Investment Under Uncertainty

(Princeton U. press, 1994)

Avinash Dixit and Robert Pindyck

(Princeton U. press, 1994)

Chapter 1

A New View of Investment

Economics defines investment as the act of incurring an immediate cost in the expectation of future rewards. Firms that construct plants and install equipment, merchants who lay in a stock of goods for sale, and persons who spend time on vocational education are all investors in this sense. Somewhat less obviously, a firm that shuts down a loss-making plant is also "investing": the payments it must make to extract itself from contractual commitments, including severance payments to labor, are the initial expenditure, and the prospective reward is the reduction in future losses.

Viewed from this perspective, investment decisions are ubiquitous. Your purchase of this book was an investment. The reward, we hope, will be an improved understanding of investment decisions if you are an economist, and an improved ability to make such decisions in the course of your future career if you are a business school student.

Most investment decisions share three important characteristics it varying degrees. First, the investment is partially or completely *irreversible*. In other words, the initial cost of investment is at least partially sunk; you cannot recover it all should you change your mind. Second, there is uncertainty over the future rewards from the investment. The best you can do is to assess the probabilities of the alternative outcomes that can mean greater or smaller profit (or loss) for your venture. Third, you have some leeway about the timing of your investment. You can postpone action to get more information (but never, of course, complete certainty) about the future.

These three characteristics interact to determine the optimal decisions of investors. This interaction is the focus of this book. We develop the theory of irreversible investment under uncertainty, and illustrate it with some practical applications.[1]

The orthodox theory of investment has not recognized the important qualitative and quantitative implications of the interaction between irreversibility, uncertainty, and the choice of timing. We will argue that this neglect explains some of the failures of that theory. For example, compared to the predictions of most earlier models, real world investment seems much less sensitive to interest rate changes and tax policy changes, and much more sensitive to volatility and uncertainty over the economic environment. We will show how the new view resolves these anomalies, and in the process offers some guidance for designing more effective public policies concerning investment.

---

[1] Some decisions that are the opposite of investment-getting an immediate benefit in return for an uncertain future cost- are also irreversible. Prominent examples include the exhaustion of natural resources and the destruction of tropical rain forests. Our methods apply to these decisions, too.

Some seemingly non-economic personal decisions also have the characteristics of an investment. To give just one example, marriage involves an up-front cost of courtship, with uncertain future happiness or misery. It may be reversed by divorce, but only at a substantial cost. Many public policy decisions also have similar features. For instance, public opinion about the relative importance of civil rights of the accused and of social order fluctuates through time, and it is costly to make or change laws that embody a particular relative weight for the two. Of course the costs and benefits of such non-economic decisions are difficult or even impossible to quantify, but our general theory will offer some qualitative insights for them, too.

1 The Orthodox Theory

How should a firm, facing uncertainty over future market conditions, decide whether to invest in a new factory? Most economics and business school students are taught a simple rule to apply to problems of this sort. First, calculate the present value of the expected stream of profits that this factory will generate. Second, calculate the present value of the stream of expenditures required to build the factory. Finally, determine whether the difference between the two - the *net present value* (NPV) of the investment - is greater than zero. If it is, go ahead and invest.

Of course, there are issues that arise in calculating this net present value Just how should the expected stream of profits from a new factory be estimated? How should inflation be treated? And what discount rate (or rates; should be used in calculating the present values? Resolving issues like these are important topics in courses in corporate finance, and especially capital budgeting, but the basic principle is fairly simple-calculate the NPV of an investment project and see whether it is positive.

The net present value rule is also the basis for the neoclassical theory of investment as taught to undergraduate and graduate students of economics, Here we find the rule expressed using the standard incremental or marginal approach of the economist: invest until the value of an incremental unit of capital is just equal to its cost. Again, issues arise in determining the value of an incremental unit of capital, and in determining its cost. For example, what production structure should be posited? How should taxes and depreciation be treated?

Much of the theoretical and empirical literature on the economics of investment deals with issues of this sort. We find two essentially equivalent approaches. One, following Jorgenson (1963), compares the per-period value of an incremental unit of capital (its marginal product) and an "equivalent per-period rental cost" or "user cost" that can be computed from the purchase price, the interest and depreciation rates, and applicable taxes. The firm's desired stock of capital is found by equating the marginal product and the user cost. The actual stock is assumed to adjust to the ideal, either as an ad hoc lag process, or as the optimal response to an explicit cost of adjustment. The book by Nickell (1978) provides a particularly good exposition of developments of this approach.

The other formulation, due to Tobin (1969), compares the capitalized value of the marginal investment to its purchase cost. The value can be observed directly if the ownership of the investment can be traded in a secondary market; otherwise it is an imputed value computed as the expected present value of the stream of profits it would yield. The ratio of this to the purchase price (replacement cost) of the unit, called Tobin's q, governs the investment decision. Investment should be undertaken or expanded if q exceeds 1; it should not be undertaken, and existing capital should be reduced, if $q < 1$. The optimal rate of expansion or contraction is found by equating the marginal cost of adjustment to its benefit, which depends on the difference between q and 1. Tax rules can alter this somewhat, but the basic principle is similar. Abel (1990) offers an excellent survey of this q-theory of investment. In all of this, the underlying principle is the basic net present value rule.

2 The Option Approach

The net present value rule, however, is based on some implicit assumptions that are often overlooked. Most important, it assumes that either the investment is reversible, that is, it can somehow be undone and the expenditures recovered should market conditions turn out to be worse than anticipated, or, if the investment is irreversible, it is a now or never proposition, that is, if the firm does not undertake the investment now, it will not be able to in the future.

Although some investments meet these conditions, most do not. Irreversibility and the possibility of delay are very important characteristics of most investments in reality. As a rapidly growing literature has shown, the ability to delay an irreversible investment expenditure can profoundly affect the decision to invest. It also undermines the simple net present value rule, and hence the theoretical foundation of standard neoclassical investment models. The reason is that a firm with an opportunity to invest is holding an "option" analogous to a financial call option-it has the right but not the obligation to buy an asset at some future time of its choosing. When a firm makes an irreversible investment expenditure, it exercises, or "kills," its option to invest. It gives up the possibility of waiting for new information to arrive that might affect the desirability or timing of the expenditure; it cannot disinvest should market conditions change adversely. This lost option value is an opportunity cost that must be included as part of the cost of the investment. As a result, the NPV rule "invest when the value of a unit of capital is at least as large as its purchase and installation cost" must be modified. The value of the unit must exceed the purchase and installation cost, by an amount equal to the value of keeping the investment option alive.

Recent studies have shown that this opportunity cost of investing can be large, and investment rules that ignore it can be grossly in error. Also, this opportunity cost is highly sensitive to uncertainty over the future value of the project, so that changing economic conditions that affect the perceived riskiness of future cash flows can have a large impact on investment spending, larger than, say, a change in interest rates. This may help to explain why neoclassical investment theory has so far failed to provide good empirical models of investment behavior, and has led to overly optimistic forecasts of effectiveness of interest' rate and tax policies in stimulating investment.

The option insight also helps explain why the actual investment behavior of firms differs from the received wisdom taught in business schools. Firms invest in projects that are expected to yield a return in excess of a required, or "hurdle," rate. Observers of business practice find that such hurdle rates are typically three or four times the cost of capital.[2] In other words, firms do not invest until price rises substantially above long-run average cost. On the downside, firms stay in business for lengthy periods while absorbing operating losses, and price can fall substantially below average variable cost without in-

[2] Summers (1987, p. 300) found hurdle rates ranging from 8 to 30 percent, with a median of 15 and a mean of 17 percent. The cost of riskless capital was much lower; allowing for the deductibility of interest expenses, the nominal interest rate was 4 percent, and the real rate close to zero. See also Dertouzas et al. (1990, p. 61). The hurdle rate appropriate for investments with systematic risk will exceed the riskless rate, but not by enough to justify the numbers used by many companies.

ducing disinvestment or exit. This also seems to conflict with standard theory, but as we will see, it can be explained once irreversibility and option value are accounted for.

Of course, one can always redefine NPV by subtracting from the conventional calculation the opportunity cost of exercising the option to invest, and then say that the rule "invest if NPV is positive" holds once this correction has been made. However, to do so is to accept our criticism. To highlight the importance of option values, in this book we prefer to keep them separate from the conventional NPV If others prefer to continue to use "positive NPV" terminology, that is fine as long as they are careful to include all relevant option values in their definition of NPV Readers who prefer that usage can readily translate our statements into that language.

In this book we develop the basic theory of irreversible investment under uncertainty, emphasizing the option-like characteristics of investment opportunities. We show how optimal investment rules can be obtained from methods that have been developed for pricing options in financial markets. We also develop an equivalent approach based on the mathematical theory of optimal sequential decisions under uncertainty-dynamic programming. We illustrate the optimal investment decisions of firms in a variety of situations-new entry, determination of the initial scale of the firm and future costly changes of scale, choice between different forms of investment that offer different degrees of flexibility to meet future conditions, completion of successive stages of a complex multistage project, temporary shutdown and restart, permanent exit, and so forth. We also analyze how the actions of such firms are aggregated to determine the dynamic equilibrium of an industry.

To stress the analogy with options on financial assets, the opportunities to acquire real assets are sometimes called "real options." Therefore this book could be titled "The Real Options Approach to Investment."

3 Irreversibility and the Ability to Wait

Before proceeding, it is important to clarify the notions of irreversibility, the ability to delay an investment, and the option to invest. Most important, what makes an investment expenditure a sunk cost and thus irreversible?

Investment expenditures are sunk costs when they are firm or industry specific. For example, most investments in marketing and advertising are firm specific and cannot be recovered. Hence they are clearly sunk costs. A steel plant, on the other hand, is industry specific-it can only be used to produce steel. One might think that because in principle the plant could be sold to another steel company, the investment expenditure is recoverable and is not a sunk cost. This is incorrect. If the industry is reasonably competitive, the value of the plant will be about the same for all firms in the industry, so there would be little to gain from selling it. For example, if the price of steel falls so that a plant turns out, ex post, to have been a "bad" investment for the firm that built it, it will also be viewed as a bad investment by other steel companies, and the ability to sell the plant will not be worth much. As a result, an investment in a steel plant (or any other industry-specific capital) should be viewed as largely a sunk cost.

Even investments that are not firm or industry specific are often partly irreversible because buyers in markets for used machines, unable to evaluate the quality of an item, will offer a price that corresponds to the average quality in the market. Sellers, who know the quality of the item they are selling, will be reluctant to sell an above-average item. This will lower the market average quality, and therefore the market price. This "lemons" problem (see Akerlof, 1970) plagues many such markets. For example, office equipment, cars, trucks, and computers are not industry specific, and although they can be sold to companies in other industries, their resale value will be well below their purchase cost, even if they are almost new.

Irreversibility can also arise because of government regulations or institutional arrangements. For example, capital controls may make it impossible for foreign (or domestic) investors to sell assets and reallocate their funds, and investments in new workers may be partly irreversible because of high costs of hiring, training, and firing. Hence most major capital investments are in large part irreversible.

Let us turn next to the possibilities for delaying investments. Of course, firms do not always have the opportunity to delay their investments. For example, there can be occasions in which strategic considerations make it imperative for a firm to invest quickly and thereby preempt investment by existing or potential competitors.[3] However, in most cases, delay is at least feasible. There may be a cost to delay-the risk of entry by other firms, or simply foregone cash flows-but this cost must be weighed against the benefits of waiting for new information. Those benefits are often large.

As we said earlier, an irreversible investment opportunity is much like a financial call option. A call option gives the holder the right, for some specified amount of time, to pay an exercise price and in return receive an asset (e.g., a share of stock) that has some value. Exercising the option is irreversible; although the asset can be sold to another investor, one cannot retrieve the option or the money that was paid to exercise it. A firm with an investment opportunity likewise has the option to spend money (the "exercise price"), now or in the future, in return for an asset (e.g., a project) of some value. Again, the asset can be sold to another firm, but the investment is irreversible. As with the financial call option, this option to invest is valuable in part because the future value of the asset obtained by investing is uncertain. If the asset rises in value, the net payoff from investing rises. If it falls in value, the firm need not invest, and will only lose what it spent to obtain the investment opportunity. The models of irreversible investment that will be developed in Chapter 2 and in later chapters will help to clarify the optionlike nature of an investment opportunity.

Finally, one might ask how firms obtain their investment opportunities, that is, options to invest, in the first place. Sometimes investment opportunities result from patents, or ownership of land or natural resources. More generally, they arise from a firm's managerial resources, technological knowledge, reputation, market position, and possible scale, all of which may have been built up over time, and which enable the firm to productively undertake investments that individuals or other firms cannot undertake. Most important, these options to invest are valuable. Indeed, for most firms, a substantial part of their market value is attributable to their options to invest and grow in the future, as opposed to the capital they already have in place.[4] Most of the economic and financial theory of investment has focused on how firms should (and do) exercise their options to invest. To better understand investment behavior it may be just as important to develop better models of how firms obtain investment opportunities, a point that we will return to in later chapters.

---

[3] See Gilbert (1989) and Tirole (1988, Chapter 8) for surveys of the literature on such strategic aspects of investment.

[4] For discussions of growth options as sources of firm value, see Myers (1977), Kester (1984), and Pindyck (1988b).

4 An Overview of the Book

In the rest of this chapter, we outline the plan of the book and give a flavor of some of the important ideas and results that emerge from the analysis.

4.A A Few Introductory Examples

The general ideas about "real options" expounded above are simple and intuitive, but they must be translated into more precise models before their quantitative significance can be assessed and their implications for firms, industries, and public policy can be obtained. Chapter 2 starts this program in a simple and gentle way. We examine a firm with a single discrete investment opportunity that can be implemented within a "window" of two decision periods. Between the two periods, the price of the output undergoes a permanent shift up or down. Suppose the investment would be profitable at the average price, and therefore a fortiori at the higher price, but not at the lower price. By postponing its decision to the second period, the firm can make it having observed the actual price movement. It invests if the price has gone up, but not if it has gone down. Thus it avoids the loss it would have made if it had invested in the first period and then seen the price go down. This value of waiting must be traded off against the loss of the period-1 profit flow. The result-the decision to invest or to wait-depends on the parameters that specify the model, most importantly the extent of the uncertainty (which determines the downside risk avoided by waiting) and the discount rate (which measures the relative importance of the future versus the present).

We carry out several numerical calculations to illustrate these effects and build intuition for real options. We also explore the analogy to financial options more closely. We introduce markets that allow individuals to shift the risk of the price going up or down, namely contingent claims that have different payoffs in the two eventualities. Then we construct a portfolio of these contingent claims that can exactly replicate the risk and return characteristics of the firm's real option to invest. The imputed value of the real option must equal that of the replicating portfolio, because otherwise there would be an arbitrage opportunity-an investor could make a pure profit by buying the cheaper and selling the dearer of the two identical assets.

We also examine some variants of the basic example. First, we expand the window of investment opportunity to three periods, where the price can go up or down between periods 2 and 3 just as it could between periods 1 and 2. We show how this changes the value of the option. Next, we examine uncertainty in the costs of the project and in the interest rate that is used to discount future profit flows. Finally, we consider choice between projects of different scales, a larger project having higher fixed costs but lower operating costs.

Even with these extensions and variations, the analysis remains at the level of an illustrative and very simple example rather than that of a theory with some claim to generality. In later chapters we proceed to develop a broader theoretical framework. But the example does yield some valuable insights that survive the generalization, and we summarize them here.

First, the example shows that the opportunity cost of the option to invest is a significant component of the firm's investment decision. The option value increases with the sunk cost of the investment and with the degree of uncertainty over the future price, the downside component of the risk being the most important aspect. These results are confirmed in more general models in Chapters 5-7.

Second, we will see that the option value is not affected if the firm is able to hedge the risk by trading in forward or futures markets. In efficient markets such risk is fairly priced, so any decrease in risk is offset by the decrease in return. The forward transaction is a financial operation that has no effect on the firm's real decisions. (This is another example of the Modigliani-Miller theorem at work.)

Third, when future costs are uncertain, their effect on the investment decision depends on the particular form of the uncertainty. If the uncertainty pertains to the price the firm must pay for an input, the effect is just like that of output price risk. The freedom not to invest if the input price turns out to have gone up is valuable, so immediate investment is less readily made. However, instead suppose that the project consists of several steps, the uncertainty pertains to the total cost of investment, and information about it will be revealed only as the first few steps of the project are undertaken. Then these steps have information value over and above their contribution to the conventionally calculated NPV Thus it may be desirable to start the project even if orthodox NPV is somewhat negative. We return to this issue in Chapter 10 and model it in a more general theoretical framework.

Fourth, we will see that investment on a smaller scale, by increasing future flexibility, may have a value that offsets to some degree the advantage that a larger investment may enjoy due to economies of scale.

4.B Some Mathematical Tools

In reality, investment projects can have different windows of opportunity, and various aspects of the future can be uncertain in different ways. Therefore the simple two-period examples of Chapter 2 must be generalized greatly before they can be applied. Chapters 3 and 4 develop the mathematical tools that are needed for such a generalization.

Chapter 3 develops more general models of uncertainty. We start by explaining the nature and properties of *stochastic processes*. These processes combine dynamics with uncertainty. In a dynamic model without uncertainty, the current state of a system will determine its future state. When uncertainty is added, the current state determines only the probability distribution of future states, not the actual value. The specification in Chapter 2, where the current price could go either up or down by a fixed percentage with known probabilities, is but the simplest example. We describe two other processes that prove especially useful in the theory of investment-Brownian motion and Poisson processes-and examine some of their properties.

Chapter 4 concerns optimal sequential decisions under uncertainty. We begin with some basic ideas of the general mathematical technique for such optimization: dynamic programming. We introduce this by recapitulating the two-period example of Chapter 2, and showing how the basic ideas extend to more general multiperiod choice problems where the uncertainty takes the form of the kinds of stochastic processes introduced in Chapter 3. We establish the fundamental equation of dynamic programming, and indicate methods of solving it for the applications of special interest here. Then we turn to a market setting, where the risk generated by the stochastic process can be traded by continuous trading of contingent claims. We show how the sequential decisions can be equivalently handled by constructing a dynamic hedging strategy-a portfolio whose composition changes over time to replicate the return and risk characteristics of the real investment.

Readers who are already familiar with these techniques can skip these chapters, except perhaps for a quick glance to get used to our notation. Others not familiar with the techniques can use the chapters as a self-contained introduction to stochastic processes and stochastic dynamic optimization-even if their interests are in applications other than those discussed in the rest of the book.

4.C The Firm's Investment Decision

These techniques are put to use in the chapters that follow. Chapters 5-7 constitute the core theory of a firm's investment decision. We begin in Chapter 5 by supposing that investment is totally irreversible. Then the value of the project in place is simply the expected present value of the stream of profits (or losses) it would generate. This can be computed in terms of the underlying uncertainty. Then the investment decision is simply the decision to pay the sunk cost and in return get an asset whose value can fluctuate. This is exactly analogous to the financial theory of call options-the right but not the obligation to purchase an asset of fluctuating value for a preset exercise price. Therefore the problem can be solved directly using the techniques developed in Chapter 4. The result is also familiar from financial theory. The option can be profitably exercised-is "in the money" when the value of the asset rises above the exercise price. However, exercise is not optimal when the option is only just in the money, because by exercising it the firm gives up the opportunity to wait and avoid the loss it would suffer should the value fall. Only when the value of the asset rises sufficiently above the exercise price-the option is sufficiently "deep in the money"-does its exercise become optimal.

An alternative formulation of this idea would help the intuition of economists who think of investment in terms of Tobin's q, the ratio of the value of a capital asset to its replacement cost. In its usual interpretation in the investment literature, the value of a capital asset is measured as the expected present value of the profit flow it will generate. Then the conventional criterion for the firm is to invest when q equals or exceeds unity. Our option value criterion is more stringent; q must exceed unity by a sufficient margin. It must equal or exceed a critical or threshold value q^*, which itself exceeds unity, before investment becomes optimal.

We perform several numerical simulations to calculate the value of the option and the optimal exercise rule, and examine how these vary with the amount of the uncertainty, the discount rate, and other parameters. We find that for plausible ranges of these parameters, the option value effect is quantitatively very important. Waiting remains optimal even though the expected rate of return on immediate investment is substantially above the interest rate or the "normal" rate of return on capital. Return multiples of as much as two or three times the normal rate are typically needed before the firm will exercise its option and make the investment.

4.D Interest Rates and Investment

Once we understand why and how firms should be cautious when deciding whether to exercise their investment options, we can also understand why interest rates seem to have so little effect on investment. Econometric tests of the orthodox theory generally find that interest rates are only a weak or insignificant determinant of investment demand. Recent history also shows that interest rate cuts tend to have only a limited stimulative effect on investment; the experience of 1991-1992 bears the latest witness to that. The options approach offers a simple explanation. A reduction in the interest rate makes the future generally more important relative to the present, but this increases the value of investing (the expected present value of the stream of profits) and the value of waiting (the ability to reduce or avoid the prospect of future losses) alike. The net effect is weak and sometimes even ambiguous.

The real options approach also suggests that various sources of uncertainty about future profits-fluctuations in product prices, input costs, exchange rates, tax and regulatory policies-have much more important effects on investment than does the overall level of interest rates. Uncertainty about the future path of interest rates may also affect investment more than the general level of the rates. Reduction or elimination of unnecessary uncertainty may be the best kind of public policy to stimulate investment. And the uncertainty generated by the very process of a lengthy policy debate on alternatives may be a serious deterrent to investment. Later, in Chapter 9, we construct specific examples that show how policy uncertainty can have a major negative effect on investment.

4.E Suspension and Abandonment

Chapters 6 and 7 extend the simple model. As future prices and/or costs fluctuate, the operating profit of a project in place may turn negative. In Chapter 5 we assumed that the investment was totally irreversible, and the firm was compelled to go on operating the project despite losses. This may be true of some public services, but most firms have some escape routes available. Chapters 6 and 7 examine some of these. In Chapter 6 we suppose that a loss-making project may be temporarily suspended, and its operation resumed later if and when it becomes profitable again. Now a project in place is a sequence of operating options; its value must be found by using the methods of Chapter 4 to value all these operating options, and then discount and add them. Then the investment opportunity is itself an option to acquire this compound asset.

In Chapter 7 we begin by ruling out temporary suspension, but we allow permanent abandonment. This is realistic if a live project has some tangible or intangible capital that disappears quickly if the project is not kept in operation-mines flood, machines rust, teams of skilled workers disband, and brand recognition is lost. If the firm decides to restart, it has to reinvest in all these assets. Abandonment may have a direct cost; for example, workers may have to be given severance payments. More importantly, however, it often has an opportunity cost-the loss of the option to preserve the capital so it can be used profitably should future circumstances improve. Therefore a firm with a project in place will tolerate some losses to keep this option alive, and only sufficiently extreme losses will induce it to abandon.

In fact we have here an interlinked pair of options. When the firm exercises its option to invest, it gets a project in place and an option to abandon If it exercises the option to abandon, it gets the option to invest again. The two options must be priced simultaneously to determine the optimal investment and abandonment policies. The linkage has important implications; for example, a higher cost of abandonment makes the firm even more cautious about investing, and vice versa. We illustrate the theory using numbers that are typical of the copper industry, and find a very wide range of fluctuation of the price of copper between the investment and abandonment thresholds.

We also consider an intermediate situation, where both suspension and abandonment are available at different costs. A project in suspension requires some ongoing expenditure, for example, keeping a ship laid up, but restarting is cheap. Abandonment saves on the maintenance cost, and may even bring some immediate scrap value, but then the full investment cost must be incurred again if profit potential recovers. Now we must determine the optimal switches between three alternatives-an idle firm, an operating project, and a suspended project. We do this, and illustrate it for the case of crude oil tankers.

4.F Temporary versus Permanent Employment

While most of our attention in this book is on firms' capital investment choices, similar considerations apply to their hiring and firing decisions. Each of these choices entails sunk costs, each decision must be made in an uncertain environment, and each allows some freedom of timing. Therefore the above ideas and results can be applied. For example, a new worker will not be hired until the value of the marginal product of labor is sufficiently above the wage rate, and the required margin or multiple above the wage will be higher the greater the sunk costs or the greater the uncertainty.

The U.S. labor market in mid-1993 offered a vivid illustration of this theory in action. As the economy emerged from the recession of the early 1990s, firms increased production by using overtime and hiring temporary workers even for highly skilled positions. But permanent new full-time hiring was very slow to increase. The current level of profitability must have been high, since firms were willing to pay wage premiums of 50 percent for overtime work and to use agencies that supplied temporary workers and charged fees of 25 percent or more of the wage. But these same firms were not willing to make the commitment involved in hiring regular new workers.[5] Our theory offers a natural explanation for these observations. A high level of uncertainty about future demand and costs prevailed at that time. The robustness and durability of the recovery were unclear; it was feared that inflation would return and lead the Federal Reserve to raise interest rates. Future tax policy was very uncertain, as was the level of future health care costs that employers would have to bear. Therefore we should have expected firms to be very cautious, and wait for greater assurance of a continued prospect of high levels of profitability before adding to their regular full-time labor force. In the meantime, they would prefer to exploit the current profit opportunities using less irreversible (even if more costly) methods of production, namely overtime and temporary work. That is exactly what we saw.

4.G Hysteresis

When we consider investment and abandonment (or entry and exit) together, a firm's optimal decision is characterized by two thresholds. A sufficiently high current level of profit, corresponding to an above-normal rate of return on the sunk cost, justifies investment or entry, while a sufficiently large level of current loss leads to abandonment or exit. Now suppose the current level of profit is somewhere between these two thresholds. Will we see an active firm? That depends on the recent history of profit fluctuations. If the profit is at its current intermediate level having most recently descended from a high level that induced entry, then there will be an active firm. However, if the intermediate level was most recently preceded by a low level that induced exit, then there will not. In other words, the current state of the underlying stochastic variable is not enough to determine the outcome in the economy; a longer history is needed. The economy is *path dependent*.

The idea of path dependence has been recently explored and illustrated, most prominently by Arthur (1986) and David (1985, 1988). They allow an even more extreme possibility: even very long-run properties of their systems are altered by slight differences in initial conditions. Here we have a more moderate kind of path dependence. The long-run distribution of the possible states of the economy is unaltered, but the short- and medium-run evolution can still be dramatically affected by initial conditions.

[5] Two articles describing these developments appeared in the New York Times on the same day, Sunday, May 16, 1993: "Fewer Jobs Filled as Factories Rely on Overtime Pay," and "In a Shaky Economy, Even Professionals Are Temps."

The path dependence can lead to the following kind of sequence o events. When the firm first arrives on the scene and contemplates investment the current profit is in the intermediate range between the two thresholds Therefore the firm decides to wait. Then profit rises past the upper threshold so the firm invests. Finally, profit falls back to its old intermediate level, but that does not take it down to the lower threshold where abandonment would occur. Thus the underlying cause (current profitability) has been restored to its old level, but its effect (investment) has not.

Similar effects have long been known in physics and other sciences. The most familiar example comes from electromagnetism. Take an iron bar an(loop an insulated wire around it. Pass an electric current through the wire the iron will become magnetized. Now switch the current off. The magnetism is not completely lost; some residual effect remains. The cause (the current was temporary, but it leaves a longer-lasting effect (the magnetized bar).

This phenomenon is called *hysteresis*, and by analogy the failure of in vestment decisions to reverse themselves when the underlying causes art fully reversed can be called economic hysteresis. A striking example occurred during the 1980s. From 1980 to 1984, the dollar rose sharply against other currencies. The cost advantage of foreign firms in US markets became verb substantial, and ultimately led to a large rise in US imports. Then the dollar fell sharply, and by 1987 was back to its 1980 level. However, the import penetration was not fully reversed; in fact it hardly decreased at all. It took larger fall in the dollar to achieve any significant reduction in imports.

4.H Industry Equilibrium

In Chapters 8 and 9 the focus turns from a firm's investment decisions to the equilibrium of a whole industry composed of many such firms. One' first reaction might be that the competition among firms will destroy any one firm's option to wait, eliminating the effects of irreversibility and uncertainty that we found in Chapters 5-7. Competition does destroy each firm's option to wait, but this does not restore the present value approach and results o: the orthodox theory. On the contrary, caution when making an irreversible decision remains important, but for somewhat different reasons.

Consider one firm contemplating its investment, knowing that the future path of industry demand and its own costs are uncertain, and knowing that there are many other firms facing similar decisions with similar uncertainty The firm is ultimately concerned with the consequences of its decision for it; own profits, but it must recognize how the similar decisions of other firms will affect it. In this respect, two types of uncertainty must be distinguished because they can have different implications for investment: aggregate uncertainty that affects all firms in the industry, and firm-specific or idiosyncratic uncertainty facing each firm.

To see this, first suppose that investment is totally irreversible, and consider an industry-wide increase in demand. Any one firm expects this to lead to a higher price, and so improve its own profit prospects, making investment more attractive. However, it also knows that several other firms are making a similar calculation. Their supply response will dampen the effect of the demand shift on the industry price. Therefore the upward shift of its own profit potential will not be quite as high as in the case where it is the only firm and has a monopoly on the investment opportunity. However, with investment being irreversible, a downward shift of industry demand has just as unfavorable an effect in the competitive case as in the monopoly case. Even though other competitive firms are just as badly affected, they cannot exit to cushion the fall in price. Thus the competitive response to uncertainty has an inherent asymmetry: the downside exerts a more potent effect than the upside. This asymmetry makes each firm cautious in making the irreversible investment. The ultimate effect is very similar, and in some models identical, to that of the option value for a firm possessing a monopoly on the investment opportunity and waiting to exercise it. In fact, the theory of a competitive industry can be formulated by giving each firm the option to invest, valuing these real options as in Chapters 5-7, and finally imposing the condition that in the competitive equilibrium this option value should be zero.

If we allow some reversibility, the exit of other firms does cushion the effect of adverse demand shocks on price. But then each firm's exit decision recognizes this asymmetric effect of demand shocks in an initially poor situation: their upside effect is more potent than the downside. Thus competitive firms are not quick to leave when they start to make operating losses; they wait a while to see if things improve or if their rivals leave. The overall effect is just like that for a single monopoly firm's abandonment decision that we found in Chapter 7. In fact, the competitive equilibrium model of joint entry and exit decisions with aggregate demand shocks that we analyze in Chapter 8 has exactly the same critical levels of high and low prices to trigger entry and exit as the corresponding monopoly model of Chapter 7.

Firm-specific uncertainty does not lead to this kind of asymmetry. If just one firm experiences a favorable shift of its demand, say some idiosyncratic switch of fashion, then it knows that this good fortune is not systematically shared by other firms, and therefore does not fear that entry of other firms will erode its profit potential in the same way. However, then the value of waiting reemerges in the older familiar form. The lucky firm does have a monopoly on the opportunity to enter with its low cost. Therefore it also has an option value of waiting; it can thereby avoid a loss if its low cost should turn out to bi transitory. Thus firm-specific uncertainty in industry equilibrium also leads tc investment decisions similar to those found in Chapters 5-7 for an isolated firm.

4.I Policy Toward Investment

Some readers might interpret the result that uncertainty makes firms less eager to invest as indicating a need for government policy intervention to stimulate investment. That would be a hasty reaction. A social planner also gets information by waiting, and therefore should also recognize the opportunity cost of sinking resources into a project. A case for policy intervention will arise only if firms face a different value of waiting than does society as , whole, in other words, if some market failure is associated with the decision process.

Chapter 9 focuses on these issues. Our first result is a confirmation o the standard theory of general equilibrium. If markets for risk are complete and if firms behave as competitive price- takers (in this stochastic dynamic context this must be interpreted to mean that each firm takes as given the stochastic process of the price and has rational expectations about it), then the equilibrium evolution of the industry is socially efficient. A social planner would show the same degree of hesitancy in making the investment decision

If markets for risk are incomplete, beneficial policy interventions do exist, but the correct policy needs some careful calculation and implementation The blunt tools that are often used for handling uncertainty can have adverse effects. We illustrate this by examining the consequences of price ceiling; and floors. For example, price supports promote investment by reducing the downside risk. However, the resulting rightward shift of the industry supply function implies lower prices in good times. Averaging over good times and bad, we find that the overall result is a lower long-run average price. In other words, the policy is harming the very group it sets out to help. Price floor of ceiling policies, for example, urban rent controls and agricultural price sup. ports, are usually criticized because they reduce overall economic efficiency Our finding is perhaps a politically more potent argument against them: their distributional effect can be perverse, too.

We also study the effect of uncertainty concerning future policy itself For example, if an investment tax credit is being discussed, firms will recognize more value in waiting, because there is a probability that the cost of investment to the firm will fall. We find that such policy uncertainty can have a powerful deterrent effect on immediate investment. If governments wish to stimulate investment, perhaps the worst thing they can do is to spend a long time discussing the right way to do so.

4.J Antitrust and Trade Policies

Chapters 8 and 9 paint a very different picture of competitive equilibrium than the one familiar from intermediate microeconomics textbooks. There we are told that firms will enter the industry if the price rises to equal the long-run average cost, and they will exit if the price falls as low as the average variable cost. Our theory implies a wider range of price variation on either side. For example, in the face of aggregate uncertainty, firms' entry as soon as the price rises to the long-run average cost will not constitute an industry equilibrium. Each firm knows that entry of other similar firms will stop the price from ever rising any higher, while future unfavorable shifts can push the price below this level. Also, a future price path that sometimes touches the long-run average cost and otherwise stays below this level can never offer a normal return on the firm's investment. Only if the price ceiling imposed by entry is strictly above the long-run average cost can the mix of intervals of supernormal profit and ones of subnormal profit average out to a normal return. Similarly, firms will exit only when the price falls sufficiently far below the average variable cost. They will tolerate some losses, knowing that the exit of other firms puts a lower bound on the price. The equilibrium level of this floor is determined by averaging out the prospects of future losses and profits to zero.

Thus we find that competitive equilibrium under uncertainty is not a stationary state even in the long run, but a dynamic process where prices can fluctuate quite widely. Periods of supernormal profits can alternate with periods of losses. A similar view of dynamic equilibrium as a stochastic process has become quite common in macroeconomics, but is surprisingly uncommon in microeconomics, particularly with regard to its implications for antitrust policy or international trade policy. The conceptual framework of such policies is generally static, and the recommendations in practice are based on observations of "snapshots" of an industry at a particular instant. We find that the dynamic view calls for a substantial rethinking of both the theory and the practice.

For example, in industrial organization theory, excess profits suggest collusion or entry barriers, calling for antitrust action. In our dynamic perspective, substantial periods of supernormal profits without new entry can occur even though all firms are small price takers. In international trade, when foreign firms continue their export operations at a loss, domestic firms allege predatory dumping and call for the standard trade policy response of countervailing import duties. However, our analysis suggests that the foreign firms may be simply and rationally keeping alive their option of future operation in our market, with no predatory intent whatever. Only a sufficiently long time series of data will allow us to test whether the supposed collusive or predator actions are merely natural phases in the evolution of a competitive industry or genuine failures of competition.

4.K Sequential and Incremental Investment

In Chapters 10 and 11 we return to a single firm's investment decision am examine some other aspects of it that are important in applications. Chapter 10 deals with investments that consist of several stages, all of which must be completed in sequence before any output or profit flow can commence The firm can constantly observe some indicator of the future profit potential, and this fluctuates stochastically. At any stage, the firm may decide to continue immediately, or wait for conditions to improve. At an early stage o the investment sequence, most of the cost remains to be sunk. Therefore the firm will go ahead with the program only if it sees a sufficiently high threshold level of the indicator of profitability. Gradually, as more steps are complete and less cost remains to be sunk, the next step is justified by an ever smaller threshold. In this sense, bygones affect the decisions to come.

In Chapter 10 we examine another effect of current decisions on the future, namely the learning curve. According to this theory, the cost of production at any instant is a decreasing function of cumulated output experience. Thus the current output flow contributes to a reduction in all future production costs. This additional value must be added to the current revenue before comparing them to the current costs of production to decide on the optimal level of production. We examine the dynamic output path under then conditions. We find that greater uncertainty lowers the value of future cost reductions, and that leads to a reduction in the pace of investment.

In Chapter 11 we turn to the study of incremental investment, where out put and profit flow are available all the time as a function of the installed stock of capital. The aim is to characterize the optimal policy for capacity expansion When production shows diminishing returns to capital, we can regard each new unit of capacity as a fresh project, which begins to contribute its marginal product from the date of its installation. Then the criterion derived in Chapter 6 for investment in such a project continues to apply. If production show increasing returns to capital over an interval, then all the units of capacity in an appropriately constructed range must be regarded as a single project, and the criterion for its installation is again a natural generalization of that for a single project described in Chapter 6.

When the firm can choose its rate of capacity expansion, we must specify how the costs of this expansion depend on its volume and pace. Different assumptions in this respect imply different optimal policies. We construct a general model that places the alternatives in context, and in particular shows the relationship between the adjustment costs models that have been the mainstay of theoretical and empirical work over the last decade and the irreversibility approach that has been the focus of our book.

4.L Empirical and Applied Research

In Chapter 12 we turn to some examples that illustrate applications and extensions of the techniques developed throughout the book. We also discuss the relevance of the theory for empirical work on investment behavior.

We begin Chapter 12 with a problem of great interest to oil companies - how to value an undeveloped offshore oil reserve, and how to decide when to invest in development and production. As we will see, an undeveloped reserve is essentially an option; it gives the owner the right to invest in development of the reserve and then produce the oil. By valuing this option we can value the reserve and determine when it should be developed. Oil companies regularly spend hundreds of millions of dollars for offshore reserves, so it is clearly important to determine how to value and best exploit them.

We then turn to an investment timing problem in the electric utility industry. The Clean Air Act calls for reductions in overall emissions of sulfur dioxide, but to minimize the cost of these reductions, it gives utilities a choice. They can invest in expensive "scrubbers" to reduce emissions to mandated levels, or they can buy tradeable "allowances" that let them pollute. There is considerable uncertainty over the future prices of allowances, and an investment in scrubbers is irreversible. The utility must decide whether to maintain flexibility by relying on allowances or invest in scrubbers. We show how this problem can be addressed using the options approach of this book.

To show how the principles and tools developed in this book have relevance beyond firms' investment decisions, we address a problem in public policy-when should the government adopt a policy in response to a threat to the environment, given that the future costs and benefits of the policy are uncertain? We will argue that the standard cost-benefit framework that economists have traditionally used to evaluate environmental policies is deficient. The reason is that there are usually important irreversibilities associated with environmental policy. These irreversibilities can arise with respect to environmental damage itself, but also with respect to the costs of adapting to policies to reduce the damage. Since the adoption of an environmental policy is rarely a now or never proposition, the same techniques used to study the optimal timing of an investment can be applied to the optimal timing of an environmental policy.

At the end of Chapter 12, we discuss some of the empirical implications of irreversibility and uncertainty for investment behavior. There is considerable anecdotal evidence that firms make investment decisions in a way that is at least roughly consistent with the theory developed in this book, for example, the use of hurdle rates that are much larger than the opportunity cost of capital as predicted by the capital asset pricing model. Some numerical simulations based on plausible parameters for costs and uncertainty are found to replicate features of reality. However, more systematic econometric testing of the theory is still at a very early stage. We review some work of this kind, point out its difficulties, and suggest ideas for future research.

5 Noneconomic Applications

Our focus in the book will be on investment decisions of firms and their implications for industry equilibrium. These matters are of primary interest to economists, their conditions (technology and resource availability) and criteria (maximization of the value of the firm) are generally agreed upon, and theory leads to quantifiable predictions about them. Investment decisions of consumers (purchase of durables) and workers (education and human capital) have obvious parallels. Research on these issues already exists, and we will sketch some of this literature in Chapter 12.

Many other personal and societal choices are made under the same basic conditions, namely irreversibility, ongoing uncertainty, and some leeway it timing. Therefore we can think about them in terms of some of the model and results outlined above, and offer some qualitative speculations. We believe there is scope for more serious research along these lines. This would amount to extending work such as that of Becker (1975, 1980) on economic approaches to social phenomena by incorporation of option values. Indeed some recognition of this aspect can be read into Becker's own remarks or the subject, for example, in Becker (1962, pp. 22-23), but a more thorough and formal analysis along these lines remains a potentially fruitful project for the future. Although we cannot spare much space in an already length; book, we believe that readers will find some speculations along these lines both interesting and thought-provoking.

5.A Marriage and Suicide

Marriage entails significant costs of courtship, and divorce has its own monetary and emotional costs. Happiness or misery within the marriage can be only imperfectly forecast in advance, and continues to fluctuate stochastically even after the event. Therefore waiting for a better match has an option value, and we should expect prospective partners to look for a sufficiently high initial threshold of compatibility to justify getting married. We should expect the option value to be higher in religions or cultures where marriage is less reversible. Therefore we should expect that, other things being equal, individuals in such societies will search more carefully (and on the average, search longer), and will insist on a higher threshold of the quality of the match. On the contrary, when divorce is easy, we should see couples entering into marriage (or equivalent arrangements) more readily.

Of course, other things are not equal. Societies that make divorce more difficult presumably attach higher value to marriage. Therefore we should expect them to counteract the rational search and delay on the part of individuals by social pressure, and provision of better "matchmaking technologies." Empirical tests of these ideas will have to be designed carefully to distinguish the separate and often opposing influences, but that only raises the interest and challenge of such research.

Perhaps the most extreme example of economics applied to sociological phenomena is the Beckerian theory of suicide developed by Hamermesh and Soss (1974). According to them, an individual will end his or her own life when the expected present value of the utility of the rest of life falls short of a benchmark or cutoff standard "zero." Most people react by saying that the model gives an excessively rational view of what is an inherently irrational action. Our theory suggests exactly the opposite. Whatever its merits or demerits as descriptive theory, the Hamermesh-Soss model is *not rational enough* from the prescriptive viewpoint, because it forgets the option value of staying alive. Suicide is the ultimate irreversible act, and the future has a lot of ongoing uncertainty. Therefore the option value of waiting to see if "something will improve" should be very large. The circumstances must be far more bleak than the cutoff standard of the Hamermesh-Soss model to justify pulling the trigger. This is true even if the expected direction of life is still downward; all that is needed is some positive probability on the upside.

Now return to the argument that most suicides are irrational, and ask exactly how they fail to be rational. There are several possibilities, but one seems especially pertinent. Suicides project the bleak present into an equally bleak future, ignoring uncertainty, and thereby ignoring the option value of life. Then religious or social proscriptions against suicide serve a useful func tion as measures to compensate for this failure of rationality. By raising thl perceived cost of the act, these taboos lower the threshold of quality of lift that leads to suicide when option value is ignored. This can correct the failure of the individuals' forethought, and bring their threshold in conformity wit] the optimal rule that recognizes the option value.

5.B Legal Reform and Constitutions

Finally, consider law reform. Different fundamental constitutional or legal principles often conflict with one another in specific contexts; for example the civil rights of the accused often stand in the way of better enforcement of law and order. Public opinion as to the appropriate relative weight to bi attached to these conflicting principles can shift over time. How quickly should the law respond to the latest swing of opinion? Given some legislative an(administrative costs of changing laws, our theory suggests that the option to wait and see if the trend of opinion will reverse itself has some value Reform should be delayed until the force of current opinion is sufficiently, overwhelming to offset this value. But there is good reason to believe that the process is subject to a kind of myopia or "political market failure." A all times, people and politicians tend to believe that at last they have got it right; the balance of opinion that has just been reached will prevail forever Therefore they will ignore future uncertainty and option values, and change the laws too frequently.

The time to anticipate and defeat this tendency is when the constitution is framed. Knowing that future generations will be too trigger-happy in changing laws, the founding fathers can artificially raise the cost of change, thereby making the politically flawed thresholds for legal change coincide with the truly optimal thresholds. Thus various supermajority requirements for constitutional change can be seen as a commitment device that corrects for the shortsightedness of future generations.

These are just a few examples of how the theory and techniques developed in this book are applicable to some fairly far-reaching problems. We have deliberately stated them in speculative and provocative ways, hoping to attract interest and research along these lines from a broader spectrum o social scientists. In the rest of the book, we will focus largely on investment but ask the reader to keep in mind the much broader potential applicability, of the theory.

Chapter 2

Developing the Concepts Through Simple Examples

Firms make, implement, and sometimes revise their investment decisions continuously through time. Hence much of this book is devoted to the analysis of investment decisions as continuous-time problems. However, it is best to begin with some simple examples, involving a minimal amount of mathematics, in which investment decisions are made at two or three discrete points in time. In this way, we can convey at the outset an intuitive understanding of the basic concepts. In particular, we want to show how the irreversibility of an investment expenditure affects the decision to invest, and how it requires modification of the standard net present value (NPV) rule that is commonly taught in business schools. We also want to show how an opportunity to invest is much like a financial option, and can be valued and analyzed according.

This chapter begins by discussing these ideas in the context of a simple example in which an investment decision can be made in only one of two possible periods-now or next year. We will use this example to see how irreversibility creates an opportunity cost of investing when the future value of the project is uncertain, and how this opportunity cost can be accounted for when making the investment decision. We will also see how the investment decision can be analyzed using basic option pricing techniques. We will examine the characteristics of the firm's option to invest in some detail, and see how the value of that option, and the investment decision, depend on the degree of uncertainty over the future value of the project. Extending the example to three periods will provide more insight into the problem of investment timing, and set the stage for modeling investment problems in continuous time. Finally, we wil briefly examine investment decisions when the cost of the investment (as op posed to its payoff) is uncertain, when future interest rates are uncertain, an(when one must choose the scale of an investment.

1 Price Uncertainty Lasting Two Periods

Consider a firm that is trying to decide whether to invest in a widget factory The investment is completely irreversible - the factory can only be used to make widgets, and should the market for widgets evaporate, the firm cannot "uninvest" and recover its expenditure. To keep matters as simple as possible we will assume that the factory can be built instantly, at a cost I, and will produce one widget per year forever, with zero operating cost. Currently the price of a widget is $200, but next year the price will change. With probability q, it will rise to $300, and with probability (1 - q), it will fall to $100. The price will then remain at this new level forever. (See Figure 2.1.)

Again, to keep things simple, we will assume that the risk over the future price of widgets is fully diversifiable (that is, it is unrelated to what happen with the overall economy). Hence the firm should discount future cash flows using the risk-free rate of interest, which we will take to be 10 percent.

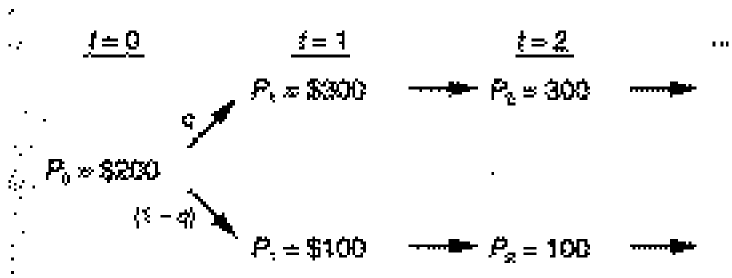

Figure 2.1. Price of Widgets

For the time being we will set I = $1600 and q = 0.5. (Later we will sec how the investment decision depends on I and q.) Given these values for I and q, is this a good investment? Should the firm invest now, or would it be better to wait a year and see whether the price of widgets goes up or down? Suppose we invest now. Calculating the net present value of this investment in the standard way (and noting that the expected future price of widgets is always $200), we get

$$NPV = -1600 + \sum_{t=0}^{\infty} \frac{200}{(1.1)^t} = -1600 + 2200 = \$600$$

It appears that the NPV of this project is positive. The current value of a widget factory, which we will denote by Vo, is equal to $2200, which exceeds the $1600 cost of the factory. Hence it would seem that we should go ahead with the investment.

This conclusion is incorrect, however, because the calculations above ignore a cost - the opportunity cost of investing now, rather than waiting and keeping open the possibility of not investing should the price fall. To see this, let us calculate the NPV of this project a second time, this time assuming that instead of investing now, we will wait one year and then invest only if the price of widgets goes up. (Investing only if the price goes up is in fact ex post optimal.) In this case the NPV is given by[1]

$$NPV = (0.5)\left[\frac{-1600}{1.1} + \sum_{t=1}^{\infty} \frac{300}{(1.1)^t}\right] = \frac{850}{1.1} = \$773\,. \qquad (2)$$

[1] In all of the calculations below, we have rounded the results to the nearest dollar.

(Note that in year 0, there is no expenditure and no revenue. In year 1, the $1600 is spent only if the price rises to $300, which will happen with probability 0.5.) If we wait a year before deciding whether to invest in the factory, the project's NPV today is $773, whereas it is only $600 if we invest in the factory now. Clearly it is better to wait than to invest right away.

Note that if our only choices were to invest today or never invest, we would invest today. In that case there is no option to wait a year, and hence no opportunity cost to killing such an option, so the standard NPV rule applies. We would likewise invest today if next year we could disinvest and recover our $1600 should the price of widgets fall. Two things are needed to introduce an opportunity cost into the NPV calculation-irreversibility, and the ability to invest in the future as an alternative to investing today. There are, of course, situations in which a firm cannot wait, or cannot wait very long, to invest. (One example is the anticipated entry of a competitor into a market that is only large enough for one firm. Another example is a patent or mineral resource lease that is about to expire.) The less time there is to delay, and the greater the cost of delaying, the less will irreversibility affect the investment decision. We will explore this point later in this book as we develop more general models of investment.

How much is it worth to have the flexibility to make the investment decision next year, rather than having to invest either now or never? (We know that having this flexibility is of some value, because we would prefer to wait rather than invest now.) The value of this "flexibility option" is easy to calculate; it is just the difference between the two NPV's, that is, $773 $600 =$173. In other words, we should be willing to pay $173 more for a investment opportunity that is flexible than one that only allows us to invest now.

Another way to look at the value of flexibility is to ask the following question: How high an investment cost *I* would we be willing to accept t have a flexible investment opportunity rather than an inflexible "now or never one? To answer this, we find the value of *I*, which we denote by $\bar{I}$, that makes the NPV of the project when we wait equal to the NPV when *I* = $1600 and we invest now, that is, equal to $600. Substituting $\bar{I}$ for the $1600 and substituting $600 for the $773 in equation (2):

$$NPV = (0.5)\left[\frac{-\bar{I}}{1.1} + \sum_{t=1}^{\infty} \frac{300}{(1.1)^t}\right] = \$600. \qquad (3)$$

Solving this for $\bar{I}$ yields $\bar{I}$ = $1980. In other words, the opportunity to build widget factory *now and only now* at a cost of $1600 has the same value as a opportunity to build the widget factory *now or next year* at a cost of $1980.

Finally, suppose there exists a futures market for widgets, with the fr tures price for delivery one year from now equal to the expected future spot price, that is, $200.[2] Would the ability to hedge on the futures market chang our investment decision? Specifically,

[2] In this example, the futures price would equal the expected future price because we assumed that the risk is fully diversifiable. (If the price of widgets were positively correlated wit the market portfolio, the futures price would be less than the expected future spot price.) Nor that if widgets were storable and aggregate storage is positive, the marginal convenience yield from holding inventory would have to be 10 percent. The reason is that since the futures price equals the current spot price, the net holding cost (the interest cost of 10 percent less the margin convenience yield) must be zero.

would it lead us to invest now, rathe than waiting a year? The answer is no. To see this, consider investing now an hedging the price risk with futures contracts. To hedge all price risk, we woul need to sell short futures for 11 widgets; this would exactly offset any fluctuE dons in the NPV of our project next year. (If the price of widgets rises to $30(our project would be worth $3300, but we would lose $1100 on the future contract. If the price fell to $100, our project would be worth only $1100, bi we would earn an additional $1100 from the futures. Either way, we end up with a net project value of $2200.) But this would also mean that the NPV of our project today is $600 ($2200 less the $1600 cost of the investment), exactly what it is without hedging. [3]

Hence there is no gain from hedging, and we are still better off waiting until next year to make our investment decision. This result is a variant of the Modigliani-Miller (1958) theorem. Operating in the futures market is only a form of financial policy, and barring the possibility of bankruptcy, it has no real consequences for investment decisions or for the value of the firm's investment opportunities.

1.A Analogy to Financial Options

Our investment opportunity is analogous to a call option on a common stock. It gives us the right (which we need not exercise) to make an investment expenditure (the exercise price of the option) and receive a project (a share of stock) the value of which fluctuates stochastically. In the case of our simple example, we have an option that is "*in the money,*" meaning that if it were exercised today it would yield a positive net payoff. (An option is said to be "*out of the money*" if exercising it today yields a negative net payoff.) We found that even though the option is "in the money," it is better to wait rather than exercise it now. Next year if the price rises to $300, we will exercise our option by paying $1600 and receive an asset which will be worth V_1 = $3300 = $\sum_0^{\infty} 300/(1.1)^t$. If the price falls to $100, this asset will be worth only $1100, and so we will not exercise the option.

We found that the value of our investment opportunity (assuming that the actual decision to invest can indeed be made next year) is $773. It will be helpful to recalculate this value using standard option pricing methods, because later we will use such methods to analyze other investment problems.

To do this, let F_0 denote the value today of the investment opportunity, that is, what we should be willing to pay today to have the option to invest in the widget factory, and let F_1 denote the value of this investment opportunity next year. Note that F_1 is a random variable; it depends on what happens to the price of widgets. If the price rises to $300, then F_1 will equal $\sum_0^{\infty} 300/(1.1)^t - 1600 = \$1700$. On the other hand, if the price falls to $100, the option to invest will go unexercised, and in that case F_1 will equal 0. Thus we know all possible values for F_1. The problem is to find F_0, the value of the option today.[4]

---

[3] Most futures markets only apply to horizons of a year or so. If there were a futures market for widgets over the indefinite future, an equivalent hedge would be to sell short one widget in each future year, that is, to sell forward the entire output stream. The result would be the same.

[4] In this example, all uncertainty is resolved next year. Therefore an option to wait has no value next year, and investment is made according to the conventional NPV criterion at that time. Waiting has relevance only this year. In later chapters we will consider more general situations where uncertainty is never fully resolved, and the option to wait retains value at all times.

To solve this problem, we will create a portfolio that has two component the investment opportunity itself, and a certain number of widgets. We will pick this number of widgets so that the portfolio is risk-free, that is, so that, its value next year is independent of whether the price of widgets goes up c down. Since the portfolio will be risk-free, we know that the rate of return on can earn from holding it must be the risk-free rate of interest. (If the return on the portfolio were higher than the risk-free rate, arbitrageurs could ear unlimited amounts of money by borrowing at the risk-free rate and buying the portfolio. If the portfolio's return were less than the risk-free rate, arbitrageurs could earn money by selling short the portfolio and investing the funds at the risk-free rate.) By setting the portfolio's return equal to the risk-free rate, we will be able to calculate the current value of the investment opportunity.

Specifically, consider a portfolio in which one holds the investment opportunity, and sells short n widgets. (If widgets were a traded commodity, such as oil, one could obtain a short position by borrowing from another producer, or by going short in the futures market. For the moment, however, we nee not be concerned with the actual implementation of this portfolio.) The value of this portfolio today is $\_{0} = F_0 - n P_0 = F_0 - 200n$. The value of this portfolio next year is $\_{1} = F_1 - n P_1$. This depends on P_1. If P_1 turns out to be \$300, so that $F_1 = 1700$, then $\_{1} = 1700 - 300n$. If P_1 turns out to be \$100, so that $F_1 = 0$, then $\_{1} = -100n$. Now, let us choose n so that the portfolio is risk-free, that is, so that $\_{1}$ is independent of what happens to the price. To do this, just set

$$1700 - 300n = -100n,$$

or, $n = 8.5$. With n chosen in this way, $\_{1} = -850$, whether the price of widgets rises to \$300 or falls to \$100.

Now let us calculate the return from holding this portfolio. That return is the capital gain, $\_{1} - \_{0}$, minus any payments that must be made to hold the short position. Since the expected rate of capital gain on a widget is zero (the expected price next year is \$200, the same as this year's price), no rational investor would hold a long position unless he or she could expect to earn at least 10 percent. Hence selling widgets short will require a payment of 0.1 Po = \$20 per widget per year. (This is analogous to selling short a dividend-paying stock; the short position requires payment of the dividend, because no rational investor will hold the offsetting long position without receiving that dividend.) Our portfolio has a short position of 8.5 widgets, so it will have to pay out a total of \$170. The return from holding this portfolio over the year is therefore

$$\begin{aligned} \_{1} - \_{0} - 170 &= \_{1} - (F_0 - n P_O) - 170 \\ &= -850 - F_0 + 1700 - 170 \\ &= 680 - F_{0.} \end{aligned}$$

Because this return is risk-free, we know that it must equal the risk-free rate, which we have assumed is 10 percent, times the initial value of the portfolio, $\_{0} = F_0 - n P_0$:

$$680 - F_0 = 0.1 (F_0 - 1700).$$

We can thus determine that F_0 = \$773. Note that this is the same value that we obtained before by calculating the NPV of the investment opportunity under the assumption that we follow the optimal strategy of waiting a year before deciding whether to invest.

We have found that the value of our investment opportunity, that is, the value of the option to invest in this project, is \$773. The payoff from investing (exercising the option) today is \$2200 - \$1600 = \$600. But once we invest, our option is gone, so the \$773 is an opportunity cost of investing. Hence the full cost of investing today is \$1600 + \$773 = \$2373 > \$2200. As a result, we should wait and keep our option alive, rather than invest today. We have thus come to the same conclusion as we did by comparing NPV's. This time, however, we calculated the value of the option to invest, and explicitly took it into account as one of the costs of investing.

Our calculation of the value of the option to invest was based on the construction of a risk-free portfolio, which requires that one can trade (hold a long or short position in) widgets. Of course, we could just as well have constructed our portfolio using some other asset, or combination of assets, the price of which is perfectly correlated with the price of widgets. But what if one cannot trade widgets, and there are no other assets that "span" the risk in a widget's price? In this case one could still calculate the value of the option to invest the way we did at the outset-by computing the NPV for each investment strategy (invest today versus wait a year and invest if the price goes up), and picking the strategy that yields the highest NPV That is essentially the dynamic programming approach. In this case it gives exactly the same answer because all price risk is diversifiable. Later in this book we will explore this connection between option pricing and dynamic programming in more detail.

1.B Characteristics of the Option to Invest

We have seen how our investment decision is analogous to the decision tc exercise an option, and we were able to value the option to invest in much the same way that financial call options are valued. To get more insight into the nature of the investment option, let us see how its value depends on various parameters. In particular, we will determine how the value of the option - and the decision to invest - depend on the direct cost of the investment, I, on the initial price of widgets, P_0, on the magnitudes of the up and down movements in price next period, and on the probability q that the price will rise next period.

Changing the Cost of the Investment

So far we have fixed the cost of the investment, I, at \$1600. How much would the option to invest be worth if I were above or below this number? We can find out by going through the same steps that we did before. Doing so, it is easy to see that the short position needed to obtain a risk-free portfolio depends on I as follows: [5]

$$n = 16.5 - 0.005\, I \qquad (6)$$

The current value of the option to invest is then given by

$$F_0 = 1500 - 0.455\, I. \qquad (7)$$

---

[5] As before, the value of the portfolio next year is (_$_1$ = $F_1 - n\, P_1$. If P_l = \$300, F_l = 3300 – I, so _$_1$ = $3300 - I - 300\, n$. If P_l = \$100, and I is not so low that we would invest anyway, F_l = 0 and _$_1$ = - $100\, n$. Setting the _$_1$'s for each price scenario equal gives the equation above for n.

Equation (7) gives the value of the investment opportunity as a function of the direct cost of the investment, I. We saw earlier that if I = \$1600, it is better to wait a year rather than invest today. Are there values of I for which investing today is the preferred strategy?
To answer this, recall that we should invest today as long as the payoff from investing is at least as large as the full cost, that is, the direct cost, I, plus the opportunity cost F_0. Since the payoff from investing today is V_0 = \$2200, we should invest today if 2200 > $I + F_0$. Substituting in equation (7) for F_0, we should invest as long as

$$I + 1500 - 0.455\, I < 2200.$$

Thus, if I < \$1284, one should invest today rather than wait. The reason is that waiting means giving up revenue in the first year, and in this case the lost revenue exceeds the opportunity cost of committing resources rather than keeping the investment option open. However, if I = \$1284, F_0 = \$916 = $V_0 - I$, and one would be indifferent between investing today and waiting until next year. (This can also be seen by comparing the project's NPV if we invest today with the NPV if we wait until next year; in either case the NPV is \$916.) And if I > \$1284, one is better off waiting.

The dependence of F_0 on I is illustrated in Figure 2.2. This graph shows the value of the option, F_0, and the net payoff from investing today, $V_0 - I$, both as functions of I. For I > \$1284, $F_0 = 1500 - 0.455\, I > V_0 - I$, so the option should be kept alive, that is, we should wait until next year before deciding whether to invest. However, if I < \$1284, $F_0 = 1500 - 0.455\, I < V_0 - I$, so the option should be exercised now, and hence the value of the option is just its net payoff, $V_0 - I$.

In the terminology of options, when I is small, the net payoff from immediate investment becomes large, or the option is "deep in the money." At a critical point or threshold where it is sufficiently deep, the cost of waiting

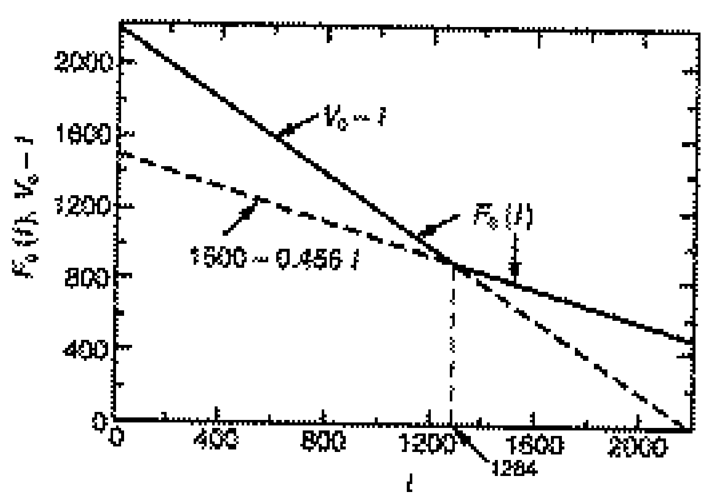

Figure 2.2. Option to Invest in Widget Factory

(the sacrifice of the immediate profit) outweighs the benefit of waiting (the value of the ability to decide optimally after observing whether the price has gone up or down), and immediate exercise becomes optimal.

Changing the Initial Price

Let us again fix the cost of the investment, I, at \$1600, but now vary the initial price of widgets, P_0. To do this, we will assume that whatever P_0 happens to be, with probability 0.5 the price next year will be 50 percent higher, and with probability 0.5 it will be 50 percent lower. (See Figure 2.3.)

To value the option to invest, we again set up a risk-free portfolio in which we hold the option and sell short some number of widgets. The value of this portfolio today is $\_{0}$ = F_0 - n P_0. Its value next year depends on P_1. The value of a widget factory next year is V_l = $\sum_0^{\infty} P_1/(1.1)^t = 11P_1$,but we will only invest in the factory if its value exceeds \$1600, the cost of the investment. Hence F_1 = max [0, 11 P_1 -1600]. Suppose P_0 is in the range where if the price goes up next year (that is, if $P_1 = 1.5\ P_0$) it will be worthwhile to invest, but if the price goes down it will not. (We will consider other possibilities shortly.) Then ($\_{1}$ = 16.5 P_0 - 1600 - 1.5 n P_0 if the price goes up, and $\_{1}$ = - 0.5 n P_0 if the price goes down. Equating the $\_{1}$'s for these two scenarios gives the value of n that makes the portfolio risk-free:

$$n = 16.5 - 1600/\ P_0$$

Note that with n chosen this way, $\_{1}$ = - 8.25 P_0 + 800 whether the price goes up or down.

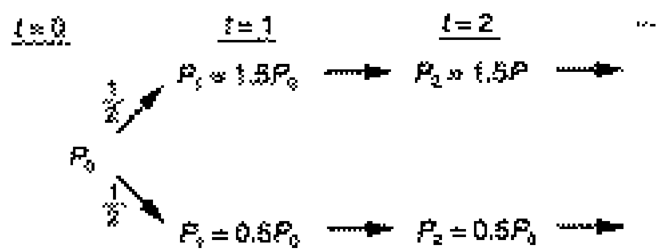

Figure 2.3. Price of Widgets

Now let us calculate the return on this portfolio, remembering that the short position will require a payment of 0.1 $n\ P_0$ = 1.65 P_0 - 160. That return is 6.60 $P_0 - F_0$ - 640. Since the return is risk-free, it must equal 0.1 $\_{0}$ = 0.1 F_0 – 1.65 P_0 + 160. Solving for F_0 gives the value of the option to invest:

$$F_0 = 7.5\ P_0 - 727. \qquad (9)$$

We have calculated the value of the investment option assuming that we would only want to invest if the price goes up next year. However, if P_0 is low enough we might never want to invest, and if P_0 is high enough it might be better to invest now rather than waiting. Below what price would we never invest? From the equation above, we see that $F_0 = 0$ when 7.5 P_0 = 727, or P_0 = \$97.

If P_0 is less than \$97, V_l will be less than the \$1600 cost of the investment even if the price rises by 50 percent next year.

For what values of P_0 should we invest now rather than wait? Once again, we should invest now if the current value of a widget factory, V_0, exceeds its total cost, \$1600 + F_0.

Hence the critical price, which we will denote by P_0, satisfies $V_0 = 1600 + F_0$, that is, 11 $P_0^* = 1600 + 7.5\ P_0^*$ - 727, or $P_0^* = \$249$. If P_0 exceeds \$249, we are better off investing today rather than waiting. The option is sufficiently deep in the money that the cost of waiting (the sacrifice of period-0 profit) outweighs the benefit. [6]

We obtained this critical price by finding the value of the investment option, but we also could have found it by calculating (as a function of P_0) the NPV of the project assuming we invest today, and equating it to the NPV assuming we wait until next year and then decide whether to invest based on the outcome for price. (We leave this for you as an exercise.)

The value of the investment option is thus a piecewise-linear function of the current price, P_0, and the optimal investment rule likewise depends on P_0. If $P_0 \le \$97$, $F_0 = 0$, and one should never invest in the factory. If $\$97 < P_0 \le \$249$, then $F_0 = 7.5\ P_0 - 727$, and one should wait a year and invest if the price goes up. If $P_0 > \$249$, then $F_0 = 11\ P_0 - 1600$, and one should invest immediately. This solution is summarized in Table 2.1.

In Figure 2.4 we have plotted F_0 as a function of P_0. Suppose our only choice were to invest today or else never invest. Then the option to invest would be worth the maximum of zero and 11 P_0 - 1600, that is, the value of a widget

Table 2.1. Value of Option to Invest and Investment Rule

Region	Option Value	Optimal Investment Rule
$P_0 \le 97$	$F_0 = 0$	Never invest.
$97 < P_0 \le 249$	$F_0 = 7.5 P_0 - 727$	Invest in period 1 only if price has gone up.
$P_0 > 249$	$F_0 = 11 P_0 - 1600$	Invest in period 0.

factory today, V_0, less the \$1600 cost of building the factory. (In financial jargon, this is the *intrinsic value* of the option - the maximum of zero and its value if it were exercised immediately.) Also, we would invest today as long as 11 P_0 > \$1600, that is, as long as P_0 > \$146. When P_0 > \$249, the option to invest is worth its intrinsic value of 11 P_0 - 1600, because then it is optimal to invest today rather than wait. Hence 11 P_0 -1600 is shown as a solid line for values of P_0 greater than \$249. When P_0 < \$249, however, the option to invest is worth more than 11 P_0 - 1600, and should be left unexercised, at least until next year.

[6] You might think that immediate investment would be justified only if we would invest next year irrespective of whether the price went up or down. In fact, the critical price for immediate investment is lower. Assuming we wait, we will invest next year if $V_1 - 1600 = 11\ P_1 - 1600 > 0$, so if the price goes down, that is, $P_1 = 0.5\ P_0$, we would invest only if $5.5\ P_0 - 1600 > 0$, or $P_0 > \$291$, which exceeds $P_0^* = \$249$. The problem with the faulty intuition is that it ignores the profit we could earn this year by investing now. Indeed, for a P_0 of, say, \$260, it is better to invest today even though if we did wait, we would *not* invest next year if it turned out that the price went down.

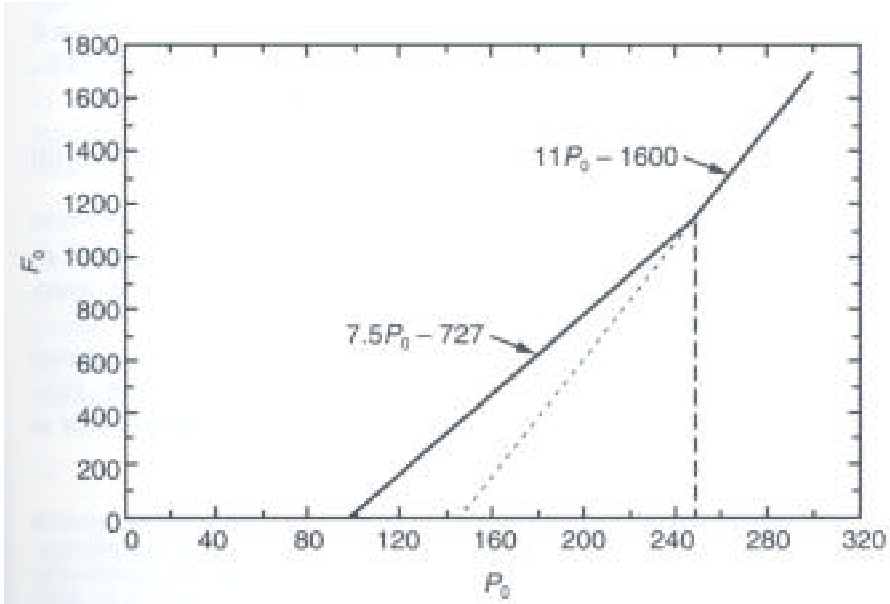

Figure 2.4. Value of Option to Invest as a Function of Initial Price

Note from Figure 2.4 that F_0 is a convex function of P_0, and that F_0 is greater than or equal to the net payoff from exercising the option today, V_0 - I, up to the optimal exercise point ($249 in this example). As we will see, the value of an option to invest typically has these characteristics.

Changing the Probabilities for Price

We can also determine how the value of the option to invest depends on q, the probability that the price of widgets will rise next year. (So far we have assumed that this probability is 0.5.) To do this, we will let the initial price, P_0, be some arbitrary value, but we will fix the cost of the investment, I, at $1600. We can then follow the same steps as before to find the value of the option and the optimal investment rule.

The reader can verify that the short position in widgets needed to construct a risk-free portfolio is n = 8.5, and is independent of q. (The reason is that n depends on the possible values for the portfolio in period 1, _$_1$, and not on the probabilities that the portfolio will take on those values.) However, the payment required for this short position does depend on q, because the expected capital gain from holding a widget depends on q. To calculate this, let _$_0$ (P_1) denote the expected price of widgets next year, calculated conditional on the knowledge of the period-0 price.[7] Thus _$_0$ (P_1) = (q + 0.5) P_1. Therefore the expected rate of capital gain on a widget is [_$_0$ (P_1) – P_0]/ P_0 = q - 0.5. Hence the required payment per widget in the short position is [0.1 - (q - 0.5)] P_0 = (0.6 - q) P_0. Setting _$_1$- _$_0$ - (0.6 - q) n P_0 = 0.1 _$_0$, and setting n = 8.5, we find that if P_0 > $97, the value of the option is

$$F_0 = 15\, q\, P_0 - 1455\, q. \qquad (10)$$

[7] We will use _ to denote the expectation (mean) of a random variable, and a time subscript on _ will indicate that the expectation is conditional on information available as of that time. Similarly, we will use _ to denote variances. These symbols are used consistently throughout the book. Other notation is specific to each chapter or even section. A symbol glossary at the end of the book collects the significant symbols and states what they denote.

Note that F_0 increases as q increases (as long as P_0 > 97). This is as we would expect, because a higher q means a higher probability that the price will go up and the option to invest will be exercised. If P_0 < \$97, we will never invest, whether or not the price goes up, and $F_0 = 0$.

How does the decision to invest depend on q? Recall that it is better to wait rather than invest today as long as $F_0 > V_0 - I$. In this case, $V_0 = P_0 + \sum_{1}^{\infty} (q + 0.5)\, P_0/(1.1)^t = (6 + 10\, q)\, P_0$. Hence it is better to wait as long as $15\, q\, P_0 - 1455\, q > (6 + 10\, q)\, P_0 - 1600$, that is, as long as $P_0 < P^*_0 = (1600 - 1455\, q)/(6 - 5\, q)$. Also, P^*_0 decreases as q increases, that is, a higher probability of a price increase induces the firm to invest more readily. Why? The cost of waiting is the revenue foregone by not selling a widget this period, which increases with P_0. Since a higher q makes a bad outcome next period less likely, it reduces the value of waiting. Thus a higher q implies that a smaller P_0 suffices to make the cost of waiting exceed the value of waiting.

Increasing the Uncertainty over Price

When we changed the probability q while keeping all the other parameters fixed, we changed the expected price in period 1. Suppose the expected price in period 1 remains fixed at the initial price level, P_0, but we increase the size of the up and down changes so that the *variance* of the period-1 price increases. What effect would such a *mean-preserving spread* in the distribution for P_1 have on the value of the investment option, F_0, and on the critical price, P^*_0, above which it is optimal to invest immediately rather than wait?

We will assume as before that q is 0.5, but now the price will either rise or fall in period 1 by 75 percent, rather than by 50 percent as before. Thus the variance of P_1 is greater, but its expected value is still P_0. To find F_0, we go through the same steps as before, creating a risk-free portfolio and equating its return to the risk-free return. Again, let the portfolio be long the investment option and short n widgets, so its value now is $\_{}_0 = F_0 - n\, P_0$. In period 1, the price will rise to $1.75\, P_0$, in which case the project will be worth $V_1 = 11\, P_1 = 19.25\, P_0$ so that $\_{}_1$ will equal $19.25\, P_0 - 1600 - 1.75\, n\, P_0$, or the price will fall to $0.25\, P_0$, in which case $\_{}_1$ will equal $- 0.25\, n\, P_0$. Equating these two possible values of $\_{}_0$ and solving for n gives

$$n = 12.83 - 1067/P_0. \qquad (11)$$

(Then, $\_{}_1 = - 3.21\, P_0 + 267$ irrespective of P_1.) Remembering that the short position will require a payment of $0.1\, n\, P_0 = 1.28\, P_0 - 107$, the return on the portfolio is $8.34\, P_0 - F_0 - 693$. Setting this equal to $0.1\, \_{}_0 = 0.1\, F_0 - 1.28\, P_0 + 107$ and solving for F_0 gives

$$F_0 = 8.75\, P_0 - 727. \qquad (12)$$

If P_0 is \$200, F_0 is \$1023, which is substantially larger than the value of \$773 that we found before when the price could only rise or fall by 50 percent. Why does an increase in uncertainty increase the value of the option to invest? Because it increases the upside potential payoff from the option, leaving the downside payoff unchanged at zero (since we will not exercise the option if the price falls).

We can also calculate the critical initial price, P_0^*, that is sufficient to warrant investing now rather than waiting. Again, just equate the current value of the widget factory, $V_0 = 11\ P_0$, to its total cost, \$1600 + F_0. Utilizing equation (12) for F_0, this gives P_0^* = \$388, which is much larger than the value of \$249 that we found before. Because the value of the option is larger, the opportunity cost of investing now rather than waiting is larger, so there is a greater incentive to wait.

A "Bad News Principle"

We can take this one step further by allowing both the probability q of an upward price move as well as the sizes of the upward and downward moves to vary. In so doing, we can determine how "good news" (an upward move) and "bad news" (a downward move) separately affect the critical price, P_0^* , that warrants immediate investment (in the above calculation the upward and downward moves had to increase or decrease together). We will see that P_0^* depends only on the size of the downward move, not the size of the upward move. The reason is that it is the ability to avoid the consequences of "bad news" that leads us to wait. [8]

Suppose that the initial price is P_0, but in period 1 the price becomes

$(1 + u)\ P_0$ with probability q,

P_1
$(1 - d)\ P_0$ with probability 1 - q.

To keep things general, we will let the cost of the investment be I. In this case, the NPV if we invest now is

$$NPV = -I + P_0 + q\sum_{t=1}^{\infty}\frac{(1+u)P_0}{(1.1)^t} + (1-q)\sum_{t=1}^{\infty}\frac{(1-d)P_0}{(1.1)^t} = -I + 10\left[1.1 + q(u+d) - d\right]P_0. \quad (13)$$

On the other hand, if we wait the NPV is

$$NPV = \frac{1}{1.1}\left\{q\max\left[0, -I + 11(1+u)P_0\right] + (1+q)\max\left[0, -I + 11(1-d)P_0\right]\right\}. \quad (14)$$

It is easy to show (and should be intuitively clear) that the point of indifference between investing now and waiting occurs in the range of P_0 where investment in period 1 is warranted if the price goes up but not if it goes down. In this case the NPV in equation (14) simplifies to

$$NPV = \frac{q}{1.1}\left[-I + 11(1+u)P_0\right] \quad (15)$$

---

[8] This "bad news principle" was first spelled out by Bernanke (1983), and the ideas can also be found in Cukierman (1980).

Equating the NPV of equation (13) for investing now with the NPV of equation (15) for waiting and solving for P_0 gives

$$P_0^* = I\left(\frac{0.1}{1.1}\right)\left(\frac{0.1+(1-q)}{0.1+(1-q)(1-d)}\right) \tag{16}$$

Equation (16) has one detail that is important to note – P_0^* does not depend in any way on *u*, the size of an upward move. It only depends on the size of a downward move, *d*, and the probability (1 - *q*) of a downward move. Also the larger is *d*, the larger is the critical price, P_0^*; it is the magnitude of the possible "bad news" that drives the incentive to wait. [9]

We can also examine the effect of a mean-preserving spread in Pl in a more general way than we did before. Suppose we set q = d/(u + d). Ther E'(Pl) = Po, and V(Pl) = u d Po . Hence if we increase a and d proportionally we keep q and F (PI) unchanged while increasing the variance of Pl. Observe from equation (16) that if d is larger but q is unchanged, Po increases. Again. we find that a mean-preserving spread increases the incentive to wait.

2 Extending the Example to Three Periods

In our example, we made the unrealistic assumption that there is no uncertainty over the price of widgets after the first year. In most markets, future prices are always uncertain, and the amount of uncertainty increases with the time horizon. In other words, while the expected future price of widgets might always equal the current price, the variance of the future price will typically be greater the farther into the future we look. Later in this book we will model the stochastic evolution of price in just this way. At this point, however, we can obtain additional insight into the nature of the investment problem by extending our example so that there are three periods in which the investment decision might be made.

We will assume as before that at $t = 0$ the price of widgets begins at some level P_0, and at $t = 1$ it will either increase or decrease by 50 percent (to $P_1 = 1.5\ P_0$ or $P_1 = 0.5\ P_0$), each with probability 0.5. Then, at $t = 2$, it will again either increase or decrease by 50 percent with equal probability. Hence there are three possible values for P_2: 2.25 P_0, 0.75 P_0, and 0.25 P_0. The price then remains at this level for all *t* ? 2. (See Figure 2.5.) We will again fix the direct cost of the investment, *I*, at $1600.

By adding one more period of price uncertainty, our investment problem becomes quite a bit more complicated. One reason is that there are now five different possible investment strategies that might make sense and must be considered. In particular, it might be optimal to (i) invest immediately; (ii) wait a year and then invest if the price has gone up, but never invest if the price has gone down; (iii) wait a year and invest if the price has gone up, but if it went down wait another year and invest if it then goes up; (iv)

[9] If current profit can be negative and the firm is contemplating a costly disinvestment or abandonment of a project, the bad news principle turns into a good news principle: the size and probability of an upturn are the driving forces behind the incentive to wait. Recalling our discussion of suicide in Chapter 1, we should emphasize this point. If the potential bad outcomes become even worse, that does not increase the incentive for immediate abandonment. However, if potential good outcomes become better, that increases the value of staying alive.

wait two years and only invest if the price has gone up both times; or (v) never invest. Which rule is optimal will depend on the initial price and the cost of the investment, and the value of the investment option must be calculated for each possible rule. The second complicating factor is that while we can still compute the value of the investment option by constructing a risk-free portfolio, the makeup of that portfolio will not be constant over the two years; we will have to change the number

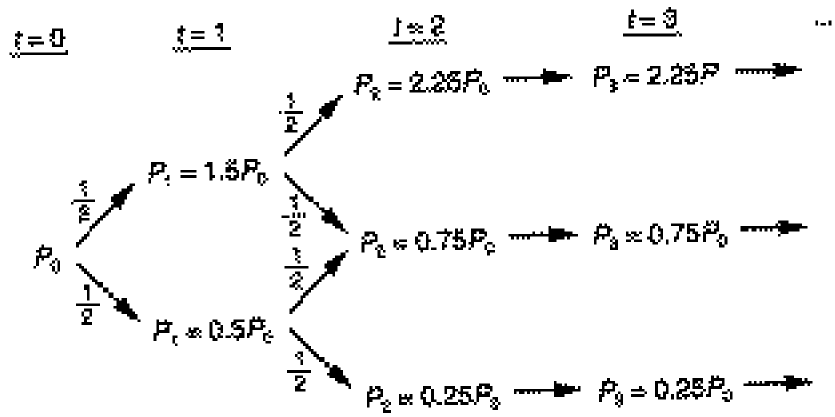

Figure 2.3. Price of Widgets

of widgets in the short position after the price changes at t = l. [10]

We will again approach this problem using option pricing methods. We want to obtain the value of the option to invest at $t = 0$, F_0, as a function of the initial price, P_0, as well as the optimal investment rule. The trick is to work backwards. We will solve two separate investment problems looking forward from $t = 1$, first for $P_1 = 0.5\ P_0$, and then for $P_1 = 1.5\ P_0$, assuming in both cases that we have not yet invested. In each case we will determine F_1, the value of the investment option at $t = 1$, by constructing a risk-free portfolio and calculating its return. Given the two possible values for F_1 (one for $P_1 = 0.5\ P_0$ and one for $P_1 = 1.5\ P_0$), we then back up to $t = 0$ and determine F_0 by again constructing a risk-free portfolio and calculating its return.

Suppose that at $t = 1$, P_1 - 0.5 P_0, and that P_0 is such that we would invest in period 2 if the price goes up, but not if it goes down. Now construct a portfolio that includes the option to invest and is short some number n_1 of widgets. The value of this portfolio is $\_{1} = F_1 - n_1 P_1$. If the price goes up in period 2 (to 0.75 P_0), we will invest, so F_2 will equal $\sum_0^\infty$ 0.75 $P_0/(1.1)^t$ -1600 = 8.25 P_0 -1600, and $\_{2}$ will equal 8.25 P_0 – 1600 - 0.75 n_1 P_0. If the price goes down in period 2 (to 0.25 P_0), we will not invest, F_2 will equal 0, and $\_{2}$ will equal -0.25 n_1 P_0. Equating the expressions for $\_{2}$ under the two scenarios, we find that the portfolio will be risk-free if $n_1 = 16.5 - 3200/P_0$; then $\_{2}$ will equal 800 - 4.125 P_0 whether the price goes up or down. Calculating the portfolio's return ($\_{2} - \_{1}$ - 0.1 n_1 P_2) and setting it equal to the risk-free return (0.1 $\_{1}$) gives us the value of the investment option: F_1 = 3.75 P_0 - 727.3. Also, note from this that $F_1 = 0$ when $P_0 = 193.94$. Hence if the price has gone down in period 1 and $P_0 < 193.94$, we will never invest.

[10] The method of keeping a portfolio riskless by changing its composition through repeated trading is called a "dynamic hedging strategy," and is of considerable importance in financial economics. We will develop it in a general setting of continuous time in Chapters 4 and 5.

We must now repeat this exercise assuming that price has gone up in period 1, that is, that P_1 = 1.5 P_0. You can verify that in this case the risk-free portfolio requires a short position of n_1 = 16.5 - 1067/ P_0 widgets, and the value of the investment option is F_1 = 11.25 P_0 - 727.3. In this case F_1 = 0 when P_0 = 64.65. Hence if P_0 < 64.65 we will never invest at all, even if the price goes up in both periods. Also, suppose we invest in period 1 rather than waiting until period 2. Then we would obtain a net value V_1 - I = 11(1.5 P_0) - 1600. Setting this equal to F_1 and solving for P_0 gives P_0 = 166.23. Hence if P_0 > 166.23, we should invest in period 1 if the price has gone up, rather than wait another year.

We now know F_1 and the optimal investment strategy (assuming we have not yet invested) for each of the two possible outcomes for P_1, so we can determine F_0 by once again constructing a risk-free portfolio and calculating its return. Since the optimal investment strategy depends on P_0, this must be done for different ranges of P_0, that is, for 64.65 < P_0 < 166.23 (in which case we would invest in period 2 only if the price goes up in both periods), for 166.23 < P_0 < 193.94 (in which case we would invest in period 1 if the price goes up, but if it goes down we would never invest), and for P_0 > 193.94 (in which case we would invest in period 1 if the price goes up, but if it goes down wait and invest if it goes up in period 2). The solution, which the reader might want to verify, is summarized in Table 2.2.

Figure 2.6 shows F_0 plotted as a function of P_0. As in the two-period model, it is a piecewise-linear function, but now there are more pieces, each corresponding to an optimal investment strategy. Also, suppose the only choice were to invest today or else never invest. Then the option to invest would be worth 11 P_0 - 1600, and we would invest today as long as 11 P_0 > \$1600, that is, as long as P_0 > \$146. When P_0 > 301.19, the option is worth this much, because then F_0 = 9.2 P_0 - 1057.9 = 11 P_0 -1600, and it is optimal to invest today rather than wait. Hence 11 P_0 -1600 is shown as a solid line for P_0 > 301.19, but is shown as a dotted line for 146 < P_0 < 301.19. Finally, note once again that F_0 is a convex function of P_0, and that F_0 is greater than or equal to the net payoff from exercising the option today, V_0 - *I*, up to the optimal exercise point (\$301.19 in this case).

As before, we could examine the dependence of F_0 and the optimal exercise point P_0 on the cost of the investment, *I*, or on the variance of price changes. (Going through these calculations is now a bit more laborious, however, because three periods are involved instead of two.) For example, we could show that if the variance of the price changes increases, while the expected price changes remain the same, the value of the option F_0 will increase, as will the critical exercise price P_0^*. As an exercise, the reader might want to check this by letting *P* increase or decrease by 75 percent in each period, instead of by 50 percent as in Figure 2.5.

Table 2.2. Value of Option to Invest and Investment Rule

Region	Option Value	Optimal Investment Rule
$P_0 \leq 64.65$	$F_0 = 0$	Never invest.
$64.65 < P_0 \leq 166.23$	$F_0 = 5.11P_0 - 330.6$	Invest in period 2 only if price goes up in period 1 *and* in period 2.
$166.23 < P_0 \leq 193.94$	$F_0 = 7.5P_0 - 727.3$	Invest in period 1 if price goes up. If price goes down in period 1, then never invest.
$193.94 < P_0 \leq 301.19$	$F_0 = 9.2P_0 - 1057.9$	Invest in period 1 if price goes up. If price goes down in period 1, then wait and invest in period 2 if price goes up.
$P_0 > 301.19$	$F_0 = 11P_0 - 1600$	Invest in period 0.

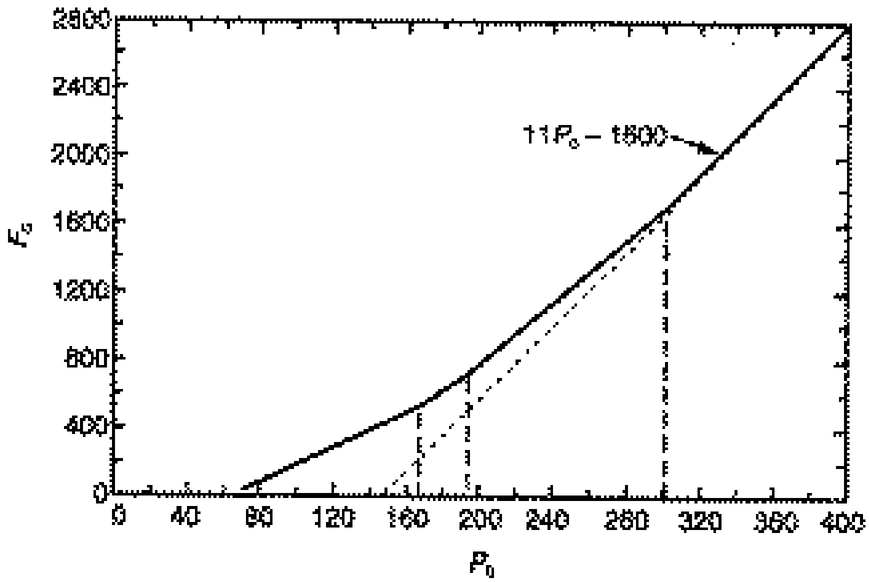

Figure 2.6. Value of Option to Invest as a Function of Initial Price

If we wanted to, we could now extend our example to four periods, allowing the price to again increase or decrease by 50 percent at $t = 3$. We could then work backwards, finding F_2 for each possible value of P_2, then finding F_1 for each possible value of P_1, and finally finding F_0. In like manner, we could then extend the example to five periods, to six periods, and so on. As we did this, we would find that the curve for F_0 would have more and more kinks. We will see in Chapter 5 that as the number of periods becomes large, the curve for F_0 will approach a smooth curve that starts at zero and rises to meet the curve showing the net payoff from immediate investment (V_0 - I). In fact the two curves meet tangentially, and the point where they meet defines the threshold P_0^* where immediate investment is optimal.

Adding more and more periods will make our example unreasonably complicated, however, and in any case would be less than satisfactory because ultimately we would like to allow the price to increase or decrease at every future time t. Thus we need a better approach to solving investment problems of this sort.

In Chapter 5 of this book we will extend our example by allowing the payoff from the investment to fluctuate *continuously* over time. As we will see, this continuous-time approach is quite powerful, and ultimately quite simple. However, it will require some understanding of stochastic processes, as well as Ito's Lemma (which is essentially a rule for differentiating and integrating functions of stochastic processes). These tools, which are becoming more and more widely used in economics and finance, provide a convenient way of analyzing a broad range of investment timing and option valuation problems. In the next two chapters we provide an introduction to these tools for readers who are unfamiliar with them.

3 Uncertainty over Cost

We now return to our simple two-period example and examine some alternative sources of uncertainty. In this section, we will consider uncertainty over the cost of the investment. Uncertainty over cost can be especially important for large projects that take time to build. Examples include nuclear power plants (where total construction costs are very hard to predict due to engineering and regulatory uncertainties), large petrochemical complexes, the development of a new line of aircraft, and large urban construction projects. Also, large size is not a requisite. Most R&D projects involve considerable cost uncertainty; the development of a new drug by a pharmaceutical company is an example.

In the context of our two-period example, suppose that the price of a widget is now \$200, and we know that it will always remain \$200. However, the direct cost of building a widget factory, I, is uncertain.

We will consider two different sources of uncertainty regarding I. The first, which we will call *input cost uncertainty*, arises because a widget factory requires steel, copper, and labor to build, and the prices of these construction inputs fluctuate stochastically over time. In addition, government regulations change unpredictably over time, and this can change the required quantities of one or more construction inputs. (For example, new safety regulations may add to labor requirements, or changing environmental regulations may require more capital.) Thus, although I might be known today, its value next year is uncertain.

As one might expect, this kind of uncertainty has the same effect on the investment decision as uncertainty over the future value of the payoff from the investment, V - it creates an opportunity cost of investing now rather than waiting for new information. As a result, a project could have a conventionally measured NPV that is positive, but it might still be uneconomical to begin investing.

As an example, suppose that I is \$1600 today, but next year it will increase to \$2400 or decrease to \$800, each with probability 0.5. As before, the interest rate is 10 percent. Should we invest today, or wait until next year? If we invest today, the NPV is again given by equation (1), that is, -1600 + 2200 = \$600. This NPV is positive, but once again it ignores an opportunity cost. To see this, let us recalculate the NPV, but this time assuming we wait until next year, in which case it will be ex post optimal to invest only if I falls to \$800. In this case the NPV is given by

$$NPV = (0.5)\left[\frac{-800}{1.1} + \sum_{t=1}^{\infty}\frac{200}{(1.1)^t}\right] = \frac{700}{1.1} = \$636. \tag{17}$$

(In year 0 there is no expenditure and no revenue. In year 1, we invest only if *I* falls to \$800, which will happen with probability 0.5.) If we wait a year before deciding whether to invest, the project's NPV today is \$636, so waiting is clearly better than investing now.

At this point it might appear that uncertainty always leads one to postpone investments, or at least increase the hurdle rate that the investment must meet, but this is not the case. If investing provides information, uncertainty can *lower* the hurdle rate for a project. Consider uncertainty over the physical difficulty of completing a project. We will call this *technical uncertainty*: assuming factor costs are known, how much time, effort, and materials will ultimately be required to complete a project? This kind of uncertainty can only be resolved by actually undertaking and completing the project. [11] One then observes actual costs (and construction time) unfold as the project proceeds. These costs may from time to time turn out to be greater or less than anticipated (as impediments arise or instead the work moves ahead faster than planned), but the total cost of the investment is only known for certain when the project is complete.

With this kind of uncertainty, a project can have an expected cost that makes its NPV negative, but if the variance of the cost is sufficiently high, it can still be economical to begin investing. The reason is that investing reveals information about cost, and thus about the expected net payoff from investing further. It therefore has a value beyond its direct contribution to the completion of the project. This additional value (called a *shadow value* because it is not a directly measurable cash flow) lowers the full expected cost of the investment.

The following simple example will help to illustrate this. Suppose that the price of a widget is and always will be \$200, but the cost of building a widget factory is uncertain. To build a factory, one must initially spend \$1000. With probability 0.5 the factory will then be complete, but with probability 0.5 an additional \$3000 will be required to complete it. Since the expected cost of the factory is 1000 + (0.5)(3000) = \$2500, and its value is \$2200, the NPV of the investment seems to be negative, suggesting that we should not invest. But this ignores the value of the information obtained from completing the first stage of the project, and the fact that we can abandon the project should a second phase costing \$3000 be required. The correct NPV is -1000 + (0.5)(2200) _ \$100. This is positive, so one should invest in the first phase of the project.

We see, then, that the uncertainty over the cost of a project can lead one to postpone investing or to speed it up. If the resolution of uncertainty is independent of what the firm does, it has much the same effect as uncertainty over the payoff from investing, and creates an incentive to wait. But if the uncertainty can be partly resolved by investing, it has the opposite effect. We will return to uncertainty over cost later in this book, and examine its implications in more detail.

4 Uncertainty over Interest Rates

Next suppose that the cost of the investment and the payoff from investing are both known with certainty, but the interest rate used to discount future cash flows changes in an unpredictable manner. How will uncertainty over the interest rate affect the decision to invest?

---

[11] This is a simplification in that, for some projects, cost uncertainty can be reduced by first performing additional engineering studies. The investment problem is then more complicated because one has three choices instead of two: start construction now, undertake an engineering study and then begin construction only if the study indicates costs are likely to be low, or abandon the project completely.

Interest rate uncertainty can have two effects on an investment decision. First, unpredictable fluctuations in interest rates can increase the *expected value* of a future payoff from investing. For example, suppose the investment yields a perpetuity that pays \$1 per year forever. The present value of this perpetuity is \$1/r, where r is the interest rate. If r = 10 percent, the value is \$1/0.10 = \$10. But suppose r is uncertain, and can equal either 5 percent or 15 percent, each with probability 0.5, so that _(r) = 10 percent. Then the expected value of the perpetuity is 0.5(1/0.05) +0.5(1/0.15) = \$13.33 > \$10. This makes the investment more attractive, and increases the incentive to invest. [12]

Nonetheless, uncertainty over future interest rates can still lead to a postponement of investment. The reason is that the second effect of interest rate uncertainty works in the opposite direction- it creates a value of waiting (to see whether interest rates rise or fall). This effect works in much the same way as uncertainty over the payoff from the investment.
To clarify this, let us again return to our two-period example. This time we will assume that the price of a widget is fixed at \$200, and that the cost of building the widget factory is fixed at \$2000. The only uncertainty is over interest rates. Today the interest rate is 10 percent, but next year it will change. There is a 0.5 probability that it will increase to 15 percent, and a 0.5 probability that it will decrease to 5 percent. It will then remain at this new level.

What will a widget factory be worth next year? If there were no uncertainty over interest rates, that is, if we knew that the interest rate will remain 10 percent, then the value of the factory next year would be

$$V_1 = \sum_{t=1}^{\infty} \frac{200}{(1.1)^t} = \$2200.$$

In our example, the expected value of the interest rate next year and thereafter is 10 percent, but the actual value is unknown. So, the value of the factory next year will be

$$V_1 = \begin{cases} \sum_{t=1}^{\infty} 200/(1.15)^t = \$1533 & \text{with probability 0.5,} \\ \sum_{t=0}^{\infty} 200/(1.05)^t = \$4200 & \text{with probability 0.5.} \end{cases}$$

Hence the expected value of V_1 is (0.5)(1533) + (0.5)(4200) = \$2867, which is higher than it is when future interest rates are certain.

We see, then, that uncertainty over the future interest rate (holding the expected value of the interest rate constant) increases the expected value of the project. But how does uncertainty affect the investment decision? First, note that if there were no uncertainty

[12] In technical language, this result is an implication of *Jensen's Inequality,* combined with the fact that the present value of a future cash flow is a convex function of the interest rate. Jensen's Inequality says that if x is a random variable and $f(x)$ is a convex function of x, then _[$f(x)$] > f(_[x]). Thus if the expected value of x remains the same but the variance of x increases, _[$f(x)$] *will* increase. In the case at hand, if the expected value of next year's interest rate remains fixed but the uncertainty around that value increases, the expected present discounted value of a payoff received next year will increase

over interest rates, we would clearly want to invest today. The NPV of the project if we invest today is

$$NPV = -2000 + \sum_{t=0}^{\infty} \frac{200}{(1.1)^t} = \$200. \tag{18}$$

whereas the NPV today if we wait until next year is only \$2200/1.1 -\$2000 = \$0. The situation is different if interest rates are uncertain. If we invest now, the NPV is

$$NPV = -2000 + 200 + \frac{\varepsilon(V_1)}{1.1} = -1800 + \frac{2867}{1.1} = \$806. \tag{19}$$

(If we invest today we can make and sell a widget now for \$200, and we will have a factory whose expected value next year is \$2867.) This NPV is positive, but suppose we wait until next year before deciding whether to invest. If the interest rate rises to 15 percent, the value of the factory will be only \$1533, which is less than the \$2000 cost of the investment. Hence we will only invest if the interest rate falls to 5 percent. Since there is a 0.5 probability that this will happen, the NPV assuming we wait is

$$NPV = (0.5)\left[\frac{-2000}{1.1} + \frac{1}{1.1}\sum_{t=0}^{\infty} \frac{200}{(1.05)^t}\right] = \$1000. \tag{20}$$

This NPV is higher, so it is better to wait than invest now.

This is a simple analysis of interest rate uncertainty, but it has some . important implications. First, mean-preserving volatility in interest rates will increase the expected value of a project, but will also create an incentive to wait , rather than invest now. The reason is that, as with uncertainty over future cash flows, uncertainty over future interest rates creates a value to waiting for new , information. Second, if an objective of public policy is to stimulate investment; the stability of interest rates may be more important than the level of interest rates. Policies that lead to lower but more volatile interest rates could end up depressing aggregate investment spending.[13] As we will illustrate in Chapter 9, the importance of stability and predictability may apply to other instruments of government policy as well, such as tax rates and trade policy.

5 Scale versus Flexibility

As students of economics or business learn early on, economies of scale can be an important source of cost savings. By building a large plant instead of two or three smaller ones, a firm might be able to reduce its average cost and increase its profitability. This suggests that firms should respond to growth in demand for their products by bunching their investments, that is, investing in new capacity only infrequently, but adding large and efficient plants each time.

What should firms do, however, when there is uncertainty over demand growth (as there usually is)? If the firm irreversibly invests in a large addition to capacity, and demand grows only slowly or even shrinks, it will find itself holding capital it does not need. Hence when the growth of demand is uncertain, there is a tradeoff between scale

[13] Ingersoll and Ross (1992) have developed a continuous-time model of investment with interest rate uncertainty, and it leads to more or less the same conclusions.

economies and the flexibility that is gained by investing more frequently in small increments to capacity as they are needed.

This problem is an important one for the electric utility industry. It is much cheaper per unit of capacity to build a large coal-fired power plant than it is to add capacity in small amounts. But at the same time, utilities face considerable uncertainty over the rate by which the demand for their electricity will grow.[14] Adding capacity in small amounts gives the utility flexibility, but is also more costly. Hence it is important to be able to value this flexibility. The options approach used in this book is well suited to do this. Here we will illustrate the basic idea with a simple example in which demand growth is certain, but there is uncertainty over relative fuel prices.[15]

Consider a utility that faces constant demand growth of 100 megawatts (MW) per year. Hence the utility must add to capacity; the question is how. There are two alternatives. The utility can build a 200-MW coal-fired plant (enough capacity for two years worth of additional demand) at a capital cost of $190 million (Plant A), or it can build a 100-MW oil-fired plant at a capital cost of $100 million (Plant B). At current coal and oil prices, the coal-fired plant is not only more economical in terms of its capital cost ($90 million per 100 MW of capacity), but also in terms of its operating cost; operating Plant A will cost $19 million per year for each 100 MW of power, whereas Plant B will cost $20 million per year. We will assume that the discount rate for the utility is 10 percent per year, and that each plant lasts forever. In this case, if fuel prices remain constant, Plant A is clearly the preferred choice.

Fuel prices are unlikely to remain constant, however. Since what matters is the relative price of oil compared to coal (and since the price of coal is in fact much less volatile than the price of oil), we will assume that the price of coal will remain fixed, but the price of oil will either rise or fall next year, with equal probability, and then remain constant. If it rises, the operating cost for Plant B will rise to $30 million per year, but if it falls, the operating cost will fall to $10 million per year. (See Figure 2.7.)

[14] The reason is not just that there is uncertainty over the growth of total electricity demand, but also because utilities now often find themselves competing with other sources of electric power (such as cogeneration).

[15] This is a modified version of an example presented in Sawhilll (1989).

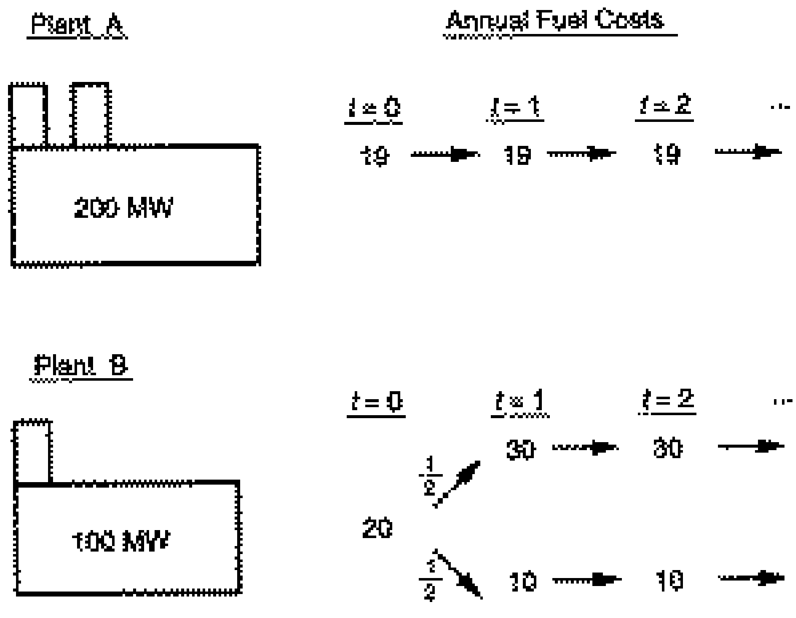

Figure 2.7. Choosing Among Electric Power Plants

The choice of plant is now more complicated. Although Plant A's capital cost is lower because of its scale, and its operating cost is lower at the current oil price, Plant B affords the utility more flexibility because it only requires a commitment to one year's worth of demand growth- if the price of oil falls, the utility will not be stuck with the extra 100 MW of coal-burning capacity in the second year. To decide which choice is best, let us calculate for each the present value of the expected cost of generating 100 MW of power per yea forever starting this year, and an additional 100 MW per year starting next year.

First, suppose we commit the full 200 MW to either coal or oil. Then, if we choose coal, the present value of the flow of cost is

$$PV_A = 180 + \sum_{t=0}^{\infty} \frac{19}{(1.1)^t} + \sum_{t=1}^{\infty} \frac{19}{(1.1)_t} = \$579. \qquad (21)$$

Note that 180 is the capital cost for the full 200 MW, and that 19 is the annual operating cost for each 100 MW, the first of which begins now and the second next year. Next, suppose we choose oil. Since the expected operating cost for oil is \$20 million per year, the present value is

$$PV_B = 100 + \frac{100}{1.1} + \sum_{t=0}^{\infty} \frac{20}{(1.1)^t} + \sum_{t=1}^{\infty} \frac{20}{(1.1)^t} = \$611. \qquad (22)$$

Thus it would seem that Plant A is preferred.

But this calculation ignores the flexibility afforded by the smaller oil fired plant. Suppose we install 100 MW of oil-fired capacity now, but then if the price of oil goes up next year, we install 200 MW of coal-fired capacity rather than another oil-fired plant. This would give us a total of 300 MW o capacity, so to make the cost comparison meaningful, we must net out the present value of the additional 100 MW, which would be utilized starting two years from now:

$$PV_B^{'} = 100 + \sum_{t=0}^{\infty} \frac{20}{(1.1)^t}$$

$$+ \frac{1}{2}\left[\frac{100}{1.1} + \sum_{t=1}^{\infty} \frac{10}{(1.1)^t}\right]$$

$$+ \frac{1}{2}\left[\frac{180}{1.1} - \frac{90}{(1.1)^2} + \sum_{t=1}^{\infty} \frac{19}{(1.1)^t}\right] = \$555. \qquad (23)$$

Note that the second line in equation (23) is the present value of the capita and operating cost for the second 100-MW oil plant (which is built only if the price of oil goes down), and the third line is the present value of the capita and operating cost of the *first 100 MW* of a 200-MW coal plant. This present value turns out to be $555 million, so installing the smaller oil-fired plant an(thereby retaining flexibility is the preferred choice.

One way to value this flexibility is to ask how much lower would the capital cost of Plant A have to be to make it the preferred choice. Let I_A be the capital cost of Plant A. Then the present value of the costs of installing and operating Plant A is

$$I_A + \sum_{t=0}^{\infty} \frac{19}{(1.1)^t} + \sum_{t=1}^{\infty} \frac{19}{(1.1)^t} = I_A + 399.$$

The present value of the cost of providing the 200 MW of power by installing Plant B now and then next year installing either Plant A or B (depending on whether the price of oil goes up or down) is

$$100 + \sum_{t=0}^{\infty} \frac{20}{(1.1)^t} + \frac{1}{2}\left[\frac{100}{1.1} + \sum_{t=1}^{\infty} \frac{10}{(1.1)^t}\right]$$

$$\frac{1}{2}\left[\frac{I_A}{1.1} - \frac{0.5 I_A}{(1.1)^2} + \sum_{t=1}^{\infty} \frac{19}{(1.1)^t}\right]$$

$$= 320 + \frac{1}{2}(90.9 + 100) + \frac{1}{2}(0.496 I_a + 190) = 510.5 + 0.248 I_{A.}$$

To find the capital cost that makes the utility indifferent between these choices, just equate these present values and solve for I_A:

$$I_A + 399 = 510.5 + 0.248\, I_A,$$

or, I_A^* = $148.3 million. Hence the economies of scale would have to be quite large (so that a 200-MW coal plant was less than 75 percent of the cost of two 100-MW oil plants) to make giving up the flexibility of the smaller plant economical.

6 Guide to the Literature

The net present value criterion and its application to investment decisions is an important topic in corporate finance courses, and is the starting point for much of what we do in this book. Readers unfamiliar with NPV calculations, including the use of the capital asset pricing model to determine risk-adjusted discount rates, may want to review a standard textbook in corporate finance. A good choice is Brealey and Myers (1992).

Although we largely ignore the implications of taxes in this book, they can affect the choice of discount rate for NPV calculations. Taggart (1991) provides a review of the various approaches to calculating discount rates (adjusted for risk and taxes) for use in the standard NPV model. Ruback (1986) shows that riskless after-tax nominal cash flows should always be discounted at the after-tax risk-free rate (for example, the Treasury bill rate times 1 minus the corporate tax rate), and Myers and Ruback (1992) derive a simple and robust rule for discounting risky cash flows in NPV calculations.

Throughout this book we will emphasize the connections between investment decisions and the valuation and exercising of financial options. Although certainly not necessary, some familiarity with options and option pricing techniques will be helpful when reading this book. Brealey and Myers (1992) provide a simple introduction; so do the expository surveys by Rubinstein (1987) and Varian (1987). For more detailed treatments, see Cox and Rubinstein (1985), Hull (1989), and Jarrow and Rudd (1983). Although somewhat dated, the survey article by Smith (1976) is also useful. Finally, for heuristic discussions of investments as options, see Kester (1984), Mason and Merton (1985), Trigeorgis and Mason (1987), and Chapter 12 of Copeland, Koller and Murrin (1991).