Sloan School of Management
Massachusetts Institute of Technology

15.013 – Economics for Strategic Decisions
Professor Robert S. Pindyck

# Lecture Notes
# on
# Intertemporal Production and Pricing

(Revised: August 2021)

Often firms must decide not simply what price to set or what quantity to produce, but how the price of a product and the firm's output should change over time. In these notes, our main concern will be with how *intertemporal production constraints* affect patterns of pricing and production. By "intertemporal production constraints," I am referring to ways in which current production can affect future costs. We will focus on two kinds of intertemporal constraints: (i) the *learning curve* (as cumulative production increases, production cost falls); and (ii) *resource depletion* (as an oil reserve is depleted, production cost rises). Before turning to these cost-related issues, however, we will begin with a brief review of how *intertemporal demand constraints* can affect pricing patterns.

## 1 Demand-Driven Dynamic Pricing

You saw an example of demand-driven dynamic pricing in 15.010 when you studied *intertemporal price discrimination*. There, you saw that book publishers typically price the hardbound and paperback versions of the book in a way that has little or nothing to do with the production costs of the book. In the case of a novel, for example, the hardbound version might sell for $25 or $30, while the paperback version, which comes out about a year later, might sell for $10 to $15. In both cases the marginal cost of printing an additional copy is likely to be only a dollar or so. The idea is that the hardbound version will be bought by

high-demand, impatient consumers, and afterwards the remaining consumers, who generally have much more elastic demands, will buy the paperback edition.

**Cream Skimming.** This strategy of setting a high initial price and lowering it over time is sometimes called "cream skimming," or "skimming pricing." It is typically used for new products that are unique, or that at least offer substantial advantages over existing ones. Besides books, examples include movies (which first play in theatres, and then later are available on DVDs), and new generations of cell phones, PDAs, music or video players, etc. In each case, some people want to own the newest and best, and other people are content to wait for prices to fall. The basic idea is to sell at a high price to low-elasticity consumers first, and then later sell at a lower price to the mass market.

A classic example of "cream skimming" is the way Polaroid priced its cameras. Table 1 shows the prices, in current and constant dollars, for Polaroid's lowest-priced color cameras. Observe the sharp decline in price during the five to ten years following the introduction of Polaroid's first color camera in 1963.[1]

| | | Price | |
|---|---|---|---|
| | Model | Current Dollars | 1975 Dollars |
| 1963-1966 | Model 100 | $164.95 | $290 |
| 1964-1967 | Model 101 | 134.95 | 234 |
| 1965-1967 | Model 104 | 59.95 | 103 |
| 1969-1972 | Color Pack II | 29.95 | 44 |
| 1971-1973 | Big Shot Portrait | 19.95 | 26 |
| 1975-1977 | Super Shooter | 25.00 | 25 |

The problem with "cream skimming" is consumer expectations. If consumers know that the price of an item will fall considerably in 6 months, they may decide to wait rather than purchase the item now. Thus producers must assess the fraction of consumers who are impatient and want the newest and best right away. Depending on the patience and

[1]The data are from *The First Thirty Years, 1948-78: A Chronology of Polaroid Photographic Products* (Cambridge, Mass: Polaroid Coproration, 1979) reproduced in T. Nagle and R. Holden, *The Strategy and Tactics of Pricing*, Prentice Hall.

rationality of most consumers, "cream skimming" may or may not be value-maximizing.

**Penetration Pricing.** In some situations, the optimal strategy is just the opposite of "cream skimming." Instead, it is best to practice *penetration pricing*—i.e., to launch a new product at a low price and then raise the price over time. As you might have seen in 15.010, and as you will see in more detail in this course, this could be optimal when there are significant positive network externalities. Obvious examples are computer operating systems (Windows) and applications software (Microsoft Office), where the positive network externalities are extremely strong. Another example is when a company wants to establish its product or system as the "standard," as with the Beta/VHS battle in videocassette recorders, and the more recent battle over the newest generation of high-definition DVD players and disks (Toshiba's HD-DVD vs. Sony's Blu-ray).

Penetration pricing can also be optimal for a monopolist facing no threat of entry. Suppose demand follows a dynamic saturation process in which the rate of growth of the number of buyers is positively related to the current number of buyers. In this case, a penetration strategy can be particularly effective because it will accelerate the rate of saturation, and thus the growth of sales.

In the remainder of these lecture notes, we will examine dynamic production and pricing in more detail, from the point of view of a firm, and also in the context of market equilibrium. Later in 15.013 we will revisit the issue of dynamic pricing, with an emphasis on gaming and competitive strategy. For example, we will examine what happens when there is a saturation process with network externalities, but there are two or three competing companies that are all trying to saturate the market simultaneously.

# 2 Opportunity Costs

A key concept that will arise repeatedly in these notes is that production today can entail an *opportunity cost* which might be positive or negative. Remember that opportunity costs are usually "invisible" to an accountant — there are no physical transfers of money that will show up on the company's books. Nonetheless, opportunity costs are true economic costs

that can affect economic decisions in important ways. An example of this (discussed in P&R, Chapter 7) is when a company makes use of an office building that it owns, but could (if management so desired) rent the building to others. Because the building is owned by the company (and might have even been fully depreciated), it is tempting to think that the cost of its office space is zero. Indeed, that is what one would see by looking at the company's books. But the office space could have been rented to others, and the forgone rent is a real cost that might well affect the company's allocation of resources. In effect, the company is renting the space to itself (at whatever is the current market rate for office space), and by taking that implicit rent into account, it might decide to make do with less office space.

How do opportunity costs arise when there are intertemporal production constraints? Consider a company that owns an exhaustible resource such as an oil reserve. Suppose that the direct cost of producing the oil is \$50 per barrel. The company should recognize that in addition to this \$50 per barrel cost, there is an additional positive opportunity cost of current production. Why? Because producing a barrel of oil today means that there will be one less barrel available to produce in the future. Thus the full economic cost to the firm of producing its oil might be much higher than \$50 per barrel.

In the case of a learning curve, the opportunity cost of current production is *negative*. Thus you might want to call this an "opportunity benefit." Why is there an opportunity benefit? Because producing one more widget today reduces the production costs for all widgets produced in the future. Suppose that currently the direct marginal cost of producing another widget is \$20, but as the firm moves down the learning curve by producing more and more widgets, this direct marginal cost is expected to fall, eventually reaching \$10. If you produce one more widget today, you incur the direct cost of \$20, and in return you get a widget (which you can sell at whatever the market price happens to be), but you also get something else — as a result of moving slightly down the learning curve, you get small reductions in all future direct costs of production. The "opportunity benefit" of producing the extra widget today is the present value of all of those small reductions in future costs. Figuring out what that number is not as difficult as it might appear, as we will see later. But suppose for the moment that the opportunity benefit (negative opportunity cost) is \$5.

In that case, the full marginal cost of producing a widget today is not $20, but instead only $15. It is this $15 number that the firm should use when deciding how much to produce and/or what price to set.

We will examine these opportunity costs and their implications for pricing and production in more detail in the following sections. We begin the next section by examining production and pricing when there is a learning curve. We will see that if the discount rate is low or if learning occurs over a short time horizon, determining the opportunity benefit and thus the full marginal cost of producing today is relatively easy — just make believe you have already reached the bottom of the learning curve, and treat the final direct marginal cost as though it were the current marginal cost. We will also see how a learning curve can increase a firm's market power.

In Section 4 we will examine the pricing and production of exhaustible resources. The owner of an exhaustible resource, such as oil, knows that the resource can be produced now or in the future, but that there is only a finite amount in the ground. Should the resource be produced now, or should the owner wait? As you might expect, the answer depends on how fast (or whether) we expect the price of the resource to rise over time. And how fast should we expect the price of the resource to rise? Thinking strategically, we might assume that other resource producers are faced with exactly the same problem we have, and know what we know, and make their decisions accordingly. We will see that this leads to some fairly simple results regarding our production decision.

Finally, we will revisit the exhaustible resource production problem in the context of uncertainty. If oil prices fluctuate unpredictably (which they in fact do), how should this affect the production decision of an owner of an oil reserve, and how does it affect the value of the reserve? We will see that oil in the ground has an *option value*, much like the value of a financial call option. Just as greater stock price volatility increases the value of a call option on the stock, so too does greater oil price volatility increase the value of an oil reserve.

# 3 The Learning Curve

Consider a market in which firms move down a learning curve: that is, as they produce, learning-by-doing reduces their average and marginal costs. In this case, the full marginal cost of current production is *less* than current marginal production cost. The reason is that an incremental unit of current production reduces future production costs by moving the firm farther down the learning curve, so that production of the unit brings a benefit (a negative opportunity cost) that partly offsets its cost.

This is illustrated in Figure 1, where marginal cost $\mathrm{MC}_t$ is constant with respect to the rate of output $Q_t$, but declines (from an initial value of $\mathrm{MC}_0$, asymptotically to $\mathrm{MC}_{\mathrm{inf}}$) as cumulative output increases. Now consider a firm that has the objective of *maximizing its value* — i.e., maximizing the present discounted value of the profits it will earn over time from producing. The firm is just starting to produce so that its marginal cost is currently $\mathrm{MC}_0$. If it faces the marginal revenue curve MR, how much output should it intially produce? If it produced where its current marginal cost intersected marginal revenue, its output would be $Q_0$ and its price $P_0$. Although this would indeed maximize the firm's *current* profit, it would *not* maximize the value of the firm. Instead the firm should take into account the negative opportunity cost (labelled MOC in the figure) of current production, and produce where its *full* marginal cost, MC*, intersects marginal revenue. As shown in the figure, its value-maximizing current output and price are $Q^*$ and $P^*$ respectively.

Note that in this particular example, $P^*$ is *less* than the firm's current marginal cost. Does this mean that the firm is engaging in some kind of predatory pricing (or worse yet, acting irrationally)? No. The firm is simply taking into account that its full marginal cost is considerably less than its current direct marginal cost. Although an accountant might point out that the firm is currently losing money, it is in fact pursuing a value-maximizing production and pricing strategy. (Make sure you can explain the virtue of pricing below current marginal cost in this situation to other people!) Of course price and marginal cost will both fall over time as the learning curve bottoms out, so that the marginal opportunity cost eventually falls to zero.

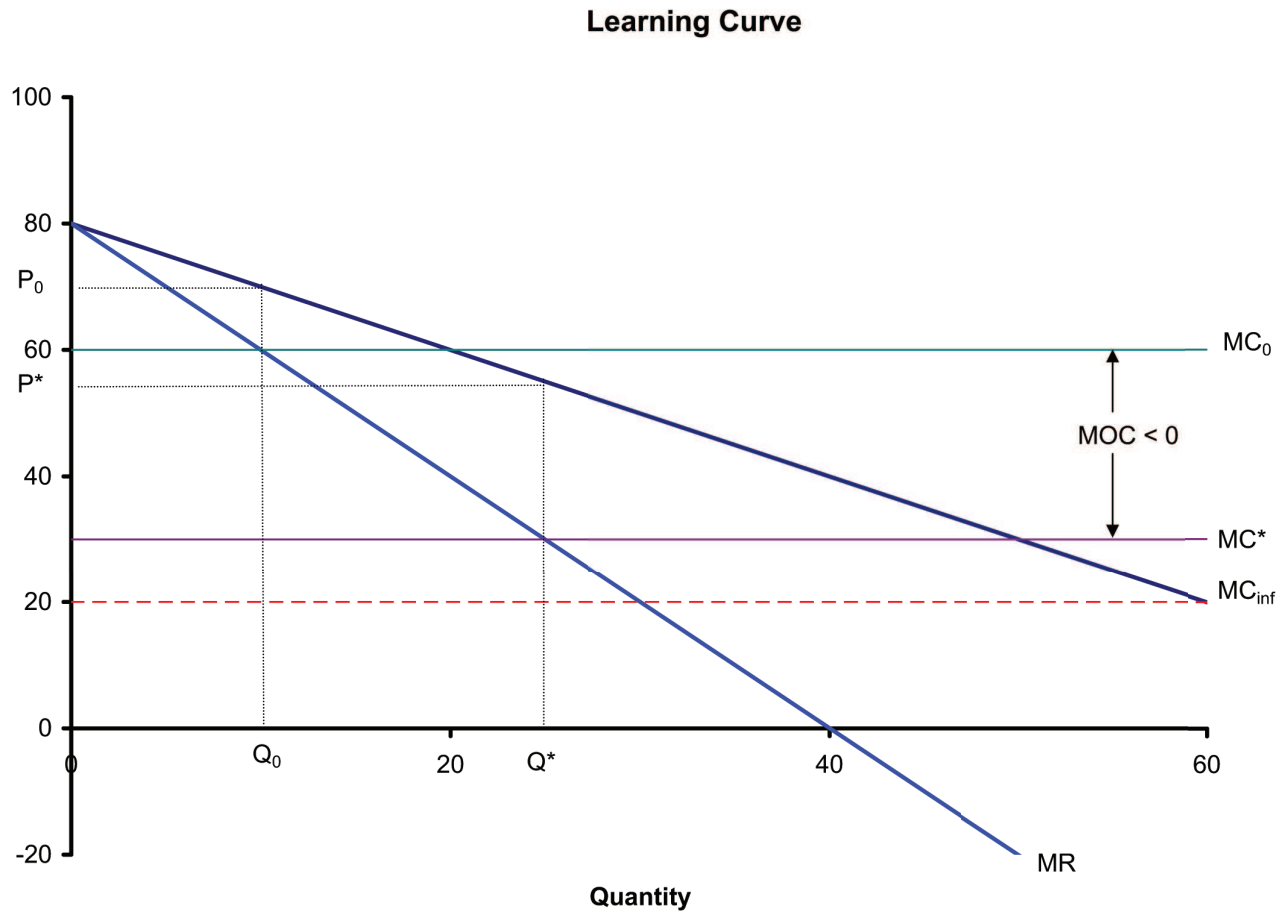


Figure 1: Learning Curve

We have still not addressed the important question of how the firm can estimate the initial value of its marginal opportunity cost MOC and thus its full marginal cost, MC*. In effect, the firm must choose the series of declining prices over time (or, equivalently, the series of increasing outputs) that maximizes the present value of the resulting series of profits. As you will see later (in the context of a more complex strategic problem), given estimates of the demand curve and learning curve, this is not all that difficult. As a practical matter, however, there is an even simpler and more straightforward way to approach this problem, as shown below.

## 3.1 Zero Discount Rate

If the firm's discount rate is very low (or if learning occurs quickly), the solution to this optimal production problem is very simple: *base current production on final marginal cost*, i.e., on the marginal cost that the firm expects to see after it has moved all the way down the learning curve. To make this clear, let's assume for now that the discount rate is zero. Why in that case is it optimal base current production on the final marginal cost? *Because if the discount rate is zero, we can collapse this month's production, next month's production, production two months from now, etc., into one great big production decision.* As a result, the marginal unit of output is the last unit of output two or three years from now after we have gotten down to the bottom of the learning curve. Hence the relevant marginal cost is the marginal cost after all learning has been completed. This is an extremely important point, and one we will come back to again and again. Make sure you understand it!

Figure 2 illustrates this simple case in which the discount rate is zero. The firm should simply ignore its current marginal cost and set output at the point where its final marginal cost (i.e., the marginal cost that will apply when the firm reaches the bottom of the learning curve) equals marginal revenue.

In practice, a firm's discount rate is rarely zero. What if the discount rate is, say, 10 percent? In that case the optimal production rule depends on the time horizon over which learning is likely to occur. For many industries, that time horizon is in fact quite short. In the case of semiconductors (e.g., microprocessors, memory chips, and application-specific integrated circuits, or ASICs), most learning-by-doing happens within a year. In the case of commercial aircraft manufacturing, most learning also happens within a year, or at most two years. (See Example 7.8 in P&R, 9th Ed.) If the time horizon is two years or less, our simple rule of basing output on final marginal cost will be very close to optimal (usually yielding a present value of profits within one or two percent of the maximum), even for moderate discount rates. And what if the discount rate is 15 percent and learning occurs over a three-year horizon? As you will see (from computer calculations in an exercise later in this course), the simple rule will still bring you to within 5 or 10 percent of the maximum

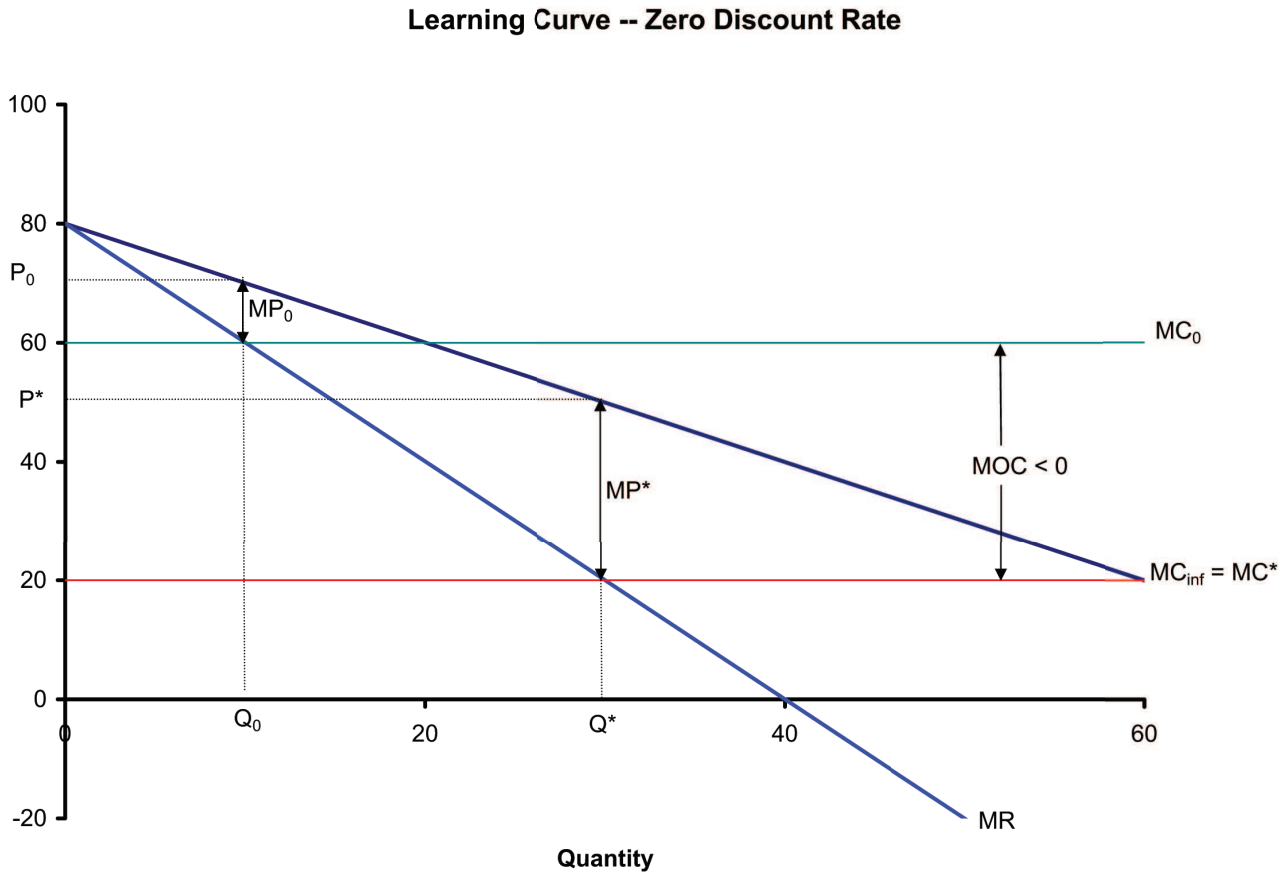


Figure 2: Learning Curve When Discount Rate is Zero

present value of profits — and will do much, much better than ignoring the learning curve and using current marginal cost to make current production decisions.

Of course to apply this simple rule, the firm needs an estimate of its final marginal cost. How can it obtain that estimate? In practice, this is not so difficult. In many cases there is historical data that can be used to estimate the learning curve, final marginal cost, and the amount of time it is likely to take to reach or get close to that final marginal cost. In the case of semiconductors, for example, Intel and National Semiconductor have been making microprocessors, memory chips, and ASICs for many years, and thus have a very good sense of what the learning curve will look like as they begin production of a new product. Likewise, Boeing and Airbus have been producing commercial aircraft for many years, and thus had good estimates of their final marginal costs of production for the 787 and A380 airplanes.

But what about a brand new product, such as Apple's iPad? Although most of an iPad's cost is for components manufactured by other companies, there is still a significant cost of final assembly and testing. Even before the first iPad was produced, Apple engineers and production planners (perhaps with the help of a consulting firm) could have estimated the iPad's final marginal cost.

### 3.2 The Learning Curve and Market Power

The firm in Figure 2 clearly has market power; it has downward-sloping average and marginal revenue curves. But is the firm's market power in any way affected by the presence of a learning curve? Compared to a hypothetical situation in which there is no learning-by-doing, does the learning curve result in more or less market power?

As Figure 2 illustrates, the learning curve *increases* the firm's degree of market power. In the figure, MP denotes the firms incremental monopoly power, measured as the difference between the price it charges and its full marginal cost. Suppose that there was no learning, so that marginal cost is expected to remain at $MC_0$. In that case the profit-maximizing (and value-maximizing) price and output would be $P_0$ and $Q_0$, and the monopoly profit would be $MP_0$. With the learning curve, the firm's value-maximizing price is lower ($P^* < P_0$), but full marginal cost is much lower ($MC^* < MC_0$), so that the monopoly profit is larger ($MP^* > MP_0$). Intuitively, the learning curve works like a reduction in marginal cost, which expands output. As output expands, the gap between average revenue (i.e., price) and marginal revenue widens, thereby increasing the firm's incremental profit.

## 4 Exhaustible Resources

We turn now to what is effectively the opposite of the learning curve. In the case of an *exhaustible resource*; pricing and production are dynamic (i.e., intertemporal) in nature because a unit of the resource produced and sold today will not be available to be produced and sold in the future. Thus, producing the unit involves not only the usual marginal production cost, but also a positive opportunity cost associated with the foregone future

production. In many textbooks, including P&R, that opportunity cost is referred to as a *user cost* of production, so I will sometimes refer to it that way here.

The nature of this user cost is easiest to understand in the context of a simple example. Suppose you own a quantity of in-ground reserves $R$ of an exhaustible resource like oil. To keep things simple, suppose you can produce the resource either this year or next year. Then how much should you produce in each year?

We will solve this problem in two ways. First, for those of you who enjoy mathematics, we will use the method of Lagrange multipliers. Those of you who are not familiar with this technique can skip this mathematical approach, because we then solve the problem again, but intuitively.

## 4.1 Mathematical Approach

Let $q_1$ be the amount produced this year, and $q_2$ the amount produced next year. Let $r$ be the market interest rate, $p_1(q_1)$ and $p_2(q_2)$ the demand functions (demand may change over time), and $C(q)$ the cost function corresponding to the physical production (extraction from the ground) of the resource. We would like to choose $q_1$ and $q_2$ so that (i) $q_1 + q_2 = R$, and (ii) we maximize the present discounted value of the flow of profits. We therefore have the following *constrained maximization problem*:

$$\max_{q_1,q_2} \; PDV = p_1(q_1)q_1 + \frac{p_2(q_2)q_2}{(1+r)} - C(q_1) - \frac{C(q_2)}{(1+r)} \tag{1}$$

subject to

$$q_1 + q_2 = R \tag{2}$$

This is a straightforward problem that can be solved using the method of Lagrange multipliers. Set up the Lagrangian, and take derivatives with respect to $q_1$ and $q_2$:

$$L = p_1 q_1 + \frac{p_2 q_2}{1+r} - C(q_1) - \frac{C(q_2)}{(1+r)} - \lambda\left(q_1 + q_2 - R\right) \tag{3}$$

$$\frac{\partial L}{\partial q_1} = \mathrm{MR}_1 - \mathrm{MC}_1 - \lambda = 0 \tag{4}$$

so

$$\lambda = \mathrm{MR}_1 - \mathrm{MC}_1 \tag{5}$$

$$\frac{\partial L}{\partial q_2} = \frac{\text{MR}_2}{(1+r)} - \frac{\text{MC}_2}{(1+r)} - \lambda = 0 \tag{6}$$

so

$$\lambda = \frac{(\text{MR}_2 - \text{MC}_2)}{(1+r)} \tag{7}$$

Combining (5) and (7) we get the following fundamental result:

$$\text{MR}_2 - \text{MC}_2 = (1+r)(\text{MR}_1 - \text{MC}_1) \tag{8}$$

or, in general:

$$\boxed{\text{MR}_{t+1} - \text{MC}_{t+1} = (1+r)(\text{MR}_t - \text{MC}_t)} \tag{9}$$

First note that MR and MC are no longer equal to each other. Why not? Because MC is not the *full* marginal cost—it is only the marginal production cost, and does not include the user cost. The *full* marginal cost is MC + $\lambda$ and $\lambda$ = MR – MC is the *user cost* associated with depleting your reserves by 1 unit.

Next, note that the user cost *rises with the rate of interest*. It should be clear to you why this must be the case. An in-ground unit is an *asset*—just like a bond—and its current value is exactly MR – MC, i.e., the extra revenue less the extra cost resulting from extracting and selling it. You would expect that asset to yield a return (in the form of capital gain—i.e., price increase) equal to other assets. And what is the return on a bond? Just the interest rate $r$. So you would expect MR – MC to increase at the rate $r$.

In a *competitive* market, $p$ = MR, so the user cost is $p$ – MC, and this rises at the rate of interest. This is the well-known "Hotelling Rule" (after Harold Hotelling) for the price of a competitively produced exhaustible resource.

## 4.2 Intuitive Approach

We have just demonstrated mathematically that in a competitive market for an exhaustible resource, price minus marginal cost should rise at the rate of interest. But this is also easy to see intuitively.

If I own reserves of oil, I own an asset, and I should expect that asset to earn a competitive return. Since reserves of oil sitting in the ground do not pay a dividend, the only way I can

earn a return is by producing and selling the oil, or by having the value of the reserves go up over time, i.e., by earning a capital gain. How will I decide whether to keep the oil in the ground or to produce it today? If I expect price net of marginal cost to rise at a rate *greater* than the rate of interest, I would keep the oil in the ground, because then the capital gain would exceed the competitive return. Likewise, if I think that price net of marginal cost is going to rise at *less* than the rate of interest, I will produce my oil and sell it immediately.

But everyone else holding reserves of oil is thinking about this in exactly the same way. Hence, if price net of marginal cost should rise at less than the rate of interest, everyone will try to sell their oil immediately, pushing the price down, so that later it will rise faster (as total reserves are more rapidly depleted). If price net of marginal cost were to rise at a rate greater than the rate of interest, everyone would try to hold their oil in the ground, so that price would immediately rise, and thereafter would rise at a slower rate (as total reserves are more slowly depleted). Hence in a competitive market equilibrium, price minus marginal cost will rise at the rate of interest.

This behavior of the competitive price is illustrated in Figure 3. Part (a) shows how price begins at a level $P_0 > c$ and rises until it reaches a point $P_T$ at which demand drops to zero (perhaps because a substitute for the resource becomes more economical). Part (b) of the figure shows how price and production move up along the demand curve during the lifetime of the resource, until finally the cut-off price $P_T$ is reached at which production ceases.

If the owner of the oil reserves is a monopolist (or has monopoly power), he or she will go through the same calculation, except with respect to *marginal revenue* net of marginal cost, rather than price net of marginal cost. Why? Because for a monopolist the value of a unit of reserves is the marginal revenue from selling the unit minus the marginal cost. Thus the owner will produce the resource at exactly that rate such that marginal revenue minus marginal cost rises at the rate of interest.

Typical price trajectories predicted by this theory for a competitively produced versus monopolistically produced exhaustible resource are shown in Figure 4. Observe that the competitive price should rise more rapidly (starting from a lower initial level). This is what you would expect, since with monopoly power price is greater than marginal revenue. Also,

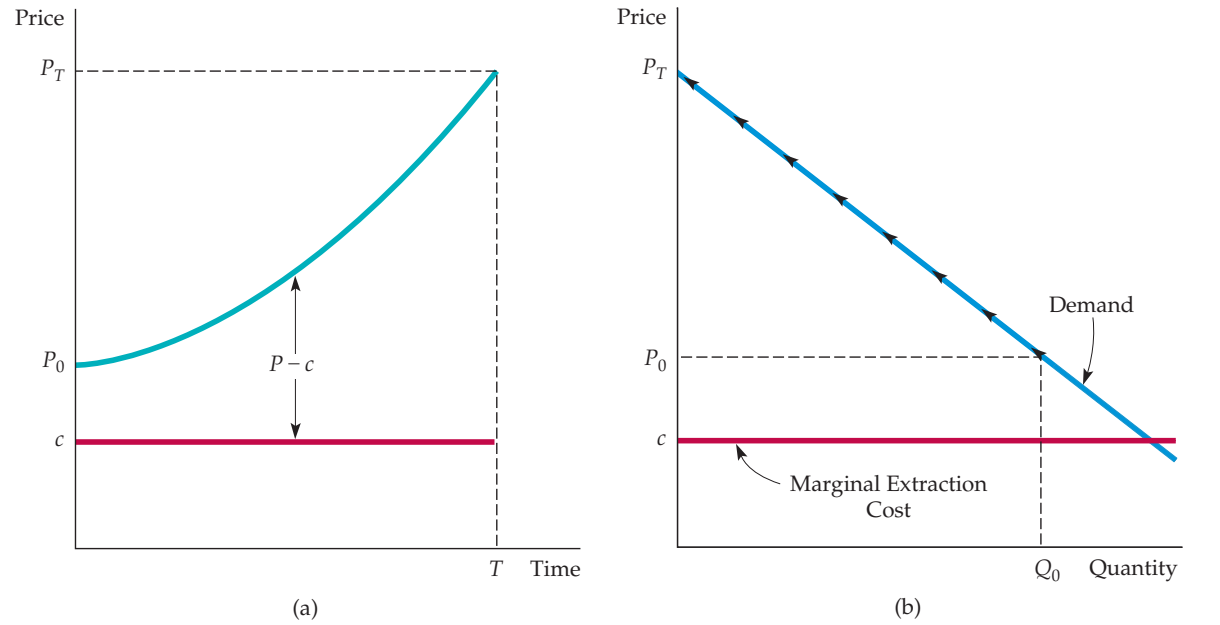


Figure 3: Price Evolution (a) and Movement Along Demand Curve (b) for an Exhaustible Resource

observe that the monopolist is more "conservationist," i.e., he/she produces the resource more slowly than a competitive market would. Finally, it can be shown that a competitive market depletes an exhaustible resource at the socially optimal rate. A monopolist, on the other hand, is overly *conservationist* (out of self-interest, of course).

Of course in the real world, prices of oil and other exhaustible resources do not follow the smooth trajectories of Figure 4. The reason is that demand and costs are constantly fluctuating as market conditions change. And in the case of oil, market structure is changing over time, as OPEC's cohesiveness and ability to affect prices waxes and wanes. In fact, these changes in market conditions are largely unpredictable, so that changes in price are partly unpredictable. Shortly we will consider the implications of this fact for production decisions and valuation.

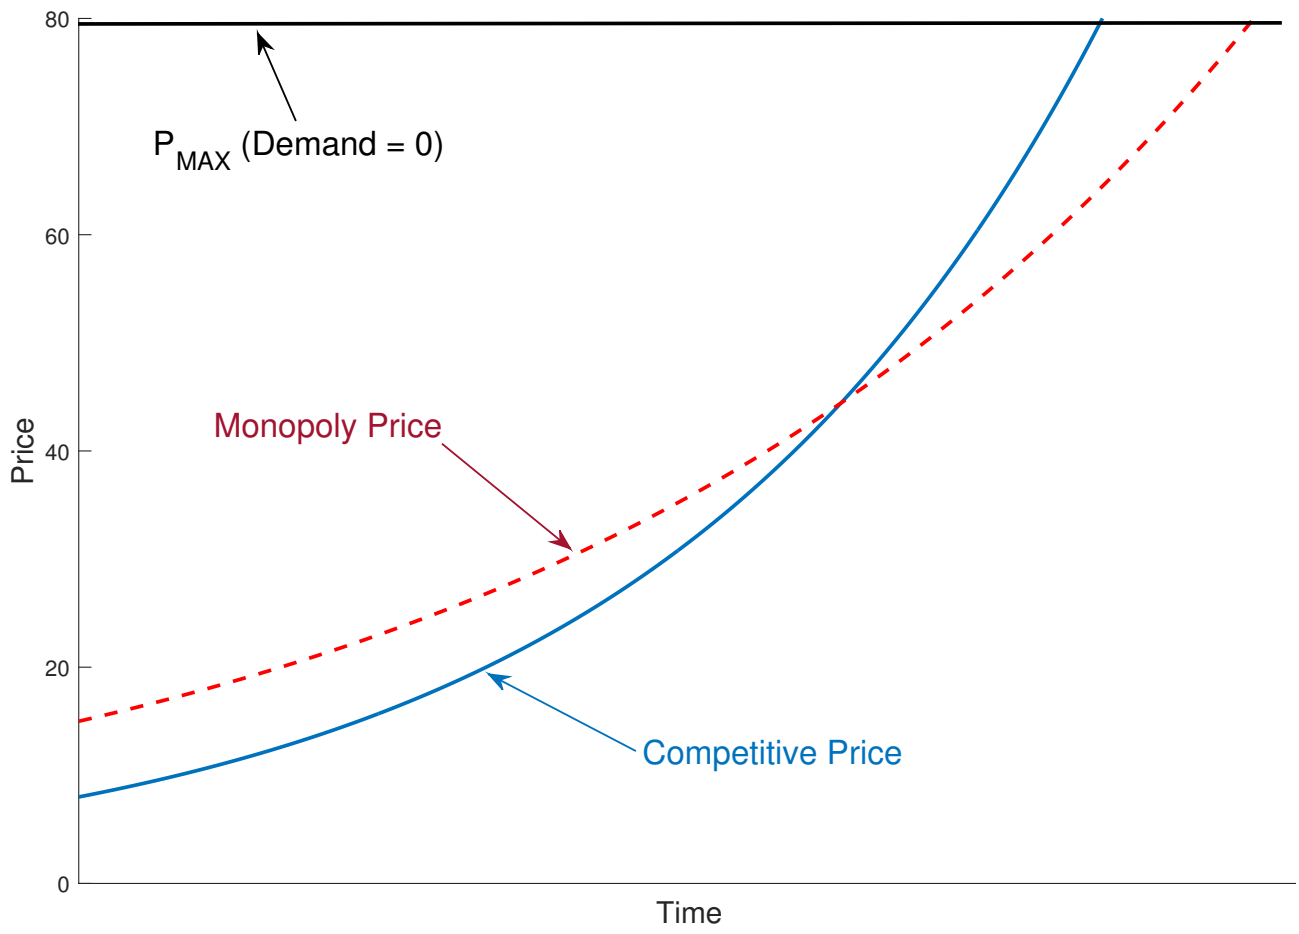


Figure 4: Price of an Exhaustible Resource: Competition versus Monopoly

## 4.3 Market Power

Other things equal, does the producer of an exhaustible resource have more or less market power than the producer of an "ordinary" good? Here we will see that the intertemporal production constraint implied by exhaustibility *reduces* market power.

### 4.3.1 Special Case: Isoelastic Demand and Zero Extraction Cost

Assume that demand is fixed and isoelastic, and marginal cost MC = 0. In this special case, it can be shown that the competitive and monopoly output paths are identical. To see this, write the demand curve as:

$$P = \text{AR} = Q^{-\eta} \tag{10}$$

Note that the price elasticity of demand, $(P/Q)dQ/dP$, is $-1/\eta$. Given this demand curve, marginal revenue is:

$$\text{MR} = (1-\eta)Q^{-\eta} \tag{11}$$

For a monopolist, MR – MC must rise at the rate of interest. Since MC = 0, MR must rise at rate of interest. The percentage rate of growth of MR is given by $(d\text{MR}/dt)/\text{MR}$. Hence:

$$\frac{d\text{MR}/dt}{\text{MR}} = \frac{-\eta(1-\eta)Q^{-\eta-1}(dQ/dt)}{(1-\eta)Q^{-\eta}}$$

$$= -\eta\frac{dQ/dt}{Q} = r \tag{12}$$

For a competitive industry, $P - \text{MC}$ must rise at the rate of interest. Since MC = 0,

$$\frac{dP/dt}{P} = \frac{-\eta Q^{-\eta-1}(dQ/dt)}{Q^{-\eta}} = -\eta\frac{dQ/dt}{Q} = r \tag{13}$$

Observe that for *both* the monopolist and the competitive industry, $(dQ/dt)/Q = -r/\eta$. Since the amount of reserves is the same in both cases (by assumption), we know that the quantity trajectories are identical. Hence at every point in time, $Q_c = Q_m$ and $P_c = P_m$. Thus, for this special case of isoelastic demand and zero extraction cost, the monopolist has *no* monopoly power!

### 4.3.2 General Case

Extraction cost is rarely zero, so that in general a monopolist would produce less than a competitive industry initially but more later (the monopolist "overconserves"). Nonetheless, the fact that resource reserves are depletable will still reduce the monopolist's degree of monopoly power. To get further insight into why depletion reduces market power, consider the case illustrated in Figure 5, where the demand curve is linear. In that figure, there is a maximum price $\overline{P}$ at which demand becomes zero.

First, suppose the interest rate $r$ were extremely high, say 10,000%. Then only current profits would matter, so future depletion could be ignored. Hence a monopolist would

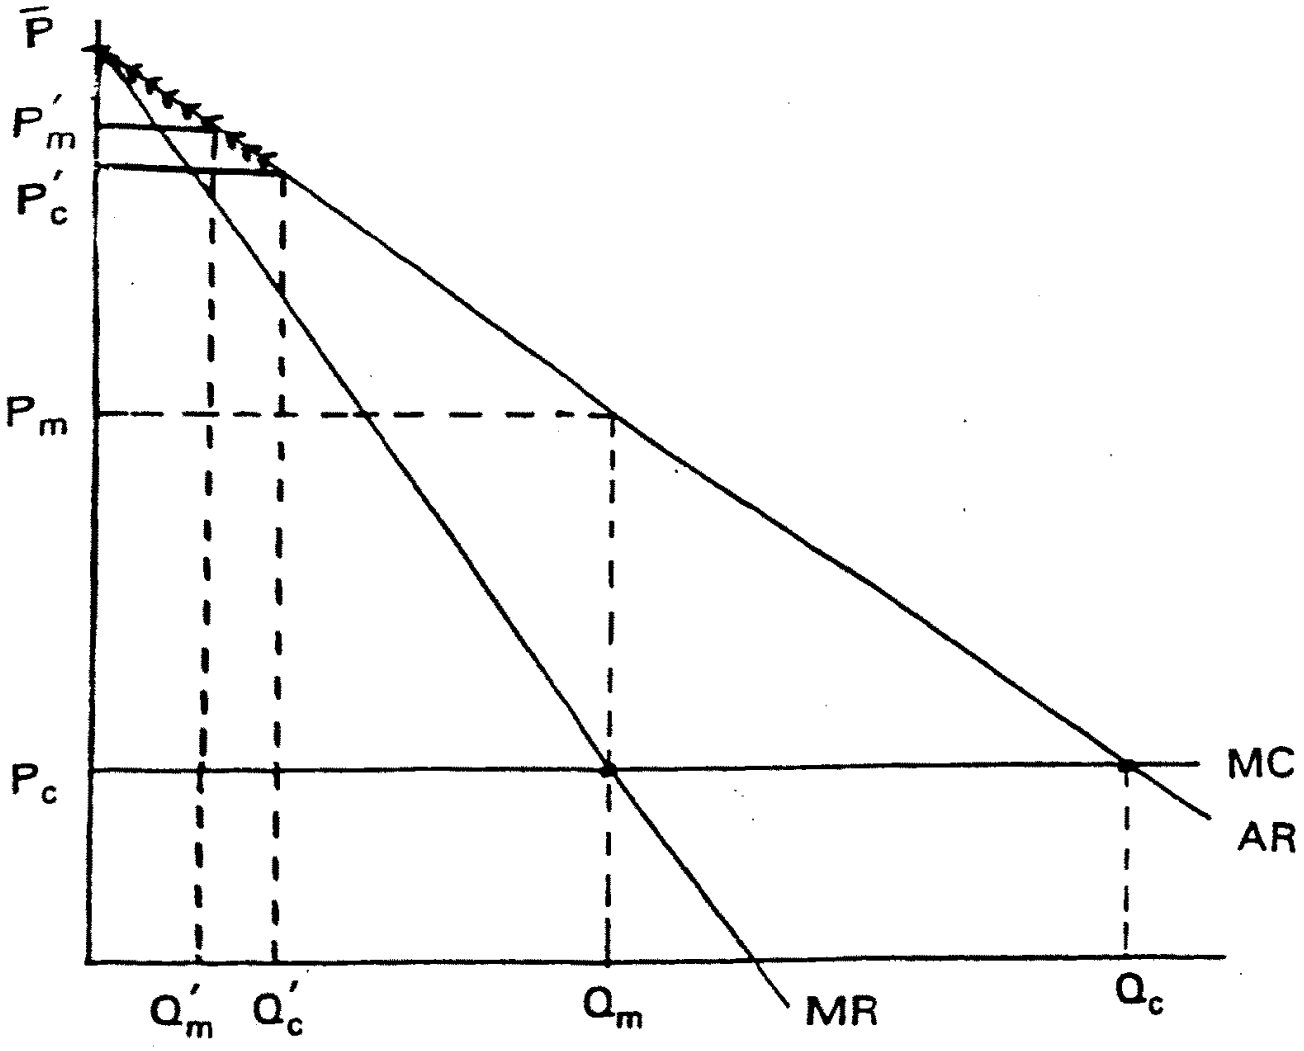


Figure 5: Depletion and Market Power

produce where marginal revenue is equal to marginal cost, i.e., $Q_m$, and a competitive industry would produce $Q_c$, where market demand intersects marginal cost. In the competitive market, price $P_c$ would equal marginal cost, and is much less than the monopoly price, $P_m$.

Next, suppose the interest rate is *zero*, so that profits in every period are worth the same. Since the last unit of the resource will be sold at the price $\overline{P}$, and since $P - \text{MC}$ does not grow (because $r = 0$), in a competitive market price is *always* equal to $\overline{P}$ (or perhaps a penny below $\overline{P}$), and $Q$ is very, very small. But if $P = \overline{P}$, $\text{MR} = \overline{P}$ as well, and since MR – MC does not grow, a monopolist would also keep price equal to $\overline{P}$ (or a penny below).

Hence the monopoly price and competitive price are the same, and there is *no* monopoly power.

If the interest rate is positive but small (e.g., $r = 1\%$), $P$ and MR both grow very slowly, and start from very high levels, not much below $\overline{P}$. Now the monopoly price ($P'_m$ in Figure 5) is above the competitive price ($P'_c$), but the difference is small. Hence there is only very little monopoly power.

## 4.4 Implications of Uncertainty

In the real world, market conditions change rapidly and unpredictably. What does this imply for production decisions, and what uncertainty over future market conditions do the the value of a firm's resource reserves?

Suppose, for example, that you owned an oil reserve that contained 1000 barrels of oil which could be produced at any time at a cost of $50 per barrel. Suppose that today the price of oil is $60 per barrel, so that if you produced all of the oil today you would earn $10 per barrel, for a total of $10,000. Suppose that your best prediction of what the price of oil will be next year and for the following several years is $60. If you were *certain* that the price would remain at $60 per barrel for the next several years you would surely produce and sell all of your oil. After all, if you don't expect the value of your oil reserve to increase at the rate of interest (and in this hypothetical you don't expect it to increase at all), you are better off selling all of the oil and investing the $10,000 that you receive in a bond.

But now suppose that although your best prediction for oil prices in the future is $60, you recognize that there is a great deal of uncertainty over the price. In other words, you believe that the probability that the price will increase by some amount is equal to the probability that it will decrease by that amount (so that the expected future price is $60), but it is quite likely that in a few years the price of oil will either be much higher or much lower than $60. Should you produce all of your oil now? And if someone offered to buy your oil reserve for $10,100, should you accept the offer?

If you think about this a little bit, you will see that owning an oil reserve is analogous to owning a call option on a stock. Figure 6 illustrates this. The straight line labelled $P - 50$

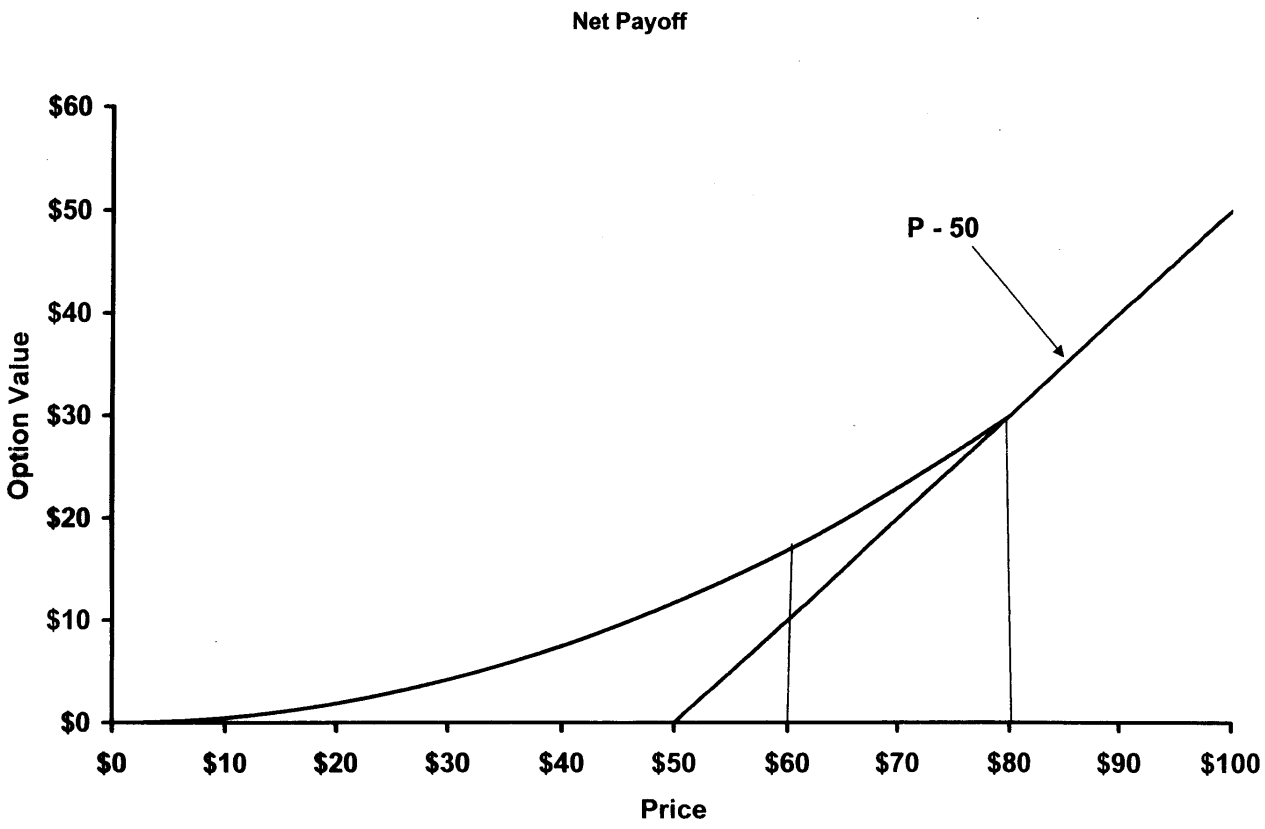


Figure 6: The Option Value of an Oil Reserve

is the payoff from producing and selling a barrel of oil as a function of the price of oil. At a price of \$50, the payoff is zero; at a price of \$60 the payoff is \$10; and at a price of \$80 the payoff is \$30. There is no limit to the payoff on the upside — but the downside is clearly limited to zero. If the price falls to \$40 you simply won't produce any oil. Likewise if the price falls to \$30, or \$20, or anything else below your production cost of \$50. This means that uncertainty over the future price of oil *increases* the value of the reserve. *Your oil reserve is equivalent to a call option on a dividend paying stock.*

The curved line in Figure 6 indicates the value of each barrel of oil as a function of the price of oil. Note that at the current price of \$60, the value of each barrel is *greater* than \$60 − \$50 = \$10. This is a result of the asymmetric payoff function combined with uncertainty over price changes. If you took a basic finance course you should immediately recognize Figure 6 as the value of a call option on a dividend paying stock. And because it

is a dividend paying stock, there is an optimal exercise point ($80 in the figure); at prices above that point, the option value is equal to the net payoff. In the case of an oil reserve, the "dividend" is the money you receive from producing and selling the oil.

Thus valuing oil reserves and deciding when to produce the oil boils down to valuing a call option and determining the optimal exercise price. In practice, "exercising" an oil reserve option means *developing* the reserve. An oil reserve that has been discovered and delineated is called an undeveloped reserve; a great deal of money must be spent to develop the reserve by sinking production wells, building pipelines, etc. Valuing the undeveloped reserve and deciding how high the price of oil should be before the reserve is developed is the fundamental option valuation problem faced by oil companies. Table 1 lays out in more detail the analogy between an undeveloped oil reserve and a financial call option.

**Table 1: Comparison of Stock Call Option and Undeveloped Petroleum Reserve**

| **Stock Call Option** | **Undeveloped Reserve** |
|---|---|
| Current stock price | Current value of developed reserve |
| Variance of rate of return on the stock | Variance of rate of change of the value of a developed resource |
| Exercise price | Development Cost |
| Time to expiration | Relinquishment requirement |
| Riskless rate of interest | Riskless rate of interest |
| Dividend | Net production revenue less depletion |

# 5 Price Forecasting

Now that you've decided what to do with your little 1000-barrel oil reserve, you want to move on to bigger things, so you become a consultant to the oil industry. Your job is to help companies figure out what to do with their *real* oil reserves. You walk in to the offices of the Acme Oil Corporation and are greeted warmly. The top executives of the company have been pulling out their hair trying to figure out whether oil prices are going to rise or fall, and they see you as their savior. So they ask you: "What will the price of oil be two years from now?" What is your answer?

In case you actually find yourself in this situation, here is the correct answer: "I don't know what the price of oil will be two years from now. Nobody knows what the price of oil will be two years from now. You should know better than to ask such as stupid question." How are the executives likely to respond to your answer? They will be impressed with your intelligence and honesty, and they will thank you profusely. Of course they may have a follow-up question: "*Why*, exactly, can't we forecast the price of oil?" Here's why.

## 5.1 Price Volatility

Over the last few decades, the price of oil (WTI crude in the U.S.) has dropped from over $140 per barrel to less than $30 per barrel. How could this be? Are price changes like this rare events?

No, the price of oil has always been extremely volatile. (The same is true for many other industrial commodities.) Figure 7 shows the price of West Texas Intermediate (WTI) crude over the past 30 years. (WTI is the benchmark price for the U.S.; Brent Crude is the benchmark price for Europe, and is often $2 or $3 higher than WTI.) Note the sharp increases — and decreases — in the price. From 2005 to 2008 the price more than tripled, going from about $45 per barrel to $140. In less than a year, it then plummeted to $40. In 2019 the price fell to just above $20 per barrel, but then in 2021 it rose to around $70. Clearly the price of oil is extremely volatile. Why is this?

The answer is that in the short run, the price elasticities of both supply and demand are

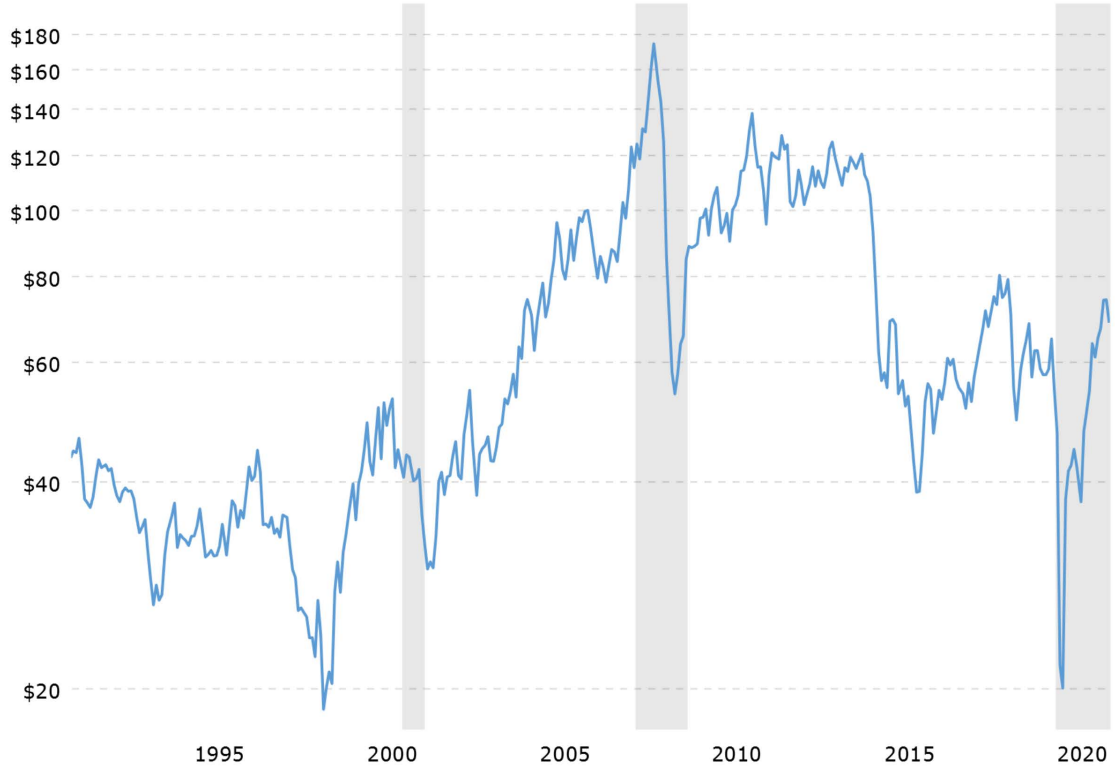


Figure 7: The Price of WTI Crude Oil (in 2021 Dollars)

extremely small. (The short-run supply elasticity is about 0.05, and the short-run demand elasticity is about $-0.05$.) As illustrated in Figure 8, this means that even a small shift in the demand or supply curve can result in a very large price change. (For a detailed explanation, see Example 2.9, "Upheaval in the World Oil Market," in Pindyck and Rubinfeld, *Microeconomics*.) In the longer run, the elasticities are larger, so both supply and demand will respond to the the higher or lower price, but over the space of a year or even several years, we will observe sharp price changes.

What happened in 2014–2015 to drive down oil prices? World oil consumption was about 94 million barrels per day (mb/d). There was a net increase in production of about 3 mb/d: U.S. production increased by 3 mb/d (mostly from fracking), Brazil and Russia each up 0.5 mb/d, Iraq up 1 mb/d, and Libya down by 2 mb/d. Consumption, however, only rose by about 1 mb/d. So either inventories would have to increase by about 2 million barrels per day on average, or the price would have to fall. The latter happened.

# WHY THE SHARP PRICE MOVEMENTS?

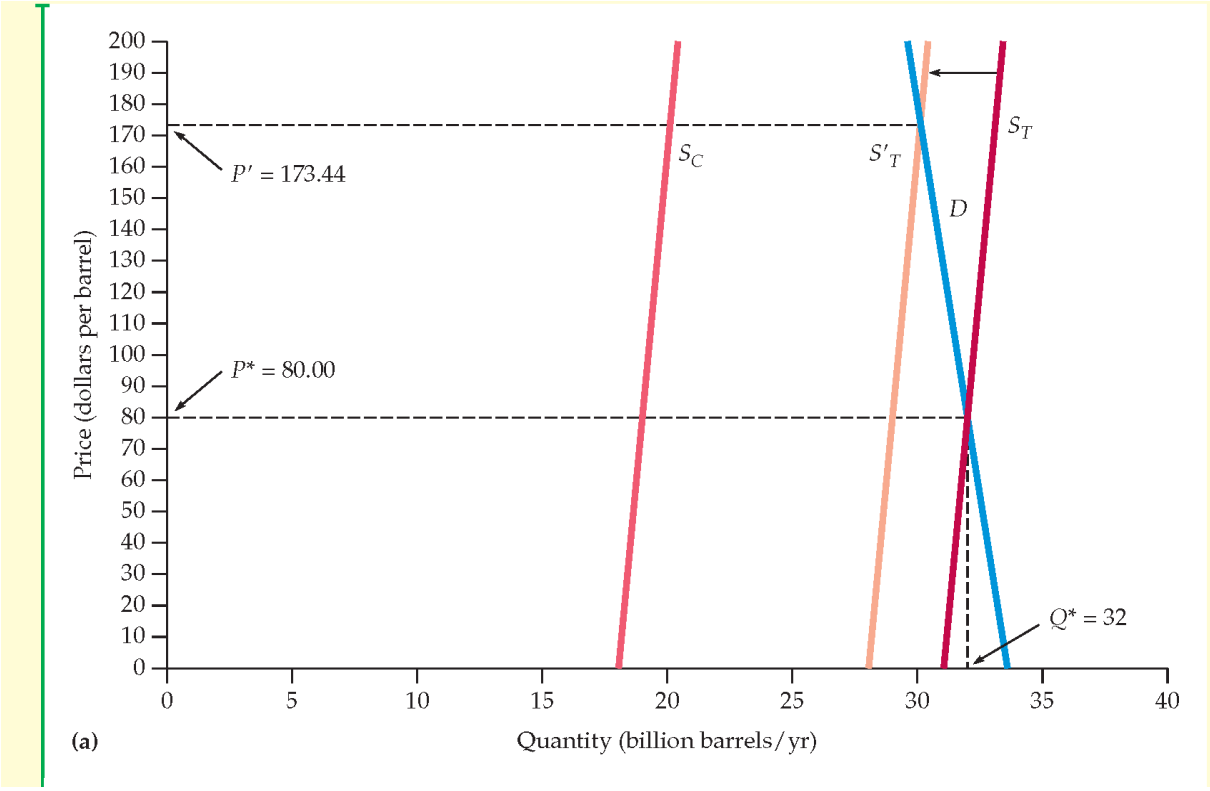


Figure 8: Short-Run Supply and Demand for Crude Oil. The short-run elasticities of supply and demand are very low, so a small shift in the supply curve or the demand curve will result in a large change in price.

Can't we forecast these kinds of changes? *No.* Take Libyan oil production as an example. Production fell in 2014 because the country did not have a functioning government, and armed groups battled for control of the country, destroying most oil production in the process. No one predicted this, and no one can predict when (or if) Libya will have a functioning government that will enable oil exports to fully resume or even expand (pushing the world price of oil down). People try to forecast oil prices, and there are plenty of forecasts published all the time. But to get a sense of how useful they are, look at Figure 10. This forecast, made in 2011 by the International Energy Agency is typical of what one can find. How useful do you think it is?

What does the futures market tell us about the future price of oil? The futures market

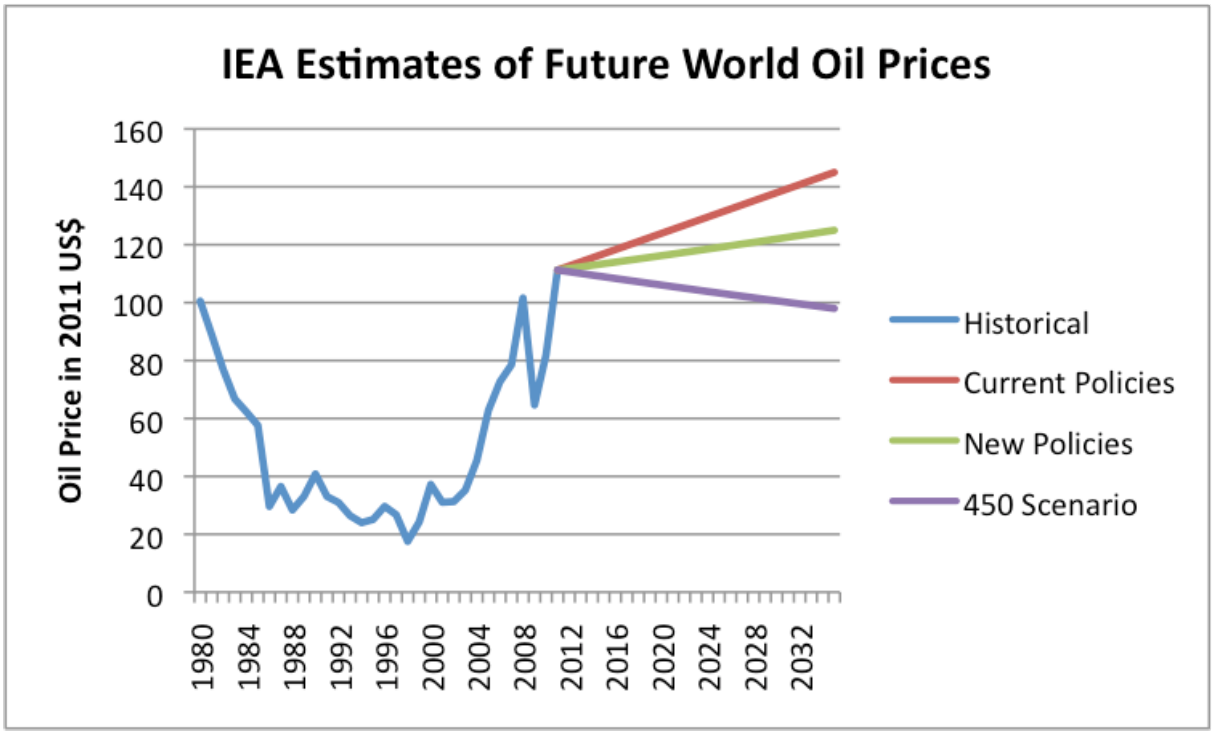


Figure 9: Forecast of Oil Price by the International Energy Agency

serves a "price discovery" role — it summarizes the beliefs of investors who hold long or short positions on futures contracts. The futures price is an estimate of the expected future spot price, but it is a biased estimate. If, in the context of the CAPM, the commodity has a positive "beta" (which an industrial commodity like oil certainly does), the expected future spot price will be higher than the futures price.[2]

Figure 10 shows the futures price for WTI crude as of August 2021, for horizons up to two years. In August 2021, the spot price was about \$70. Note that the futures price declines slightly to about \$64 in mid-2023, so that the forecasted spot price also decreases (to about \$60 at the end of the two year horizon). That price forecast is interesting, but what is much more interesting is the 95% confidence interval, which at the end of the two-year horizon ranges from \$30 to nearly \$120 — a very large spread. (The confidence interval is calculated from the concurrent prices of options on the futures.) Put simply, this means that *there*

[2]Specifically, the bias is equal to the risk premium applied to the time horizon of the futures contract.

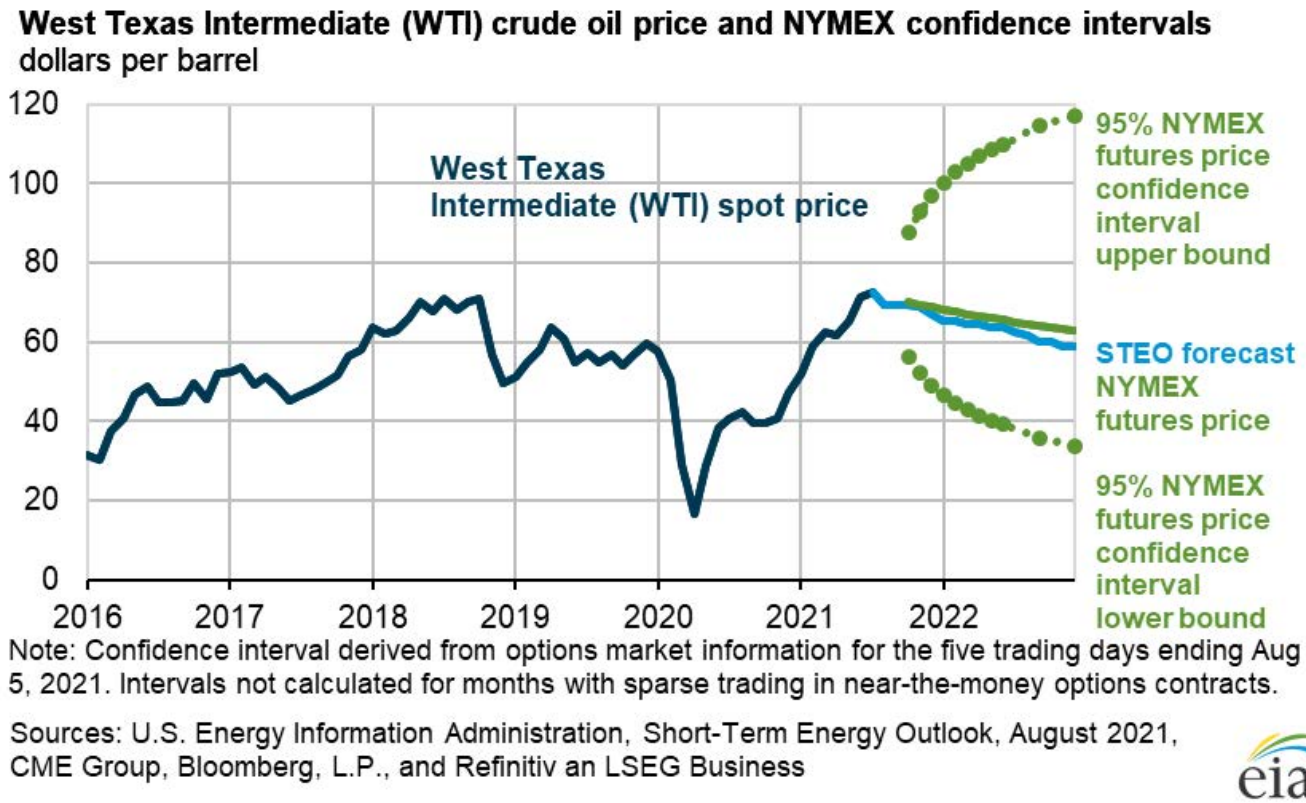


Figure 10: What Does the Futures Market Tell Us? This figure shows the futures price for WTI crude oil as of August 2021.

*a huge amount of uncertainty over the future price of oil.* In effect, the forecast from the futures market doesn't tell us much beyond the fact that there is a great deal of uncertainty.

Oil is not the only commodity with a volatile price. In fact, the prices of most commodities — copper, gold, silver, lumber, orange juice, coffee, ... — are extremely volatile. For example, Figure 11 shows how volatile the price of copper has been over the past few decades. The reason for the volatility is the same as for oil: in the short run, elasticities of demand and supply are very low.

Asking what the price of oil (or the price of copper or gold) will be in two years is a silly question. But is there a better question that oil company executives might ask? Yes. A better question is: "How much volatility is there in the oil market? What is the range of uncertainty over the future price of oil?" Answers to these questions can help you determine the option value of oil in the ground, and the kinds of risks your company might face.

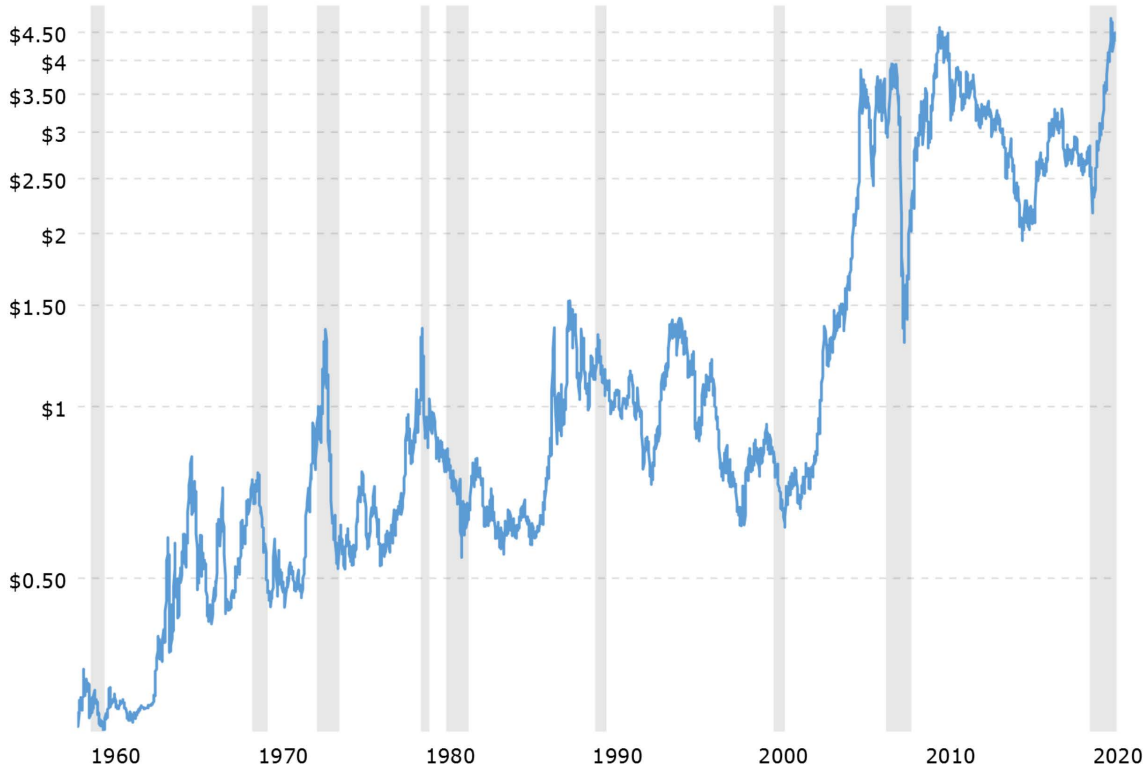


Figure 11: The Price of Copper in 2021 Dollars per Pound.

# 6 Related Problems of Intertemporal Pricing and Production

The production and pricing of an exhaustible resource or a good for which there is a learning cure are just special cases of the general problem of *intertemporal pricing and production*. In other words, pricing and production decisions today are interrelated with pricing and production decisions in the future. In the case of an exhaustible resource, the intertemporal aspect of the problem comes about because of depletion, i.e., the "using up" of in-ground reserves. You have seen many other ways in which pricing and production decisions have intertemporal aspects. For example, the demand for most products is inherently dynamic (recall the saturation curve model for the demand for computers, or the dynamic demand characteristics in the case of automobiles or gasoline).

Basically, the solution of intertemporal pricing and production problems usually boils

down to maximizing the present discounted value of a net cash flow over time. This is illustrated once again in the example that follows.

**Example: Investing in Wine.** Consider a decision to invest in cases of wine. Suppose that cases of freshly pressed wine cost \$100 each, and from experience you know that the value of a case of wine held for $t$ years is $(100)t^{1/2}$. There are 100 cases available. Suppose that the (continuously compounded) rate of interest per year is $r = 0.10$.

1. How many cases should you buy, how long should you wait to sell them, and how much money will you receive at the time of their sale?

2. Suppose that at the time of purchase, someone offers you \$130 per case immediately. Should you take it?

3. Now suppose that the inflation rate is $g = 0.05$. How many cases should you buy, and how long should you hold them?

**Solution.**

1. Buying wine is a good investment if the net present value is positive. If we buy a case and sell it after $t$ years, we pay \$100 now and receive $\$100t^{1/2}$ at the end. The NPV of this investment is

$$\begin{aligned} \text{NPV} &= -100 + e^{-rt}100t^{1/2} \\ &= -100 + e^{-0.1t}100t^{1/2} \end{aligned}$$

If we do buy a case, we will choose $t$ to maximize the NPV:

$$\frac{d\text{NPV}}{dt} = e^{-0.1t}\left(50t^{-1/2}\right) - 0.1e^{-0.1t}100t^{1/2} = 0$$

$$50t^{-1/2} - 10t^{1/2} = 0$$

so $t^* = 5$ years. Then, if we hold the case 5 years, the NPV is:

$$\text{NPV} = -100 + e^{-0.1(5)}100(5)^{1/2} = 35.670$$

Therefore, we *should* buy the wine and hold it for 5 years. Since each case is a good investment, we should buy all 100.

Another approach is to compare holding wine to putting your \$100 in the bank. The bank pays you 10 percent interest, while the wine increases in value at the rate

$$\frac{d(\text{value})/dt}{\text{value}} = \frac{50t^{-1/2}}{100t^{1/2}} = \frac{1}{2t}$$

As long as $t < 5$, the return on wine is greater than 10 percent, the rate of interest; after $t = 5$, the return on wine drops below the rate of interest. Therefore, $t = 5$ is the time to switch your wealth from wine to the bank. As for whether to buy the wine at all, if we put \$100 in the bank, we will have $100(e^{rt}) = 100(e^{0.5}) = \$164.87$ after 5 years, whereas if we spend the \$100 on wine, we will have $100t^{1/2} = 100(5^{1/2}) = \$223.61$.

2. You just bought the wine and are offered \$130 for resale. You should accept if the NPV of the deal is positive. You get \$130 now, but lose the $\$100(5^{1/2})$ you would get if you hold the wine and sell it in 5 years. The NPV is

$$\text{NPV} = 130 - e^{-0.1(5)}100(5^{1/2}) = -5.6 < 0.$$

Therefore, you should *not* sell. Another way to see this is to note that the \$130 could be put in the bank and would grow to \$214.33 ($\$130e^{0.5}$) in 5 years. This is still less than \$223.61 ($\$100(5^{1/2})$), the value of the wine in 5 years.

3. Assume that wine prices keep up with inflation, i.e., the *real* price of wine is $(\$100)t^{1/2}$. If the nominal interest rate is $r$, the *real interest rate* is $r - g$. Thus, if $r = 0.1$ and $g = 0.05$, the NPV of wine sold at $t$ is

$$\text{NPV} = -100 + e^{-0.05t}100t^{1/2}$$

Maximizing this expression as before implies

$$\begin{aligned}
\frac{d\text{NPV}}{dt} = e^{-0.05t}(50t^{-1/2}) - 0.05e^{-0.05t}100t^{1/2} &= 0 \\
50t^{-1/2} - (0.05)100t^{1/2} &= 0 \\
50 - (0.05)100t &= 0 \\
t &= 10
\end{aligned}$$

Inflation reduces the real interest rate (for a fixed nominal rate) and so it pays to hold onto wine longer before selling it and putting the money in the bank. (In reality, inflation may cause investors to look for real assets like wine with prices that keep up with inflation.) You should again buy all the cases. (Notice that if the nominal interest rate $r$ rises when inflation $g$ increases, so that $r - g$ is constant, inflation does *not* change the results of part 1.)