Intermediate Microeconomics — Week 15

Professor Boyd **November 29 & December 1, 2022**

12.4 Repeated Games

The Final is in DM-110 at 12 noon on Thursday, Dec. 8.

A **repeated game** is one that we play more than once, obtaining a payoff every time we play it.

12.4.1 Repeated Prisoner's Dilemma

We start with the repeated Prisoner's Dilemma. The Prisoner's Dilemma refers to any game with the following payoff pattern, where

$$t_i > g_i > b_i > s_i$$

for $i = 1, 2$.

Generic Prisoner's Dilemma

Player Two \ Player One	Cooperate	Defect
Cooperate	g_1, g_2	t_1, s_2
Defect	s_1, t_2	b_1, b_2

12.4.2 Payoffs in the Prisoner's Dilemma

The diagonal payoffs are called the good payoff (g_i) and the bad payoff (b_i). Of course, the good payoff is better than the bad payoff, $g_i > b_i$.

Off the diagonal we have two other payoffs, the temptation (t_i) and the sucker's payoff (s_i). You get the temptation if you defect and the other player cooperates. It's even better than the good payoff ($t_i > g_i$). However, if you let the other player take advantage this way, you get the sucker's payoff, which is the worst of all ($s_i < b_i$)

Given this, the best responses are marked below. For each $i = 1, 2$, $t_i > g_i > b_i > s_i$.

Generic Prisoner's Dilemma: Best Responses

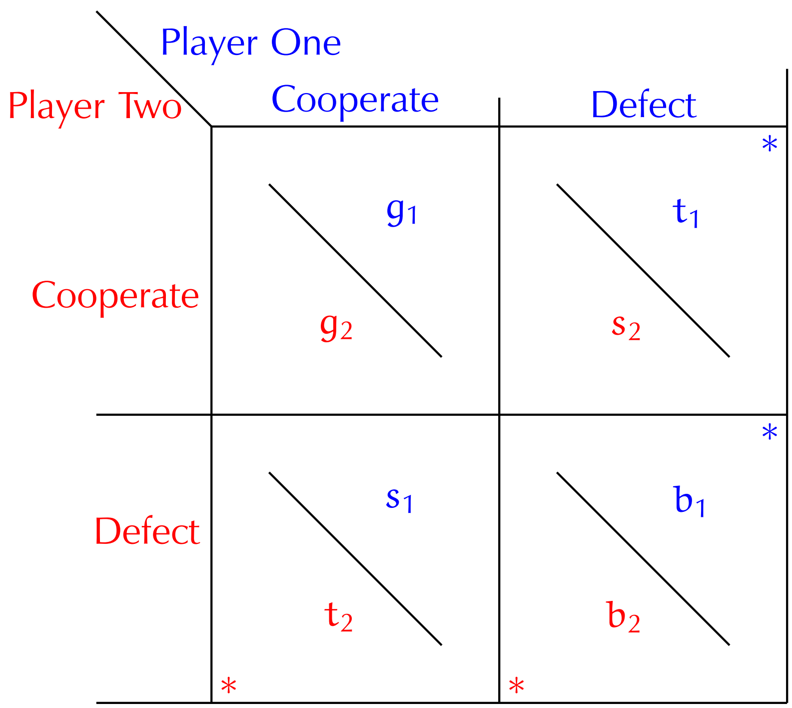

12.4.3 Dominant Strategy Equilibrium

Both player's have a dominant strategy, **Defect**. As a result, the only equilibrium is for both players to defect and receive the bad payoff, even though cooperation is better for both of them.

The problem with cooperation is that the temptation lures them both away from it.

Generic Prisoner's Dilemma: Dominant Strategy Equilibrium

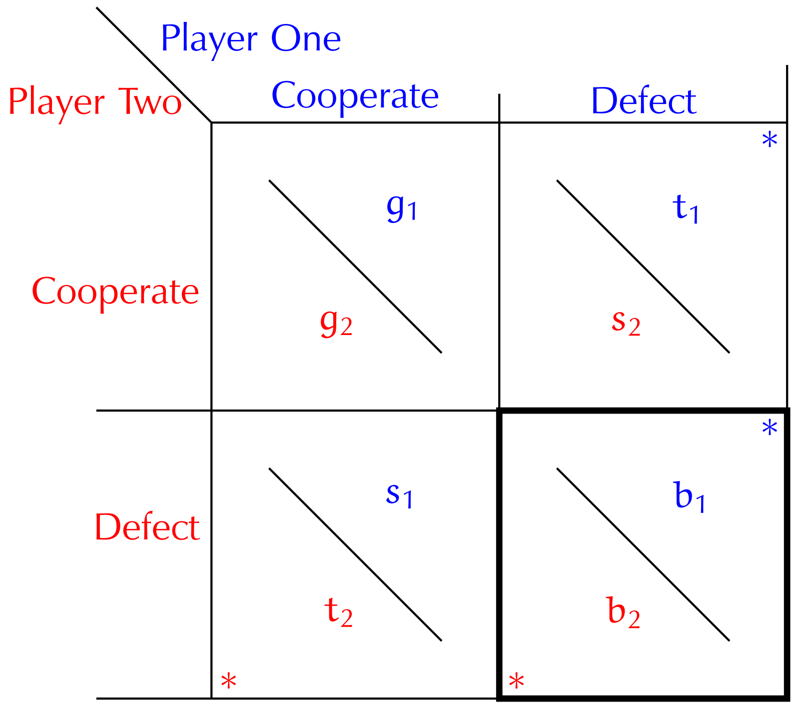

12.4.4 Cartels as Prisoner's Dilemmas

Suppose we have two firms in Cournot competition. They face an inverse demand curve of $p = 106 - Q/2$, and both have constant marginal costs of $MC = c = 10$.

The firms consider forming a cartel. For the cartel, marginal revenue is $MR = 106 - Q$. The resulting cartel or monopoly output is $q_M = 106 - 10 = 96$ and the monopoly price is $p_M = (106 + 10)/2 = 58$.

Producer's surplus for the monopolist would be $(p_M - c)q_M = (58 - 10)96 = 4608$. With $FC = 0$, this is equal to monopoly profit. Each firm produces half, 48, and earns a profit of 2304.

If they are in Cournot competition, they each produce 2/3 the monopoly output, 64, for a total of 128. The resulting price is 42 and the profit of each firm is $(42 - 10)64 = 2048$.

Finally, if one firm produces the cartel amount of 48 and the other produces the Cournot quantity of 64, total output is 112. The resulting price is 50. The producer's surplus for the firm producing the cartel amount is $(50 - 10)48 = 1920$, the sucker's payoff. The producer's surplus for the firm producing the Cournot amount is $(50-10)64 = 2560$, the temptation.

12.4.5 Cartels as Prisoner's Dilemmas

Here's the game matrix. **Cooperate** means to produce the cartel amount. **Defect** means to act as a Cournot competitor and produce the Cournot equilibrium amount. To avoid complicating the model, we have only considered these two choices.

Cournot Cartel

Player Two \ Player One	Cooperate	Defect
Cooperate	2304 \ 2304	1920 \ 2560 *
Defect	* 2560 \ 1920	* 2048 \ 2048 *

12.4.6 Can we Punish Defectors? Reward Cooperators?

In the standard prisoner's dilemma, defection is a dominant strategy. It's in both sides' interests to defect. They end up in the bad outcome, even though both would be better off in the good outcome. However, as long as the rules of the game are followed, there is no escape

The players might try to change the game to make a cartel possible, but anything of that sort that they would try is illegal in the United States. The Justice Department actively looks for such shenanigans.

There are some cases where the government will enforce a cartel, or something simliar. An example would be zoning, where it would be illegal for homeowners to make an agreement to limit their properties to single-family homes, but it is perfectly legal for the local government to do so via zoning.

Repeated prisoner's dilemmas suggest yet another possibility. Can the players reward cooperators or retaliate against current defectors in the future? Can one hold out the carrot of future cooperation (good payoff) and stick of further defection (bad payoff) to ensure cooperation now?

12.4.7 Repeating the Prisoner's Dilemma

Suppose we repeat the above prisoner's dilemma once. What happens in equilibrium?

We use backwards induction. In round #2, we have already received our payoffs from round #1. They don't change. Only the round 2 payoffs matter. We know that defecting is the dominant strategy, and must use it. The payoffs are $(2048, 2048)$ in round #2.

Now, working backwards, knowing both players defect in round #2, we consider round #1. Only the first round payoffs are at issue, and defecting is again dominant. Both players defect. The only Nash equilibrium is $\big((D,D),(D,D)\big)$. The fact that no punishment is possible after the last round means that the players use their dominant strategies then. But that means there's no punishment after the first round either. The whole project unravels from the end.

The story doesn't materially change if the are 3 rounds, 4 rounds, ..., or any number of rounds. We use backwards induction to find that both players defect in the last round, then the next-to-last, etc., all the way to the first round. Any attempt to set up an alternative outcome unravels from the future to the present.

Nonetheless, it is a mistake to think the repeated prisoner's dilemma is as simple as that.

12.4.8 What if the End is Unknown?

What if we don't know when the last round is? Then we can't use backwards induction from the last round. Suppose there is a probability p that the game continues, the **continuation probability**. Then at the end of round 1 there is a chance $(1 - p)$ that the game is over, and a chance p that we play again.

The inability to unravel the game means we can consider some new strategies based on game history. We can attempt to reward players when they cooperate and punish them when they don't.

The diagram below indicates how the continuation probability works. A prisoner's dilemma is played at each node on the main diagonal. With probability $(1 - p)$, the game end. With probability p it continues to the next node where the prisoner's dilemma is repeated. This continues in same fashion indefinitely. There is probability p the game is played twice, probability p^2 it's played a third time, etc.

Repeated Game Continuation

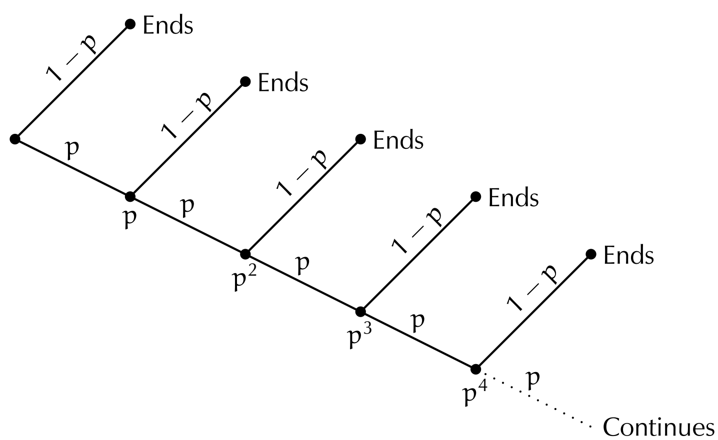

12.4.9 Trigger Strategies

We will consider a new strategy, the **grim trigger strategy**. Players using the grim trigger cooperate in the first round. As long as their opponent has always cooperated, they continue to cooperate. However, if the opponent ever defects, they defect forever after in response.

The term **trigger** refers to the fact that defection on our part is **triggered** by defection on the part of your opponent. The strategy is **grim** because we show no mercy once the opponent has defected.

12.4.10 The Grim Trigger in Action

Suppose player one uses the grim trigger strategy. Player two now considers whether to always cooperate, or to defect in the first round. We will hold off considering later defections for a bit.

Suppose the continuation probability p with $0 < p < 1$. If player two always cooperates, player one, playing grim trigger, will also always cooperate. Then player two gets a payoff of \$2304 in every round that actually occurs. The probability that round t occurs is p^{t-1}. Player two's expected payoff is

$$
\begin{aligned}
&2304 + 2304p + 2304p^2 + 2304p^3 + \cdots \\
&= 2304\bigl(1 + p + p^2 + p^3 + \cdots\bigr)
\end{aligned}
$$

12.4.11 Summation Trick

There is an easy trick for computing such sums. Define S by

$$S = 1 + p + p^2 + p^3 + p^4 + \cdots$$

Then

$$\begin{aligned} S &= 1 + p + p^2 + p^3 + p^4 + \cdots \\ pS &= \quad\;\; p + p^2 + p^3 + p^4 + p^5 + \cdots \\ S - pS &= 1 \\ (1-p)S &= 1. \end{aligned}$$

This means $S = (1-p)^{-1}$. The payoff to always cooperating is

$$2304S = \frac{2304}{1-p}$$

12.4.12 Defection Against Grim Trigger

What if player two defects in the first round. Then the payoff is 2560 in the first round. Player one defects in subsequent rounds (grim trigger), and the best player two can do is also defect, obtaining a payoff of 2048 in subsequent rounds. The expected value is

$$\begin{aligned}2560+&2048p + 2048p^2 + 2048p^3 + 2048p^4 + \cdots \\ &= 2560 + 2048p\big(1 + p + p^2 + p^3 + \cdots\big) \\ &= 2560 + 2048\frac{p}{1-p}\end{aligned}$$

12.4.13 Payoff Comparison

We now ask whether there is a p between 0 and 1 where those two payoffs are equal. So we solve

$$\begin{aligned}
\frac{2304}{1-p} &= 2560 + 2048\frac{p}{1-p} \\
2304 &= 2560(1-p) + 2048p \\
2304 &= 2560 + (2048 - 2560)p \\
2304 &= 2560 - 512p \\
512p &= 2560 - 2304 \\
p &= \frac{256}{512} = \frac{1}{2}.
\end{aligned}$$

If $p > 0.5$, it is better to always cooperate against the grim trigger rather than defect in the first round. If $p < 0.5$, it is better to defect immediately rather than always cooperate.

12.4.14 What about Late Defection?

What if player two defects in a later round against the grim trigger strategy. If the defection is in round four, the payoff streams for player two look like this.

$$\begin{aligned}\text{Defect :} & 2304, 2304, 2304, 2560, 2048, 2048, 2048, \cdots \\ \text{Cooperate :} & 2304, 2304, 2304, 2304, 2304, 2304, 2304, \cdots\end{aligned}$$

The first three payoffs are the same. We need only compare the remainder

$$\begin{aligned}\text{Defect :} & 2560, 2048, 2048, 2048, \cdots \\ \text{Cooperate :} & 2304, 2304, 2304, 2304, \cdots\end{aligned}$$

But the remainder is the same as with defection in round one. In terms of the differences in expected payoffs, the only change is that both expected payoffs have been multiplied by p^3. That doesn't affect which is bigger.

Defection in round four is better than always cooperate if and only if $p < 1/2$, in which case defection in round one is even better. Its expected value is higher by a factor of $1/p^3 > 1$.

12.4.15 An Equilibrium in the Repeated Cartel Game

We have established that if $p > 1/2$, always cooperating is a best response to the grim trigger. But then, it is also a best response to use the grim trigger itself since that leads to permanent cooperation.

That means that for $p > 1/2$, one equilibrium in the repeated cartel game is for both sides to use the grim trigger strategy.

It is possible to get cooperation in the repeated cartel game.

12.4.16 Additional Equilibria in the Repeated Prisoner's Dilemmas

The grim trigger strategy is not the only way to encourage cooperation in repeated prisoner's dilemmas.

Another useful strategy is called **Tit-for-Tat**. In the first round, you cooperate. In subsequent rounds, you copy the other player's last move.

Unlike Grim Trigger, Tit-for-Tat is a forgiving strategy. If the other player learns his lesson, and attempts to cooperation, you're willing to return to cooperation. Generally speaking, this requires a higher continuation probability in order to work.[1]

[1] The political scientist Robert Axelrod speculated on how these ideas might apply to cooperation in general in his book, Robert Axelrod (1984), "The Evolution of Cooperation", Basic Books, New York.

12.4.17 Historical Cartels

The study of repeated games tells us that a cartel can be a Nash equilibrium, contrary to our earlier analysis. It also tells us that such cartels depend on the cartel members continuing to be in competition.

Anything that disrupts this can cause the cartel to fail. In the 19th century, before antitrust legislation, railroads would often attempt to form cartels (then called "pools"). Some were successful for a while, other quickly dissolved.

If you read their history with repeated games in mind, it's clear that pools had the most problems when economic conditions were changing, or when their membership was changing. These both correspond to changes in the continuation probability. Interestingly enough, they were sometimes threatened by good economic conditions, not just potential failure of the railroads involved.

12.5 Another Take on Repeated Games: Discounting

Continuation probabilities are not the only factor affecting cooperation in repeated prisoner's dilemmas. Another factor is discounting. To keep the model simple, we'll assume that the continuation probability is $p = 1$. However, future payments will be discounted.[2]

12.5.1 Valuing Treasury Bills

Suppose you are certain to receive a payment of $1000 one year from today. E.g., you have a $1000 treasury bill that matures then. The key fact here is that a treasury bill that matures in 52 weeks pays its face value on its maturity day.[3] How much is such a T-bill worth today? One way to answer this is to ask how much money you would need today to guarantee a payment of $1000 in one year.

[2] We make a brief foray into Chapter 14.

[3] You can buy T-bills from the federal government with various maturity dates of 4, 8, 13, 26, or 52 weeks in the future. You can also purchase T-bills on the secondary market. Those may even be purchased if they are going to mature the next day. Treasury bonds are government debt with longer maturities. They also make coupon payments to the owner every six months.

12.5.2 Interest Rates and Discounting

Suppose a bank will pay you an annual interest rate $r > 0$, written either as a percentage or a fraction. You could put amount A (the principal) into the bank now, and withdraw your principal plus accumulated interest in one year. Such deposits are guaranteed by the federal government, and are as safe an investment as a T-bill.

In one year's time, you would receive

$$A + rA = (1 + r)A.$$

For your savings to pay \$1000, A must solve

$$1000 = (1 + r)A \qquad \text{or} \qquad A = \frac{1000}{1 + r}$$

12.5.3 Discounted Value

We refer to A as the **discounted value** or **present value** of \$1000 received one year from now. In fact, treasury bills **pay no interest**. Rather, they are sold at a discounted price over the face value. The discount determines the interest rate that they pay.

If A is the discounted price and \$1000 the face value, the discount rate r is found by rearranging the equation above.

$$1 + r = \frac{1000}{A}.$$

This is how interest rates on T-bills are computed. The price A of T-bills is determined at auction. Then A is compared to the face value to find the implied interest rate. Older T-bills still in circulation are actively traded, and a discount rate is computed for them based on their current market price.

12.5.4 Multi-year Discounting

What happens if we have money being paid 2, 3, 4, or more years in the future?

It gets discounted once for each year. A 2-year security paying $1000 is worth

$$\frac{1000}{1+r}$$

one year from now (discounted from 2 years to one year). We discount it a second time to bring it to the present, so its discounted value today is

$$\frac{1000}{1+r} \times \frac{1}{1+r} = \frac{1000}{(1+r)^2}.$$

The interest rates are **compounded**, not added as with simple interest.

12.5.5 Frequency of Compounding

A \$1000 payment received n years in the future is discounted n times:

$$\frac{1000}{(1+r)^n}.$$

One way to get a weekly interest rate is to divide the interest rate by 52. However, an interest rate compounded weekly is different from one compounded annually. For example, if the annual interest rate is 5.2%. We divide by 52 to get a weekly interest rate of 0.1%. Compounding weekly for 52 weeks gives us a discount factor of $1/(1+0.001)^{52} = 1/1.05335$ for a rate of 5.335%. Larger interest rates or more frequent compounding will lead to a bigger difference between the two rates.

The extreme is continuous compounding. In that case we compute e^{-rt}, where t is the length of time the interest accumulates. A continuously compounded rate of 0.1% for 52 weeks yields discount factor $e^{-(0.001)52} = 1/1.05338$, for an annual rate of 5.338%.

12.5.6 Present Value

The concept of present value is closely connected with discounting. Suppose we buy a 5-year bond with a face value of \$1000 that makes a \$50 coupon payment every year. This gives us the following income stream over the next 5 years.

$$-1000, +50, +50, +50, +50, 1050$$

To find the present value, we discount each of the payments/receipts, and add them up. We will use an interest of 5% for this. Then

$$\begin{aligned}
PV &= -1000 + \frac{50}{1.05} + \frac{50}{1.05^2} + \frac{50}{1.05^3} + \frac{50}{1.05^4} + \frac{1050}{1.05^5} \\
&= -1000 + \frac{50}{1.05} + \frac{50}{1.1025} + \frac{50}{1.157625} \\
&\quad + \frac{50}{1.21550625} + \frac{1050}{1.2762815625} \\
&= -1000 + 47.62 + 45.35 + 43.19 + 41.14 + 822.70. \\
&= 0
\end{aligned}$$

In this case, the present value is zero because the bond pays 5% interest every year (the coupon payment).

12.5.7 More on Present Value

If the interest rate were different from 5%, say 4%, we can use the income stream from the bond (leaving out its cost), to compute what the current price should be.

$$
\begin{aligned}
PV &= \frac{50}{1.04} + \frac{50}{1.04^2} + \frac{50}{1.04^3} + \frac{50}{1.04^4} + \frac{1050}{1.04^5} \\
&= +\frac{50}{1.04} + \frac{50}{1.0816} + \frac{50}{1.124864} \\
&\quad + \frac{50}{1.16985856} + \frac{1050}{1.2166529024} \\
&= 48.08 + 46.23 + 44.45 + 42.74 + 863.02. \\
&= 1044.52.
\end{aligned}
$$

When the current interest rate is less than bond is paying, its present value exceeds it face value. If the interest rate were higher than 5%, the bond would sell at a discount.

12.5.8 Consols

The British government once issued bonds called **consols**. They paid interest forever. If the interest rate is constant, it is easy to value consols. Just divide by the interest rate.

If C is the annual payment, the consols present value is

$$\begin{aligned}
PV &= \frac{C}{1+r} + \frac{C}{(1+r)^2} + \frac{C}{(1+r)^3} + \cdots \\
&= \frac{C}{1+r} \times \left(1 + \frac{1}{1+r} + \frac{1}{(1+r)^2} + \frac{1}{(1+r)^3} + \cdots\right) \\
&= \frac{C}{1+r}\left(\frac{1}{1-\frac{1}{1+r}}\right) \\
&= C\left(\frac{1}{1+r-1}\right) \\
&= \frac{C}{r}.
\end{aligned}$$

12.5.9 Varying Interest Rates

The valuation of financial assets becomes more complex because the interest rate varies over time, and still more complex due to uncertainty about what the interest rate will be.

For example if the annual interest is r_1 this this year, r_2 next year, and r_3 the year after, and we have a payment stream of

$$-1000, 50, 50, 1050,$$

the present value is

$$PV = -1000 + \frac{50}{1 + r_1} + \frac{50}{(1 + r_1)(1 + r_2)} + \frac{1050}{(1 + r_1)(1 + r_2)(1 + r_3)}.$$

12.5.10 The Power of Compound Interest

One thing that can be a shock is the power of compound interest—the power of exponential growth.

Suppose you invest \$10,000 at age 20, planning for retirement. A typical real return for the stock market would be 6% per annum. In 47 years, that \$10,000 will grow to

$$\$10,000 \times (1.06)^{47} = \$154,659.17$$

giving you a healthy start on retirement. If you did this for 10 years, you'd have around a million dollars saved.

12.5.11 The Rule of 72

It's hard to get intuition about exponential growth, growth at a constant rate. Fortunately, there's a rule of thumb that works well here—the rule of 72.

The rule of 72 is helpful for estimating exponential growth (or decline). If you divide 72 by a small interest rate (72/6 = 12 in our example) you get the approximate time it takes to double your money (or cut it in half if it's shrinking by 6% per year).

So in 12 years, you'd have $20,000. In 24 years it would double again, to $40,000. By 36 years (56 years old), you'd have $80,000 and if you waited until 48 years (68 years old instead of 65), you'd have about $160,000.

Of course, this is only approximate. The actual number at 48 years is $163,938.70. Still, a 2.5% error at 48 years is not bad. You shouldn't use the rule of 72 for precise numbers, only for rough estimates.

12.5.12 Coronavirus Doubling

Dec. 1, 2022

NB: The Final is in DM-110 at 12 noon on Thursday, Dec. 8.

In April 2022, the number of new US cases of the Omicron-2 variant (BA.2 line) of the coronavirus was growing at about 30% per week. Based on the rule of 72, that would mean that the new cases doubled roughly every 2.4 weeks. That would mean quadrupling in a bit over a month. Fortunately, the rate of growth slowed down.

These numbers are large enough to make the rule of 72 a very rough approximation. Let's be more precise. If it grows at 30% per week, it actually doubles every $\ln 2/\ln 1.3 = 2.64+$ weeks (or $\log 2/\log 1.3$, which is still $2.64+$, the logarithm you use doesn't matter). That's about $18\frac{1}{2}$ days. It quadruples in 37 days, is 8 times as large in $55\frac{1}{2}$, etc.

The calculation is based on setting $2 = (1.3)^t$ where t is time in weeks. This can be written $e^{\ln 2} = e^{t \ln 1.3}$. Taking the natural logarithm yields $\ln 2 = t \ln 1.3$, so $t = \ln 2/\ln 1.3$.

12.5.13 "A Trend that Can't Continue, Won't" — after Herb Stein

The heading paraphrases Herb Stein (1916–1999), head of the Council of Economic Advisors in the early 70's. Stein's observation was based on studying economic trends. But his point applies quite generally. Growth is almost always limited by resource constraints.

In the case of the coronavirus, growth slowed down a bit, and in 37 days, the number of new cases "only" increased by a factor of 3.37 rather than 4. The actual quadrupling took 42 days to occur (a weekly growth rate of 26%).

Amazingly, in another week it fell almost to zero! This illustrates how growth can come to a sudden stop, making it difficult to make reliable predictions.

With coronavirus variants, the successful varients all spend some time growing at a nearly constant rate (exponential growth), and slow down as they come to a peak. This slowdown happens because more and more of the potential victims are immune due to previous infection or vaccination. Sometimes, as happened in early May, growth stops quite suddenly.

This often occurs in fast-growing systems, where it is just impossible to maintain that rate of growth due to resource constraints (e.g., potential infectees). This makes reliable prediction for COVID-19 impossible. Economic systems don't grow as fast, and economic predictions are not as bad as the coronavirus predictions were.

12.5.14 Repeated Prisoner's Dilemma with Discounting

Suppose we play our repeated prisoner's dilemma annually. This time, the game never ends (the continuation probability is $p = 1$), but we discount the future at annual interest rate r.

If player one uses grim trigger and player two defects in the first round, it is again best for both to defect in the second and later rounds. The income stream for player two is

$$2560, 2048, 2048, 2048, 2048, \ldots$$

with present value

$$\begin{aligned} PV &= 2560 + \frac{2048}{1+r} + \frac{2048}{(1+r)^2} + \frac{2048}{(1+r)^3} + \frac{2048}{(1+r)^4} + \frac{2048}{(1+r)^5} + \cdots \\ &= 2560 + \frac{2048}{1+r}\left(1 + p + p^2 + \cdots\right) \end{aligned}$$

where $p = 1/(1+r)$. That is

$$PV = 2560 + 2048\frac{p}{1-p}.$$

It is the same as the expected value we found in subsection 12.4.12.

We can conclude that the grim trigger leads to the Nash equilibrium where both players cooperate if and only if $p = 1/(1+r) > 1/2$, which means $1 + r < 2$ or $r < 100\%$.

12.6 Sequential Games

A **sequential game** is a game where moves are made in sequence. Repeated games are a particular type of sequential game, but not all sequential games are repeated games.

We've seen two other sequential games already, the full information version of Escape! and Stackelberg competition.

We'll examine one more such game, the Entry Game, based on the idea of predatory pricing. In the Entry Game, subgame perfection, which we introduced in section 12.3.15, plays a key role.

12.6.1 Predatory Pricing

One long-standing issue in economics has been predatory pricing, where a monopolist will extend its market by selling at a price below cost to drive competitors out of business. Due to economies of scale, it is sometimes enough to price the product below their competitor's cost. Rockefeller's Standard Oil reputedly did this, buying up its competitors as they were forced to sell by the threat of bankruptcy.

This was an important component of the Standard Oil antitrust case of 1911. According to McGee (1958), "Historians tell us that the facts revealed in *Standard Oil* were in good part responsible for the emphasis that the antitrust laws came to place upon unfair and monopolistic business practices. Perhaps the most famous of all the monopolizing techniques that Standard is supposed to have used is local price cutting."[4]

[4] John McGee (1958), Predatory price cutting: The Standard Oil (N.J.) case, *Journal of Law and Economics*, vol. **1**: 137–169.

12.6.2 Predatory Pricing and Antitrust

Some such practices were outlawed by the Clayton Act (1914), and more additional practices were outlawed by the Robinson-Patman Act of 1936. As McGee puts it, "According to most accounts, the Standard Oil Co. of New Jersey established an oil refining monopoly in the United States, in large part through the systematic use of predatory price discrimination Standard struck down its competitors, in one market at a time, until it enjoyed a monopoly position wherever competitors dared enter. Price discrimination, so the story goes, was both the technique by which it obtained its dominance and the device with which it maintained it."

McGee (1958) then argues there is little evidence to support this account. He later returned to this issue in McGee (1980).[5] If you are curious, you'll find at least the main body of both McGee articles need no more than intermediate micro.

[5] J. McGee (1980), Predatory pricing revisited, *Journal of Law and Economics*, vol. **23**, 289–330.

12.6.3 The Entry Game

We can also think about the predatory pricing issue using game theory. The Entry Game is a stylized version of the predatory pricing problem. The game pits an existing monopoly (the incumbent) against a would-be entrant into the same market.

The incumbent uses the following strategy: Reap monopoly profit if there is no competition, and if there is competition, use economies of scale to sell at a very low price, lower than is profitable for the smaller entrant. A variant of the game, which has the same solution, even has the monopolist selling below its own cost in order to drive the other firm into bankruptcy sooner.

12.6.4 Entry Game: Matrix Form

Here's a game matrix for the Entry Game:

Entry Game

Entrant \ Incumbent	Contest	Allow
Enter	-10, 20	20, 40
Stay Out	0, 100	0, 100

12.6.5 Best Responses in the Entry Game

It's clear that the incumbent's best response to entry is to allow it, and that either allowing entry or contesting it is a best response if the entrant does not entry. These moves would not actually be made in such a case, but the threat of contesting the entry still plays a role in determining equilibria.

The entrant's best response to an incumbent that allows entry is to enter, while the best response to an incumbent that contests entry is to stay out.

Entry Game

Entrant \ Incumbent	Contest	Allow
Enter	20, -10	* 40, 20 *
Stay Out	* 100, 0 *	* 100, 0

12.6.6 Nash Equilibria in the Entry Game

There are two Nash equilibria: (Allow, Enter) and (Contest, Stay Out). They are indicated on the game matrix below.

There are no mixed strategy equilibria. One way to see this is to recall that for a mixed strategy to be viable in equilibrium, it must make the opponent indifferent between at least two different moves. Yet if the entrant chooses to enter with probability p, $0 < p < 1$, the incumbent has expected payoff $20p + 100(1-p)$ if he contests entry, and $40p + 100(1-p)$ from allowing entry. Since $p > 0$, the incumbent's best response is to allow entry.

This is a rather curious result, and suggests that maybe the (Allow, Enter) equilibrium is somehow more likely.

Entry Game

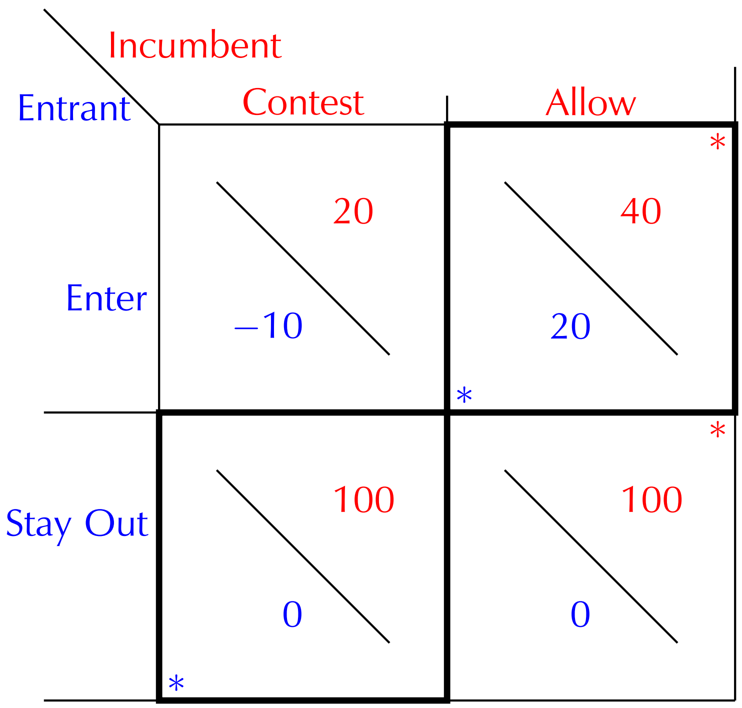

12.6.7 The Entry Game: Extensive Form

Let's take a look at the extensive form of the entry game. The incumbent knows whether the entrant has entered the market when he decides to contest. He may threaten earlier, but his ultimate decision is only made after entry.

If we try using backwards induction to solve this game, we find that it is not optimal to contest if the entrant has entered. The only equilibrium that backwards induction gives us is (Enter, Allow).

Entry Game

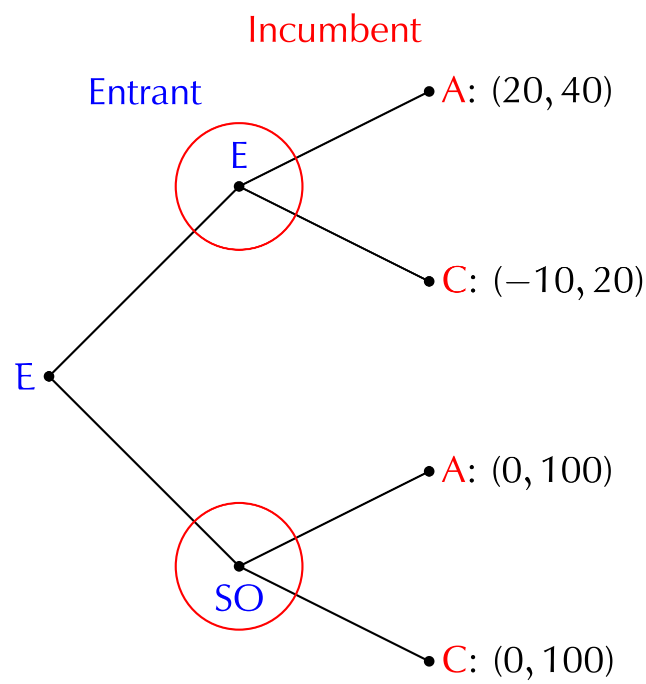

12.6.8 Not Subgame Perfect

The other equilibrium is not subgame perfect. Consider the (one-player) subgame that occurs after entry. It's illustrated below.

Contesting once the entrant has entered the market is not a best response. That means its not an equilibrium is the subgame below. Although the incumbent can threaten to contest the entry, it is not a credible threat.

As such, it's not part of any subgame perfect equilibrium.

An Entry Game Subgame

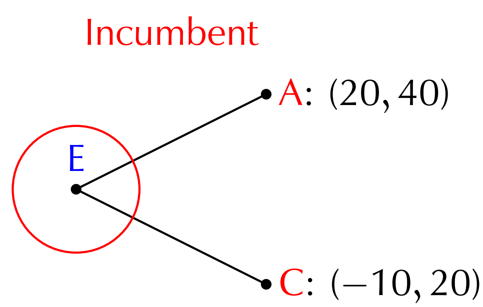

12.7 Monopsony

A **monopsony** or **buyer's monopoly** is a market with a single buyer. Just as a monopolist faces a negatively sloped demand curve, the monopsonist faces a positively sloped supply curve. We'll consider the case of a firm that is a monopsonist in one of its input markets. It uses a specialized input that no one else uses and has no close substitutes.

12.7.1 Buy More, Pay More

Let's consider the marginal cost of increasing the use of this input. A monopsonist that wants to use more of an input must pay a higher price.

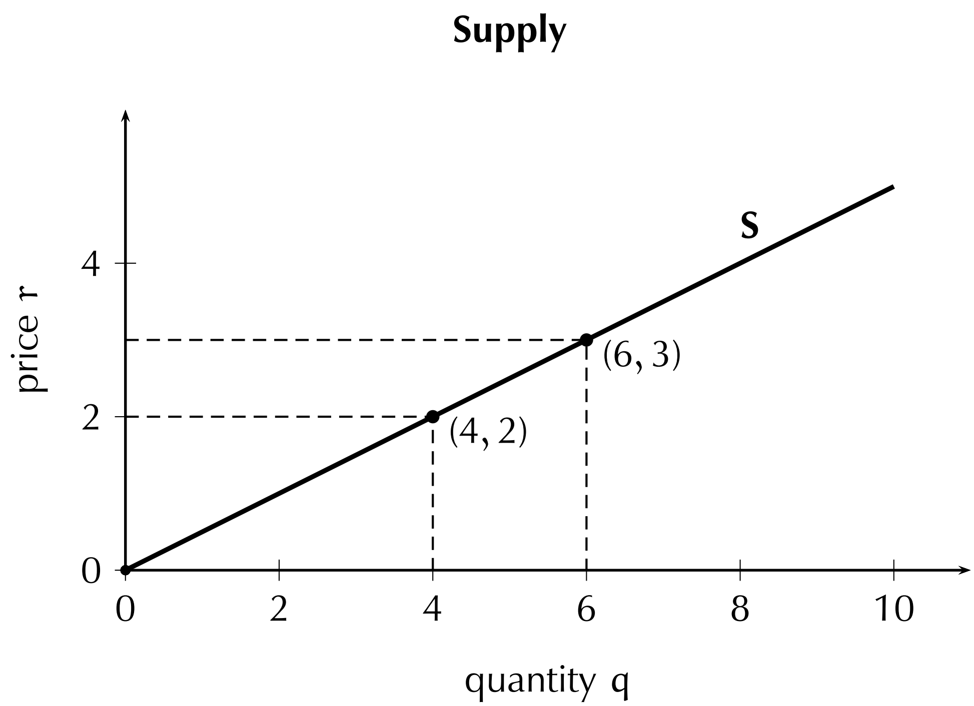

12.7.2 Marginal Factor Cost

If we buy 4 units, we pay \$2 each, for a cost of \$8. If we buy 6 units, we pay \$3 each, for a cost of \$18. Those two extra units cost us \$10, or \$5 each.

That \$5 is called the **marginal factor cost** and can be quite a bit more than the price.

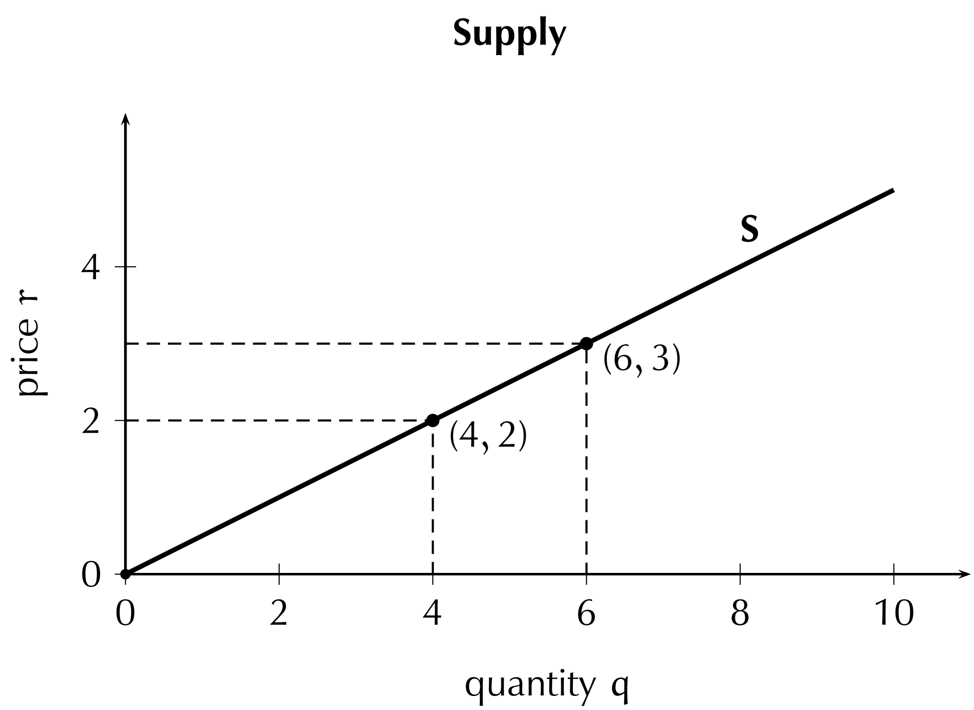

12.7.3 Marginal Factor Cost with Linear Supply

Suppose the supply curve has is described by the inverse supply function

$$p(q) = a + bq$$

where $a \geq 0$, $b > 0$. The total cost of purchasing q units is

$$p(q)q = aq + bq^2.$$

Its derivative is the marginal factor cost, MFC.

Then

$$\text{MFC}(q) = a + 2bq.$$

Like marginal revenue, it has the same vertical intercept and twice the slope. But now it's relative to the supply curve, not the demand curve.

To maximize profit, we must set marginal factor cost equal to the value of marginal profit (factor demand), solve for q, and then use the supply curve to determine the price. Here $\text{MFC}(q) > p(q)$ for $q > 0$, so the price will be less than the value of marginal product, which is equal to MFC.

12.7.4 Graphing Marginal Factor Cost with Linear Supply

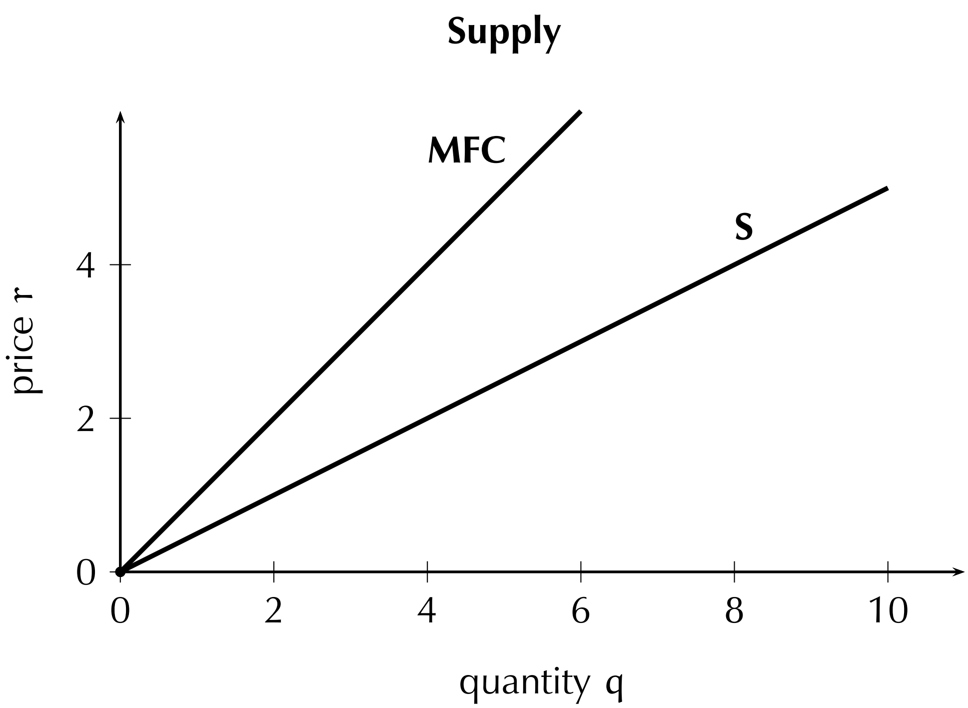

12.7.5 Marginal Factor Cost: Elasticity Formula

Another approach to marginal factor cost is via calculus. Given inverse supply function $p(q)$, the cost of q units is $p(q)q$. We differentiate this with respect to q to find the marginal factor cost.

There is an elasticity formula for marginal factor cost that is very similar to the elasticity formula marginal revenue.

$$\begin{aligned} MFC(q) &= \frac{d}{dq}\big[p(q)q\big] \\ &= p(q) + q\frac{dp}{dq} \\ &= p(q)\left[1 + \frac{q}{p}\frac{dp}{dq}\right] \\ &= p(q)\left[1 + \frac{1}{e_s}\right] \\ &> p \end{aligned}$$

The last line follows because $e_s > 0$. In comparison, $MR = p[1 + 1/e_d] < p$ as $e_d < 0$.

12.7.6 Reduced Quantity Demanded by Monopsonist

The monopsonist uses their own demand curve (for a competitive firm, this is given by the value of marginal product, $VMP_K = p\,MP_K$ for capital or $VMP_L = p\,MP_L$ for labor. The monopsonist sets MFC = VMP to maximize profit. This determines the quantity to purchase.

Compared with a competitive firm, which would produce at the intersection of supply and demand, both price and quantity sold are reduced.

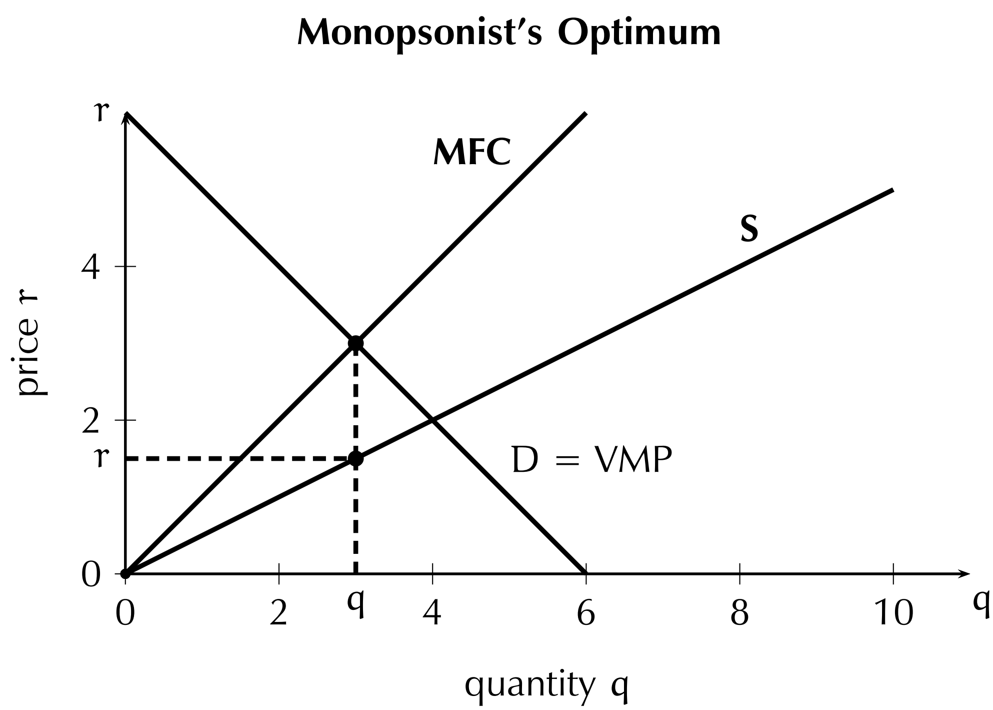

12.7.7 Surplus and Deadweight Loss of Monopsony

The buyer's (monopsonist's) surplus is light blue, the seller's surplus is green, and deadweight loss is red.

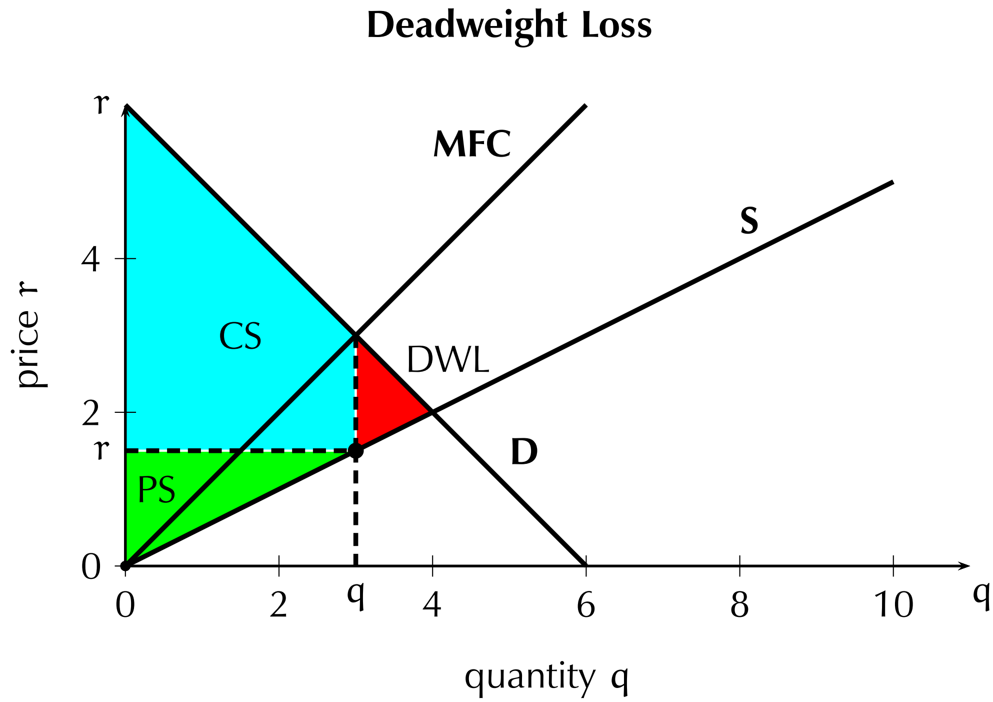

12.7.8 Pricing Strategies

Monopolists and other sellers with market power can employ alternate pricing strategies. Such tactics are also available to monopsonists.

For example, the monopsonist of the previous diagram could use a form of two-part tariff. The efficient price and quantity are determined by the intersection of supply and demand at $(r_e, q_e) = (2, 4)$. Set the per unit price to \$2, and the entry fee in order to sell to the monopsonist at the producer's surplus of $\frac{1}{2} \times 4 \times \$2 = \$4$.

Two-part Tariff

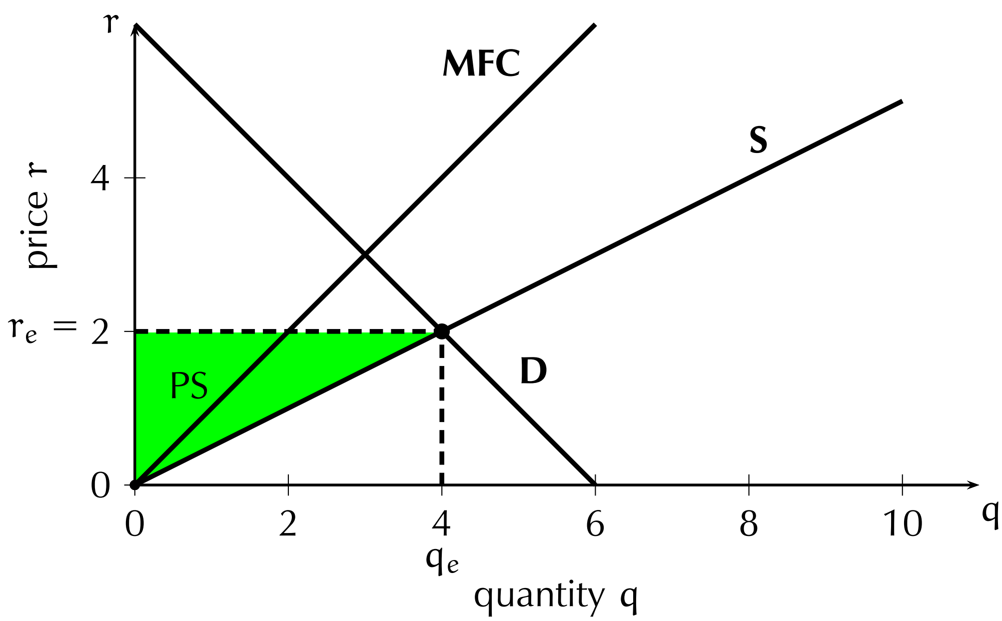

12.7.9 Monopoly vs. Monopsony

What if a monopoly faces a monopsony, where there is only one buyer and one seller of a particular good? This is also called **bilateral monopoly**. The upshot is that the players have to bargain. Such problems were considered by Edgeworth (1881), who took them as crucial for an understanding of price formation and economic equilibrium.[6] John Nash later wrote about bargaining solutions.[7]

The two sides both want the maximum possible for themselves, and obtaining the maximum possible surplus facilitates that. We can regard the problem as deciding how to divide the maximum possible total surplus—a type of Nash bargaining game.

[6] Francis Y. Edgeworth (1845–1926) was an economist and statistician. His major contribution to economics was his 1881 book, "Mathematical Psychics", C. Kegan and Paul, London.

[7] J. Nash (1950), The bargaining problem, *Econometrica*, vol. **18**, 155–162, and J. Nash (1953), Two-person cooperative games, *Econometrica*, vol. **21**, 128–140.

12.7.10 Bilateral Monopoly in Health Care

This sort of thing happens in health care, in particular, with one of my doctors. His practice is part of a larger group that periodically negotiates with various insurance companies such as FloridaBlue. Negotiations often run to the last moment, with notices being sent to customers that they may no longer be in-network.

I spoke with one of my doctors shortly after receiving such a notice. His comment was that the relationship was too valuable to not come to an agreement. All that was at issue was how much each side got. Soon afterward I received a notice that a deal had made.

12.7.11 Two Part Tariff Again!

What can happen is similar to a two-part tariff, with the provider making a payment (rebate) to the buyer (FloridaBlue). I do not know the details, but from former students in the industry, I gather they usually involve different levels of rebate if various targets are met.

The actual amount changed per visit or procedure (unit prices) is set with an eye toward how the patients will react. This all interacts with the copays and coinsurance payments made by the patients. In the end it is a rather complex problem. For reasons I don't understand, all this is often ignored in discussions of medical care pricing, even by health economists.

12.8 The Economics of Life and Death

Many economists are interested in economic policy, and economic policy often involves trade-offs between life and death. Indeed, life itself involves trade-offs between life and death!

Do you drive to the market to buy your groceries, or is it too dangerous? If you don't, you'd better find another way to get your groceries, perhaps they will deliver.

Maybe you can get someone else to take that risk for you by paying a fee. Is it worth paying? If it's \$5 is it worth paying? What about \$50? Or \$500?

12.8.1 Some Numbers on Traffic Accidents

First, we need some data. Just how dangerous is it to drive?

2020 Motor Vehicle Fatalities

	Total	Per Million
All	38680	118
Pedestrian	6236	19
Bicyclists	891	3
Motorcyclists	5458	17
Vehicle Occupants	26095	79

Source: NHSTA Safety Facts, June 2021

The average American drives about 14,250 miles per year. Suppose a trip to the grocery and back is 4 miles. Then the risk is about

$$\frac{4}{14250} \times 79 \text{ per million} \approx 22.2 \text{ per billion}$$

It's not much, but it's not zero either. How much are you willing to pay to avoid such a risk?

12.8.2 The Value of a Statistical Life Defined

Economists have addressed this very question. How much are people willing to pay to avoid small risks. From this, we can calculate the **value of a statistical life**. We multiply it by the small probabilities of death to find how much the average person is willing to pay to avoid that risk.

If you multiply by one, you have a measure of the value of a single statistical life. However, it's not a really a whole life, but merely an extrapolation from small risks. In policy analyses these small risks can add up to many lives, and this is how we value them. It's also why we call it the value of a statistical life.

Kip Viscusi is the leading researcher in this area. Much of his work is summed up for a general audience in W. Kip Viscusi (2018), "Pricing Lives: Guideposts for a Safer Society", Princeton University Press, Princeton, NJ.

12.8.3 Measuring the Value of a Statistical Life

In New York state, there is a method for valuing a lost life for legal purposes. It depends on estimating future income. This is not the sort of method economists use as it is not based on utility.

One method that is theoretically sound is to measure willingness to pay to avoid risk by looking at wages paid for jobs that are equivalent in terms of skills, difficulty, etc., but differ in the amount of risk to life and limb.

To make such measurements is not easy, and is somewhat less precise than we would like due to the various statistical corrections for other factors affecting wages. However, it ultimately produces consistent, but noisy, answers.

12.8.4 The Value of a Statistical Life

Currently, a statistical life in the US is worth about \$10,000,000, at least for people in the prime of life. It's less for both young and old people. However, although the statistical value of a life declines for older people, the statistical value of a life-year remains about the same. The increasing death rate is what reduces the statistical value of their lives.

So \$10 million times 22.2 per billion gives us 222/1000, which is a little over 22 cents.

Driving for a whole years poses a risk to vehicles occupants costing about

$$(79 \times 10^{-6}) \times (\$10 \times 10^{6}) = \$790.$$

12.8.5 COVID-19 Risk by Age

Much like deaths from all causes, deaths from COVID-19 increase rapidly with age. In fact, COVID-19 death rates increase even more rapidly than deaths from all causes up to around age 40. The table below shows the COVID-19 death rate per million in 2022 (annualized), and the fraction of deaths due to the coronavirus in each age group.

COVID-19 Deaths by Age, 2022

Age Group	Per Million	Relative to All
Infants	47.9	0.99%
1-4	8.5	3.46%
5-14	5.8	3.39%
15-24	18.3	1.95%
25-34	65.3	3.45%
35-44	151.6	5.08%
45-54	386.5	7.31%
55-64	922.6	8.03%
65-74	1609.3	8.59%
75-84	4142.7	8.62%
85 and up	11976.3	8.11%

As you can see, deaths rise steeply with age, and the risks to older people can be quite high. However, vaccination can do a lot to tame those risks. The vaccine with bivalent booster reduces the risk of death more than 14-fold according to the CDC. In my case, the risk after the bivalent booster is about 115 per million, about 50% higher than my risk of dying in a car accident while inside a car.

In dollar terms, it's about $1,150 per year. In comparison, the risk of the average 15-24 year old is 18.3/115 of that, about $183 if unvaccinated, or $13 if you've had the bivalent booster.

December 1, 2022