individual is taken and given to another individual. Also note that the second fundamental theorem relies on the same assumptions as the first fundamental theorem. There are times where individuals might (rightly) criticize the assumptions of the first fundamental theorem as unreasonable, but then invoke the second fundamental theorem as an argument for redistribution. But if one criticizes the assumptions used as the basis for the first fundamental theorem, one cannot then invoke the second theorem without criticizing those same assumptions.

2 Equity

While the First and Second Fundamental Theorems of Welfare Economics are important in establishing results that competitive equilibria are Pareto optimal, they provide no foundation for choosing one equilibrium over another. A Pareto optimal outcome could be that all individuals have equal amounts of all goods or it could be that all individuals but one are living at a subsistence level and the one individual has everything beyond what is needed for everyone else to be at subsistence level. Neither of those is likely to be an outcome desired by society; the latter because there is a very unequal distribution of goods, the former because that assumes everyone desires the exact same amount of each of the goods. How then do we, as a society, rank outcomes, regardless of whether they are Pareto optimal or not?

2.1 Social Welfare Functions

One method of ranking outcomes is to use a social welfare function. A social welfare function assigns a value, call it W, to an allocation of goods based on some underlying criteria of how individual utility levels map into societal welfare. There are three commonly used social welfare functions: utilitarian, weighted utilitarian, and Rawlsian.

Assume there are N individuals in society. A simple utilitarian approach would sum the utilities of all individuals in society:

$$W_{utilitarian} = U_1 + U_2 + ... + U_N$$

where the U_i are the utility levels of each individual. Whichever allocation yields the highest level of $W_{utilitarian}$ would be the preferred allocation.

A weighted utilitarian social welfare function is slightly more sophisticated as it sums weighted utilities of individuals:

$$W_{weighted} = \alpha_1 U_1 + \alpha_2 U_2 + ... + \alpha_N U_N$$

where the α_i are the weights assigned to each individual's utility. Whichever allocation yields the highest level of $W_{weighted}$ would be the preferred allocation.

A Rawlsian social welfare function defines the social welfare of society as the minimum utility level of any individual in society:

$$W_{Rawlsian} = \min \{U_1, U_2, ..., U_N\}$$

Under this social welfare function, society's welfare is determined by the least well-off individual in society. The idea is based John Rawls and his veil of ignorance concept, which asks what social welfare function an individual would choose if the individual was completely ignorant of all other information about society and what position they would have in that society.

There are challenges with all these approaches. One challenge underlying all three approaches is that we initially defined utility as an ordinal measure, such that the utility number itself did not matter, and only the ranking of one allocation of goods to another was important. We used $U(Q_1, Q_2) = \sqrt{Q_1 Q_2}$ as our utility function in earlier notes. Holding income and prices of Goods 1 and 2 constant, a consumer with that utility function will choose the same optimal bundle as a consumer with $\widehat{U}(Q_1, Q_2) = 10\sqrt{Q_1 Q_2}$ or $\widetilde{U}(Q_1, Q_2) = \sqrt{Q_1 Q_2} + 92.4$ or $\overline{U}(Q_1, Q_2) = \sqrt{7 Q_1 Q_2} + 14$. When we work through the math, the indifference curves will all be the same – the number associated with each indifference curve will be different, but the indifference curves themselves will be the same.[11] With a social welfare function, the actual utility values are important, so if we want to make comparisons across individuals we need to turn utility into a cardinal measure.

The utilitarian approach may not lead to an equitable distribution of goods, as goods will likely be transferred to those who receive more utility from consuming them. With a weighted utilitarian approach, the first question to ask is: Who determines the weights? The Rawlsian approach should lead to an outcome that is fairly egalitarian, but it assumes an individual is extremely risk averse. If individuals maximize expected utility, then the choice from behind the veil of ignorance could lead to a social welfare function that has a wider distribution of outcomes than one that maximizes the minimum possible utility.

The difficulty with the social welfare function approach is that it relies on value judgements – which individual's or groups of individuals' utilities are more important than others? Depending upon who makes that decision, the outcome could be egalitarian or it could be very inegalitarian.

2.1.1 Arrow's Impossibility Theorem

Another critique of the social welfare function approach comes from Kenneth Arrow, who proved[12] that it is impossible to have a social welfare function that satisfies the following four criteria if there are more than two possible allocations:

- Unrestricted Domain: Individuals can have whatever preferences they want over the allocations, provided they are complete and transitive.
- Weak Pareto: If all individuals prefer allocation a to allocation b, then the social welfare function should prefer allocation a to allocation b.
- Independence of Irrelevant Alternatives: If society prefers allocation a to allocation b, then introducing allocation c should not affect society's preference of allocation a over allocation b.
- Non-dictatorship: Society's choices do not reflect those of one single individual for every single social choice, regardless of the views of everyone else.

We will not work through the mathematical proof, but these seem like reasonable criteria. If there are only two possible allocations (or candidates, if we want to apply this theorem to voting), then we can find a social welfare function that satisfies these criteria. If we only want to consider three criteria, then we can find a social welfare function that satisfies three of the criteria but not the fourth.

2.2 Other Considerations

There are other equity considerations beyond the social welfare function. One is to return to the concept of gains from trade. Gains from trade are, at a minimum, denominated in dollars, which are more tangible than utility. If we focus on maximizing the gains from trade, we could then examine possibilities for redistributing those gains. One could be to attempt to equalize consumer and producer surplus in a market. Another, which we will discuss in more detail later in the course, could be to have individuals who receive a utility or monetary improvement as a result of some change in the market (possibly due to a policy change) compensate individuals who have a utility or monetary reduction.

[11] Technically, utility functions are identical up to any positive monotonic transformation, which means as long as we preserve the ordering of the preferences over the bundles of goods the utility functions represent the same preferences.

[12] The word "proved" is a really strong word which I reserve for mathematical proofs. I will not use the word "proved" if someone conducts an empirical study and finds a result. "Shows" or "finds" is a much better word for empirical studies that more accurately conveys the results of the work.

Our focus has primarily been on outcomes, not process. Equitable access to markets is important. If individuals are excluded from markets, they cannot share in the gains from trade. Allowing individuals to participate in markets and share in the gains from trade is one method of redistributing the gains from trade that does not require any individual or agency to impose a redistribution requirement.

3 Criticisms

Many results established in the general equilibrium framework rely on what some view as unrealistic assumptions (and some of them are). However, these results and the assumptions need to be considered in the context of the time at which they were developed. The efficiency properties of market processes had not been established, so researchers were searching for some minimal assumptions that would generate equilibrium outcomes without a central authority dictating what needed to be produced and consumed. There was (and still is) a debate over whether a market-based economy or a command economy would be better for society. These basic models were established during that time and are best considered as a benchmark economy. They have been modified by countless researchers throughout the years, looking to make the models more closely resemble reality.

I have already provided some criticisms of the social welfare function approach. Constructing a social welfare function is a difficult enough task for a set of immediate family members or close friends; it is an incredibly difficult task to make judgements about individual's utilities twenty miles away, let alone two hundred or two thousand miles away.

4 Appendix

4.1 Calculating gains from trade

Figure 2 illustrates the gains from trade in a market. With the linear supply and demand functions, we can calculate the gains from trade using geometry. Figure 13 is just Figure 2 rotated so that the price axis is on the bottom. I have done this rotation because it is easy to see that the gains from trade is simply a triangle when supply and demand are linear. The height of the triangle is given by the equilibrium quantity (3), and the base of the triangle is the difference between the price intercept for the demand curve and the price intercept for the supply curve. With $P = 5Q_S + 7$ the inverse supply function and $P = 37 - 5Q_D$ the inverse demand function, those numbers are 37 and 7, respectively, so the base is 30. Then we use the area of a triangle to calculate the gains from trade, which is $\frac{1}{2} base * height = \frac{1}{2} * 30 * 3 = 45$.

If supply and demand are curved, and not linear, then we would need to use calculus to find the gains from trade. We would need to find the area between the demand and supply function between 0 and Q^*, so we would need the following integral:

$$Gains = \int_0^{Q^*} \left(D\left(Q\right) - S\left(Q\right)\right) dQ$$

where $D(Q)$ is the demand for the good as a function of quantity and $S(Q)$ is the supply of the good as a function of quantity. We can use this method for our linear functions as well. Finding $D(Q) - S(Q)$ first:

$$\begin{aligned} D\left(Q\right) - S\left(Q\right) &= 37 - 5Q - (5Q + 7) \\ D\left(Q\right) - S\left(Q\right) &= 30 - 10Q \end{aligned}$$

and substituting in $Q^* = 3$, we now have:

$$Gains = \int_0^{3} (30 - 10Q)\, dQ$$

Solving the integral[13] we get:

$$Gains = 30Q - 5Q^2 + C|_0^3$$

[13] To solve the integral you want to find a function such that the derivative of that function is the term in the integral. We want to find a function such that its derivative is $30 - 10Q$. A function that meets that criteria is $30Q - 5Q^2 + C$, where C is a constant of integration. That C could be any number as the derivative of a constant is zero.

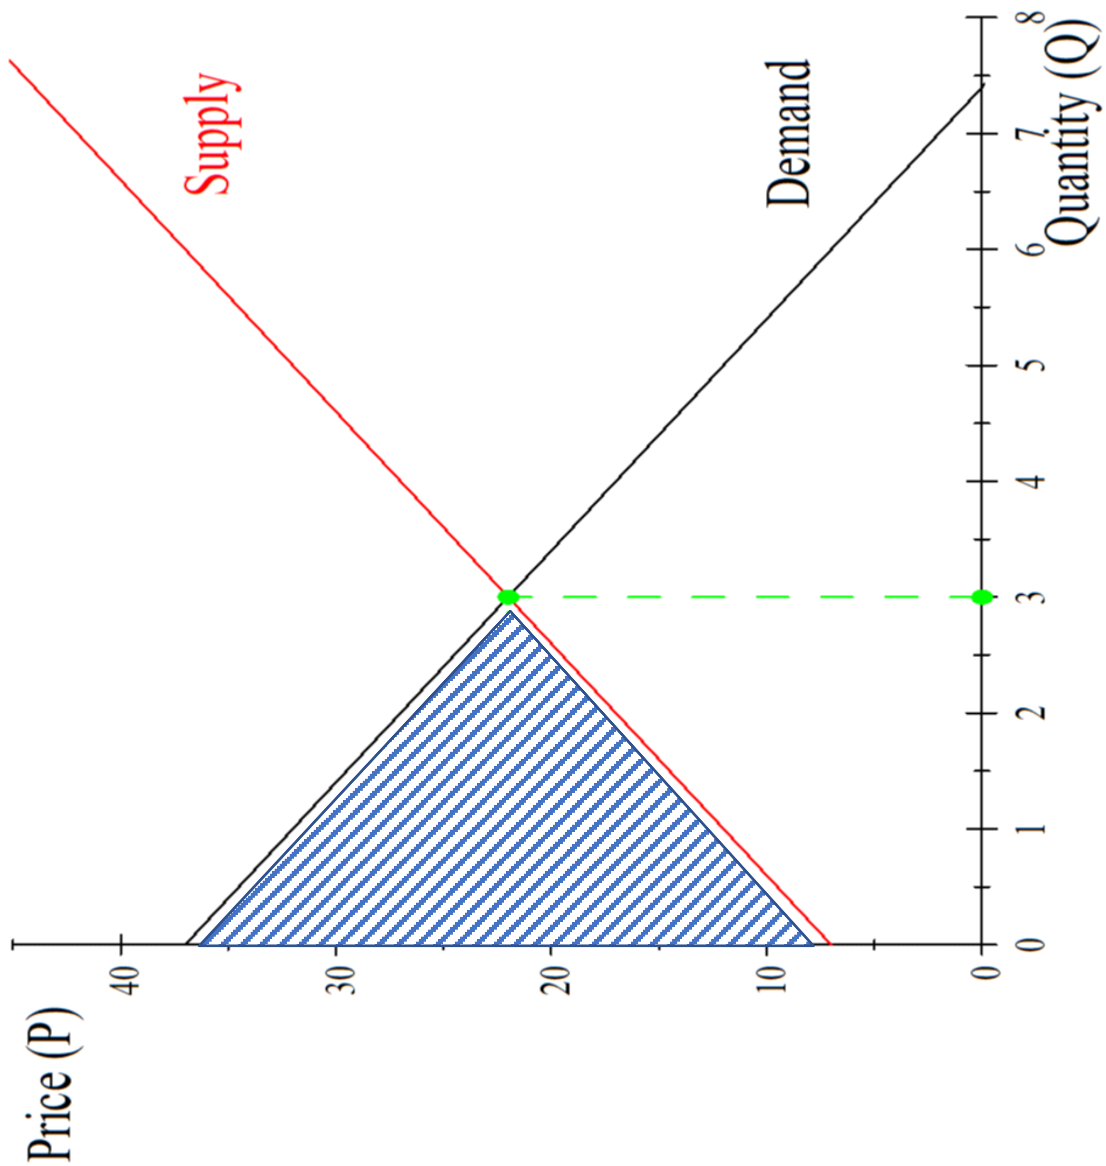

Figure 13: A supply and demand picture with the price axis on the bottom. Used to clearly show that the gains from trade, when supply and demand are linear, is a triangle.

where C is just a constant of integration. First we substitute in $Q = 3$ and calculate that number, which is $45 + C$. Then we substitute in $Q = 0$ and calculate that number, which is C. Then we find the difference between the two, which is $45 + C - C = 45$. It is (and should be) the same answer from the geometric calculation.

I do not expect you all to calculate integrals (it is highly unlikely I will ask you to calculate the area of the triangle, let alone an integral), but I want to at least mention why and how they are used, particularly for those students who may not have a strong calculus background. In reading economics papers you will eventually come across papers that use integrals, and you should be aware that, at their core, integrals are just a method for calculating an area under a curve or for calculating a sum.