Principles of Macroeconomics–Honors

PPF & Trade Notes

Instructor: Dr. Kevin Kuruc

January 27, 2021

1 Production Functions, PPFs & Gains From Trade

Economics is the study of allocating (and producing) resources. When deciding how to use some resource, we necessarily face trade-offs because no resource is ever in infinite supply. The class begins by studying a basic model meant to formalize our thinking about such trade-offs, or **opportunity costs**, of different goods/services. We'll then see how the cost of obtaining certain goods can fall through the magic of exchange between economies.

1.1 Single Economy PPFs

To formalize these concepts, we're going to introduce our first **economic model** of the course.

Economic Models

An economic model is a simplification of a complex economic phenomenon into a mathematical (equation or graphical) structure. Making such simplifications ensures that the model will be "wrong," but turns the phenomenon into something that can be usefully analyzed.
Further reading: https://www.imf.org/external/pubs/ft/fandd/basics/models.htm.

The models we introduce here are called production *functions* and production *possibilities frontiers*. Why is production going to be the primary focus at the start? *Because consumption—broadly defined as all the things that make us happy—is restricted by production.* That's your first big lesson of macroeconomics and it will come up repeatedly throughout the course. Why is this true? Well, if something is never made, it can't ever be used or consumed! And if we want to buy goods from other countries, we're going to have to send them something in return. So even when our consumption is not strictly limited to our production, we can stick with the rule of thumb that $C \leq Y$; the value of consumption, C, is limited by the value of output, Y.

The concept of a production possibility frontier (PPF) is a way to graph our production options—sometimes called our "choice set"—as a society.

Production Possibility Frontier (PPF)

A graphical representation of the combinations of goods that an economy can produce.

By illustrating the options the economy can choose between, we can infer the trade-offs faced by that economy. As an example, imagine there is a factory that can produce cars (also C here; sorry, notation will have to get re-used a lot throughout the course) and trucks (T). A car can be produced in 1 hour, a truck in 2. In Figure 1 we see the graph of the possible combinations of output in a 24 hour day.

Figure 1: Car-Truck PPF

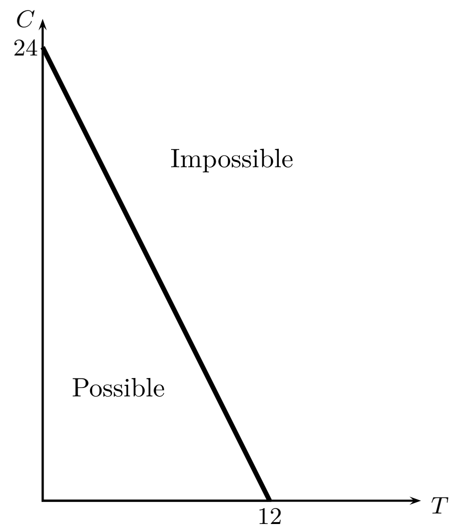

These points are plotted by asking "if we wanted to make T trucks, how many cars are attainable?" The point reaching up to 24 on the C-axis indicates that when 0 T are produced, 24 C are possible: since all 24 hours are used to make cars, and cars are produced at a rate of 1 per hour, we can have 24 C when we spend no time making T. Alternatively, if we make 12 T, we can't make any C since all 24 hours are used making 12 T. **Exercise: Verify that the point (4, 16) is on that line.** Any combination on or below that line is possible. The points within the "Possible" region, but *not* on the line, are considered inefficient; it is possible to make 3 T and 3 C. But if that point was chosen more C could be produced without sacrificing any T. Clearly, there is unused capacity at points within that set. Outside of that line are points that require the plant operate more than 24 hours a day, this is why it is labeled Impossible. The points that fall on that line are what we call **efficient**. All productive capacity is used making C and/or T.

The PPF allows us to see the trade-offs, or **opportunity costs**, of these goods.

Opportunity Cost

Whatever is sacrificed to obtain some good, service, or other resource.

An example will help illustrate the point. Looking at that graph, you might ask what a T "costs." Well, we have no information about the monetary costs of this factory. What we do know is that we're running this factory for 24 hours a day, and therefore if we want more T, *we have to give up some* C. In fact, we can say precisely that the **opportunity cost** of the first truck is 2 C: if we want to move rightward from (0,24) we go to (1,22), a loss of 2 C. This factory could be a solar-powered, completely automated machine with no monetary costs to run, and still, producing that T is not free. Something is given up to produce it: the opportunity to produce another good.

So we've established that the first T costs 2 C; how about the 6th? It turns out, the 6th also costs 2 C.[1] Every time we want to add a T, we give up 2 C, because we shift 1 hour of production towards T. The opportunity cost is constant in this example. For those of you who remember your algebra, what else is constant for a line like we have plotted above? The slope! That's right, the opportunity cost of the good on the x-axis is the slope of the PPF at that point. Its how much you need to reduce whatever is on the y-axis to add one more of whatever is on the x-axis. What do we think the opportunity cost of a C is here? It's going to be $\frac{1}{2}T$ because a car only takes 1 hour to make, which means we only forego half of a T. The opportunity cost of the good on the y-axis is the reciprocal of the slope (remember: reciprocal means flip the fraction).

This car-truck example is fairly simple because we've got a linear PPF which implies there are constant opportunity costs. This is not usually the case. Imagine the city of Norman provides you data for points on what they believe is their PPF. The city only produces fancy restaurant meals M and smartphone apps A. The following combination of output are possible based on the data provided.

The PPF here is not a line, and the trade-off is not constant between these goods. In particular, the first app we want to produce only costs us 5 meals (on the graph this is the far upper-left corner—notice that the slope is flatter there). When we're moving from 9 to 10 apps we are giving up 23 meals. Why would this shape arise?

Specialization of inputs is why economy-wide PPFs usually have a bowed shape. Specialization means that those inputs which are relatively better at producing some good will be the first resources we employ in that activity. In this example, we know that the city of Norman is made up of many young, college-educated individuals who are creative, talented, and know what would make for a good smartphone app—but are terrible at cooking. The other residents of Norman are older folks who have fantastic cooking skills but no idea how to make smartphone apps. Imagine we're starting at the far upper-left point on that PPF; the

[1] Notice I asked about the 6th T, not 6 T; we typically care about the costs of individual goods in economics.

Figure 2: Norman PPF for Apps and Meals

Meals	Apps
100	0
95	1
70	5
40	8
23	9
0	10

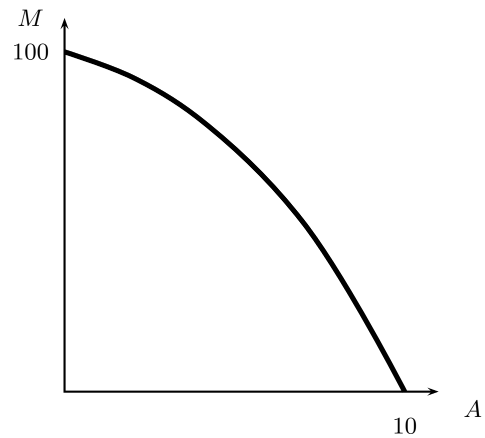

city is only producing meals here. All of Norman's resources are used making meals which means that both the old folks of Norman and the college students with no idea how to cook are working in this sector. If we wanted to move from that far upper-left point slightly to the right so that Norman produces 1 app, who would we move from cooking → coding? It would be the individuals who are the best at coding relative to their cooking skills—we'd find the computer science student that subsists on ramen noodles and move him to app making. With this reallocation, almost nothing is given up in the kitchen—this person is terrible at cooking anyway—but a lot would be gained in app production because they're great at coding! As we keep moving along there are less and less people fitting this description. By the time we get near 10 apps, the only people left in the kitchen are skilled cooks without any idea how to make an app. If we move any of them to app making, we gain very little in app production but lose a skilled cook in the kitchen. Its very costly (in terms of how many meals we give up) to make those last apps!

Suppose now that as time goes on the University realizes that cooking and app creation are what Norman thrives on, so it decides to offer a mandatory cooking class as part of its general education requirements. How would this affect the PPF, and hence Norman's choice set? Norman residents have become more productive in the cooking sector, but no more productive in app creation. To talk yourself through this, it will help to look at the extreme points and just fill in the general shape. The point (0,100) would shift: if we are all better at cooking, and we spend all of our time cooking, we'll produce more meals. Let's say our capability improves to 120 meals for the sake of this example. The point (10,0) would **not** change: we've gotten no better at app creation, so if we dedicate all of our times to apps we will still only make 10. Figure 3 demonstrates such a shift.

In these exercises there are really only 3 possible shifts. Either we become more productive for the good on the y-axis (which looks like Figure 3); we become more productive for the good on the x-axis (which

Figure 3: An Increase in Cooking Productivity

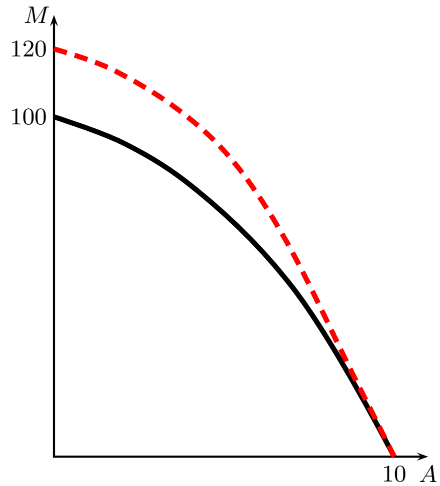

would look like the example just rotated); or there can be an increase in productivity for both goods. This last example would shift out both axes; the PPF would retain its shape. Increases in productivity, either through technological increases or increases in the proficiency of the workforce, are good **because these expand our *choices*.**

1.2 Production Functions → PPFs

These PPFs we've talked about implicitly come from objects we call *production functions.* Production functions are equations that link inputs to output for specific goods; we put x inputs into our production function and $f(x)$ units are produced. The PPF, on the other hand, was a relationship between possible combinations of outputs. As you can imagine, these will be closely linked. If scarce inputs (the x's) are used to produce one good, these is necessarily less that can be used to produce some other good.

For example—and I wont walk through it here—you should verify that the "car-truck" PPF from above comes from the following 3 equations.

$$C = h_C \tag{1}$$

$$T = \frac{1}{2} h_T \tag{2}$$

$$24 = h_C + h_T \tag{3}$$

Here h_C, h_T are the hours spent producing cars and trucks, respectively. Equation (1) is the production function for cars; Equation (2) is the production function for trucks; and Equation (3) illustrates the "resource constraint." This resource constraint tells us that total hours can only add up to 24, so any hour used for trucks necessarily reduces what is available for case, and *vice-versa.*

To introduce a new and important concept to production functions, we'll proceed with a more interesting example. In the pacific northwest (Washington State, for example) the agrarian economy produces both lentils L, and grains, G. Equation 4 is the production function for lentils using the input "land dedicated to lentil growing," l_L. This function takes the land used for lentils, l_L, as an input and determines how many lentils, L, can be grown.

$$L = \sqrt{l_L} \tag{4}$$

Try plotting the square root function for these different values for land used to grow lentils: $l_L = \{0, 1, 4, 9, 16, 25, \ldots\}$. First you'll notice that the slope is not constant—it is *decreasing*. The curve gets flatter and flatter the fur-

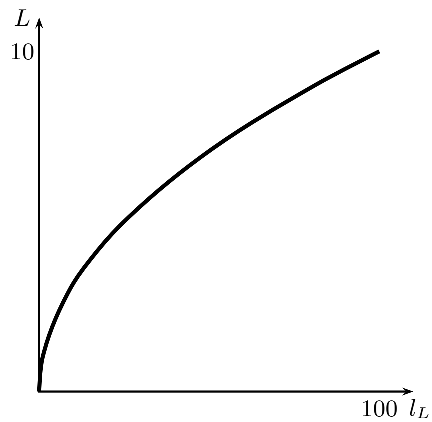

ther out you go (notice the axes run to 10 and 100; this would look extremely flat on "normal" axes). This is a representation of the common economic concept of **diminishing marginal returns**.

Diminishing Marginal Returns

As the use of a resource increases, the additional (**marginal**) return (benefit) of that input decreases.

In this example, going from 0 to 1 increases production by 1; going from 100 to 101 increases production by only 0.05 (0.05=$\sqrt{101} - \sqrt{100}$)! The same one-unit increase in land has very different returns depending on how much land is laready dedicated to this activity. Many resources work this way, in many different settings, so we will see this concept throughout the course. This is one of the key reasons that economists **think at the margin**. When determining how much of a resource to allocate to a specific use, it shouldn't matter how productive the earlier inputs were. What matters is how much adding *another* unit of this resource will contribute to this output, compared to the (opportunity) cost of deploying that additional unit here. If the unit in question passes this test and is deemed useful, you move on to the next and ask the same question. This process continues until the marginal (additional) cost of using the next unit of that resource exceeds the marginal benefit.

The reason we think diminishing returns are common to many economic processes resembles the specialization reasoning for why the society-wide PPF has an outward bow. Imagine we were planning to grow lentils (i.e, moving from 0 to 1 acre dedicated to lentil farming). We would first find the land and soil best suited for producing lentils. Therefore that first l_L input we use is going to be very productive. As we keep expanding production we are forced to use less and less lentil-friendly land so our overall productivity is decreasing (even though *total* production always increases when l_L increases—the rate of increase (slope) declines).

Suppose that grains, G, have a similar production function, but scaled by $\frac{1}{2}$, and there are 100 acres that can be farmed for either good. Then we have the following system of equations that describes our farmers' choices.

$$L = \sqrt{l_L}$$
$$G = \frac{1}{2}\sqrt{l_G}$$
$$100 = l_L + l_G$$

We can plot the PPF by:

i. Pick an l_L

ii. This determines l_G since $l_G = 100 - l_L$

iii. Plug in l_L, l_G to their respective production functions

iv. Plot the combination

That produces a table and corresponding PPF that look like the ones below.

l_L	l_G	L	G
100	0	10	0
81	19	9	$\frac{1}{2}\sqrt{19} \approx 2.17$
49	51	7	$\frac{1}{2}\sqrt{51} \approx 3.57$
25	75	5	$\frac{1}{2}\sqrt{75} \approx 4.33$
9	91	3	4.78
0	100	0	5

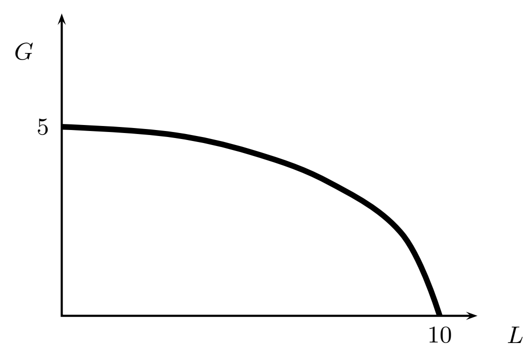

Pacific NW PPF for Lentils and Grains

1.3 Multi-Economy PPFs

We've seen so far where an individual PPF comes from, but more interesting things start happening when we consider multiple economic units and how they aggregate and interact. Let's start by considering an example from South America. The Andean mountains of Peru are remarkably steep—geographically close villages can be at very different elevations. Different elevetions imply different production functions; they're more or less suited to produce certain goods so they would need more or less inputs for the same amount of output. Two important things ancient Peruvians needed to survive were clothing, C (courtesy of the alpaca they domesticated), and food from their staple crop quinoa, Q. Let's consider a high- and low- elevation community with 10 people each.

In the high-elevation economy, each member's labor can produce 4 units of Q or 4 units of C. In the low-elevation economy each member can produce 8 units of Q or 2 units of C. It's easier to raise alpaca in the highlands; harder to grow quinoa up there. Below is a sketch of their individual PPFs.

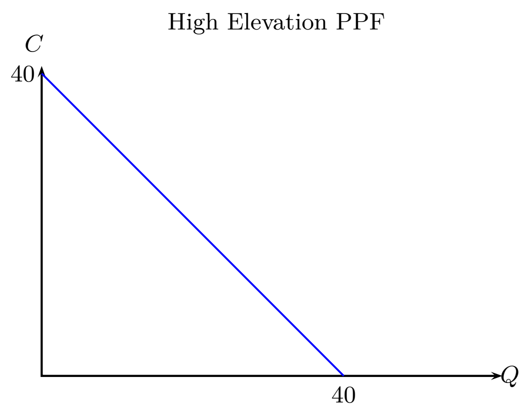

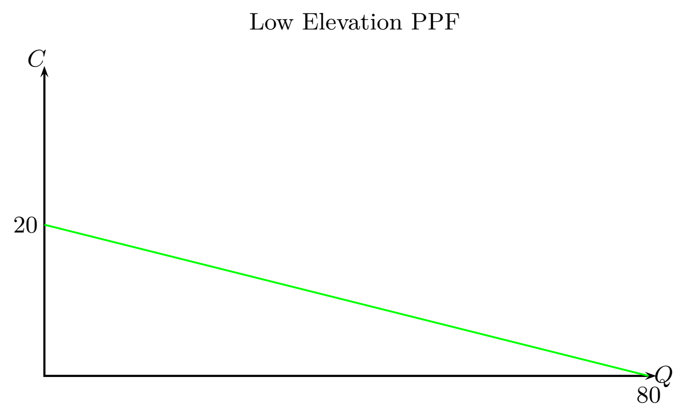

Let's take this opportunity to introduce two new key terms: **absolute and comparative advantage**.

Absolute and Comparative Advantage

If economy A has an **absolute** advantage over economy B in some good, it can produce more of that good. If economy A has a **comparative** advantage in producing some good, it can produce that good at a lower opportunity cost.

Having an **absolute advantage** in a good implies that the economy can produce more of that good than the economy we are comparing it with. You can think of this as their raw productivity for that good. Here we see that if their resources were put to it, the high-elevation economy could produce more clothing; they have the absolute advantage in that good. On the other hand, the low-elevation economy can produce more quinoa if their resources were spent doing so; they have the absolute advantage in quinoa. It need not be the case that economies flip absolute advantages between goods. The United States, for example, has the absolute advantage in almost all goods compared to Nepal—we have the technology to make more of nearly any good were we to hypothetically face off.

Having the **comparative advantage** in a good implies that the economy produces that good at a lower (opportunity) cost. That is, it is *relatively* better at making the good. The opportunity cost for Q is 1 C for the high-elevation economy, but it is only $\frac{1}{4}C$ for the low-elevation economy (remember: compute the slopes of the PPF). The low-elevation economy therefore has the comparative advantage in quinoa. They give up less clothing to produce quinoa. The case is flipped for C. *Comparative advantages always flip in a 2-good, 2-economy setting.*[2] An economy cannot have a comparative advantage in *both* goods. Comparative advantage determines relative strength, it is impossible to be "relatively" better at both actions being considered.

Before moving on, let's fix ideas and assume that because each economy needs both Q and C they each devote half of their resources to producing each good. The high-elevation economy would then produce (and consume) (20,20); the low-elevation economy would produce and consume (40,10).

These economies, up until this point, do not know the other exists. But if you oversaw this land, it might make sense to think of them as one big economy. Let's plot the PPF you would be facing as the King/Queen of the Inca's (the most famous historical Andean society). To plot this PPF, begin with the extreme points: if you wanted to produce all C or all Q you could now produce 60 C or 120 Q, respectively (add the capabilities of the two economies).

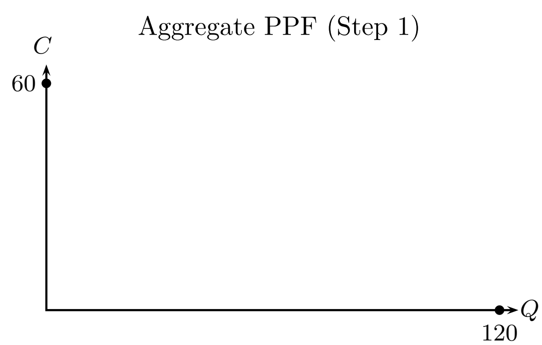

But how do we get from that far left point (0,60) to (120,0)? If you're a wise leader, you'd begin by utilizing the economy that is *relatively better* at producing Q. This is just like in the Norman example above (Figure 2) where we got our best coder/worst cook out of the kitchen when we first begin making apps. Here, when we want to begin making quinoa, we assign some individuals from the low-elevation economy to produce it. This is the efficient decision because they are (relatively) better at producing quinoa. As we continue expanding quinoa production, we continue assigning low-elevation economy resources for this same reason. This continues until we run out of them; when the low-elevation puts all of their resources towards Q, we can no longer assign more of their resources there. Then, should we want more quinoa than that economy can produce, we allocate the people and resources that are relatively worse at producing Q (the high-elevation group).

[2] See section 1.5 for a proof of this; you won't be responsible for the math proof, just remember that it is a true statement.

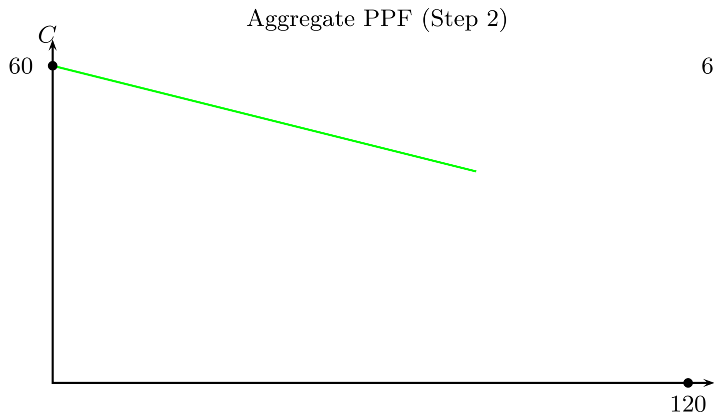

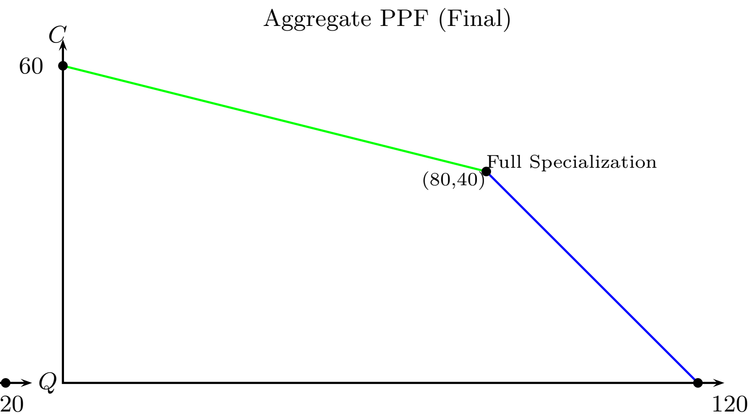

Something amazing happens here: despite both PPFs being linear, the aggregate PPF resembles a bowed out PPF. This is because the specialization issue here is analogous to the Norman apps and meals example. We send the relatively better economy to work on producing Q when we start, so the (opportunity) cost is low at first. Once we run out of their resources we need to send some people in the high-elevation economy to produce Q; the cost increases at that point because they are worse at this activity. Imagine we had a 3rd group—we'd have a third kink. The more groups we add, the more kinks (and slope changes) we add. If we imagine the whole U.S. economy being made up of every household's PPF, we'd have (literally) 100 million kinks—100 million kinks would make such a curve look pretty smooth!

Another paradoxical thing to notice is where their current production falls on this aggregate PPF. We assumed above that the high-elevation economy produces (20,20) and the low-elevation group was producing (40,10); that's total joint production of (60,30). That point is inefficient (see figure below) when looking at the aggregate PPF, despite being efficient within each economy. This phenomenon arises because—despite them using their individual resources efficiently—they haven't used the resource of "specialization" properly.

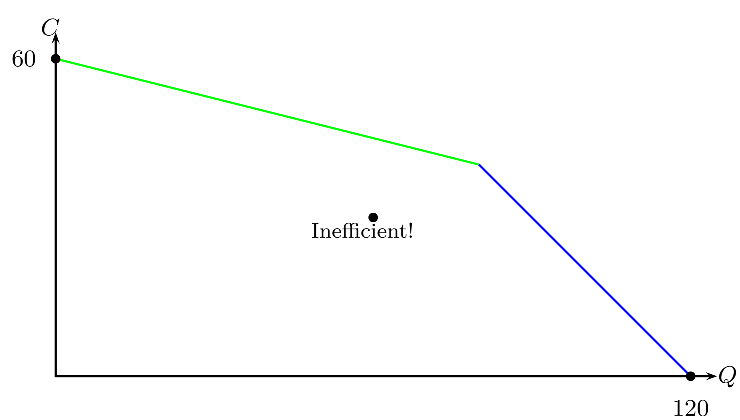

1.4 Gains From Trade

To see how powerful specialization can be, let's analyze a potential trade between these groups. Let's say the high-elevation economy hatches a plan where they would *only* make clothing, and would offer to send down 20 alpaca shawls (the clothing unit) for 40 sacks of quinoa. As we will see, if the low-elevation economy *only* produces quinoa this can be a win-win. Consider their PPFs again, with the dots marking their production, and the star marking their consumption—with trade these points are no longer the same. The high-elevation economy is producing 40 shawls, but is sending 20 away for 40 sacks of Q; so their production is (0,40), but their consumption is (40,20).

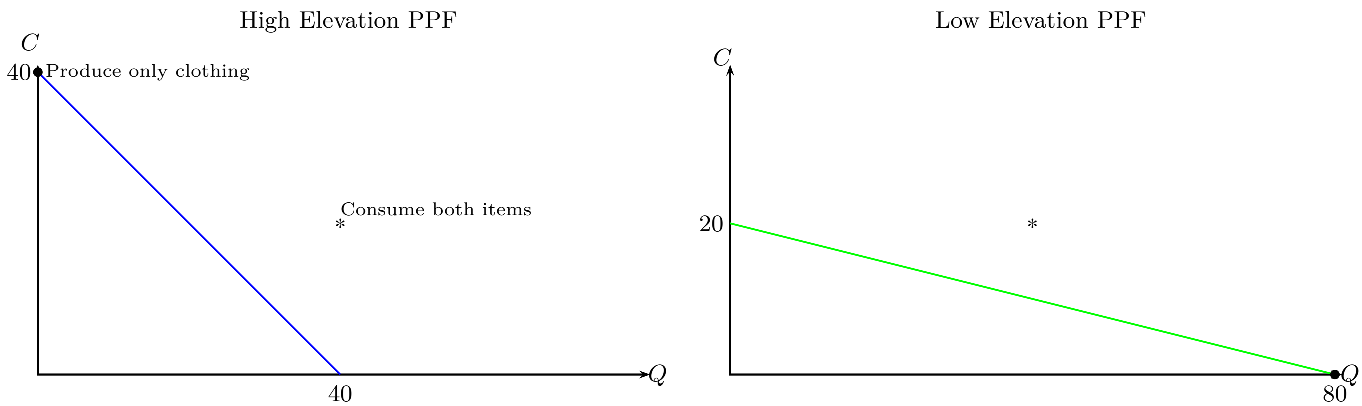

Trade allows both economies to consume *outside of their PPFs.* Leveraging specialization allows these economies to access consumption combinations they could not produce on their own. Its like an extra "resource" they can take advantage of when operating in a coordinated fashion. It was a failure to take advantage of this resource that resulted in them being within the joint-PPF above [verify that specializing in this way puts them on their joint PPF]. In fact, there is so much winning going on that these mutual-gains are possible in tons of potential trades. Try this exercise again if the high-elevation offers only 15 C for 40 Q. You'll see that it's still beneficial for both groups!

It isn't until the high-elevation economy offers only 10 C for 40 Q that the low-elevation economy ought to tell them to get lost. Why is that an important point? Let's look at the economy *importing*—purchasing from another economy—C. In a 40 : 10 trade, the low-elevation economy pays $4Q$ for each C they buy.[3] What is their cost of making it? Their opportunity cost is $4Q$ for every $1C$. They can tell the high-elevation economy to get lost, they can make it just as cheap as the deal being offered. Therefore, the *highest* price that can be charged for $1C$ in this trade is $4Q$. How about the lowest? Here we need to look at the *exporting* economy. Their cost of making $1C$ is $1Q$—if the price falls below $1Q$ its no longer "profitable" for them to make C and sell it for only $1Q$. **Therefore, the trade price for a good must fall between the**

[3]Divide each side of the trade by 10, so it becomes a "for each" statement. To figure out how much each Quinoa is sold for, divide each side of the trade by 40: each Q sells for $\frac{1}{4}C$.

opportunity costs of the two economies. [Do this for the price in reverse; what are the bounds for prices on $1Q$?]

The real takeaway is even stronger than the bold statement above. Its not just that the price must be between their opportunity costs, for the reasons laid out above **if the trade rate for a good is between two economy's opportunity costs, both economies can gain from that trade.**

David Ricardo's Trade Claim

As long as opportunity costs differ between economies, these economies can gain from trading with one another. Comparative advantage determines mutually-profitable trading patterns, *not* absolute advantage.

This is a remarkable statement and it follows exactly from our claims about comparative advantages above. Economies can't be relatively better at produce all goods, so there is always an opportunity for specialization. Because specialization is a new resource these economies take advantage of, they both gain. This is true even when an economy has an absolute advantage in all goods. Our common sense might suggest that economies can only gain from trade if absolute advantages differ: the U.S. is better at producing everything than some poorer countries, how could it be that we can gain from trading with them? This common sense is wrong. What matters is if they have different comparative advantages, and this will **always** be true; therefore there are **always** gains from trade. [Note that economies don't need to fully specialize as in the Andean example here; the larger economy might end up making some of the good they are not specialized in if the smaller economy can't satisfy the needs of both groups.]

Takeaway: specialization leads to situations in which there are always gains from trade. Trading allows for more joy to be created with the same resources! In the next section of the course we will move to determining—in more realistic settings—who produces/purchases what goods at what prices.

Professor Kuruc's Real World Takeaways

What you think these theories mean for the real world depends on combining them with other facts and your values; and then your reasoning gets all muddled by cognitive biases. So take these with a grain of salt.

1. In your own life, specializing in something (especially in something not many are good at) is an economically wise choice.

2. Free trade and globalization are pretty great! Rich and poor countries both can benefit from specializing and exporting in exchange for things they want but cannot produce cheaply.

   - This is why almost all economists think President's Trump tariff policies designed to reduce our imports are bad ideas. His biggest complaint is that other countries undercut American companies. But think about what "undercutting" means. It means these countries are sending us their products and asking for very little in exchange. This is a good thing! If some country gifts us product x at very low prices, we should graciously accept this, stop producing x, and move our resources to producing things we can't easily get on global markets. Trade creates win-wins, we should aim for more, not less, of these.

1.5 Math Appendix

Here I'll prove that that in a 2-good, 2-country setting one country can't have a comparative advantage in both goods. This is the main claim that gives David Ricardo's trade result force.

Suppose Country A and Country B produce good "1" and "2". If Country A has some opportunity cost x for good 1, then it must have opportunity $\frac{1}{x}$ for good 2 (recall our reciprocal rule). We can call Country B's opportunity cost y for good 1, and $\frac{1}{y}$ for good 2. Then it must be the case that:

1. $x > y$; Country B has the comparative advantage in good 1. If $x > y$ it must also be true that $\frac{1}{x} < \frac{1}{y}$; the denominator x is bigger, so the fraction is smaller. Therefore: if Country B has the comparative advantage in good 1, Country A has the comparative advantage in good 2.

2. $x < y$; Country A has the comparative advantage in good 1 $\Rightarrow \frac{1}{y} < \frac{1}{x}$. Therefore: if Country A has the comparative advantage in good 1, Country B has the comparative advantage in good 2.

3. $x = y$; neither country has a comparative advantage. This is the rare case where countries cannot gain from trade: they can't specialize if they have exactly the same relative skills.

Principles of Macroeconomics–Honors

Supply, Demand, & Markets

Instructor: Dr. Kevin Kuruc

February 10, 2021

In the first section of the course we analyzed production possibility frontiers and trade. We saw, analytically, the gains that are possible from trade. This shouldn't be surprising. We're all trading all the time—when you purchase groceries or prepared food, you're implicitly trading some work that you (or your parents) have done for the work that went into producing that food. It's far more efficient to do the thing you are relatively best at in return for the output of others, rather than trying to build your own shelter, grow your own food, etc., yourself. In this section we learn how market economies—like the United States and most of the world utilize—organize these trades and incentivize production.

1 Markets

1.1 Market Economy vs. Command Economy

From a bird's eye view, market economies rely on *individual* decisions to guide the allocation and production of resources. The government's role is to ensure that property rights are respected (i.e., no stealing) and that trades/contracts are enforced, but the government does not tell producers what to produce, or consumers what to purchase. Having lived in a market economy your entire life this may seem like the only way things could be, but this is not the structure of command economies once popular with communist governments. In those systems the government decides what is produced and what individuals receive for their consumption. Market economies instead rely on the combination of individual self-interest and prices to determine what is produced and who gets what.

> *It is not through the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard to their own self-interest*
> —Adam Smith

Market forces organize incredibly complicated chains of production to get consumers what they want at the lowest sustainable prices. Pencils, for example, are demanded by consumers. The pencil manufacturer, after doing her researcher, realizes that she can sell pencils for more than the cost of the inputs. So she outbids a company making wooden spoons for some of the lumber, outbids a hose company for some rubber to make erasers, and outbids somebody else for a bit of graphite. Her new demand in this market induces slightly more future lumber, rubber, and graphite production because these industries are now slightly more profitable. She then combines the inputs to sell you a pencil. Her sale price is then limited by her competition—if she charges too much, consumer's will leave her and purchase from someone else (or buy pens). Amazingly, your demand for a pencil creates pencil companies that then demand the raw inputs for pencils; competition between sellers keeps prices low.

Generally, when too much of something is produced there is a surplus → fierce competition drive prices down → self-interested sellers realize that perhaps this isn't the best market to operate in and instead find one where supply doesn't quite reach demand. If too little of something is produced there is a shortage → producers can charge a high price because of the competition among buyers for their product → self-interested sellers produce more (or enter this market) because they can make a high profit. **Prices incentivize buyers and sellers to make the decisions that are in the best interest of society.** For a government to try and emulate this process without markets, they would need an unrealistic amount of information. Markets preserve our economic freedom to choose what we do and what to purchase, while generating outcomes that are almost always better than what a government could hope to organize. Of course, there will be exceptions to this, but the fundamental takeaway is incredibly surprising and powerful. By the end of this section you'll understand these forces at an analytical level and why there is so much faith in markets worldwide.

1.2 Demand

Because market forces derive from individual decisions, our analysis will start with individuals. On the consumer side we have what we will call a demand curve.

> **Demand Curve**
>
> A relationship between the price of a good/service/resource and the amount that would be purchased at that price. In other words: at some price P how much quantity is demanded for purchase, Q_d?

Let's take some individual's demand for pizza slices. Our hypothetical consumer is walking to a pizza place after a late night of studying and is very hungry. As he's gaming out his decisions, he thinks to himself: I could eat an entire pizza right now. From that we can learn that *if pizza were free, his quantity demanded would be 8* (row 1 of his demand table). Even though he *could* eat an entire pizza, once he has 6 slices he'd be pretty full, so if the price is \$2 he's going to stop at 6. If the price is \$4, he's going to stick just to 4 slices because that's a bit much and 4 would still leave him pretty satisfied. And if the price is something crazy like \$6 a slice he's going to just get 2 slices to hold him over until he's home to make himself a sandwich. If the price is as high as \$8 he's not going to purchase any.

Figure 1: Individual Pizza Demand

Price	Q_d
0	8
2	6
4	4
6	2
8	0

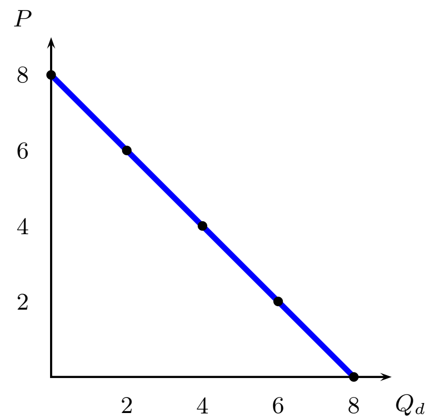

A few key concepts jump out of this example. The first is diminishing marginal utility.

Diminishing Marginal Utility

The more of some good/resource is consumed, the less benefits consumer's receive from an additional unit.

This is very similar to diminishing marginal product, and it works here both at the individual and social level. At the individual level, our pizza consumer got less satisfaction from every additional slice; once he's satiated his worst hunger pains, the next slice adds slightly less happiness. At the social level, if we only produced 1 pizza, we'd hope that the person who likes pizza the best gets it. The second pizza goes to someone who likes it a bit less, leading to the same phenomenon. The fact that marginal utility diminishes gives rise to the Law of Demand.

Law of Demand

Other things equal, if the price of a good falls, more will be demanded for purchase. If the price rises, less will be demanded for purchase.

This is a statement about the *inverse* relationship between price and quantity demanded we expect. It manifests as a negative slope on the graph above: when the price is high, quantity demanded is low; when prices are low, quantity demanded is high. The "other things equal" qualifier is important here. You might think "when prices are low, that's because not many people want to purchase that good." As we will see, that is true, but it is different from the claim here. We're looking at a product, assuming nothing changes about that product except for the price, and asking whether more or less will be purchased. If the product is identical, but it is half-off, we'd expect more quantity demanded.

One last important feature of this graph to note are the axes. Something is a bit weird: we're asking how Q_d responds to changes in P—but we've put Q_d on the x-axis and P on the y-axis. Usually our *independent* variable goes on the x-axis and we see how the y evolves. Here we've done the opposite. As far as you need to be concerned, there's not a good reason for this. But you must remember to do it this way! Whenever we graph supply and demand P goes on the y-axis. That is especially important when dealing with demand equations, rather than a chart. For instance, the equation that most naturally describes this scenario is $Q_d = -P + 8$. However, to graph this we would want to rearrange and put P on the left-hand side.

With these individual demand curves we can build a market demand curve. This is a simple process: we add all quantities demanded at a given price. If 10 individuals demand 1 slice at $7, then the market demand at a price of $7 is 10 total slices. As macroeconomists what we care about is market demand, not individual demand. Graphically, this is a *horizontal* summation (as opposed to vertical). In a microeconomics class we would spend more time on the individual demand curves and their properties. Here we're going to just assume we can build up to a market demand curve.

Market demand curves inherent the Law of Demand describing individual demand curves. Before moving onto supply curves we should discuss changes—or shifts—in demand. We are going to want to analyze both (i) how markets allocate goods for some given demand curve and (ii) how markets respond to *changes* in demand. Therefore, we need to analyze what a change in demand looks like graphically.

Suppose we have our pizza market with some demand curve, denoted by the column D. If something

Figure 2: Market Pizza Demand

Price	D	D'
0	800	1000
2	600	800
4	400	600
6	200	400
8	0	200
10	0	0

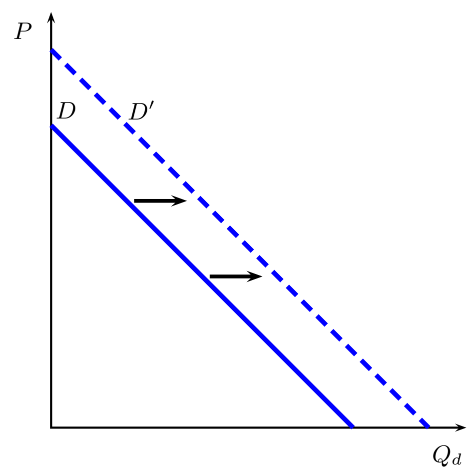

happens that increases the demand for pizza—let's say we find out that tomato sauce is extremely healthy, or some viral internet challenge involves pizza—what we would expect is that **for any given price, more**

pizza will be purchased. This is a rightward shift, not a shift "up." You'll be tempted to see this as an upward shift. Don't do it. An increase in demand is a rightward shift. A decrease in demand causes a leftward shift. This will be an important distinction once we bring supply into the picture. Also note the terminology: this is an increase in demand. If more of some good is purchased because the price falls, that is an increase in *quantity demanded*. Keeping this distinction straight will also be important once we have supply in the picture.

A list of factors that could shift demand are as follows:

- Change in number of buyers
- Change in consumer preferences
- Change in incomes
- Change in the price of a related good
  - Complementary Goods
  - Substitute Goods
- Change in expectations

1.3 Supply

The supply side of the economy has a similar structure to the demand side—the only difference is the picture will be flipped a bit. We continue to study the relationship between quantities and prices, but the supply curve depicts how much is produced and sold at a given price.

Supply Curve

A relationship between the price of a good/service/resource and the amount that would be produced and brought to market at that price. In other words: at some price P how much quantity is made available for purchase, Q_s?

Suppose we are interested in the market for babysitting. This is a simple market to think about because we can think through what its like to be on the supply-side of a labor market—individuals sell their time in return for some hourly wage. Technically, in a given weekend individuals have 48 hours to supply for baby-sitting. However, at different hourly pay rates, they will choose to supply different numbers of hours. Suppose we have two potential teens in this town, Adam and Betty. Adam doesn't have much going on this weekend, so his (opportunity) cost of baby-sitting is lower than Betty's; she has concert tickets for her favorite band Saturday, so she would need an incredibly high wage to get her to work more than 15-20 hours.

The supply curves of Adam (S_A) and Betty (S_B) are given below along with their graphs and the market supply curve S_{Mrkt}.

Figure 3: Baby-sitting Supply

Price	S_A	S_B	S_{Mrkt}
0	0	0	0
5	0	0	0
10	5	0	5
20	20	10	30
40	28	20	48
50	30	22	52

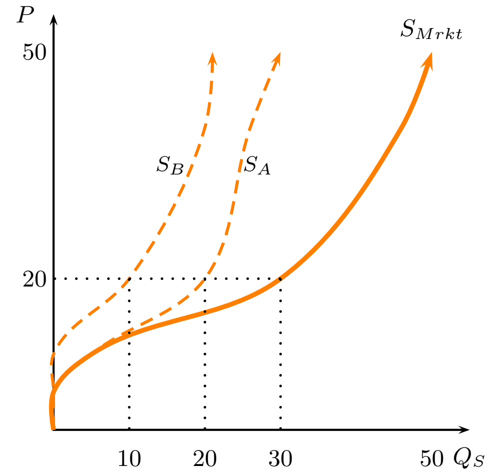

The supply curves here—despite their whacky, wandering shape—follow what we call the Law of Supply.

Law of Supply

Other things equal, if the price of a good rises, more will be made available for purchase. If the price falls, less will be available for purchase.

The Law of Supply tells us that these curves should have positive slopes at all points. These supply curves meander, but always towards the north-east of this graph. This law comes from the concept of increasing marginal costs.

Increasing Marginal Cost

As more of a good/service/resource is supplied, it becomes increasingly expensive to supply an additional unit.

Increasing marginal costs directly follow from our earlier concept of diminishing marginal returns. Recall, the more of a good we want to produce, the less efficient we become at producing it—this implies that the *cost* of supplying that additional unit should rise. These concepts might make more sense in our baby sitting example. When the price offered is low, neither teen is willing to supply much, if any, babysitting hours. When the price increases to \$10 per hour, Adam is now willing to supply some hours (Betty still is not). Why might Adam only be willing to supply 5 hours? His weekend is relatively free and he realizes that if he plays video games all weekend he'll get pretty bored. So, at a low price of \$10, he's willing to give up 5 of those hours, since by that point he'll be pretty bored with these games. Because Betty has a lot going on this weekend, even her last few hours free she would rather nap or catch up with friends than work, so \$10 doesn't get any time out of her.

When the price rises up to \$20 an hour, however, things change. Although Adam wants to play his video games, for \$20 he's willing to give up a lot more of his time. We can infer from this information that the opportunity cost of these next 15 hours was higher than \$10 because he wouldn't have supplied them when that was the price; the cost must not be higher than \$20 per hour, though. Once he gets to 20 hours, he's given up most his weekend, and being able to relax at least a bit is important to him. These next hours are even higher opportunity cost hours. As the price rises, he's willing to supply more and more of these high cost hours.

Betty, in contrast, has very high opportunity costs this weekend. She's got a concert to get to! Notice, her supply curve lies to the left of Adam's. She is the high-cost (or inefficient) supplier, so **at any given price, she is willing to supply less of this service.** Being high cost isn't necessarily a bad thing! In some markets this is due to being an inefficient producer, in markets like this its because you have better things to be doing. Regardless, as we can see at a price of \$20 clearly, she is willing to supply less time than Adam (10 vs. 20 hours).

The market supply curve, like the market demand curve, is the horizontal addition of these supply curves. At \$20, Adam supplies 10 hours, Betty supplies 20, so the market supply at \$20 is 30 hours. Because we add horizontally, notice that as supply *increases* this curve shifts to the *right*. This is unituitive, so take careful note. If you think "increasing supply means shifting the curve up" you would get this wrong. This is why I stressed in the demand section to think right or left.

Let's look at an example in this market where the supply curve shifts. One easy example would be something that changes the cost of supplying this service. Because their costs are only the opportunity cost of their time in this example, we would need something to alter the time-costs. So, suppose a big storm rolls in, and the outdoor concert Betty was going to attend is canceled. Her weekend has become free, so the opportunity cost of her time has fallen: **at any given price she is now willing to supply more hours.** Because one of our suppliers increased their supply, the market supply curve will increase. Any factor that changes how much is supplied *at a given price* will increase the supply curve in some market. Taking a more common "goods" market, let's come up with some examples of what could shift the supply of coffee.

- Change in producer conditions: if the weather is very bad in Latin America (a large coffee exporter) the supply of coffee will be reduced at any given price.
- Change in input prices: if the cost of labor goes up in Latin America, it is more expensive to produce coffee, so less will be supplied at a given price.
- Change in technology: if the machines that process coffee beans become more efficient, more will be supplied at a given price

Figure 4: Increase (Shift) in Supply of Baby-Sitting

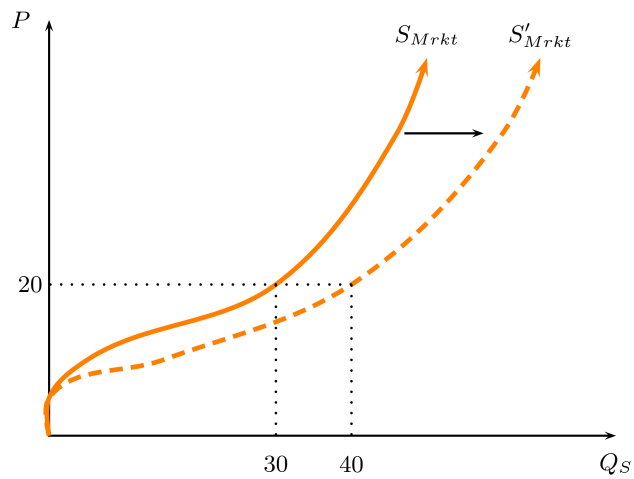

- Number of Sellers: if the market for selling beef collapses, some ranchers may become coffee farmers, increasing the supply at any price

- Expectations: If sellers expect the price of coffee to rise next year, they may reduce supply today in order to have more to sell next year

1.4 Equilibrium

With supply and demand within a market, we can analyze where that market will operate when left unregulated: the **market equilibrium**.

Market Equilibrium

The price & quantity combination in a market such that buyers and sellers are in harmony. That is, the number of goods sellers want to produce at some price equals the number of products wished to be purchased at that price.

In physics, the concept of an equilibrium refers to the state where a system settles down and becomes stable. In economics, we have a very similar concept of an equilibrium; we're describing where this market will become stable. The ability to analyze markets in this way ultimately allows us to describe why some markets operate as they do (why is the wage for baby sitting so much lower than for physics professors?) and predict how a market will respond to certain changes (what will happen to the coffee market if the weather in Latin America is bad next year?).

Let's first begin by graphing supply and demand together for some arbitrary good. The equilibrium—labeled in Figure 5 as $E^{\star}$—is the price ($P^{\star}$) such that the quantity demanded and the quantity supplied are

Figure 5: A Market with Supply and Demand

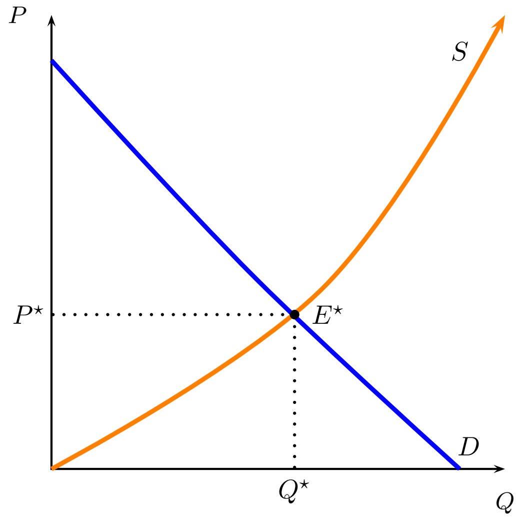

equal ($Q_d = Q_S = Q^\star$). At the equilibrium price, $P^\star$, all sellers who want to sell their product can, and all buyers who want to buy that product can. This is a bit tricky: yes, there may be some consumers who wish the price were lower, but because we need someone selling these goods, if the price we're lower there would be less available. When we say "want to buy this product" we mean "want to buy this product, *at the price its being sold for*."

The intersection of these curves *feels* like a natural place for the equilibrium, but why is the market stable here and not somewhere else? As I will demonstrate: at any other price, there are opposing forces of competition and profit motive that are not in balance. Only at the equilibrium does neither force kick in. The absence of either force leaves the price where it is.

Figure 6 describes what happens in each of two possible scenarios. Either the price is higher than the equilibrium (panel a) or the price is below the equilibrium (panel b). If the price is too high, the number of sellers exceeds the number of buyers: a surplus. In a state of surplus, sellers realize that they may be left with their product. Rather than take this risk, they can slightly lower their prices to ensure their products sell.[1] Because the loss of failing to sell is worse than dropping your price a bit and ensuring a sale, every seller has the self-interested incentive to lower their price. Hence, when the price is above the equilibrium, competition among sellers drives the price down as they would all prefer to ensure they get the sale. This is true for any price that remains above $P^\star$. Once the equilibrium is reached from this direction, no seller wants to lower their price any further—they realize there are now enough buyers, so why would they?

If the price is instead too low (panel b), sellers realize there are more buyers than sellers: a shortage.

[1] We're assuming a market with perfect knowledge right now. If I drop my price a few cents lower than yours, customers will purchase from me instead.

Figure 6: Dynamics towards equilibrium

(a) Price too high (b) Price too low

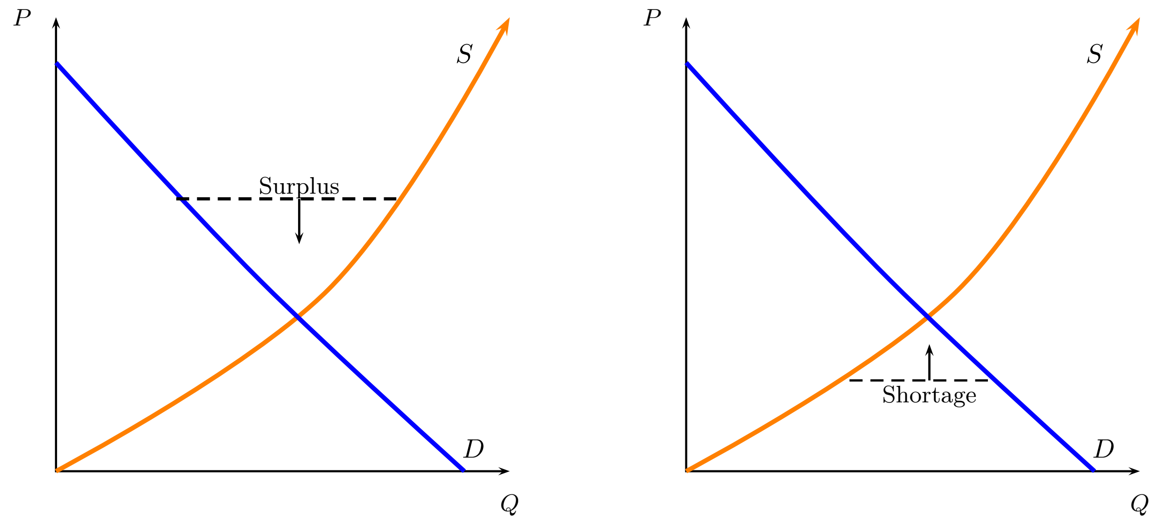

In this case, buyers would rather pay a few extra cents to ensure they receive the product in this state of shortage. Sellers are happy—because of the profit motive—to increase their prices. Since all sellers can increase their prices a bit and still sell their products, they do. This higher price induces more sellers to provide products and we move towards the equilibrium. Again, once we reach the equilibrium (after this process repeats a bit) this pressure is alleviated. If there are no extra buyers hanging around, sellers no longer have the liberty to increase their prices without losing customers. **Models of supply and demand, therefore, allow us to predict where (and understand why) markets operate at the size and prices that they do.**

These models also give us the tools to predict what will happen when there are changes in this market. I didn't have time to diagram these shifts, but if you try them yourself for studying, make sure you get the following results:

- Demand increase: $P\uparrow, Q\uparrow$
- Demand decrease: $P\downarrow, Q\downarrow$
- Supply increase: $P\downarrow, Q\uparrow$
- Supply decrease: $P\downarrow, Q\downarrow$
- Demand, supply both increase: $P?, Q\uparrow$
- Demand increase, supply decrease: $P\downarrow, Q?$

- If supply changes in one market, think about how that will impact the production of related goods
  - In class we discussed coffee supply reducing → coffee prices rising → which pushed some consumers towards tea (increasing the demand) → which increased the quantity of tea supplied (as it shifted along its supply curve). So price signals running through this chain get more tea producers to make tea in response to something bad happening in the coffee market.

We've hinted at the idea that there is something "good" about markets reaching their equilibrium. Shortages (surpluses) are eliminated, which was a rough indication that too few (many) resources were allocated towards a specific good/service. However, there is a more formal analysis we can undertake to see why markets operating at their equilibrium is the "optimal," or "efficient" point. The analysis will concern itself with a concept called **social surplus.**

Social Surplus

The total gains from trade in a market.

Pretty simple concept, but not so obvious how we'll go about analyze something as abstract as gains from trade. We'll start with the consumer side.

Consumer Surplus

The (total) difference between "willingess to pay" and actual price paid in some market.

Figure 7: Consumer Surplus

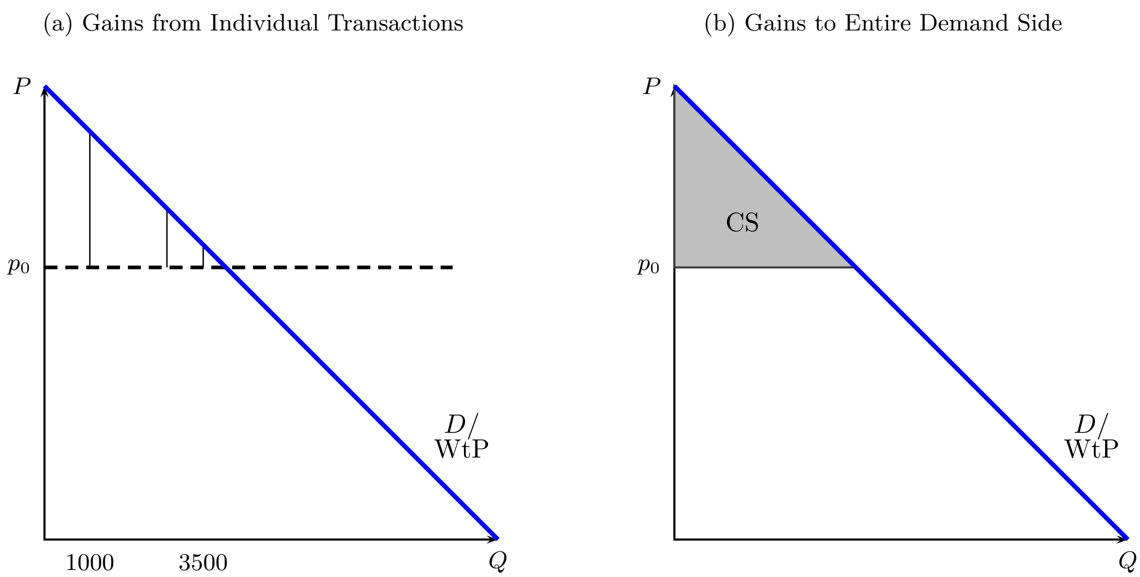

Figure 7 depicts consumer surplus for a given demand curve and price p_0. Looking first to the panel (a), I've picked out a few arbitrary purchases in this market: the 1000'th purchase and the 3500'th purchase. When lined up by their demand curve in this way, we know the 1000th purchase came from a higher willingness to pay: that is the 1000th purchaser would've been willing to pay more than whoever made that 3500th purchase. Taking the (vertical) difference between the willingness to pay for that item, and the actual price (p_0 here), gets you the surplus (or gain from purchase) for that individual transaction in dollar terms. Notice that I've not drawn any transactions beyond where the price and demand curve intersect—the willingness to pay is below the price, so those transactions will not happen at a price of p_0. Those trades generate *negative* gains to the buyers, so they'd never willingly purchase.

To compute this surplus for the entire demand side of the market you would need to add up all of those tiny little lines. Thankfully, using the power of math, we can notice that if we really drew all those little lines and wanted to add them, it'd look a lot like a shaded triangle (panel b). Therefore: **we can conceptualize and compute consumer surplus as the *area* beneath the demand curve but above the price.**

The same logic holds for the supply side, but flipped: the gains from selling is the price received minus the supply curve (a willingness to produce/sell curve). If a producer receives p_0 and they'd be willing to sell at some lower price, then their surplus is this difference. We again add them all for the result of a triangle-like shape (here a triangle, but in examples with curve-y demand/supply curves its not exactly a triangle).

Figure 8: Consumer Surplus

(a) Producer Surplus

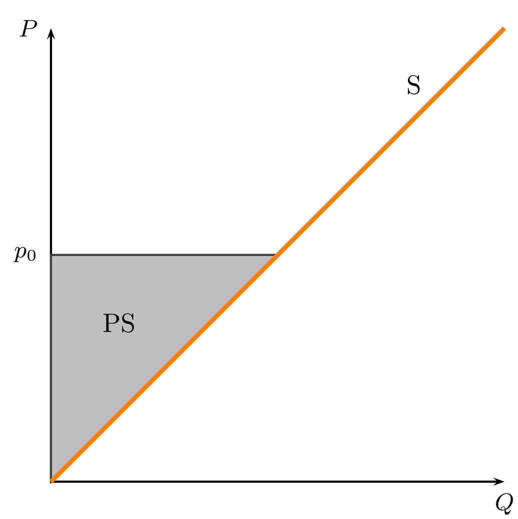

(b) Total Surplus in Equilibrium

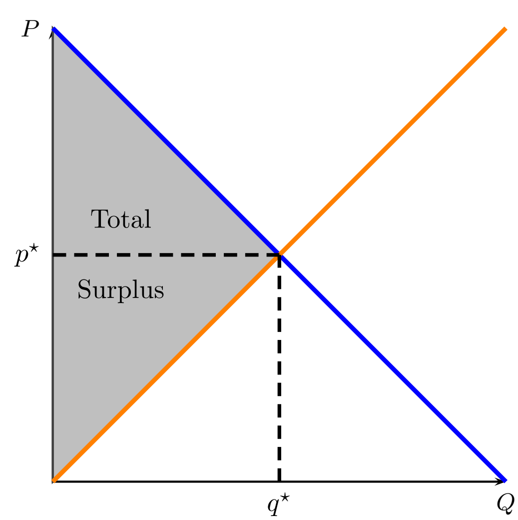

We can then put supply and demand curves into a market to determine the "total surplus" (consumer + producer); this is possible because our equilibrium predicts where the market will operate without regulations. For the most part we will jump straight to analyzing total surplus, because our aggregate analysis doesn't care much about consumers versus producers. This total surplus graph could have been directly constructed by thinking about trades instead of a market. The very first trade in this market creates a ton of value: there are very high willingness to pay buyers, and some very low cost sellers. This could be imagined as one very long line going all the way from the demand to supply curve. As more and more trades are made, less and less surplus is created by each. Once we hit the equilibrium, no surplus is available. This is the miracle of unregulated markets: we know they get to our equilibrium, and at our equilibrium all trades have been made that generate positive surplus, and not a single one more. There will be cases where our model here is a bad approximation of the real world, but this logic should increase your confidence a bit that markets are getting us to a useful operating point on their own, in many cases.

1.5 Government Policy: Price controls & taxes

This framework can now let us make both positive and normative statements about government intervention in markets. Positive economics is a description of the effects; normative economics is a statement about whether these effects would be good or bad. Using our equilibrium analysis we can predict what will happen; and our tool of "total surplus" will help us determine whether this is good or bad.

Let's start with a simple policy: a price control. Imagine, for political purposes, we prohibited the price of smartphones from exceeding $400; consumer's complained that these necessary objects were too expensive. The equilibrium price of smartphones is above $400—the newest iPhone goes for around $1000 dollars. This law would prevent the market from reaching its equilibrium price, and hence is called a "price ceiling." Figure 9 depicts the effects of this policy.

Because the market price is too low, we have a *shortage* of smartphones—notice where this $400 intersects the demand curve, many people want to purchase at this price. Far fewer, however, are willing to produce/sell at that price (only 400, as I've labeled it). So the overall effect of this price ceiling on our market outcome is to lower both the price and the number of phones available. Our equilibrium analysis allows us easily recognize that when prices change we need to think about how all market participants respond.

With our concept of social surplus, we can do more than just predict the new outcome. We can ask whether this was good. After all, prices fell—good for consumers—but the quantity available for purchase fell—bad for consumers. The way we make such a judgment is to compare the social surplus with a price ceiling to the social surplus without a price ceiling. We know that without a price ceiling we would get to our equilibrium, and therefore the social surplus is the big triangle stretching from demand (top) to supply (bottom) out to the equilibrium quantity (remind yourself what this looks like in Figure 8b). Here, we still get those first 400 sales, and each still generates a lot of surplus.[2] However, once the 400th is sold, we've

[2] Although notice that we can no longer even be sure that the 400 individuals who want these phones most get them, when there are shortages sometimes folks just get lucky in who gets the product—so the depiction here is the most social surplus possible generated and it relies on sellers somehow getting these phones to the individuals with the highest willingness to pay.

Figure 9: Effects of Price Control

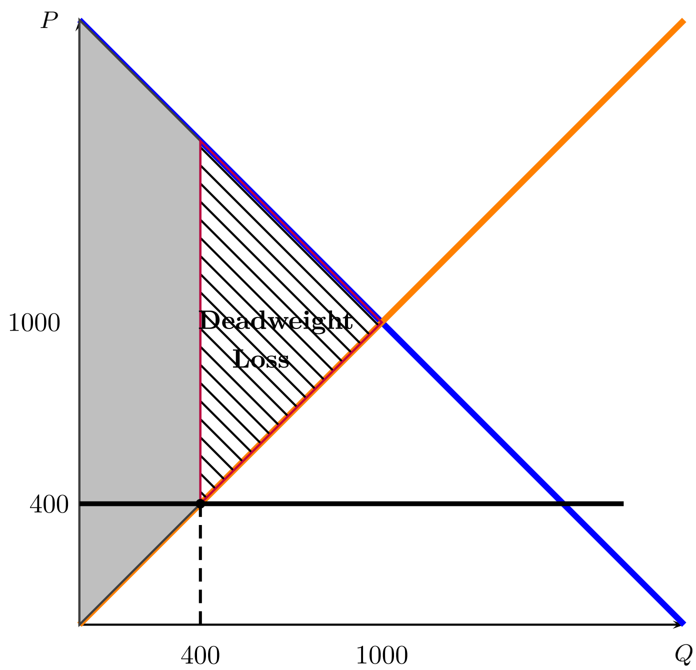

run out of producer's who find it worth their while to sell in this market. The surplus from trades stops here, since no more trades are made. There is a big triangle missing relative to the surplus the unregulated markets get: this is called "deadweight loss." We give it such a demeaning term because its additional surplus we could have had if we just let this market achieve its equilibrium (which is easy, since it would naturally head that way!).

Think about where such losses come from: after the 400th sale, the 401 phone could happily be sold for \$450, lets say (if the price increased above \$400 we know from our supply curve we'll draw at least a few more producers in). Based on our demand curve we know that there would be plenty of people willing to pay more than \$450 (at \$400 tons of folks are left without phones, but willing to pay that amount). So there is a seller that'd be happy to sell at \$450, and many buyers who would line up to buy that product. But the price ceiling makes such transactions *illegal*! Price controls are indirectly quantity controls, and when we shrink the size of our market we are artificially restricting trades that individuals are happy to make (otherwise they wouldn't make them in the first place!).

As an exercise you should flip this graph and study a price floor (a price above the equilibrium). Show that this leads to a product surplus. Notice that even though many are supplied, because there are only a small number of buyers at high prices, the market quantity—the actual number of trades—will still be

If this point confuses you, you should ignore it, its not critical.

below the equilibrium. Therefore, the logic that we shrink the size of the market holds even if we impose the opposite price restriction.

Taxes are a bit more difficult to analyze—however you should have a hunch that because they mess with prices in a well-functioning market, they certainly can't help. What does a tax do? If its on the sellers, it raises their effective costs by $$t$—if previously 100 were supplied at $2, now those 100 producers would need $(2+$t$). This can be represented by a shift of the supply curve up by $t.

Figure 10: Taxing Sellers

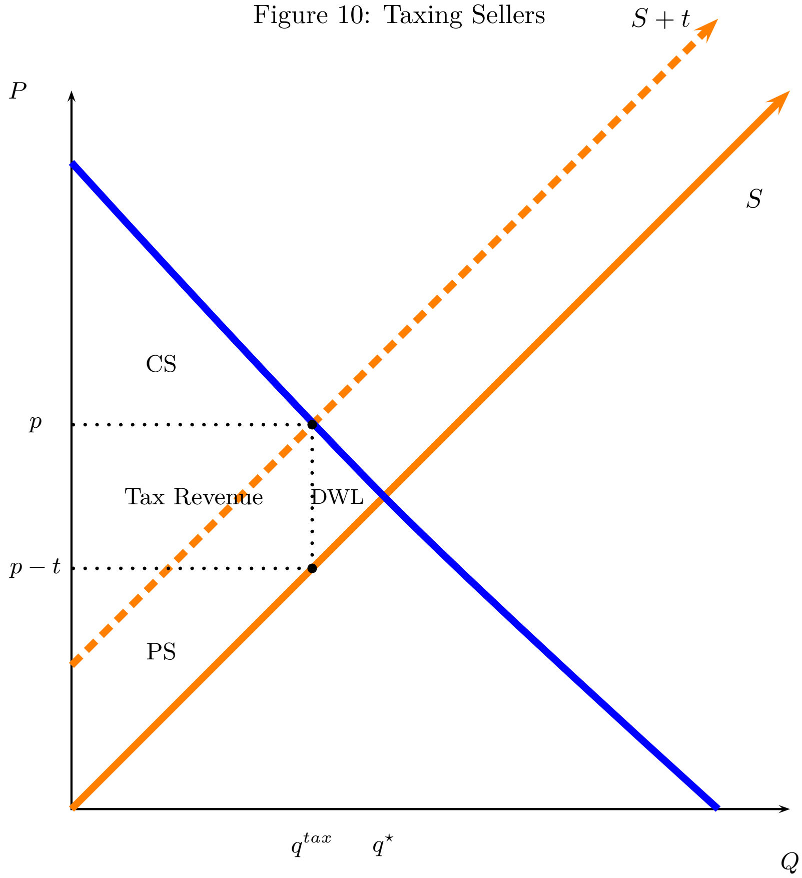

Takeaways:

1. The quantity in this market reduces (relative to the untaxed outcome, $q^{\star}$)

2. The price consumers pay rises; the price producers recieves falls (each by less than the tax)

   - We call this "sharing" the tax burden

3. Consumer (CS) and Producer (PS) surpluses shrink; we introduce a new area: government revenue

   - The base of the "revenue" rectange is the quantity bought/sold; the height is the tax.
     Base × height = area of a rectangle = revenue

4. Deadweight loss is introduced because we produce too little of this good

This analysis will have all of the same takeaways if you taxed buyers instead. Rather than shift the supply curve up, you shift the demand curve down by the amount of the tax. If consumers used to be willing to give sellers x for a good, they're now only willingly to give $x - t$ because they need to pay that t in taxes. Try this exercise in reverse: see that you get the same after-tax price paid by consumers, the same take-home pay for producers, and the same q^{tax}.

Externalities Externalities are a case where some of our main normative assumptions break down.

> **Externality**
>
> An effect that a transaction has on individuals who do not participate in that transaction.

A classic example of an externality is pollution. If when a good is produced some pollution is produced as a by-product, then my decision to pay you to produce that good for me has an impact on individuals outside of that transaction. I get no say as to whether you buy a gallon of gas and emit a bit of pollution into the Norman atmosphere. The situation I'm describing is one where *private* costs do not equally *social* costs. In all of our examples earlier we assumed the only costs were the ones producers bore. Importantly, because producers/consumers respond to their incentives, these external costs have no force to keep production low. This is in contrast to costs the seller's do in fact face—if those go up we expect the supply curve to shift left.

To represent this difference, we draw 2 supply(/cost) curves: the standard private one, and a social cost curve that includes both private costs and external costs. Specifically, if the externality (such as pollution costs) can be represented monetarily as being worth \$$e$, for example, that means the social cost curve will be the private cost curve shifted up by \$$e$. The market operates at a point, q^{Mrkt}, that has *more* production than is socially optimal ($q^\star$). We reason that the social optimum is $q^\star$ because any trade after $q^\star$ has total costs that exceed benefits. In this case the deadweight loss comes from producing too much of this good—hence it has a slightly different shape than the deadweight loss we saw earlier from underproduction.

The stunning takeaway we finished class with was to notice that if we tax suppliers by \$$e$ it will shift the private supply curve to overlap with the social cost curve! So the adverse effect of a tax (to shrink the market size) can be used to correct markets that operate at an inefficiently large size. This requires estimates of the externality, which is much easier than the government needing to estimate $q^\star$. This strategy is called "internalizing" the externality—if your activity causes harms to others (social costs) we can get suppliers to account for this by taxing them at this level. Plus, this raises revenue; a win-win.

Figure 11: Production Externality

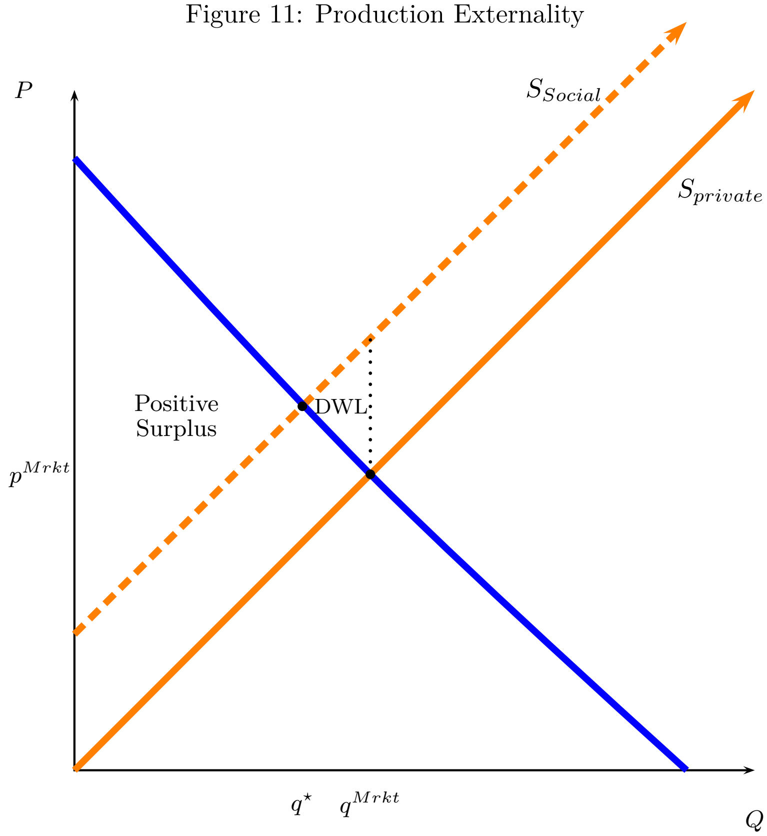

Principles of Macroeconomics–Honors

GDP, Inflation, and Living Standards

Instructor: Dr. Kevin Kuruc

December 10, 2020

Sections I and II of the course laid out two of the most important tools for macroeconomists: production possibilities frontiers and a method for analyzing markets. We'll now jump more thoroughly into a study of macroeconomic questions. To do so we must get our hands dirty with some data concepts.

1 Gross Domestic Product

Questions surrounding economic health, both for the individual and the economy at large, are typically concerned with **Gross Domestic Product**. For example, when we say an economy is doing well, or growing fast, we almost always mean "the gross domestic product has increased." When we say that a policy will be "good for the economy," likewise, we are usually suggesting that it will increase gross domestic product. An intimate understanding of this concept is therefore necessary for assessing most macroeconomic claims.

> **Gross Domestic Product**
>
> The total market value of all final goods/services produced within an economy's borders, within a given time period.

Gross Domestic Product, or **GDP**, captures all production within our national borders within some time period. Because GDP asks us to add up different goods and services, we need a common yardstick by which to compare products. It would be obviously silly to just add up the number of apples, oranges, and computers produced. One slightly more reasonable, but admittedly still bizarre, proposal would be to add up the total weight of everything produced so that refrigerators and cars count more than pillows and paper. Or maybe, if we were in a strictly agragian economy, it wouldn't be crazy add up total calories produced. In our economy, however, it is most sensible to measure products by their **market value,** or prices. Prices are driven, in part, by household's valuation of a product (i.e., demand); this method therefore counts products more or less based on their consumer-driven value.

A second important feature of GDP—as noted in the definition—is that we only count "final" goods. Final goods are purchases made with the intention of use, not re-selling. For example, imagine Amelia and Matt are in the business of selling chocolate milk. Amelia is a producer of chocolate milk, Matt is a retail seller of chocolate milk. The following chain of events transpires in my pursuit of chocolate milk (non-dairy, of course):

1. Amelia produces chocolate milk (assume she can do this at no monetary cost)
2. Matt buys chocolate milk from Amelia for $1.
3. I buy chocolate milk from Matt for $2.
4. I consume it.

When the chocolate milk was sold from Matt to Amelia, this was not a final good—Matt resold it to me. When I bought it, I consumed it, making it a final good. The reason we only count final goods is to avoid double counting. If we had counted the sale between Matt and Amelia we would have counted two sales of the same chocolate milk; even though it was only produced once. This chain of transactions should only be worth $2. One product was produced that a consumer valued at $2. This makes counting GDP much easier than you might have first thought. We don't have to keep track of all the interconnected links within an economy: we can just add up everything bought and sold as a final good because that final transaction embodies the sum of production and sales within that chain.

Explicitly, the method I'm referencing is called the **expenditure approach** to computing GDP. This method sums all expenditure on final goods in an economy. Because, in theory, market prices equilibrate to ensure that everything we've produced is sold, our total expenditure on products will equal the total value of everything produced.[1] The most straightforward way to compute GDP is to look at the purchases of all final goods and add them together.

Alternatively, we can add up the value of all income. Naturally, this is called the **income approach** to computing GDP. Income is defined in a very broad way here: it includes all wages, interest income, rental income, profits, and any other income source you can think of. This isn't just a boring statistical point. Understanding that, at the macroeconomic level, **our collective income is our collective spending** is far from trivial. As nobel prize winning macroeconomist and New York Times columnist, Paul Krugman, likes to remind his readership "my income is your spending, and my spending is your income." To see this equivalence in our simple chocolate milk example, try using the income approach to compute what this transaction contributes to GDP. Matt's income (profit) in his trade with me is $1 ($2-$1); Amelia's income

[1] In reality, some adjustments are made for unpurchased goods if they become "inventory." Inventory messes with our GDP calculation because the products are made in one year, but sold in another year.

(profit) is \$1; total income generated from this transaction is \$2! Not coincidentally, the total income generated was equal to the money the final consumer injected into the process.

A simple economic model called the **circular flow diagram** formalizes this relationship. Rather than add my own diagram here, I'll direct you to the first 3 minutes of this corny video. We skipped adding the government, so I'll focus on the first simple diagram (which he calls a matrix). Notice that money flows in a complete circle: the consumer spending (bottom right) flows to businesses as revenue, which flows out of businesses as costs, which flows into households pockets as income. If you're thinking "businesses don't pay out *all* of their revenue as costs—some is left as profit, right?" Right! Except that profit gets paid to households too. The business owner(s) is part of the household side of the market too, so yes, *all* firm revenue flows out to households.

The takeaway from this model is that total expenditure from households circulates around to become their total income. That's the importance of looking at things from a macroeconomic perspective, rather than a microeconomic perspective considering only one household, business, or even industry. Based on this model, it doesn't make much sense to say something is "good for businesses, but bad for households" or vice-versa. We're linked in this deep way, such that anything good for businesses becomes good for households, and vice versa. What people might mean when they complain in this way is that the policy in question is good for *business owners* and bad for typical employees. Therefore, their concern must stem from inequality, which is an important distinction. It is the case that those who own businesses (including stock holders) are wealthier, on average, than those who do not. It would make our public debates much clearer if we acknowledged this fact, and asked what we think about inequality and whether some policy would promote it; rather than frame it around firms vs. households.

Moving past this simplified diagram of our economy, we can be comfortable that **total income = total spending on final goods.** Therefore, we can pick either method for computing GDP. It is, in my opinion, easier to conceptualize this from the spending or production side. To that end, we can separate GDP, Y, into 3 main expenditure categories.

$$Y = C + I + G \tag{1}$$

- C = Private Consumer Spending
- I = Private Investment
- G = Government Spending (on final goods!)

C is straightforward. It represents total consumer spending and makes up the largest share of our consumption/production. **Investment** is a little trickier. The way economists use the term investment is different than how financial markets, or normal people, use the term.

Economic Investment

A purchase of *real resources* for the objective of more production/consumption/income in the future.

Economic investment has to involve the allocation of resources away from consumption today, for the benefit of future production/consumption. One example is a taxi company purchasing a fleet of new cars: this is spending that allows the company to produce more taxi-rides in the following month/quarter/year. Another example is a family building a new home: they are purchasing the production of a house that they can live in next quarter/year/decade. **Economic investment does not count merely financial transactions.** When I buy a stock or a bond, for example, I am buying a financial asset that is already owned by someone. Nothing new was produced. Financial markets help companies raise money to make economic investments, but these "financial investments" themselves do not count towards I in GDP because nothing was produced.

G requires a similar level of nuance. Most things we call government spending in our public debates do not count in this G. Namely, transfer payments do not count. For example, unemployment insurance, social security, food stamps, etc. are not G because they aren't purchasing final goods, these programs merely transfer funds. They go to some household that then either consumes or saves this money, so these transfers get included indirectly once they are spent or saved. This G only represents government expenditures on real goods or services: providing a police force; maintaining roads; upkeep of parks; Covid testing sites; etc.

Despite some of those nuances, this accounting identity is nice because we can easily imagine our own household income being split between these sources. When income is earned, some is taxed for the government (G), the rest is either spent (C) or saved (I).[2] That's it—there's no where else your money can go. Because GDP is the sum of all expenditures, if we account for everything we can spend money on, we've gotten GDP.

There is one small "fudge factor" we need to add in the real world when we perform our expenditure approach to GDP. One major simplification we've made so far is that we've assumed the economy is "closed." That is, we've assumed our economy does not trade with outside sources. Because GDP is supposed to capture our production, it would be a mistake to count expenditure on foreign goods to GDP; that expenditure **does not** circulate back to our household income. The term we add to this equation is exports minus imports: "Net Exports," NX. Anything we export should be added—its production we don't count when we look at U.S. spending, since foreign consumers purchased it—and we should subtract imports because they were produced abroad but part of our expenditure.

$$Y = C + I + G + \underbrace{(EX - IM)}_{NX} \tag{2}$$

[2]Technically, the investment your household does may not count as "economic investment," but money you save is available for firms to borrow and make these investments, so its a good approximation.

It's conceptually simple to compute GDP now that we recognize it is merely the total expenditure in our domestic economy on final goods/services. Let's go ahead and try actually doing this for a very-very small simple economic unit.

The small economic unit we'll study here is a hamburger-stand-nation-state, Burgtopia. The only final goods it produces and sells are Burgers, B, and milkshakes, M. We have data on the prices and quantities sold of these goods for the years 1950 and 2020.

Table 1: Sales Data for Burgtopia

Year	P_B	Q_B	P_M	Q_M
1950	\$0.50	100	\$0.33	75
2020	\$5.00	150	\$5.00	150

To compute the GDP for Burgtopia in each of these years, we want to sum the **market value of all transactions.** In 1950 the market value of all transactions was $\$0.50 \times 100 + \$0.33 \times 75 = \$75$. In 2020 the market value of all transactions was $\$5.00 \times 150 + \$5.00 \times 150 = \$1500$. We can add this new column to the table.

Table 2: Nominal GDP for Burgtopia

Year	P_B	Q_B	P_M	Q_M	GDP
1950	\$0.50	100	\$0.33	75	**\$75**
2020	\$5.00	150	\$5.00	150	**\$1500**

There are a few things to notice here. First, as we've stated many times now: computing GDP is very simple, conceptually. Once we identify final goods sales, we just total up their expenditure—nothing fancy. In practice it can be challenging to count all of these transactions, but the point here is that once we've identified them we don't need any complex math. Second: according to Table 2 GDP went up a ton over this period. A change from $\$75 \rightarrow \$1500$ is a 20x increase! And because we're hoping that GDP serves as a measure of production, that would seem to imply that production has increased by 20x. But…burger production only went up 1.5x ($100 \rightarrow 150$) and milkshake production only went up 2x. What's going on? GDP seems to be failing us as a measure of production here.

Over time *prices* change, which makes GDP difficult to compare across years. In 1950, we are assigning the value of a hamburger \$0.50; in 2020 we assign the same item a value of \$5.00. Like the United States economy, Burgtopia prices have risen over time. That's a problem if we want to count hamburgers equally (as we should!) when comparing production in different years. To solve this we're going to define two types of GDP: real and nominal.

Real/Nominal GDP

Nominal GDP is the market value of production in a given year, *using that year's prices.* Real GDP is the market value of production, using a *fixed set of prices.*

Nominal GDP is what we've already compute above: the market value of production using the prices in that year. My helpful trick for remembering this is that **N**ominal GDP is **N**aive GDP—we unthinkingly use the prices in that year without adjustment. **Real GDP**, on the other hand, takes a fixed set of prices—so that hamburgers are counted the same in every year—and computes what the market value of that year's production would be under these fixed prices. These fixed prices usually come from one of the years we have data from, we call this the "base year." Because we already have nominal GDP, lets just go ahead and add a Real GDP column to this table and compute it. For this example we will call 2020 the base year[3], and so we will use the 2020 prices to compute the market value in both years.

One thing you should quickly notice is that for 2020 we've already computed the market value using 2020 prices: that's what nominal GDP was. **In the base year real GDP = nominal GDP** because the base year prices *are* that year's prices. For 1950, we need to perform the following calculation:

$$\underbrace{\$5.00}_{2020 P_B} \times \underbrace{100}_{1950 Q_B} + \underbrace{\$5.00}_{2020 P_M} \times \underbrace{75}_{1950 Q_M} = \underbrace{\$875}_{1950\text{Real GDP}}$$

We've taken the prices from 2020 and asked "how much would the production in 1950 be worth, if prices had not changed." This number now represents something we can compare between 1950 and 2020.

Table 3: Real GDP (using 2020 prices) for Burgtopia

Year	P_B	Q_B	P_M	Q_M	Nom GDP	Real GDP
1950	\$0.50	100	\$0.33	75	\$75	**\$875**
2020	\$5.00	150	\$5.00	150	\$1500	**\$1500**

The increase in Real GDP is now just under 2x ($1500/875 \approx 1.7$). This makes much more sense. Burger production went up 1.5x, milkshake production went up 2x; the economy's growth across all industries was something in between these two numbers. These calculations would look different had we chosen 1950 as the base year—try this for practice! In fact, even the growth rate of real GDP would be different: you won't get exactly 1.7x as your increase.[4] Picking the base year can therefore be consequential for making claims about how much the economy grew, but the differences will be rather small and the reasons for this discrepancy are beyond the scope of our course. What you need to know is that comparing nominal GDP across time captures two phenomenon: the fact that production changes over time *and* the fact that prices change over

[3]I need to give you the base year for any problems, there's no "right" way to pick one that you need to know.

[4]Think about why this would be the case: it has to do with the *relative* value of goods.

time. Because GDP is designed to measure production, we've focused here on eliminating the effects of price changes to compute Real GDP. In the next subsection we will focus on the changes in prices, or **inflation.**

1.1 Shortcomings of GDP

Before that, however, we should conclude our discussion of GDP by noting some of its shortcomings. GDP, in reality, is only supposed to measure economic activity, but many link Real GDP with economic "well-being." Is this reasonable? The "GDP haters" will rightly point out that GDP fails in a number of ways:

1. Population sizes make GDP bigger, without raising well-being!
    - Can easily correct this by using per capita GDP (GDP divided by population).
2. Non-market activities not counted
    i. Non-traded activities like child-care
    ii. Black market trades that are not recorded
    iii. Leisure
3. Environmental degradation not counted
4. Individual conditions that make life worth living: literacy, political freedom, health, etc. not counted

In class we saw that when we adjust for population, GDP per capita is highly correlated with all of these measures. Rich places have cleaner environments, and more leisure time, and better health, and higher literacy, etc. So as an approximation, despite these short-comings, GDP per capita is a useful indicator of well-being.

In light of our earlier discussions on Production Possibility Frontiers, this shouldn't be terribly surprising. When GDP increases, that necessarily means our societal set of choices must be increasing. It would be very odd if, when our choices increased, we didn't chooes for life to get better in multiple dimensions at once.

2 Inflation

Returning to our question of how Nominal GDP increases, we saw that rising prices–or inflation—is the other way in which this could rise.

Inflation

The rate of increase (in percent terms) of average prices in the economy. It is typically denoted by the variable π.

We can even decompose the growth rate of nominal GDP into the growth rate of real GDP, g_{RGDP} and the growth rate of prices, π. Yes, that's the same symbol that's usually used to represent $\pi = 3.146\ldots$, but economists use it as a variable to denote inflation. Don't get confused. Because its challenging to measure real GDP, in practice we measure nominal GDP and then correct it for inflation.

Inflation measures the changes in the overall price level in the economy, so we are going to reverse the process we used to compute real GDP. For real GDP we held prices fixed, and only studied how quantities changed. To compute inflation, we're going to ask how the prices changed for a fixed set of goods/services. We'll refer to this fixed set of goods/services as the "consumer basket" and ask how much this identical hypothetical basket would cost in different years.

Formally we follow the steps below:

1. Choose some basket of goods, the "consumer basket"

2. Choose a base year

3. Compute the price of the basket in the base year

4. Compute the Consumer Price Index (CPI) in all years: 100 $\times$ the price of the basket in a given year $\div$ the price of the basket in the base year

5. Compute inflation: percent change in the CPI

For a concrete example let's use the same economy of Burgtopia and compute the CPIs and the inflation rate. First: I have to tell you both the basket and the base year. Let's call the basket: 1 burger and 1 milkshake. This is a reasonable choice since it reflects a standard consumer purchase in 2020 (the same number of each were sold). Let's call the base year 2020 since that's the year we're in (but, again, there's really no good reason to choose a base year so that's information I'd have to give you).

- The cost of the basket in 2020 was: $\$10 = 1 \times \$5 + 1 \times \$5$

- The CPI in 1950 was: $8.3 = 100 \times \frac{\$0.83}{\$10}$

- The CPI in 2020 was: $100 = 100 \times \frac{\$10}{\$10}$

  - The CPI in the base year is **always** 100!

- Total inflation between 2020 and 1950 was $100 \times \frac{100-8.3}{8.3} = 1,105\%$ (Prices went up about 12x)

That's really all there is to computing inflation. Solve for CPIs using the price of a basket in different years, and use those CPIs to compute how much the price of a basket has changed.

3 Real Wages

We've seen so far that one way to compare living standards across time (and places) is to (i) measure nominal GDP; (ii) correct it for price differences, like inflation (or different currencies); and (iii) correct it for population size. That gives us a broad sense of how much stuff was produced, and therefore available, per person. This is an imperfect statement because goods/services within an economy are not equally distributed, but for the most part provides a pretty good measure of living standards. However, one of the features this continues to miss out on is leisure time—an economy can increase its GDP per capita merely by increasing the hours we all work. Because we like both our economic output *and* our leisure time, its not obvious that this is improves welfare, even though it increases real GDP per capita.

This is where my favorite measure of economic well-being comes in: the concept of a **real wage**.

> **Real Wages**
>
> How much stuff that can be purchased for each hour of work.

Real wages are a particularly intuitive concept: they ask how much you can buy (the "real" part) for each hour you work (the "wage" part). In the early 1900s, for example, almost all time worked translated to purchasing necessities (food, shelter, etc.). Now, for the average American, almost all of our necessities are covered by a few hours of work. This fact represents an expansion of our *choices*. We could choose to have the same material lives as our great-grandparents and enjoy 10 hour work weeks; or we could work hours like they did and live far better lives materially. These choices are what economic growth is all about!

When comparing real wages over time, the question we're asking is whether hourly wages have increased more than the general price level. Equivalently, we're interested in the ratio between pay and the CPI. In equation form, we can relate the nominal wage W (your wage in dollar terms) to the real wage w, in the following way.

$$w = \frac{W}{CPI} \tag{3}$$

We can use Equation 3 to ask a few questions. First: you could ask whether real wages have increased over time by knowing W in both years and CPI in both years. If $\frac{W_{2020}}{CPI_{2020}} > \frac{W_{2000}}{CPI_{2000}}$ that means wages—relative to the price level—have increased. The dollars you earn at work can now purchase more of these "baskets." We could be even more exact. Suppose I knew that wages in 1950 in Burgtopia were \$0.75 per hour and I wanted to know how much someone would need to earn now to have the same purchasing power as someone in 1950 Burgtopia. Well, the CPI in 1950 was 8.3; the CPI in 2020 was 100, so the ratio we want to solve is below.

$$\begin{aligned} \frac{\$0.75}{8.3} &= \frac{\$x}{100} \Rightarrow \\ \$x &= \$0.75 \times \frac{100}{8.3} \approx \$9.00 \end{aligned} \tag{4}$$

It would take \$9.00 to buy the same amount as \$0.75 bought in 1950. So, if that's what wages were now, citizens of Burgtopia would be just as well off. If we surveyed the economy and found wages were actually \$18 per hour, then we'd know that they are twice as well off! They only need a wage of \$9.00 to earn—in real terms—what their grandparents did, but they earn twice that amount meaning their salaries now go twice as far.

We can use Equation 4 to ask about moving other prices through time. In general, if something cost $\$p_t$ in year t and we want to know how much we would expect it to cost in year t' we just convert it using the CPI ratio.

$$p_{t'} = p_t \frac{CPI_{t'}}{CPI_t} \tag{5}$$

To put this in more concrete terms: if the average prices in the economy doubled over some period ($\frac{CPI_{t'}}{CPI_t} = 2$) then we'd expect this item to cost twice as much now as it did ($p_{t'} = \frac{CPI_{t'}}{CPI_t} p_t = 2p_t$). Because the CPI computes average changes in prices, this won't always be dead-on, but it gives a good approximation of how much we'd expect it to cost.

4 Takeaways

1. GDP is a measure of our production over time, since our production is what we are able to consume, this is an approximation for well-being over time once we look at per capita versions of GDP.
2. We compute GDP as the market value of all goods and services purchased in a given year
   - But we could just as well ask about total income, because this is equal to total expenditures
3. Because prices are changing over time, we need to correct for this in our computation
   - Real GDP corrects for these prices
   - Changes in nominal GDP can be split between real GDP growth and inflation
4. We compute inflation by asking how the price of a representative market basket changes over time; this is an index of how far a dollar goes
   - When the CPI is low; the basket costs little; so a dollar goes far
5. We then saw we can use the CPI concept to ask how much our labor buys us
   - These "real wages" may be the best measure of economic well-being, it measures whether an hour of work affords more or less than it used to; when this is rising, our choices for how to spend our lives expands.

- Notice: inflation makes no one worse off if wages rise too! As this onion article demonstrates, it makes little sense for people to complain about how cheap things used to be if they don't account for the fact that wages were also low!

Principles of Macroeconomics—Honors
Theories of Economic Growth

Instructor: Dr. Kevin Kuruc

March 14, 2021

Section III of the course had us thinking about how to measure the economic well-being at both the national level (real GDP) and individual level (real GDP per capita, real wages, etc.). Using these tools, we then looked over time and across the world at living standards. Massive differences in GDP per capita stood out. Not only are Americans much much wealthier than South Asians and Sub-Saharan Africans, but Americans are also much much wealthier than their grandparents. Because this wealth brings with it so many other things that make life worth living—health, literacy, culture—there may be no more important area of study than the causes of economic growth. This section of the course introduces you to the field of growth economics.

1 Theories of Growth

1.1 Malthus Model

Before getting to theories of growth, its important to cover the most famous theory of *non-growth*. This was, after all, the norm for thousands of years prior to the mid-1700's. Thomas Malthus (1766-1834) proposed one such theory as to why humanity had never—and likely would never—better its lot. His theory rests on principles that make quite a bit of sense, so despite that it has been spectacularly wrong over the last 200+ years, it continues to have many proponents in the modern day.

To introduce any of these theories, we need to define the **aggregate production function.**

Aggregate Production Function

A mathematical description of how total societal inputs determine total output. Generally:

$$Y = F(tech, capital, labor, land, \ldots) \tag{1}$$

This function works like any other function. Some combination of inputs—here technology, capital, labor, land—determine total output (or Real GDP) in that economy. The following theories will begin by simplifying this production function to capture important features of the economy, and then use that simple function to study how our economy might operate.

The insight that Malthus' began with was an uncontroversial one: there are diminishing returns to labor in the aggregate production function. Specifically, he had in mind natural resource constraints, as many do today. The economy was mostly agricultural in Malthus' time, so his thinking was along the following lines: there is only so much good land to farm, so the more people we add to the population, the further onto the fringes of arable land they need to work. Therefore those last workers don't raise as many crops for themselves as the original folks working the better land. A sketch of this situation can be seen in Figure 1.

Here Y is total output, L is the labor supply, or total population. In these long-run growth theories, a simplifying strategy is to ignore unemployment because we trust that labor markets will roughly sort this problem out. $100\sqrt{L}$ is chosen to bring some concrete math to this using our standard diminishing returns

Figure 1: Output vs. Population in Malthusian Model

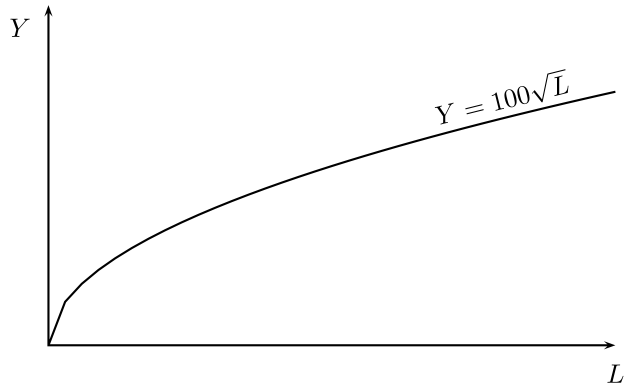

function ($\sqrt{L}$). Additionally, it captures Malthus' observation that the only "variable" was L; he lived in a time where technology, capital, and natural resources were unchanging. The shortcut mathematically is to lump those unchanging factors into a constant multiplicative term. As is common by this point in the course, our graph here formalizes that when L is large, the slope of this curve is falling—an additional person at this end of the curve produces less output than when L is small.

Uncontroversially, output rises with each additional person. But what's happening to output per person is less obvious. Our macroeconomic theories will, hopefully, have something to say about both total production *and* living standards. Define output per person (or GDP per capita) as $y = \frac{Y}{L}$. We can go ahead and chart out the relationship by picking different values for L, computing Y, and then computing $\frac{Y}{L}$. This is the approach in Table 1. Alternatively, we can solve directly for y in the following steps.

$$\begin{aligned} y =& \frac{Y}{L} = \frac{100\sqrt{L}}{L} \\ y =& 100\frac{L^{\frac{1}{2}}}{L} = 100L^{-\frac{1}{2}} \\ y =& \frac{100}{\sqrt{L}} \end{aligned} \tag{2}$$

The following table documents the numerical results.

Table 1: L, Y, y in Malthusian Model

L	Y	$\frac{Y}{L} = y$
1	100	100
4	200	50
9	300	33.33
25	500	20
49	700	14.28

The first two columns there produce Figure 1—output increasing when L increases, but at a diminishing rate. When L moves from 1 (billion) to 9, we add 8 billion people and 200 (trillion) dollars of output. Moving from the row with 25 to 39 was an addition of 19 billion people, but still only that 200 trillion in output.

The last column, plotted against L though, tells a very different story: GDP per capita is actually *shrinking* as the population grows.

Key Point #1 in Malthus Model: Because our natural resources are fixed ⇒ output has *diminishing returns with respect to labor* ⇒ standards of living *fall* with population growth.

Figure 2: Output *per capita* vs Population in Malthusian Model

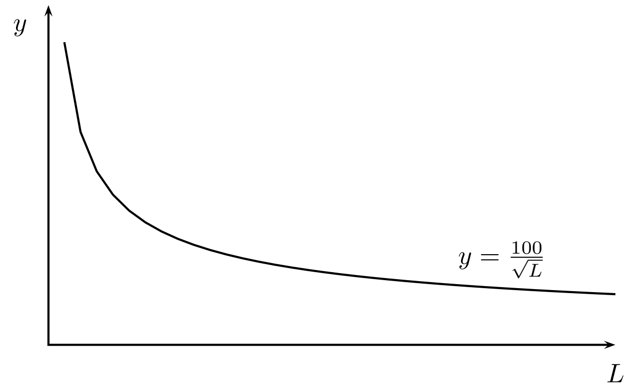

This takeaway is fairly intuitive. If we only have so many resources, of course adding people to split them between decreases the amount any individual can have. Malthus' broad points have therefore remained popular for 200 years. This is in spite of the theory's spectacular failure for the years following his conjectures.

Malthus made a second prediction about the dynamics of this model. He assumed that when people have extra wealth, they use these resources to raise more children. Hypothetically, if you were looking at your finances and saw you had money left over after supplying enough food for your current family, rather than say "great, lets eat a little more next year!" you'd instead say "great, we've got enough to have another baby!"[1] When you combine this statement about human nature with the reasonable claim that there are diminishing returns to labor, you end up with a gravitational pull towards poverty. Any time you have enough output per person to have an extra child (increasing L), you've drawn yourself further down that y curve, thus erasing the standard of living you just had. This will continue until some "subsistence level" of output—a level where you just have enough to survive, but no more—$\underline{y}$. I've labeled the maximum sustainable population as L^{max}—if more people than that existed, y would be too low to even keep everyone alive.

Figure 3: Subsistence Dynamics in Malthus Model

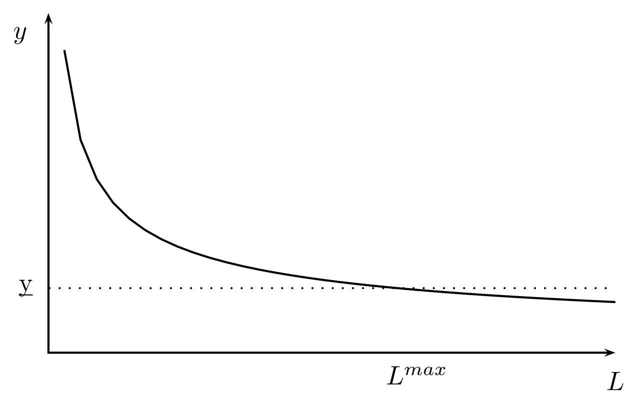

Key Point #2 in Malthus Model: We are drawn from either direction towards L^{max}. If the population is smaller than L^{max}, our standard of living is high, so we have babies, which lowers our standard of living next year. If we're still below L^{max}, there will still be extra wealth, which will still draw us towards that

[1] This claim sounds a bit ridiculous, but it actually has sound biological backing—we've evolved over millions of years to really, really want to produce off-spring.

intersection. If we had too many people and managed to exceed L^{max} somehow, there'd be too little food to go around, and we'd starve ourselves back to L^{max}. **Therefore, humanity is doomed to perpetually toil near the subsistence level.**

Quite the optimist Malthus was! Luckiliy for us, something happened that nullified Malthus' key premise. We began expanded both our technological capabilities and our physical capital which has allowed us to produce more even with fixed natural resources. The poor guy wrote down a good model based on the last 1000 years of economic life, but as the ink was drying on his *An Essay on the Principle of Population* the world changed. The 200 years since Malthus wrote have been characterized by an explosion in both world population levels *and* standards of living—exactly counter to what his model claimed was possible. The pessimists (and environmental scientists) among us might say something like "well now its all great, but Malthus's key principle was right—we are exhausting our fixed resources which will cause future generations to lead miserable lives!" That may be right; only time will tell. We do seem to be seriously mismanaging our fossil fuel use, and this will harm many future people. But the view of most economists is that life will continue to get better and better, year after year. To describe this formally, we'll need to enrich the Malthus model.

1.2 Solow Model

To enrich the Malthus model, the Solow Model formalizes long-run economic outcomes in quite broad terms. This does not prevent the model from leading to surprisingly detailed and insightful conclusions. The Solow model is premised on the idea that GDP (also referred to here simply as "output") can be described as some combination of capital, K, labor, L and technology A. We'll denote this formulation of the aggregate production function as:

$$Y = AF(K, L) \tag{3}$$

A stands in for anything that would result in two economies that have the same capital and labor inputs having different GDPs—how efficient workers are, how politically stable the economy is, how open it is to trade, how fertile the soil is, etc. For the purposes of our class, we will assume that $F(K, L) = \sqrt{K}\sqrt{L}$.

So far there has been nothing ground breaking—all we have done is written out a fairly general way in which output is produced. And in fact, this function is not even nice to work with since it has 2 non-linear terms (K, L) which makes it difficult to even think about how to start graphing this. The beauty of the Solow model is that it allows for a very nice mathematical description of GDP per capita, $\frac{Y}{L}$, which is ultimately what determines an individuals standard of living.[2] To see this, start by dividing both sides of Equation 3 by L and then working through the resulting expression. [**You do not need to know how to replicate these steps for an exam.**]

$$\frac{Y}{L} = \frac{A\sqrt{K}\sqrt{L}}{L} \Rightarrow \qquad \frac{Y}{L} = \frac{A\sqrt{K}\sqrt{L}}{\sqrt{L}\sqrt{L}} \Rightarrow$$

$$\frac{Y}{L} = \frac{A\sqrt{K}}{\sqrt{L}} \Rightarrow \qquad \frac{Y}{L} = A\sqrt{\frac{K}{L}}$$

Using lower case letters to denote per capita terms (so that $y = \frac{Y}{L}$), we can rewrite that last equation as:

$$y = A\sqrt{k} \tag{4}$$

Equation 4 makes it clear that GDP per capita, y, is now only a function of technology, A, and capital per worker, k. Importantly, it is subject to diminishing returns to k, as can be seen in Figure 4.

There are two ways to get richer in this economy. Either A can increase—that is the economy finds more efficient ways to turn k into y. Or k can increase—that is, capital accumulates. These cases are depicted in

[2] Recall our China vs. US example. Chinese GDP is nearing US GDP, but US GDP per capita still is far larger than Chinese GDP per capita.

Figure 4: Function for per capita GDP

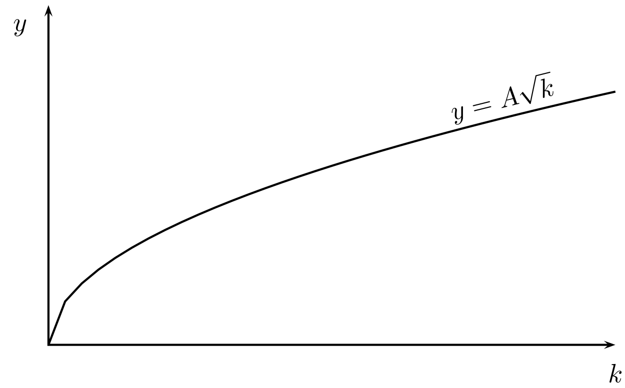

Figure 5. In 5a technology has increased, and so the same level of capital produces more GDP per capita; in 5b capital per person has increased and so there is more output per worker.

It may seem challenging to increase A. If A represents how fertile an economy's soil is, for example, that would be very challenging to improve. It is therefore much more straightforward to increase the capital stock; we know how to build more buildings, or cars, or machines, etc. So this simple model suggests a powerful way to increase our wealth: save a fraction of GDP and use it to add to the capital stock.

Figure 5: Getting Wealthier in the Solow Model

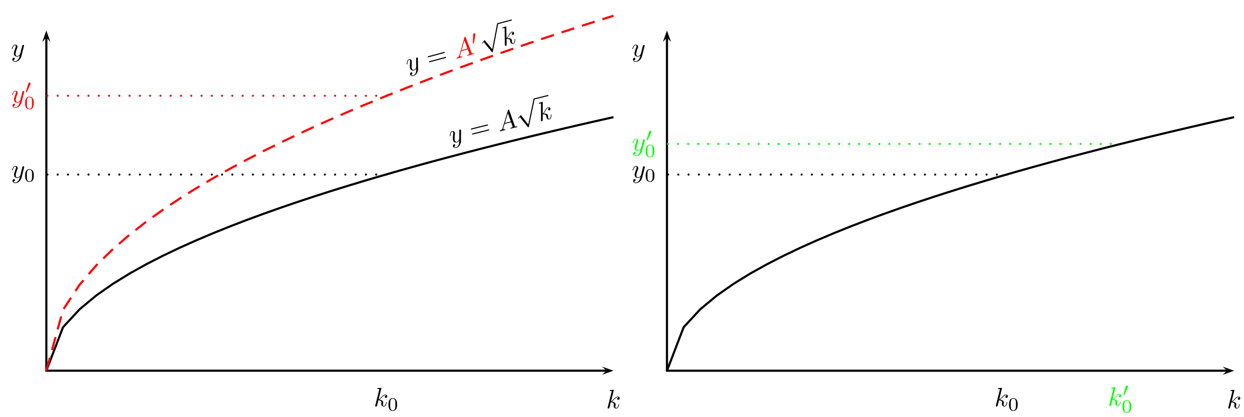

(a) A increases to A'

(b) k_0 increases to k'_0

In order to analyze how savings can increase our wealth, we need to introduce two new terms. s is the savings rate of the economy: a number between 0 and 1 that describes what fraction of GDP is saved for improving the capital stock. δ is depreciation rate in the economy: a number between 0 and 1 that describes how quickly capital wears down if nothing is spent to keep it up. Imagine what would happen to a building if you spent nothing year after year on up-keep—it deteriorates until it is no longer useable. As you can probably infer here, some savings will be necessary *just to keep up our current standard of living.*

Formally, we will bridge these concepts for an equation known as the "law of motion" for capital. This relationship describes how capital today evolves into capital tomorrow.

$$K_{t+1} = \underbrace{(1-\delta)K_t}_{\text{Remaining } K} + \underbrace{sY_t}_{\text{Total Saving}} \tag{5}$$

Equation 5 uses t subscripts to denote how variables are related over time. This is called a **dynamic** relationship. Equation 5 is read as "capital **tomorrow** equals the undepreciated capital **today** plus what we save (our savings rate × production **today**) to restore our old capital and build new capital." It may help to keep in mind that (δ, s) are numbers close to zero, like 0.1; so $1-\delta$ will be something close to 1, like 0.9 or 90%. $(1-0.1)K_t$ would therefore be "90% of the capital stock in t." This equation can be divided through by L to translate these variables into their per capita counterparts, k_{t+1}, k_t, y_t. These are the equations we will work with.

$$k_{t+1} = (1-\delta)k_t + sy_t \tag{6}$$

$$\text{or:}$$

$$k_{t+1} = k_t + \underbrace{sy_t - \delta k_t}_{\text{net savings}} \tag{7}$$

Equation 7 rearranges Equation 6 to draw attention to an important concept: **net savings**. It is the amount that total (per person) savings, sy_t, exceeds total (per person) depreciation, δk_t. **Net savings is, therefore, how much we *add* to the capital stock in a given period.** As Equation 7 makes clear, if net savings are positive then $k_{t+1} > k_t$, if net savings are negative then $k_{t+1} < k_t$, if net savings are zero then $k_{t+1} = k_t$.

Net savings, therefore, deserves special attention as it determines how capital tomorrow will relate to capital today—and recall, more capital means the economy can produce more GDP per capita tomorrow. Net savings is the difference between sy_t and δk_t, so it might help to put these on a common graph. Before looking at the figure, think about how we would do this. First, notice that we need a common variable to graph them against (in this case k_t), so we need to rewrite $sy_t = sA\sqrt{k_t}$. This is a function with diminishing returns so it shares the shape of the production function in Figure 4. In other words, because s is a number smaller than 1, and is constant, its as if we're just shifting down the curve in figure 4 by s% at each point. However, δk_t is going to be linear in k_t. How do we know this? Every time I add a unit to k, I increase depreciation by δ—two houses means twice as much depreciation as one house. *Therefore, the depreciation curve will not have a diminishing shape.*

Figure 6: Gross Savings vs. Depreciation

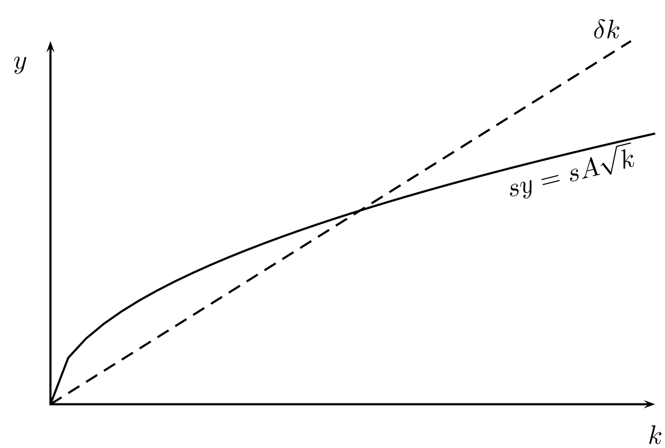

This combination of shapes tells us that when k_t is small—look at the left half of Figure 6—$sy_t > \delta k_t$. When this is true, net savings is positive. When net savings is positive $k_{t+1} > k_t$; the economy is growing. If

k is large (right half of graph) $\delta k_t > sy_t$, so net savings is negative. There is so much capital in this economy that the amount saved is not even enough to keep up with repairs on all current capital. When net savings is negative, $k_{t+1} < k_t$; the economy is shrinking.

Understanding this allows us to analyze the dynamics of an economy that starts with some initial level of capital k_0. Figure 7 plays forward an economy with very little initial capital, k_0. The distance between total savings (sy_t) and depreciation (δk) are the net savings, and so this vertical distance is the amount of capital this economy adds. At first net savings are large, so the jump from $k_0 \to k_1$ is large. Later this gap shrinks, so the capital increase is small.

Figure 7: Dynamics of an Economy Starting with Low Capital

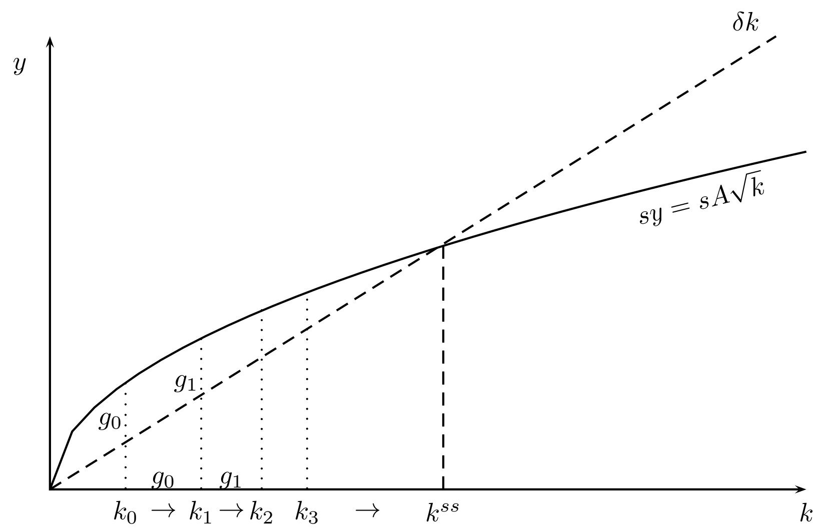

The capital is headed for that intersection point—but what happens when it reaches there? Net savings shrinks to 0 so this economy *stops growing.* This is called the "steady state" level of capital since net savings is 0 which implies k is steady; $k_t = k_{t+1}$.

The powerful lesson of this simple diagram is that **by itself, accumulating capital can only lead to temporary, not permanent, growth.** When this growth runs out, the economy is as rich as it can sustainably be with its current level of A and s. Keeping in mind that k^{ss} tells us about the steady-state level of GDP per capita, y^{ss}, it might be useful to solve for what k^{ss} is.

Start with the idea that the steady-state is where net savings is 0, which allows us to start with the following equation.

$$\begin{aligned} sA\sqrt{k^{ss}} &= \delta k^{ss} \Rightarrow \\ s^2A^2k^{ss} &= \delta^2(k^{ss})^2 \Rightarrow \\ k^{ss} &= \frac{s^2A^2}{\delta^2} \end{aligned} \tag{8}$$

The steady state level of capital in this economy is $\frac{s^2A^2}{\delta^2}$. It is increasing in A, s since these terms are in the numerator, but is decreasing in δ. This all should make sense: if A is higher, the economy is productive and can afford to maintain more capital, if s is higher the economy is saving more and can afford to maintain more capital, but if δ is higher, capital depreciates faster and is more costly to maintain.

How would shifting one of these variables look on our graph? Consider Figure 8 where the savings rate in the economy has increased from $s \to s'$.

The gross savings curve increases since our savings are higher, and accordingly the intersection has shifted up. Will the economy immediately jump to the new steady state? No. When the savings rate first increases,

Figure 8: Steady State Change as s Changes

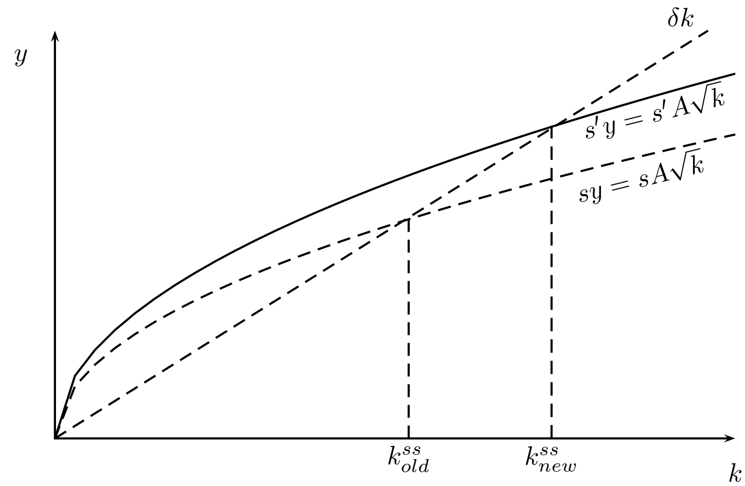

capital is still at its old level. Since savings are much higher than depreciation, the capital stock increases that period and begins following the dynamics traced out in Figure 7. Its as if the economy has been artificially reset to a lower than appropriate level of capital.

We can go one step further; we can plot how economic growth, g evolves in response to this increase in saving rates. Look back to Figure 7. When an economy is below its steady-state, as it will temporarily be when savings rates jump, growth was high (g_0 in Figure 7). As the economy gets closer to its steady-state, growth rates fall ($g_0 > g_1 > g_2$ in Figure 7). The same pattern will hold here: the vertical gap between the savings curve starts large and shrinks to zero—because this gap represents growth, we know that growth will start large and shrink to zero as well.

Figure 9: Growth Rates in Solow Model

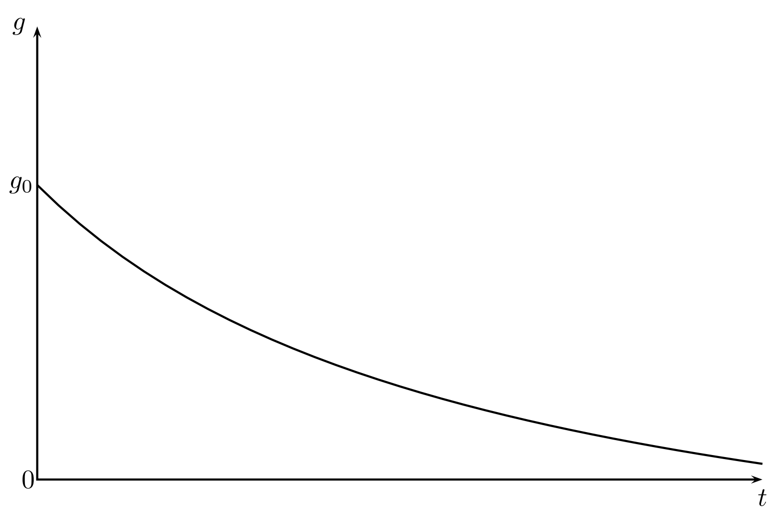

Figure 9 plots this evolution. When savings jump, growth immediately jumps. Then it fades back to zero, because once we get to the new steady-state growth will be zero again. Notice that *time* is on the x-axis. Like in the two bonus questions on the homework, this asks you to graph the evolution of these varialbes as time passes—forcing you to change your mindset a bit.

Principles of Macroeconomics—Honors
Business Cycles

Instructor: Dr. Kevin Kuruc

April 28, 2021

The Solow Model and our study of long-run growth lead to rather sensible conclusions: economic growth is driven by (i) capital accumulation and (ii) technological progress. Capital accumulation cannot continue indefinitely on its own, so we're left rely on innovation and increases to our economic efficiency to explain sustained long-run growth. But everything in Section IV of the course was described as being smooth—recall the graph of what growth looks like following a change in A or s: jumps up, decays smoothly back to zero. And the Solow Model allowed for *negative* growth only if we found ourselves above the steady-state; a scenario that was all but impossible in the first place.

But even looking at modern U.S. history—one of the most stable economies to ever exist—and we see the economy contract in the late 1970s, early 2000s (briefly), late 2000s, and as recently as March 2020! Our Solow Model is inadequate for providing insight into these important economic events. In this last section of the course we will examine the facts, causes, consequences, and policy responses to **business cycles**.

1 Business Cycles

See powerpoint lecture.

2 Aggregate Demand and Supply

Here we are going to build a model that links total output, or GDP, Y, with the price level in the economy, P. This model will fit well with what we know about the long-run: that the only way to produce economic growth is through changes in A or K; but it will give us better insights into the short-run. Namely, that changes in consumer/investor confidence and other non-A, K factors *can* (temporarily) influence economic output. As we saw in the Business Cycle Powerpoint, the key to these different effects of changes to the "demand side" of the economy will run through a channel where prices P are temporarily distorted in the short-run and correct themselves only over time.

2.1 Aggregate Demand

Our model of aggregate supply and demand is going to look very similar to a model of supply and demand for an individual market. But because we are now representing the entire economy, the variables have a new interpretation. P, on the y-axis as always, represents the total price level in the economy; something like the Consumer Price Index we studied in Section III of this course. When this P rises or falls that is what we called inflation or deflation. Y, on the x-axis (I know, confusing), represents total production, or real GDP. In this section of the course we will stick with real GDP, Y, rather than per capita real GDP, y. In the short-run populations are fairly stable, so we can ignore that nuance here.

The **aggregate demand curve** (Figure 1) describes how much total expenditure is desired at any given price level, P. Recall that our total expenditure in the economy can be decomposed as: $Y = C + I + G$. Like demand curves in individual markets, the aggregate demand curve is characterized by a negative relationship between P and Y. The reasons for this, however, are quite different than the reasons we get a negative relationship in individual markets. In individual markets, we typically think of substituting across goods; when Nikes are expensive some consumers switch to Reeboks. That can't happen here. If P goes up, that's a situation where the *economy-wide* price level rises. If prices are all drifting up, substituting across goods isn't possible. And even if it were, switching goods can't decrease Y because the new good is also counted in GDP.

Figure 1: Aggregate Demand with an Increase

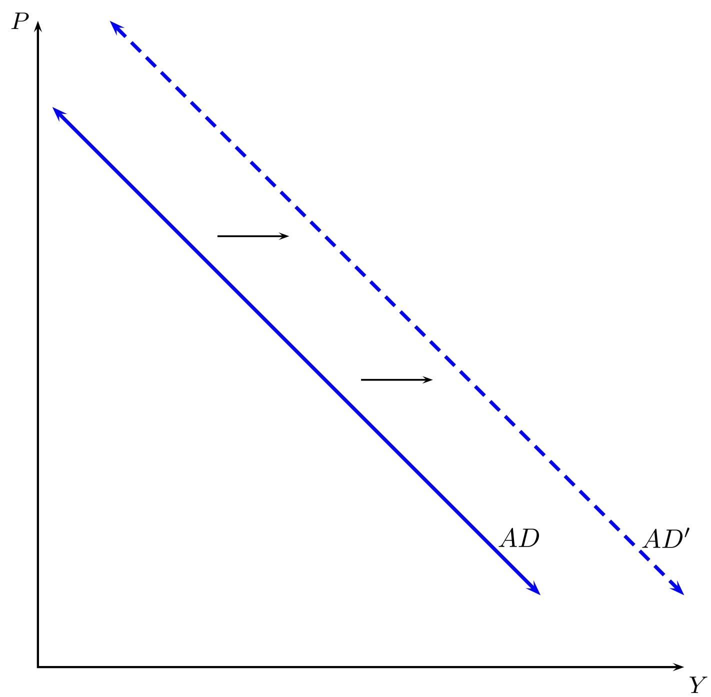

The inverse relationship in the AD curve can best be conceptualized as the **real balances effect**.

Real Balances Effect

When P falls, household *real* wealth increases; this wealth effect causes individuals to demand more goods/services.

Household wealth was accumulated in past months and years, and in the present moment sits in accounts as a **nominal** value. If I look at my own bank account, it tells me how many \$ I have (a nominal variable), not how many Beyond Burgers I can afford (a "real" variable). If the price level in the economy falls, the number in my bank account does not change! All the sudden that means my bank account is now worth more Beyond Burgers—I've become instantly wealthier because price have fallen. Wealthier people buy more stuff, so if the whole economy gets a bit wealthier because prices fall, we all demand more stuff in aggregate.

An alternative way to understand this negative slope of the aggregate demand curve is through the concept of a **pre-determined nominal budget.**

Pre-Determined Nominal Budget

Consumers decide at the period's start how many dollars ($P \times Y$) to spend. If P rises $\Rightarrow Y$ falls; and vice-versa.

Households are busy and, at least in the short-run, may pre-determine their spending rather than rethink their choices every time prices change. If this is the case, P and Y will be inversely related because $P \times Y$ is already determined by consumers. If P rises, Y has to fall for them to hit their budgetary objective.

Investment, I, will also respond positively to falls in the price level through the **real return effect.**

Real Returns Effect

The real return on a project is determined by the expected profits $\div$ current costs. If P falls, current costs fall, so the real return of projects rise.

Recall from our loanable funds market, the demand for investment is determined by how many projects have a return on investment that is higher than the interest rate. If costs today fall, more projects clear that threshold because costs are in the denominator of an ROI calcuation. If more projects have an ROI > the interest rate, investment increases.

Government spending is the only component of total expenditures that we do not expect to be sensitive to prices. The government picks their actions ahead of time, and if prices rise, they happily spend more; if prices fall they happily spend less. However, because we expect C and I to respond inversely to P, total demand for goods/services responds inversely to P.

Also like our individual demand curves, we might expect this aggregate demand curve to increase (rightward shift) or decrease (leftward shift). Because this curve describes the relationship between the price level and total expenditure, anything other than the price level that influences total expenditure would shift this curve. Here are examples of scenarios where each of the respective components of expenditure ($C + I + G$) change for reasons other than the price level:

- Consumers may increase their budgets/expenditure if confidence rises or personal patience falls
- Firms may increase their investment expenditure if interest rates fall, or if their outlooks for the future increase
- The government may increase its expenditure when parties turnover, or when there is a military conflict

When any of these happens we get rightward shifts as depicted in Figure 1. If any of those components decrease we would instead see a leftward shift.

2.2 Aggregate Supply: A Tale of Two Horizons

The aggregate supply side of the model is where things get interesting, and conceptually a bit tricky. It's already been hinted at that prices will respond differently in the short and long run. And since businesses ultimately set prices, this must enter on the supply side of our analysis. Therefore, the relationship between prices in the economy and production decisions—our aggregate supply curve—will differ in the short and long run.

Long Run Aggregate Supply (LRAS)

We'll start by looking at the long-run, because its extremely simple mathematically and builds from what we learned in the Solow model. The Solow model was built on the premise that our aggregate production function was only a function of A, K, L—specifically, $Y = AF(K, L)$. Stare deeply into that equation. Is there a P anywhere in it? No! **In the long-run, the amount our economy is capable of producing is independent of the price level.** We can't change our sustainable level of production just by lowering or raising prices in the economy—prices are a way to count things, how could that impact what we're capable of producing?!

Because Y is completely unaffected by P, our LRAS curve will be perfectly vertical (see Figure 2). No matter the price level, we're capable of producing $Y^N = AF(K, L)$. Y^N here is what we call the "natural" rate of output. It represents where the economy would operate once we've settled (back) into our productive capacity.

Figure 2: Long-Run Aggregate Supply

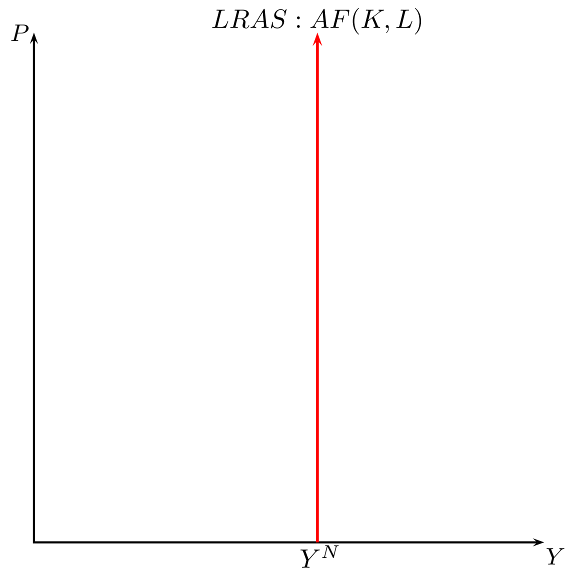

This may seem like a paradox: here I'm claiming that the changes in P can't impact output in the long-run; but before I claimed that prices need to be "right" to organize us efficiently. How can they both not matter and matter? The conceptual difference is between the *level* of prices and the *relative* prices across goods/services. **Relative prices**—that is, how prices compare to one another—determine how we organize ourselves and our resources. The **price level**, P, does not. For example, suppose only the price of coffee rises: we'd drink less coffee. That's because the relative price of coffee has increased. If instead the price of coffee went up at the same rate that the prices of everything else was going up, there wouldn't be cheaper alternatives to substite towards. Additionally, if *all* prices really did go up, our wages would rise as well, so the "true" (or opportunity) cost of coffee wouldn't have risen.[1] To summarize: P is the price level, which we have good reason to believe cannot impact production in the long-run.

With this long-run aggregate supply curve, we can examine the aggregate equilibrium of this economy. Figure 3 plots the LRAS and the AD where point 'A' is the equilibrium. Just like in individual markets, the equilibrium is where these curves intersect. We can then analyze how shifts in both aggregate supply and aggregate demand impact total output, Y, and the price level, P. Changes in aggregate demand, lets say from a boost in consumer confidence that leads to more desired spending, shifts us from point A→B in panel (a) of this figure. There has been **no change** in output, prices just rise. How could this be? We hear all the time on the news how important consumer confidence is. The best way to see this is that in the long-run, because relative prices are "right", there are few resources sitting around unused ('idle' resources we can call them). If there are few idle resources, then a boost in demand—which doesn't change the availability of resources—can't expand production. In panel (b) we see that a shift in Y^N can change output. Just as in

[1] Remember, the opportunity cost is "what we give up to obtain something"; if the price of coffee remains the same as tea, then we still give up one cup of coffee for one cup of tea.

the Solow model, long-run increases in living standards come from changes in A or K; in this model L could cause this shift as well because this line represents total production, $Y^N = AF(K, L)$.

Figure 3: Aggregate Demand and Supply Shifts in the Long-Run

(a) AD Increase (b) AS Increase

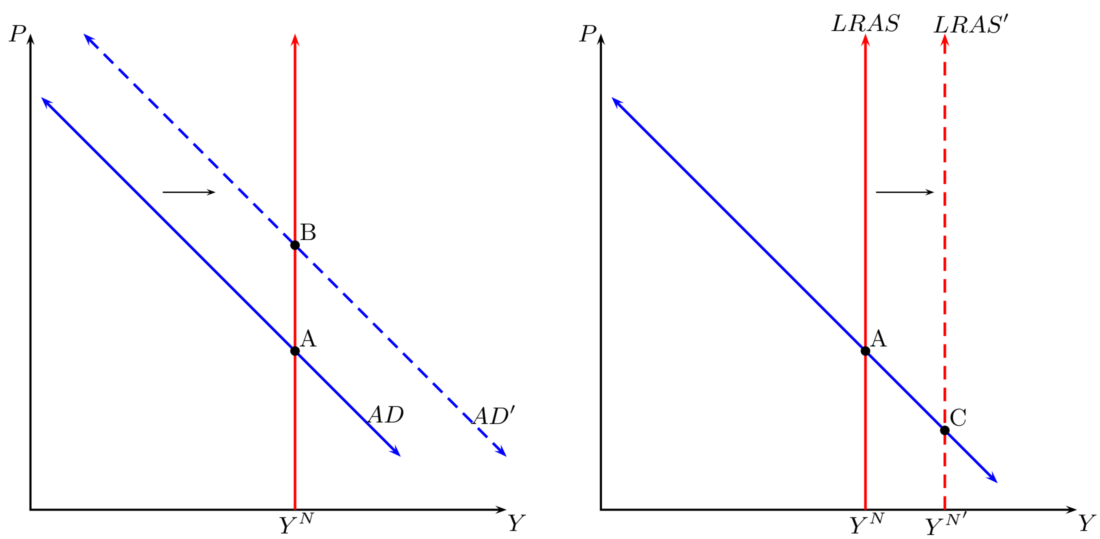

To summarize the long-run of our macroeconomy: relative prices are right (that's our long-run assumption), therefore the price level P cannot impact how we organize ourselves, and therefore cannot change Y. Because of this our economy can only expand or contract from expansions or contractions to A, K, or L.

Short Run Aggregate Supply (SRAS)

If we want any hope of changes in demand flowing through to changes in output, something must be different in the short-run. The assumption we will add to our short-run representation of the macroeconomy is that some (or all) prices are "sticky."

Sticky Price Assumptions

Some (or all) prices do not change in the short-run, regardless of economic conditions.

If prices don't respond to changes in economic conditions in the short-run, they will not continually organize us in a way that our resources are used efficiently. We will represent "sticky prices" in two distinct ways. First, we will consider a situation where prices do not change at all in the short-run. This is an extreme assumption, but it may not be terribly unrealistic in the very short-run. If aggregate demand increased today, its possible prices wouldn't adjust for a month or two. Businesses need some time to learn that demand is truly up before they find it in their best interest to increase their prices. If they've only seen slightly increased consumer demand for a week or two, they may believe this is fleeting rather than persistent. If demand decreases, business owners may be even more stubborn. There is some psychological evidence that when demand falls business owners can be very reluctant to lower their prices, even if this would in theory increase their profits by restoring sales. In any case, it appears to be a feature of the real world that prices are slow to respond to economic conditions, which we will formalize by saying "in the short-run, prices do not respond to economic conditions."

This representation of the SRAS can be seen in panel (a) of Figure 4. If prices truly do not respond to economic conditions, then prices must always stay fixed at their original level (here I've called that P_0). This form of "independence" is exactly the opposite of what we see in the LRAS.

Figure 4: Short-Run Aggregate Supply

(a) Case 1: Extreme Sticky Prices

(b) Case 2: Sticky Wages

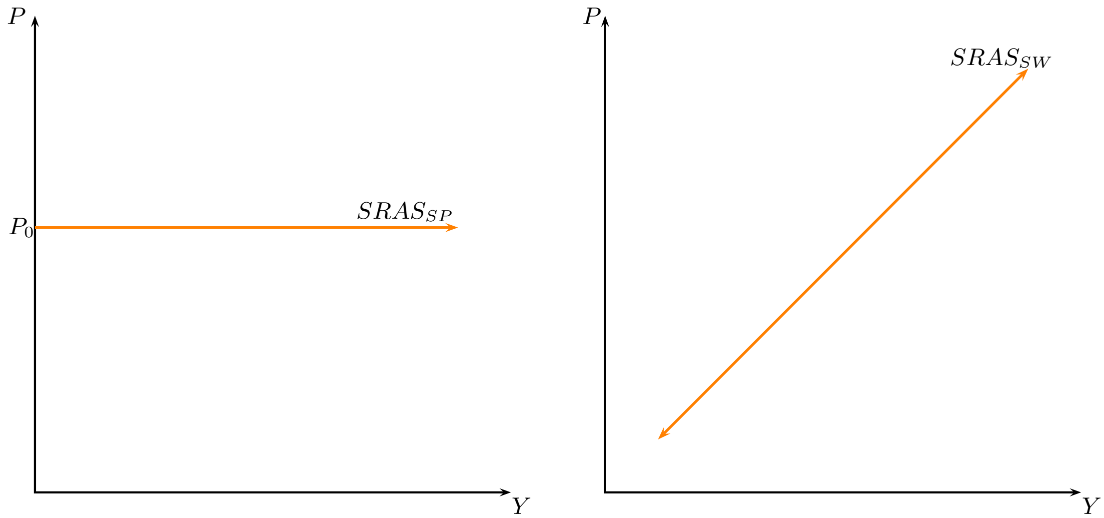

However, some economists think this assumption of sticky prices is so extreme that it is worthless. These economists instead believe that a very special form of sticky prices can better describe the short-run; these economists instead employ a **sticky wage** theory of short-run aggregate supply. The sticky wage theory states that, while prices for goods can change in the short run, wages cannot change. This certainly strikes me as more believable. Wages are negotiated at infrequent intervals, and typically these negotiations are anchored the wage of previous years. Wages really do seem to adjust very-very slowly, perhaps even on the timescale of years.

This assumption leads to an aggregate supply curve that is upward sloping. The reason for this is conceptually more challenging. The key is how profitability changes as P changes now that wages—a primary input price—are fixed. Imagine a scenario where the price you as a business can charge for your product rises. If wages are fixed, that means each product you sell has a temporarily higher profit margin; therefore, you are excited to expand production. If instead prices fell in the economy, but your wages were fixed, now each product is much less profitable—reducing production is the prudent decision. So, P and Y would have a positive relationship on the supply side.

With these descriptions of how supply looks in the short-run, we can analyze how shifts in demand impact the aggregate economy immediately after they occur. Figure 5 depicts a demand shift in the short-run under either of these assumptions. In both cases we see that output increases in response to the increase in aggregate demand. News commentary about how consumer confidence drives the economy may not be so wrong after all—at least if what we care about is the short-run. What happens with inflation is less straightfoward. If prices are perfectly stuck, there is no change in prices *by definition*. If instead we assume that wage stickiness is the defining characteristic of the short-run, this shift in demand moves us up and along the SRAS to the upper right. In either case, we would say there is "inflationary pressure" in this economy because prices have risen or they want to rise to restore Y^N.

In a scenario where AD falls instead, we would see output fall and deflationary pressure. This matches what was seen in the Great Depression (1929) and the Great Recession (2007)—output fell and inflation was extremely low. This is a key reason why many economists believe that these recessions are best described as "demand-driven" rather than supply driven.

Figure 5: Demand Shifts Change Y in Short-Run

(a) Extreme Sticky Prices

(b) Sticky Wages

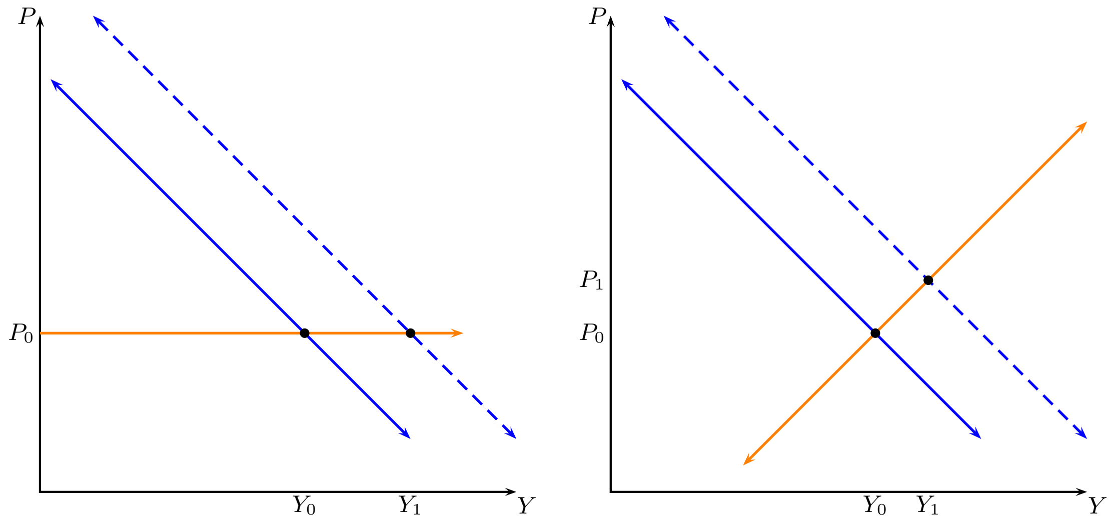

2.3 Linking the Short-run & Long-run

Obvious questions that arise from this short-run/long-run distinction are (i) "how short is the short-run?" and (ii) "what happens to get from the short-run to the long-run?" The answer to the first question is: maybe years! The Great Depression lasted quite a long time, as did the Great Recession. Before the Great Depression economists really only had a long-run model; John Maynard Keynes came up with the short-run model to try and understand the Great Depression. When his opponents were annoyed at him for claiming their ideas were irrelevant in the here and now, they responded that his theory wasn't important because the economy would adjust itself in the long-run so we shouldn't worry too much. To this Keynes made his famous reply "Yes, but in the long run, we're all dead."

More interesting from a modeling perspective is how we link these in a unified framework. These two models give such different predictions, how could we possibly bring them together in a nice way? Let's follow the progression depicted in Figure 6. Imagine we start as it looks in panel (a). Our "long-run equilibrium" is determined by where AD and the LRAS intersect—point 'A' here. If we assume that we're in an economy that has settled into its long-run, then prices will be at P_0. The SRAS curve says "because we're at P_0 right now, if something changes in the economy we will initially remain at P_0." Therefore, we overlay the SRAS where P_0 is. Things are then in a sustainable harmony, all three curves overlap so that short-run and long-run equilibrium are the same. This occurs when the economy is on a smooth trajectory for a long period of time.

If something unexpectedly changes—here the AD curve shifts right, perhaps because of an increase in consumer confidence—this stability is shattered. The short-run equilibrium is now at point 'B'; output has increased above its natural rate to Y_{SR}. The problem is that 'B' cannot be a long-run equilibrium, because the LRAS and AD curve do not intersect at 'B.' If it cannot be a long-run equilibrium, we cannot stay here indefinitely. Something's gotta give.

To see how this adjustment happens, consider what makes the short-run supply curve flat in the first place: we've assumed prices can't adjust in the short run. But eventually prices *will* adjust! And in fact, we know that because demand is up, the pressure on prices is also upwards. When prices finally "unstick" the SRAS curve shifts up to the new long-run equilibrium (panel c).

This entire process will look very similar with a sticky wage theory. The big difference comes from wages

Figure 6: Short-run → Long-Run

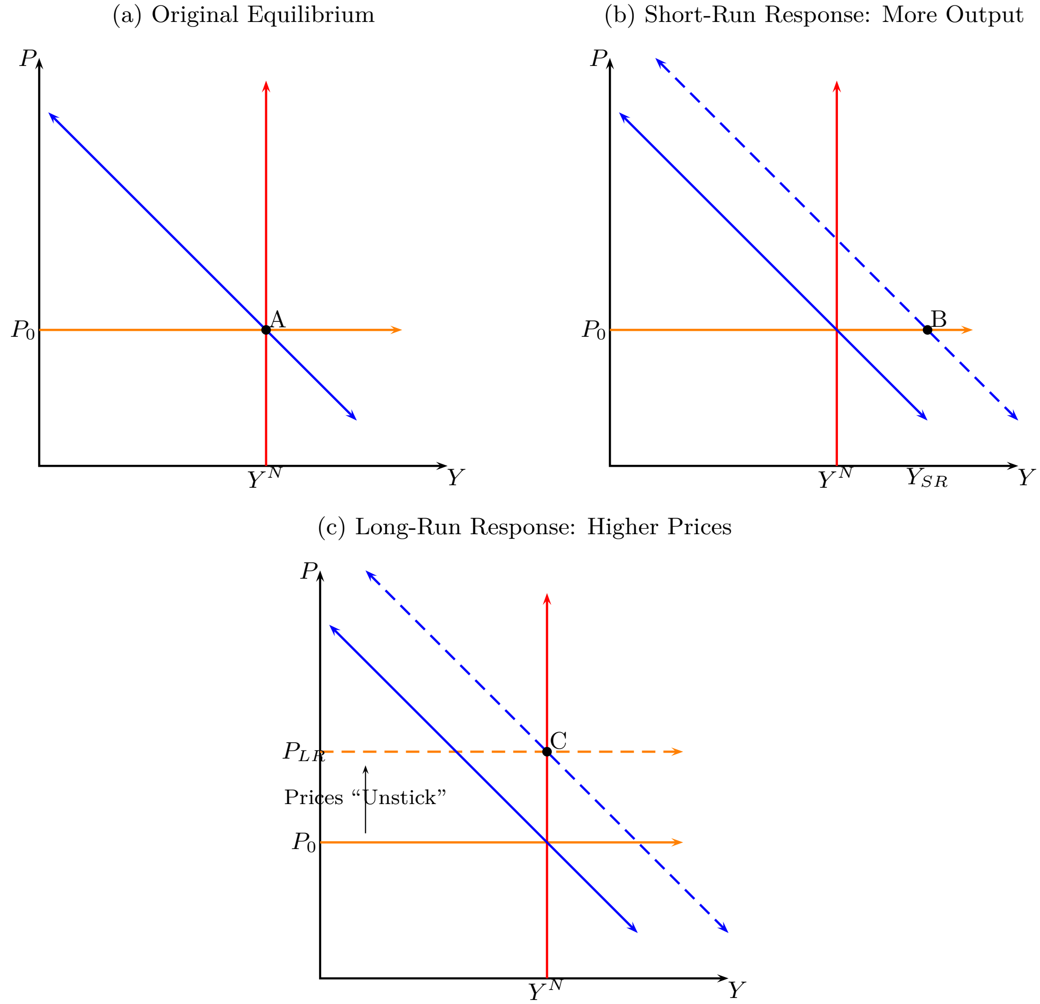

unsticking—you should try this. An increase in AD increases output and prices in the short run. To get back to a long-run equilibrium, wages unstick and increase. An increase in wages shifts the SRAS to the left. The reason for this is the same as why an increase in wages would shift the supply curve in a market to the left. Wages are a primary input price, if they rise it reduces the amount businesses are willing and able to supply at any price. That corresponds to a leftward shift in the SRAS.

Before moving to policy, it will be useful to explicitly state the main takeaway of this model. In the long-run, the supply side determines output. In the short-run, the demand side determines output. The transition between these phases involves prices eventually readjusting to bring us back to Y^N.

3 Fiscal Policy

With a unified framework for thinking of the short- and long-run, we will now turn to the question of macroeconomic policy. When short-run economic crises strike—such as a standard recession—we expect the government to step in with some sort of action. What might good government policy look like?

First, we will study what is called **fiscal policy** before moving onto monetary policy.

Fiscal Policy

Government policy regarding spending, transfers, taxes, and (as a result) the national debt.

The fiscal policy we will focus on will be broad; we'll say things like "government spending increases" or "taxes increased" without making any statement about which programs are funded or who faces the burdens of a tax. We're interested here in the effects of a bout of government spending or changes in tax policy on the business cycle. The answer to this question will be surprisingly simple now that we've mastered the AD-AS model.

To see why it will be conceptually simple how the government can influence the macroeconomic equilibrium, remind yourself about how we separated aggregate expenditures when we computed GDP all the way back in Section III of the course.

$$Y = C + I + G \tag{1}$$

Total aggregate expenditures are the sum of consumer spending, household and business investment, and whatever the government spends. Because our AD curve tracks the relationship between real spending Y and the price level P, our AD curve will shift around as G changes. If the government undertakes a large spending project, such as a Green New Deal (or the Original New Deal in the 1930s), that's an increse in G $\Rightarrow$ increases aggregate expenditures $\Rightarrow$ increase AD. Alternatively, the government can increase C and I by reducing taxes. C will directly respond to reduced taxes, but remember that I is determined by the rate of returns on projects versus the interest rate in the economy, so tax policy will be less effective in changing I. Those concerns are beyond the scope of this class; what you should remember is the takeaway that C and I will roughly increase when taxes decrease. All of these effects can work in the opposite direction as well: falls in G or increases in taxes will decrease AD.

That leaves the question of when the government should attempt to increase or decrease aggregate demand. Two points will help us answer this.

1. We've learned that changes in AD do not promote long-run changes in output. In our context, this implies that there is little long-run benefit from increasing G or lowering taxes.

2. Being below *or above* Y^N is not good, so the goal of fiscal policy should be "stabilization."

The first point is a nuanced one—there *are* government policies that can create long-run growth. But to do so, based on our LRAS curve (and the Solow Model), we must see increases in either investment (so that K increases) or A. So government investment in research will hopefully promote A, which will promote long-run economic output. The fiscal policy we will discuss here are claims that large increases in government spending, on their own, create jobs. This is only true in the short-run, because outward shifts in the AD can only keep us off our LRAS for so long.

The second point is also difficult to grasp for many. We are used to thinking that more GDP is better in all cases. But remember our Keynes article about work vs. leisure from the start of the semester—it's not always true that producing more is better. If these increases don't come from increases in A or K it just means we're working harder and stretching our existing resources further. As an economy we could, after all, increase GDP a lot by increasing working hours to 80 a week instead of 40. That wouldn't necessarily be good! The point here is: **we don't want output to exceed Y^N**. Additionally, this has the problem that we will return to our long-run equilibrium through a combination of falling Y and increasing inflation, a combination nobody likes. We will refer to episodes when short-run Y exceeds Y^N as an "overheating" economy to remind ourselves that it is not a good state of affairs.

What this leads us to, then, is the takeaway that:

i. If AD decreases so that $Y^{SR} < Y^N$, the government should either increase G or decrease taxes to undo this decrease in AD.

ii. If AD increases so that $Y^{SR} > Y^N$, the government should either decrease G or increase taxes to undo this increase in AD.

To stabilize the macroeconomy, the government should push in the opposite direction of the private sector. This is called counter-cyclical fiscal policy.

Counter-cyclical Policy

Policy that moves counter to the "business cycle." It is expansionary when the economy is down; contractionary when the economy is up.

In theory, either taxing or spending decisions of the government can get the job done, as long as they don't offset one another. For example, an increase in G *shouldn't* be offset by raising taxes—the point is to inject new funds into the economy. This requires the government to run a budget deficit in recessions to push AD up. When the economy is overheated the government should instead run a budget surplus. The more general result is then: **our federal deficit should be counter-cyclical; it should be large and positive during downturns and negative during economic booms.**

An Aside: Debt vs Deficit

The government's deficit is the year-to-year shortfall in revenue; the National Debt is the accumulation of these deficits. If a \$10 deficit is run every year for 10 years, a \$100 debt is built up. Our current debt is large because we've run deficits for the last 20 years.

This theory is all well and good, but there are a few big practical problems to running such wise fiscal policy.

1. Knowledge Lag: Takes time for Congress to learn about the severity of a crash, or how overheated we are. This is very hard to measure in real time.
2. Implementation Lag: Once Congress knows what's going on, they need to write bills, compromise, and get legislation passed. This takes a very long time.
3. Never Running Surpluses: Its easier to boost spending or cut taxes to stave off collapse, but its very hard politically to rain on the economic parade when we're overheating. In this way, we only do one half of the recommended policy. This results in an ever growing national debt.

So, within our model fiscal policy seems like a fantastic tool to stabilize the economy. In practice it is very hard to do correctly for both knowledge reasons and political reasons. That might sound depressing—we have an excellent tool at our disposal that we just can't quite use right. However, in a fit of accidental brilliance, we've built instantaneously responding counter-cyclical fiscal policy right into our laws. These go by the name automatic stabilizers.

Automatic Stabilizers

Fiscal policy tools that automatically and immediately push government funds into the economy during recessions, and remove them during booms.

What's more, you've probably heard of many of these without even realizing it. Take unemployment insurance. When the economy suffers a downturn, by definition people are getting laid off. When they get laid off, they call the government and file for unemployment payments. The government then sends these folks checks immediately. The government never needs to go out and measure how bad the economy is, or pass any new legislation; money starts automatically getting dispensed to households. Even better: when the economy is booming, unemployment benefits automatically fall since very few people are requesting them, and the program naturally gets cheaper. Its a program that works to automatically stabilize the economy. A similar story goes for welfare programs like food stamps. On the tax side, the same thing happens. Our government taxes income/profit/sales/etc. When the economy is crashing, they automatically reduce how much they collect in taxes; when the economy booms tax collection increases. Its a beautiful thing!

The takeaway then is:

1. Fiscal policy (in particular our deficits) should be counter-cyclical.
2. Aside from very big crises (Great Depression, Great Recession, Covid-recession) its almost politically impossible to get Congress to run fiscal policy correctly.
3. Our automatic stabilizers do most of the heavy lifting and do indeed push fiscal policy in the right direction over the business cycle.

4 Monetary Policy

Aside from Fiscal Policy, the monetary system may impact real GDP - and will certainly impact the price level. The monetary system is closely tied to the financial system, so we will discuss these in parallel. Overseeing this system we have a public, but independent, institution called The Federal Reserve System. The monetary system is so critical to the functioning of our economy, and the chair of the Federal Reserve System is so independent from political checks and balances, that some argue this individual is the 2nd most powerful person in the United States! When Janet Yellen made history by becoming the first female Fed Chair, she may have been the most powerful woman in the history of the United States; even more so than Kamala Harris is now.

The monetary system is fascinating. In the U.S. our monetary system is one of **fiat currency.**

Fiat Currency

A currency that has no *intrinsic* value.

The U.S. dollar is, intrinsically, worth nothing. Our paper money has no value aside from the value we all claim/believe it has. It used to be the case that you could turn in a US dollar for a fixed quantity of gold promised by the government, but under Franklin Roosevelt (FDR) we abandoned the so-called Gold Standard. **Money is only worth the value that others place on it.** We are inclined to think that money makes us wealthy, but it is purchasing power that makes us wealthy. If, during some national calamity, businesses and other individuals stopped wanting the US Dollar, it would make no sense for you to accumulate any. **Money is only worth the value others place on it.** This seems like a terribly fragile set-up; the only reason I want money is because I know that y'all want money. And the only reason y'all want money is because you know I want it. If either of our beliefs falter—say because I suspect the US government will collapse—the entire monetary system collapses. We refer to this as a **reflective equilibrium**.

Reflective Equilibrium

A context where the possible outcomes are characterized by everyone reflecting one another's beliefs. In the case of $, either we can all value it, or none of us value it.

This concept of a reflective equilibrium is best seen in a silly two-person example. Suppose you and I have a fiat currency that we use to make trades. I bother to collect it because I know you want it. And you do actually want it, but only because I want it. Our beliefs are self-reinforcing. If I stopped valuing cash, your only rational response is to also stop valuing it (why would you keep collecting it if I no longer want it?). So my belief that it is worthless is also self-reinforcing. In our real world you need more than a single person to start this chain, but the concept is nonetheless the same: our collective beliefs are self-reinforcing, making the system suprising vulnerable.

If our system was really as fragile as I am claiming, how could it be that we've gone for nearly 100 years off the Gold Standard without any destabilizing episodes? That's a great question. American democracy and political institutions were seen throughout the 1900s as the most longstanding and stable institutions in the world. Further, the Federal Reserve System that manages our currency is free from political influence. This ensures that the Fed does not bend to political whims; when we need to make difficult choices to keep our system stabilized the Fed can do so without worrying about angering voters. Paradoxically, because the

Fed *can* act in these unpopular ways means it rarely needs to. Just the credible promise that it will act responsibly is enough to keep the system from getting into trouble.

As a proof of concept that things can indeed unravel incredibly fast, we can look at periods of **hyperinflation.**

Hyperinflation

Episodes where entire economies lose faith in the purchasing power of their currency. Prices have been known to rise by *billions* (with a B!) percent per year, rendering the currency worthless in a very short period of time.

Before discussing these incredibly interesting episodes, we should formalize a link that has thus far been informal. **The value of money is equal to the inverse of the price level.**

Relationship between money & prices

The value of money, V_M is equal to the inverse of the price level, P.

$$V_M = \frac{1}{P} \tag{2}$$

If the price of goods in an economy doubles, the value of that currency has been sliced in half. If, as in hyperinflationary cases, prices spiral out of control, the value of money falls to zero. Money is worthless if it can't buy anything! Hyperinflationary episodes are unique and accompanied by both (i) political turmoil and (ii) mass printing of money by the central governments. Because the citizens lose all faith that the currency will be worth anything, no one wants to hold any of it, which means this prophecy becomes self-fulfilling and reinforcing. Linked here is a video describing Zimbabwe's hyperinflation as an illustration of how extreme these events can be.

Moving forward we are going to work under the asumption that we stay in the "good" reflective equilibrium; there's no use in studying money in a scenario where it is worthless (and these episodes are rare). We can the describe the relationship between the value of a currency and the quantity of money circulating using the **quantity theory of money.**

$$M \times V = P \times Y \tag{3}$$

The 4 terms in the quantity theory are as follows:

- M is the supply of money (currency) circulating in the economy; this quantity is determined by the Central Bank (Federal Reserve in the United States)
- V is the velocity of money. This represents how much the average dollar changes hands in a given year.
- P, Y represent the price level and the level of real GDP like they do in the AD-AS model

Equation 3 is what we call an identity; it is not a theory about a relationship, its an equation that must always hold, no matter what. To see why, consider what the right hand side of that equation equals. It is real GDP (all the stuff that is produced/purchased in the economy) times P an index of the price of all that stuff. If you think back to section III of the course, you'll remember that we call this **Nominal GDP**. Its the value of everything produced and sold in a given year. The left hand side is the amount of money circulating times the number of times each dollar changes hands. If there is \$10 circulating, and it has a velocity of 3 (its spent 3 times) that must mean there was \$30 of total spending in that economy. Therefore, the left-hand side is equal to total expenditures, or **Nominal GDP**. This is just like our earlier result that total expenditure = Nominal GDP = total value of production; same thing going on here.

With two important further facts layered into this equation, it can tell us about the relationship between M and P. First, we will assume that V is fixed: the velocity of money doesn't change over time. The second fact we will add here is that—in the long-run—the amount of money in circulation does not impact

Y. This is argument is again premised on our production function. The amount of money circulating cannot change real facts about the resources available for production. Finally, lets turn Equation 3 into one that is represented in percent changes. We reviewed this earlier, but when an equation is multiplicative, it can be represented in growth rates using addition.

$$g_M + g_V = \underbrace{g_P}_{\text{inflation}} + g_Y \tag{4}$$

Because $g_V = 0$ (we assumed that V never changes) we can rewrite this equation in the following way (now using the term π for inflation).

$$\begin{aligned} g_M &= \pi + g_Y \\ \pi &= g_M - g_Y \end{aligned} \tag{5}$$

The rate of inflation (or how much prices are rising) is determined by the rate of money printing (g_M) minus the rate of real economic growth, g_Y. The reason we subtract the real rate of economic growth is because prices would remain stable if the balance between currency and goods is stable. If the amount of stuff we have grows by 2% we need 2% more money to keep these quantities balanced. For technical reasons we want a slightly positive rate of inflation[2], the Fed targets 2% for example, so they increase the money supply each year by slightly more than the rate of economic growth. If the real economy grows at 1% and the Fed wants a 2% rate of inflation, it must increase the money supply by 3%.

This equation also lets us see the harm that can be caused by large increases in M. If M grows by 20%, and this money printing doesnt change that Y grows by 1-2% each year, then the inflation rate will be nearly 20%. In fact, because g_Y at the yearly level is fairly small, any episode of high inflation must come from excessive money printing. To quote Milton Friedman who won the Nobel Prize for this discovery: "Inflation is always and everywhere a monetary phenomenon."

Bringing this back to our original question: inflation is the rising of prices, and rising prices implies a decreasing currency value. So, **the rate at which money is printed is the only important factor determining the value of a currency.**

This result relies on the assumption that printing money cannot impact g_Y, which is a good assumption for the long-run. But think back to our study of fiscal policy and government spending, G. On its own, it increases AD which doesn't have long-run effects either, but we saw that it can be quite powerful in the short-run. The same thing will be true of printing money, because printing money will also operate by shifting the AD curve. It's long-run effects will be inflationary, but its short-run effects will help drive/stabilize short-run output.

Before getting to the $AD - AS$ model, its important to understand how the Federal Reserve "prints money", how that impacts interest rates, and then how the change in interest rates impacts the AD curve. The Fed influences the money supply through **open market operations**.

Open Market Operations

The act of the Federal Reserve buying (selling) bonds on open markets to increase (decrease) the money supply.

When the Fed wants to inject money into the economy, it purchases bonds on an open market that financial firms are trading on. In this transaction they take money they've had hidden in a virtual vault and distribute it to households/firms in exchange for bonds. Because they are paying the market rate for these bonds this never includes a "handout" of any sort. When the Fed instead wants to reduce the money supply, it sells bonds in this open market and collects money to put back into this virtual vault.

This process, rather than spurring economic activity by handing cash out, spurs economic activity by influencing the interest rate in the economy. By influencing the interest rate the Fed is able to increase

[2] Take more macroeconomic courses if you'd like to know why!

investment I and disincentivize household savings, which increases C. We'll revisit loanable funds market to study this.

The loanable funds market was our market for savings and investment that we studied right after the Solow Model. To refresh your memory: savers supply the funds businesses want access to for investment; banks are the intermediaries. The interest rate is the price in this market that balances savings and investment. What the Fed does is enter this market as a saver (or an anti-saver). If the Fed enters as a saver, it pushes the savings curve rightward. This reduces the interest rate. When the interest rate is reduced because of the Fed's entry into the market, some households no longer find it in their best interest to save, and so they spend instead. In total though, savings and investment increase. Therefore, the Feds entry increases I and C, which increases AD.

At this stage, its clear the Fed's monetary injections control the AD curve just as the Federal Government's spending and taxing can. The analysis on the $AD - AS$ curve will be identical. If the fed contracts the money supply by selling bonds, they enter as "anti-savers" (shifting the savings curve leftward). This raises interest rates, leading to an increase in household savings and a decrease in I; both reducing the AD curve. In the long-run printing money, which shifts the AD to the right, will only raise prices. This is just as we saw predicted by the quantity theory. In the short-run however, monetary policy can stabilize the AD curve, just as fiscal policy can.

The advantages of monetary policy rather than fiscal policy is (i) the level of expertise and control at the Fed and (ii) the lack of partisan politics. The Fed has hundreds of PhD economists monitoring the economy and informing the board of directors how to move interest rates. And the board of directors do not need to feel any political pressure, nor do they partake in partisan fighting like Congress does. For these reasons monetary policy tends to dominate fiscal policy in terms of how we fight recession and control inflation in practice. It is only in extreme circumstances that fiscal policy comes to the aid of monetary policy. Again highlighting the power the Fed Chair in the US economy has—they not only control the long-run stability of our currency, but they control the short-run stability of the entire macroeconomy.

At this stage we can summarize a few important points.

1. Fiscal policy and monetary policy—our two primary sources of economic policy—operate primarily on the demand side of the model $\Rightarrow$
2. They will not influence long run economic well-being (roughly measured as Y^N here).
3. However, because most recessions and overheatings happen because of demand shifts, these policies are (in theory) great for stabilization (assuming they operate "countercylically").
4. In reality, fiscal policy suffers from more than a few implementation problems and so is typically less useful (though automatic stabilizers are an exception to this).
5. Therefore, countercyclical monetary policy (expanding the money supply when the economy contracts) has been our primary stabilization tool of the last few decades.

There are two last policy relevant points to make about the Fed before wrapping up. First, the Fed has a dual mandate to keep both unemployment and inflation low. However, when trying to accomplish these goals simulatenously they run into a relationship known as the **Phillips Curve**.

Phillips Curve

The inverse relationship between unempolyment and inflation that policymakers can choose between arising from their control of the Aggregate Demand curve.

The Phillips Curve is the formalization that the Fed (and the fiscal policy) don't have good tools to get both unemployement and inflation down; they are in a bind where they need to choose between the two. To see this, look back at the AS-AD model with sticky wages. If the Fed can move the AD curve, but the AS curve is immovable, you can imagine the Fed as having the choice between any given point on the SRAS curve. If they choose high AD, they get high output but high prices (i.e., low unemployment paired with high inflation). If they don't like inflation, their option is to turn AD down to reduce output and prices (pairing high unemployment with low inflation). So, while the Fed is pretty darn good at stabilization, demand side policy has its limits because it is constrained to choices on the SRAS.

Everything we've covered so far captures what was true in US economic policy from about 1980-2007. The Fed got very very good at tinkering with interest rates to stabilize AD and therefore our entire economy. If they foresaw a slow down, they expand the money supply (and nudge interest rates down) to nudge spending back up before a recession hits. Likewise, they (mostly) increased interest rates as things overheated to incentivize households and firms to save, slowing economic activity back to its proper level. It seemed like this could go on indefinitely! Bob Lucas, then president of the American Economics Association, was so impressed he declared the problem of recessions basically solved: the bigger the initial scare, the bigger the interst rate drop; no problem!

However, something happened in 2008 that very few foresaw: the initial demand shock (i.e., leftward shift) was so big that the Fed dropped interest rates all the way to zero and that wasn't enough to restore economic activity. They reached what we know call the **Zero Lower Bound (ZLB)**.

Zero Lower Bound

The Fed's open market operations cannot push interest rates into negative territory, so they face what is known as a Zero Lower Bound on interest rates. If a recession is so big that its not fixed by the time interest rates are at zero, our primary stabilization tool can no longer help.

The reason the ZLB exists is intuitive. Think about what would happen if interest rates fell below zero: you'd be paying the financial system to hold your cash. Instead of paything them, you could just keep some duffle bags in your storm shelter, so there's an alternative option once banks pay a negative interest rate.

For reasons we won't get into now, the world has naturally drifted towards lower long term interest rates and this trend doesnt seem to be reversing. This means that even small recessions have the chance to push us to the ZLB; if interest rates are normally 1%, it doesn't take much expansion to push them to zero. This isn't just a fun technical problem, its really how monetary policy has operated since 2007 and will likely continue operating into the future. If this is your only macroeconomic course, you better know something about it!

In the face of the ZLB there are two options we have a society. One should be obvious based our prior discussion: if monetary and fiscal policy can do the same thing—influence aggregate demand—then fiscal policy better step up when monetary policy hits its boundary. This is imperfect for all the reasons we stated before. But ZLB episodes, for now, have been relegated to very big recessions (2007 + March 2020). In very big recessions we have some hope of Congress getting their stuff together and passing fiscal stimulus; and indeed they did with the American Reinvestment and Recovery Act (ARRA) in 2009 as well as the covid stimulus we've seen. But if the ZLB is going to be a common phenomenon, we'd prefer having monetary policy tools that can deal with this reality rather than kicking the issue to fiscal policy.

The Fed has come up with a very clever tool that we'll refer to broadly as **expectations management**, or more specifically **forward guidance**. To see how these tools work, we need a small detour towards something called the **Fischer Equation**.

Fischer Equation

The relationship between the real and nominal interest rate, as it is mediated by inflation (Equation 6).

$$r = i - \pi \tag{6}$$

The Fischer Equation tells us that the real interest rate r (how your savings' purchasing power changes) varies depending on both the nominal interest rate i (how your dollars saved grow) minus the rate of inflation

(π). To understand how this equation works, imagine the bank promises you a 2% nominal interest rate: for every \$100 you have saved they give you \$2. But there is also 2% inflation in the year we're considering that erodes the value of your dollars by 2%. Then your gains are cancelled out by your losses, your "real" returns were zero. In general, if you earn nominal interest that is larger than inflation, your purchasing power grows; if you earn less than the rate of inflation, your purchasing power falls. That's why, if you save in a 0% checking account, you're really losing purchasing power over time (since inflation erodes that value).

The Fed's objective is actually to lower r, which they influence indirectly through i. The ZLB is a statement on nominal interest rates; i, what the bank pays you in dollars, can't go negative otherwise you could store your cash in your cellar. But when the Fed gets i to zero, they've really gotten r below zero, to $-\pi$ precisely. Even so, sometimes this isn't enough to generate enough spending. Now look back at the Fisher Equation, if the Fed can't get i any lower, but it wants to get r lower, which variable is it going to have to push, and in what direction?

The Fed can secondarily push π higher, which further reduces r into negative territory, and inducing more households and firms to spend now rather than lose purchasing power. For example: imagine today the Fed announced that they were going to try and drive π up to 10%; all of the sudden I worry that money I hold in the bank is going to lose a ton of value. Might as well spend it now and turn it into real things while it still holds its current value. This is precisely what the Fed has tried doing in a policy it calls **forward guidance**.

Forward Guidance

An inflation expecations method for pushing down real interest rates. If the Fed convinces households and firms inflation is coming, this is a further motive to spend now, and hence just the changing of expectations can act as a type of monetary stimulus.

This is a brand new tool, and its long run success has yet to be tested. It seems to have worked well when they've tried it so far, but they need to follow through and actually abide by their threats, otherwise households won't believe them next time they make this claim. Because this is part of the monetary system, expectations matter a ton. The jury is out on whether the Fed can reliably manage this dimensions of expectations, but it seems a worthwhile policy to try given the low interest world we may exist in for the coming years.