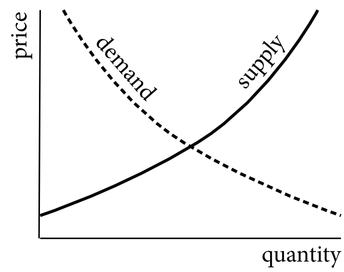

Fig. 15.1. Supply and demand curves (1).

The Ideology of Supply and Demand

Mainstream, neoclassical economics is built on a model of individual economic decisions by consumers and producers. The ideological core of that model is a set of curves that relate demand for some commodity by consumers and supply by producers to price, shown in figure 15.1. These supply and demand curves are used by economists to make economic predictions and give advice, and to argue against government intervention in the economy, among other things.

The model and its two curves are derived by a number of steps from certain assumptions about human economic behavior, preferences, production costs, and so forth. Most of those assumptions are patently false and many are even absurd, and some of the steps in the argument are dubious at best. By far the most rigorous analysis of these problems can be found in Steve Keen's *Debunking Economics*,[14] but given the central role of figure 15.1 in mainstream economics and capitalist ideology, as well as for reasons mentioned above, I believe it is essential to spend a few pages scrutinizing the ideas behind it.

The theory starts with a model of an individual consumer. That individual consumer is perfectly rational, perfectly informed, and perfectly selfish. Furthermore, if that consumer derives satisfaction or "utility," measured in "utils," from the consumption of one unit of some good, then he will get more satisfaction and thus more utils from consuming more, but for every additional good the increase of satisfaction will be smaller. This is called "marginal utility." Thus, consuming one piece of cheesecake will produce, for example, 10 utils, the second piece 8, the third piece 6, and so forth. And by implication, consuming one piece produces 10 utils in total, consuming two pieces 10 + 8 = 18 utils, three pieces 10 + 8 + 6 = 24, and so forth.

This is, of course, absurd. Consuming many pieces of cheesecake will make the consumer feel sick, resulting in dissatisfaction (i.e., negative utils). If nausea starts to kick in after the fourth piece and overwhelms the enjoyment of eating cheesecake by the fifth, then the fifth piece would carry a satisfaction or utility of approximately –25 or even lower because that piece would cancel out all previous enjoyment, and every next piece would just lead to more nausea and thus more negative utils (but not as low as –25 as these following pieces would only increase nausea and not cancel out the previous enjoyment because the fifth piece already did that). Figure 15.2 compares the utility curve according to orthodoxy, the dotted line, with a corrected, more realistic utility curve, the continuous line.

14 Steve Keen, *Debunking Economics*, rev. and exp. edn. (London: Zed Books, 2011).

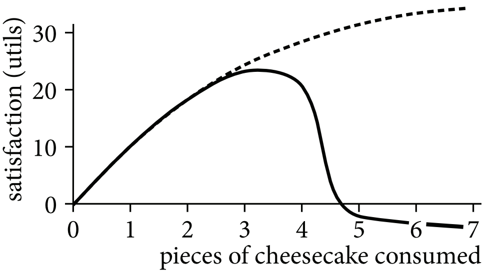

Fig. 15.2. Cheesecake utility.

Utility curves can have considerably more irregular shapes than the cheesecake curve in this figure, however. Let's assume that some consumer's utility of one new, white car of some particular brand "Acme" and model "Coyote" is a certain number of utils.[15] Owning another white Acme Coyote would not increase the consumer's satisfaction by much, however, and might actually decrease it if parking space would be rare or costly, which it is in many cities. From the third, white Acme Coyotes may very well start becoming a nuisance. Having to own three of them might, especially in a city, be so expensive and troublesome that it would be better to own none at all. Consequently, the utility curve for white Acme Coyotes looks somewhat similar to that of pieces of cheesecake, but with two important differences. First, the line for cars peaks and sinks much earlier than that for pieces of cheesecake and its peak is (probably) much higher. Second, and much more importantly, half a piece of cheesecake will produce some satisfaction (or some nausea), but half a car is most likely just an annoyingly big piece of trash. Therefore, while the line for cheesecake is smooth, the line for white Acme Coyotes will be flat or even decreasing until the first whole car where it peaks, after which it sinks because of the increasingly large part of unusable half-a-car, and suddenly peaks again at the second whole car, and then sinks below zero because of the trouble and expenses associated with owning too many (pieces of) cars.[16]

The important point here is that utility curves do not have the shape mainstream economists suppose they have. They will not continue to rise infinitely and they may have flat starts, valleys and peaks, and other odd and irregular features. And therefore, the "law of diminishing marginal utility," which holds that all utility curves have shapes like the dotted line in figure 15.2 is false.[17] Furthermore, the theory cannot be rescued by claiming that this single consumer could sell off the additional pieces of cheesecake or the additional white Acme Coyotes, and that this would cause the curve to continue to climb. This is a model of a *single consumer*. He is a consumer, and not a supplier or trader, and moreover, if he is single (not in the sense of

15 And let's assume that Acme Coyotes do not explode and actually do have a positive utility.

16 Utility curves with a flat start or peaks and valleys can also occur for other kinds of commodities. Let's say that our consumer plays a 47-string concert harp. He's been using strings of type *A*, but now wants to try strings of type *B*. He needs to replace all of the strings to do that, however, so the utility of 1 to 46 strings is pretty much zero. Then, at 47 strings the utility suddenly jumps, after which it continues to climb very gradually because it is nice to have a few spare strings, until the number of strings gets so large that issues of storage start to play a role and utility starts decreasing again.

17 More accurately, the "law of diminishing marginal utility" holds that every next unit of a commodity will have a positive value in utils but a smaller value than the previous unit. But it is still false.

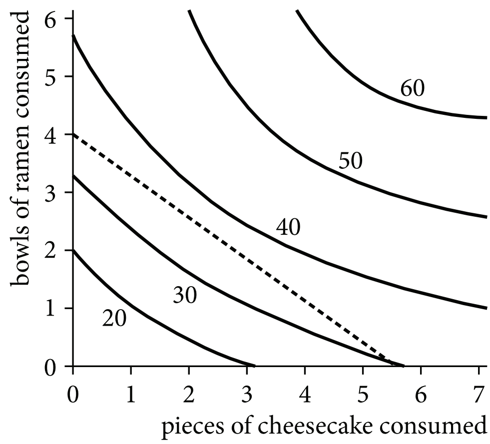

Fig. 15.3. Indifference curves (1).

"unmarried" but in the sense of "alone"), then there is no one else to buy the surplus from him. But even if that is ignored, it does not really change much. There is a limit to how many pieces of cheesecake this consumer or trader could sell. One million pieces of cheesecake would quickly turn into a rotting mess with a very large negative utility, and thus the utility decline would merely be postponed. And half cars are still useless, so the utility curve for Acme Coyotes would still have peaks and valleys. In other words, the "law of diminishing marginal utility" is still false. But let's ignore this, and move on.

Let's say that the same consumer likes bowls of ramen slightly more than pieces of cheesecake. Every combination of pieces of cheesecake and bowls of ramen also gives him a certain satisfaction or utility, shown in figure 15.3. The lines in the figure are called "indifference curves." They connect all combinations of goods with the same total utility. The line marked "30" in the figure marks all combinations that add up to 30 utils: 5.6 pieces of cheesecake and no ramen, 4 pieces and half a bowl, 3 pieces and 1 bowl, and so forth.

If bowls of ramen cost 550 yen each and pieces of cheesecake cost 400 yen, and the consumer has 2200 yen in his pocket, then he can buy either 4 bowls of ramen or 5.5 pieces of cheesecake or any other combination of ramen and cheesecake on the dotted line or below in figure 15.3. The highest total satisfaction the consumer can reach with his money is the point closest to the 40 utils line. And because our consumer is perfectly rational, perfectly informed, and perfectly selfish, he will try to achieve that maximum satisfaction by purchasing approximately 2.4 pieces of cheesecake and 2.25 bowls of ramen.

By varying the price of ramen and adding a few more indifference curves in between 30 and 40 to the diagram, the effect of price on the consumer's consumption can be shown. The thick black lines in figure 15.4 show the amounts of ramen and cheesecake available for different prices of ramen. Because the price of cheesecake is constant, all of these lines go through the same point on the *y*-axis: 0 bowls of ramen and 5.5 pieces of cheesecake.

The thick black line crossing the *y*-axis just below 5 is the line for a ramen price of 450 yen per bowl. At that price, and a fixed price of cheesecake, the highest indifference curve he can reach is 40 utils. That point corresponds to 3.3 bowls of ramen and 1.9 piece of cheesecake, so at a price of 450 per bowl, the consumer would buy

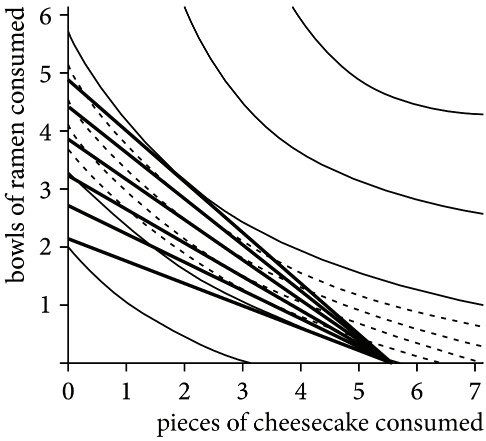

Fig. 15.4. Changing the price of ramen.

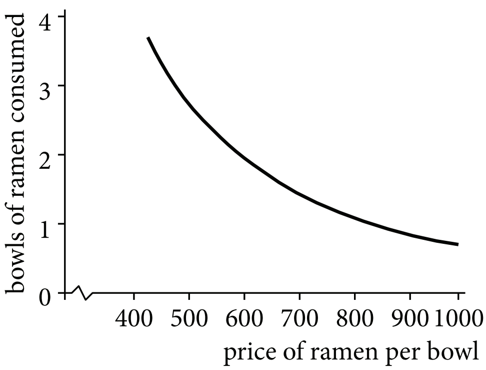

Fig. 15.5. The demand curve of ramen for a single consumer.

3.3 bowls of ramen. Similarly, at a price of 500 yen per bowl, he would buy 2.7 bowls, at a price of 580 2.1, and so forth. With this data a new diagram can be drawn that shows the relation between price and consumption. This is the "demand curve" for an individual, single consumer in figure 15.5.

The demand curve in figure 15.5 looks much like the one in figure 15.1 (with the axes switched!), but remember that it depends on fictional utility curves and several other assumptions. That utility curves do not have the shape they are supposed to have was already shown above, but many of these other assumptions are almost equally nonsensical. For example, the consumer must always be able to make a choice between different combinations of goods or commodities. And his preferences must be transitive — if he prefers *A* over *B* and *B* over *C*, then he must prefer *A* over *C*. Both these assumptions may be acceptable for an idealized, perfectly rational consumer, but psychologists and experimental economists have shown that they do not always apply to real people. And the assumption that a consumer can always compare all possible combinations of goods in terms of their utility becomes especially ludicrous if it is taken into consideration that real consumers do not compare two commodities but possibly thousands.

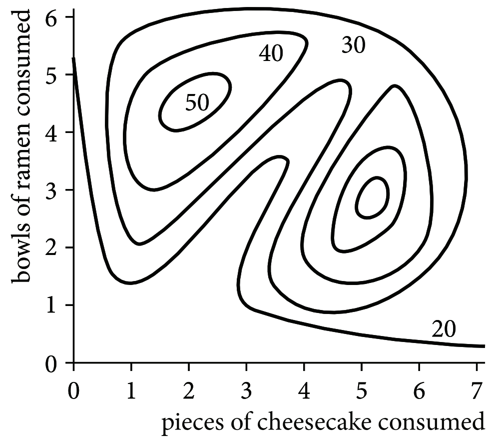

Fig. 15.6. Indifference curves (2).

It is further assumed that the satisfaction derived from consuming one good is completely independent from the satisfaction derived from consuming any other good and that there are no interaction effects between consumption of various goods. However, it might be the case that the consumer is fine if he eats lots of ramen and a little cheesecake or the other way around, but gets sick if he eats roughly equal amounts of both, for example. If this kind of interdependencies are not ignored, and if it is also taken into account that eating too much ramen or cheesecake also makes our consumer nauseated and thus leads to dissatisfaction, the picture starts to change rather drastically. Then, even with all of the other aforementioned assumptions, a rational consumer's indifference curves would look something like figure 15.6.

Obviously, this seriously messes up the demand curve. Now, at some price levels, there is not a single point at which this consumer can maximize his utility, but two, which means that the supposed "curve" is not a curve at all. Furthermore, if some of the other aforementioned assumptions are relaxed as well, indifference curves cannot be drawn anymore at all, not even weird or irregular ones, and thus, no demand curve could be derived either. But, again, let us ignore all that and move on.

Thus far, it was assumed that the consumer's budget is fixed at 2200 yen, but obviously, if his income rises or falls, so does the budget he has available for ramen and cheesecake. With a smaller budget he would end up with a different combination of goods to maximize his utility. Keeping prizes constant, a number of different lines can be drawn in figure 15.3 parallel to the dotted line to represent different budgets. These lines each touch different indifference curves. Connecting the points at which such budget lines reach maximum utility results in the thick black line in figure 15.7. This line is called an "Engel curve." It has a slightly peculiar shape in the figure, which is largely due to the fact that the indifference curves in my figures are not as nice and smooth as the indifference curves of a perfectly rational, perfectly informed, and perfectly selfish consumer that can (and prefers to) consume infinite amounts of anything are supposed to be. On the other hand, as Steve Keen pointed out, "Engel curves can take almost any shape at all,"[18] so there is no reason why this oddly shaped

18 Keen, *Debunking Economics*, 50.

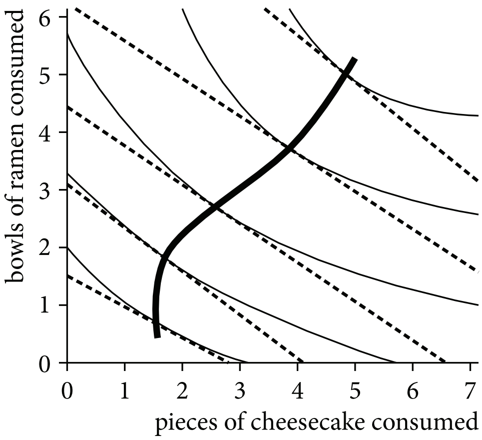

Fig. 15.7. Indifference curves and the Engel curve.

curve would be impossible. And besides, the Engel curve in this figure is rather pedestrian compared to what you'd get if you'd try to fit an Engel curve into figure 15.6.

An Engel curve shows how a consumer's spending pattern changes with a change in income. For reasons to be explained below, mainstream economists assume that Engel curves are straight lines unlike the line in the figure, which implies that spending patterns do not change when income changes. If the consumer gets richer, he just gets more of the same things and in the exact same proportion (and he gets less, but still in the same proportion, if he gets poorer). This is obviously absurd — spending patterns change very much with a change in income or budget, but this too we will have to ignore to move on.

No economy consist of a single consumer, so a demand curve for a single consumer is rather useless. We need a demand curve for an economy as a whole, or in other words, for all the consumers in that economy together. To get that demand curve, you would have to find the individual demand curves of all consumers in the economy, and add up the total number of bowls of ramen consumed at each price level. If there are a thousand consumers in the economy and they have similar preferences, then the market demand curve could, supposedly, look something like figure 15.8.

There is a rather nasty complication, however. The derivation of the individual demand curve depends on a fixed price of cheesecake and a fixed budget. Regardless of the price of ramen, the single consumer still has only 2,200 yen to spend. That is fine in the single-consumer model, but that becomes very implausible in the case of whole economies. If the price of ramen goes up or down or if the consumption of ramen rises or falls, that will influence the income, and thus the budget, of some people in the economy and those people are consumers too. In other words, changes in price and changes in consumption change the incomes and budgets of some consumers, and therefore, in an economy with multiple consumers, budgets *cannot* be fixed.

Furthermore, there is another kind of income effect that comes into play when three or more commodities are taken into account. For example, if the consumption of one of those three commodities cannot easily change but its price can, then a change in the price of that commodity will effectively raise or lower the consumer's income, and thereby his budget available for the other two commodities. Think of a raise in rent, for example. Moving house may not be a short-term option, so such a change would effectively just decrease a consumer's remaining income and budget.

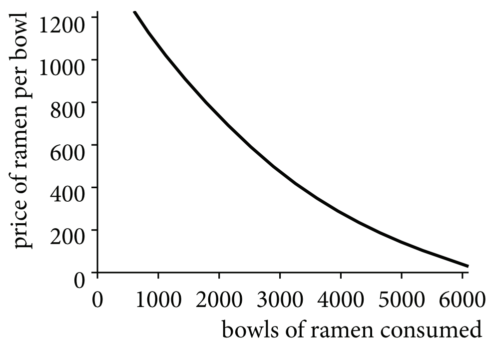

Fig. 15.8. The demand curve of ramen.

Or if a commodity takes up much of a consumer's budget but is really considered inferior by that consumer, then a decrease in the price of that commodity may make it possible for the consumer to buy better alternatives and thus lead to a decrease in consumption of the inferior good. For example, if due to poverty all you can afford is one kind of bread, if that bread becomes cheaper, you have to spend less on it, but because of that you can buy less of it and spend what you save on better food. Hence, while a demand curve is supposed to slope downwards (i.e., rising prices mean falling consumption and falling prices mean rising consumption) it can be the other way around.

When such effects and the role of income are properly taken into account it turns out that a demand curve "can take any shape at all — except one that doubles back on itself."[19] In a paper published in 1953, W.M. Gorman proved that the only way to get demand curves with the shape mainstream economists believe they have — downward sloping lines similar to the curve in figure 15.8 — is by assuming that the Engel curves of all consumers are parallel, straight lines.[20] This is an interesting assumption. As mentioned above, assuming that Engel curves are straight lines is assuming that spending patterns do not change with income, which really only could be the case if there is just one commodity available. Assuming that they are parallel for all consumers is assuming that all consumers have identical preferences, which really only could be the case if there is just one consumer. These are obviously absurd assumptions. If the only way a continuously downward sloping demand curve or line can be derived is by assuming that an economy consists of a single consumer consuming as single commodity, then this demand curve has little if anything to do with a real economy. Even ignoring most of the problems mentioned in previous sections, in a real economy, a demand curve can have almost any shape. But let us ignore all that as well and move on. IIII

Assuming a downward sloping demand curve, if there is only one producer, then there are only two things that matter: the market demand curve for whatever it is producing and the costs of production per unit. (But keep in mind that every step in the derivation of that demand curve turned out to be invalid and that every important assumption it is based on is false.)

19 Ibid., 52.
20 W.M. Gorman, "Community Preference Fields," *Econometrica* 21, no. 1 (1953): 63–80.

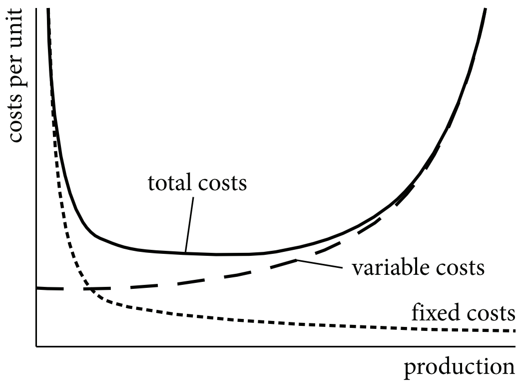

Fig. 15.9. Production costs.

Economists divide production costs into two different kinds: fixed costs and variable costs. Fixed costs do not depend on how many units are produced. They include the costs of buildings, machinery, tools, and so forth. Variable costs depend on the level of production and include labor and resources. It is assumed that variable costs rise with the level of production: the more units one attempts to produce at the same fixed costs, the more labor or other variable cost factors are needed. Therefore, the curve of variable costs per unit of production rises. The curve of fixed costs, on the other hand, starts very high, then drops steeply, and almost flattens out. The first units (bowls of ramen, pieces of cheesecake, cars, or whatever) are expensive to make due to fixed costs (buildings, machinery, and so forth), but the more units are made, the less these fixed costs matter. In figure 15.9, the dotted line represents fixed costs, the dashed line variable costs, and the continuous line the total average costs of production per unit.

It was already shown in the 1950s, however, that this figure does not look anything like a real average cost curve.[21] 95 percent of managers reported that there is no significant rise in the variable costs. In fact, for the vast majority of firms, the variable cost curve appears to be nearly or completely flat. This will turn out to have important implications.

Profits are total income minus total costs. Total costs are obtained by multiplying the number of units produced with the average production costs (represented by the continuous line in figure 15.9). Total income is the number of units produced multiplied by the price at which they can be sold. In case of the demand curve of ramen derived above, if the producer sets the price at 900, it can sell 1,000 units. The more units are produced, the lower the price has to be to sell all of them. Figure 15.10 shows total costs (thick dotted line) and total income (very thick gray line). The difference between those two is total profit (continuous black line). The two thinner lines show what happens if the nonsensical assumption of rising average variable costs is discarded. Then total production costs (thin dashed line) are more or less linear and the profit peak (thin continuous line) occurs later (i.e., at a higher level of production).

Supposedly, the situation changes drastically if there are many producers of the exact same good. Economist assume that in typical markets there will be very many producers, and that, because of that, none of them can influence the total supply or

21 W.J. Eiteman and G.E. Guthrie, "The Shape of the Average Cost Curve," *American Economic Review* 42, no. 5 (1952): 832–38.

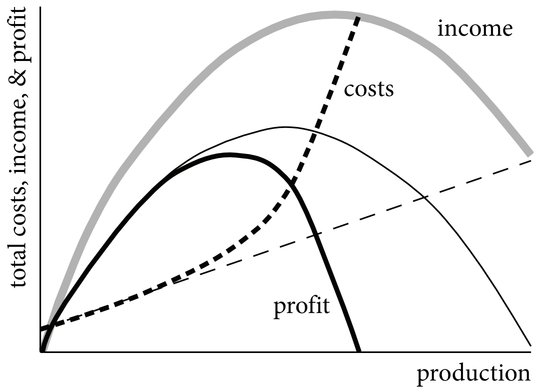

Fig. 15.10. Profit.

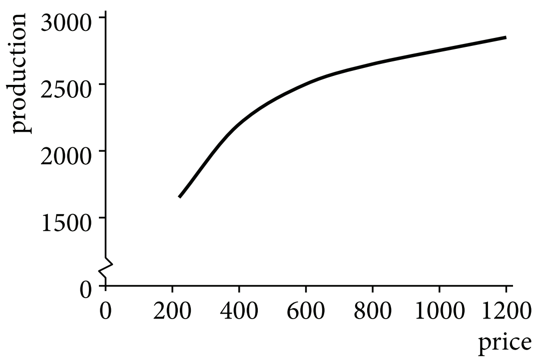

Fig. 15.11. Supply curve for a single producer.

price. All of them, therefore, have to accept the market price, because if one of them would sell above the market price, it would lose all its customers to its competitors and consequently, go bankrupt. And if it would sell below the market price it would decrease its own profits. (There's something very fishy about this, but we will get to that soon.) If the production cost curves for ramen would be like figure 15.9 with the *y*-axis (i.e., costs per bowl) ranging from 0 to 1,000 and the *x*-axis (production) from 0 to 4,000, then the maximum profits a producer can reach can be calculated for every price level. Figure 15.11 shows these maximum profits. This is the supply curve for an individual producer. If the market price is 600, the producer can reach the maximum profit by producing and selling 2,458 units. The supply curve starts at a price of 230. Below that price, the producer cannot make a profit at any production level.

To derive this supply curve, it is essential to assume that the producer can sell everything it makes because that is what calculated income depends on. This assumption is defended in the same way as the assumption that individual producers cannot influence price or total supply: there are so many producers in the market that the production of each of them is just a drop in the ocean of total supply, and therefore, every firm can sell everything without influencing prices or total supply. This assumption is absurd from a mathematical point of view, but we'll get to that shortly.

The total supply is simply the sum total of individual supply curves of all producers in the market, which obviously implies that total supply at each price level depends on the number of producers. In fact, given that it is assumed that all pro-

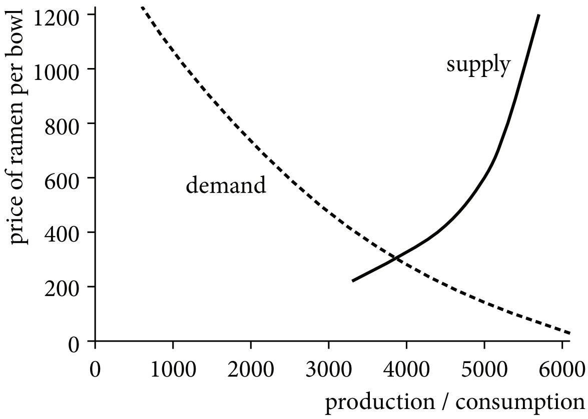

Fig. 15.12. Supply and demand curves (2).

ducers are identical, the total supply at each price level is just the individual supply multiplied by the number of producers. Hence, the total supply curve for a market with 10 producers would look exactly like figure 15.11, but with one zero added to each number on the *y*-axis.

If the total supply curve is added to figure 15.8, which showed the demand curve, then the result is figure 15.12 (but note that the axes in figure 15.12, relative to 15.11, are switched). The point where the two lines cross is where the market is in equilibrium. At that price level all units produced will indeed be sold. According to the figure, slightly less than 4,000 units will be produced and sold at a price of 327 each.

What the figure does not show is the number of producers in this market. That number is 2. If there would be 3 producers, then the two lines would not cross, because the three of them would together produce almost 5,000 units, which would reduce the market price to well below the minimum they need to cover their costs.

Recall that there were supposed to be many producers in this market. Two is not very many. This is an obvious problem for the theory indeed, but there is an easy way to address this problem, just lower fixed costs. If fixed costs are extremely low, then even small production volumes relative to the market are profitable and there is room for very many producers. Unfortunately, for the vast majority of industries, the fixed costs are much larger (rather than smaller) than in the example thus far, so we would have to, again, sacrifice realism to save the theory.

The assumption, then, that there are very many producers is not a plausible assumption in most industries. (And keep in mind that they must be producers of the *exact* same commodity.) But the derivation of the total supply curve depended on some other assumptions as well: because of the large number of producers, none of them can influence price or total supply, and rising average variable costs of production apply. The importance of the latter assumption cannot be overstated. If it is assumed that a producer can sell whatever it produces, then there is a maximum profit only if average production costs rise and at some point overtake income. If average production costs flatten out — as they do in reality; see above — then profits keep rising with production. And if there are no maximum profit levels at various price levels, then there is no supply curve. Then, regardless of the price level and still under the assumption that all production is sold, the profit-maximizing producer sets its production level at *infinite*. This is obviously absurd.

Furthermore, even if there are very many producers, it is mathematically impossible that a change in production level of one of them without compensation by the others would not affect total supply and price. Even if there are 100 producers, and one of them increases its production only slightly, that changes total supply by a tiny little bit, and thereby price by a tiny little bit, and that tiny little bit affects everyone in the market. Economists argue that these "tiny little bits" are in fact so tiny that they can be ignored, but that is a rather dubious claim (and mathematical nonsense), especially if it is taken into consideration that in most markets there will not be very many producers.

Given the model we have now it is rather easy to simulate the effect a single producer could have on the market. Let's assume that in this market somehow a situation has evolved in which there are 10 producers, each making 250 units (2,500 total) and selling them for 585 per unit, which is the price level at which all 2,500 units will be sold according to the demand curve. All 10 of them make a rather modest profit of 8,850 (or 8,710 if rising average variable costs are assumed). Now, if one of them would suddenly decide to quadruple its production, then total supply would rise from 2,500 to 3,250, which would reduce the price to 458. The firm that raised its production now makes a profit of 208,000 (or 203,000 if rising average variable costs are assumed), but the other 9 firms (that did not change their production level but that are affected by the same change in price) now *lose* approximately 23,000 each. (And even a much smaller production increase would increase profits of that firm, while reducing profits of all others.) Hence, a perfectly rational and omniscient, profit-maximizing firm would raise its production to put all of its competitors out of business and then establish a monopoly in which it could have maximum profits. But because, supposedly, all firms in the market are perfectly rational, omniscient, and profit-maximizing, they will all try to do that, leading to massive overproduction, a collapse of the price, and all of them going bankrupt. Except, of course, if they all predict that consequence as well.

That's not what happens, however, and the main, but not only, reason for that is that production costs work a little bit differently; or actually, not a *little* bit. Variable costs are more or less flat, but fixed costs are not as "fixed" as they are supposed to be either. So-called fixed costs are the costs of buildings, machinery, and so forth needed for a certain maximum production level, that is, fixed costs are related to a certain production capacity. If a producer runs at 100 percent capacity, it produces the maximum number of units that can be made with the buildings, machinery, and so forth available. There is no point in adding more labor if you are at 100 percent capacity because those extra workers would just stand by and watch.[22] And there is no point in trying to put more resources or inputs in the machines if those are already at 100 percent capacity either — you cannot make a machine run faster than its maximum speed. Typically, fixed costs are very high and are not earned back until a producer is well over 50 percent capacity, and it may be much closer to maximum capacity, and because of that, the vast majority of firms run at capacities between 80 percent and 100 percent. If you are already close to 100 percent capacity, then

22 Recall that it is assumed that variable costs rise with the level of production because if one tries to produce more units at the same fixed costs, more labor or other variable cost factors are needed. This effectively means adding more labor (etc.) to a machine or process that already runs at (close to) full capacity, which is useless. This just is not how production works and it is a mystery to me how mainstream economics can seriously believe that their conception of production costs makes sense.

a production expansion is possible only by spending "fixed costs" to open another production line, which also would have to run at close-to-full capacity immediately to be profitable. That means that increasing production requires a large investment, which can be earned back only if the full jump in production can be profitably sold. Profit levels are generally insufficient to make such investments, and banks and investors are only willing to lend the money if they are sufficiently confident that the production expansion will pay off (or even not at all; see the section "Playing with FIRE" below).

There are various other complexities and complications, but those matter little right now. The important thing is that every assumption that underlies the derivation of a supply curve has been shown to be false. Production curves are not even remotely similar to what economists imagine them to be. There is not enough room for many producers in most markets. And even if there would be many producers, the actions by any one of them would affect total supply and price. The implication is that there is no way to derive supply from price, and therefore, that there is no such thing as a supply curve. But let us ignore all of that and move on.

So here we are. We have consistently ignored reality by making absurd assumptions, by confusing small amounts with zero (in the derivation of the supply curve), and by making one invalid step after the other (in the step-wise derivation of the demand curve, especially),[23] but it has paid off: we now have our two curves. So, what can we do with them? *Nothing*, seems to be the obvious answer, as these curves have nothing to do with reality, but economists see things a bit differently.

Faced with criticism like the foregoing, many economists would protest and claim that their "science" is really much more sophisticated than the sketch I'm providing here, which to some extent is true. However, as David Orrell pointed out, "what counts is less what economists say — they are skilled at deflecting criticism, and have plenty of practice — than what kinds of calculations they actually perform."[24] Economists might accept that some of their assumptions are false and point out that there are sophisticated theories and models taking this into account, but in practice, all of their calculations and all of their predictions and recommendations remain based on the same problematic assumptions. Hence, even if the problem is recognized, that recognition is mere empty rhetoric and does not lead to any significant change. More often the problem is brushed aside. It does not matter that all of the assumptions that the theory is based on are false because all scientific theories are based on false assumptions, the argument goes, and all that matters is whether the resulting theory is a useful tool. This is really mainstream economics' last line of defense, and it fails as miserably as everything that came before.

A physicist might assume that friction does not matter when predicting the effects of gravity on a cannon ball. In that case, she is making two kinds of assumptions at once. First, she restricts her theory to cannon balls and similar objects. And second, she assumes that for those objects the effects of friction are negligible. Economists make similar assumptions, but what they do not seem to realize is that those assumptions imply that its theories only apply to perfectly rational, perfectly informed, and perfectly selfish, profit-utility-maximizing beings (in the same way

23 Presumably, you have noticed the tally marks in the margins. They count variants of the phrase "but let us ignore that and move on." The tally mark on this page shows that reality had to be ignored in order to accept a bunch of absurd assumptions five times to get to the point where we are now.

24 David Orrell, *Economyths: Ten Ways Economics Gets It Wrong* (Ontario: John Wiley, 2010), 6.

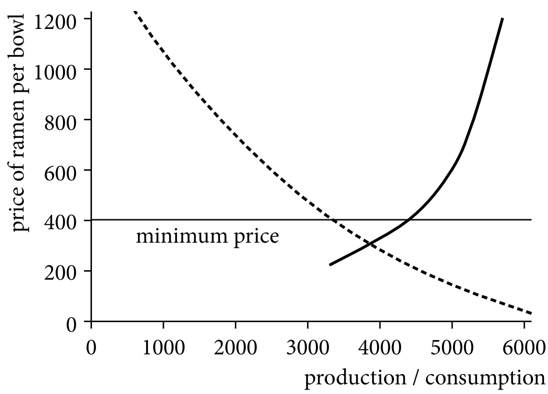

Fig. 15.13. Effects of a minimum price.

that the physicist's theory is restricted to cannon balls), and that for such beings certain effects can be ignored (like friction for cannon balls). But perfectly rational, perfectly informed, and perfectly selfish, profit-utility-maximizing beings do not exist (while cannon balls do), and economists thus restrict the domain of their theories to *nothing* — their theories have no application.[25]

Whether mainstream economic theories are useful is debatable as well. They are not useful to understand how an economy works because they assume things that do not and cannot exist in a real economy, such as a supply curve. And neither are they useful as tools for prediction, as the predictions of mainstream economists are consistently wrong. For example, the consensus among mainstream economists just a few months before the Great Recession of 2007 to 2009 erupted was that there would be continuous economic growth and no more crises. Economic theories may be "useful" in another way, but let us return to our question first: what can we do with these theories? What can we do with our fictitious supply and demand curves?

Rather neat stuff, supposedly. We can "prove," for example, that the market is always right and that governments can only get it wrong. Here is how that is supposed to work. Suppose that the government decides that ramen shops need to be supported and towards that end implements a minimum price for ramen of 400 yen per bowl. If the supply and demand curves in the previous figure apply to the ramen market, then we can add a third line to that figure showing this minimum price, as in figure 15.13. The total supply is where that line crosses the supply curve: 4370 bowls of ramen. The total demand is where the minimum price line crosses the demand curve: 3600 bowls of ramen. Hence, we now have an overproduction of 770 bowls (if those extra bowls are actually produced). Consumers have to pay more for their ramen than before and are thus unhappy. The two producers see their profits rise from 183 thousand to 306 thousand each and are very pleased. But because there are very many more consumers than producers, overall welfare decreases. The same applies to virtually any other kind of government interference in the market: whatever the government does, it will decrease overall welfare. By treating labor as a commodity,

25 There are other major disanalogies between abstraction in economics and abstraction in physics, as I pointed out in *The Hegemony of Psychopathy* (Earth: punctum books, 2017), 46ff. See also the section "Ideal Theory, Utopia, and Ideology" in chapter 16.

specialization and trade		no		yes	
		A	B	A	B
labor	corn	10	10	6	14
	cars	10	10	14	6
production	corn	40	45	24	63
	cars	40	35	56	21
trade	corn			17	−17
	cars			−15	15
consumption	corn	40	45	41	46
	cars	40	35	41	36
	total	80	80	82	82

Table 15.1. The supposed benefits of free trade.

for example, it can be shown in the same way that a minimum wage increases unemployment. And so forth.

All of this is nonsense, of course. It depends on fictitious supply and demand curves that bear little if any resemblance to reality. But this also reveals why these fictitious curves are so pervasive anyway: they serve an ideological agenda. They serve a pro-market and anti-government agenda that favors deregulation and small governments. They serve an agenda that benefits the financial sector (FIRE; see below) and large corporations and that harms almost everyone else. It is in this sense — and only in this sense — that mainstream economic theories are "useful": they are useful to serve the interests of the global financial and industrial elite, the Transnational Capitalist Class.[26]

Free Trade Ideology

Free market ideology often comes together with free trade ideology. According to conventional "wisdom" free trade leads to prosperity. There are various more or less sophisticated theories making this claim, but they all go back to David Ricardo's theory of "comparative advantage."[27] There is a fundamental problem with that theory, as was shown by Frank Graham in 1923, but unfortunately, that problem tends to be ignored.[28]

Assume a world with two countries, A and B, and two commodities, corn and cars. Both country A and country B produce corn and cars, and both have a labor force of 20 "units." These units might be 10,000 workers or 1 million work hours — this does not matter. In both countries, half the labor force produces corn and the other half

26 Leslie Sklair, *The Transnational Capitalist Class* (New York: Wiley, 2000), and William Robinson, *Global Capitalism and the Crisis of Humanity*, new edn. (Cambridge: Cambridge University Press, 2014). See also the section "Ideology" in chapter 4.

27 David Ricardo, *On the Principles of Political Economy and Taxation* (London: Murray, 1817).

28 Frank Graham, "Some Aspects of Protection Further Considered," *The Quarterly Journal of Economics* 37, no. 2 (1923): 199–227.