Journal of Economic Perspectives—Volume 6, Number 1—Winter 1992—Pages 107–132

Investment and Hysteresis

Avinash Dixit

The economic theory of investment under competitive conditions rests on the foundations of Marshall's analysis of long and short run equilibria. If price exceeds long run average cost, this induces existing firms to expand, and new ones to enter. If price falls below average variable cost, then firms suspend operations or even exit from the market.

Reality is very different. Firms invest in projects that they expect to yield a return in excess of a required or "hurdle" rate. Observers of business practice find that such hurdle rates are three or four times the cost of capital.[1] In other words, firms do not invest until price rises substantially above long run average cost. The hurdle rate appropriate for investment with systematic risk will exceed the riskless rate, but it seems hard to justify the large discrepancies observed. On the downside, firms stay in business for lengthy periods while absorbing operating losses, and price can fall substantially below average variable cost without inducing disinvestment or exit. Many U.S. farmers in the mid-1980s were in this situation.[2]

[1]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 was close to zero. Summers' concern was the discount rate applied to depreciation allowances. But he found that almost all firms used the same rate to discount all components of cash flow. See also Dertouzas et al. (1990, p. 61).

[2]In 1983, average net income per farm operator was $6,000. Even if rent and mortgage payments on land are excluded from costs on the theory that the land had no alternative use, the figure rises

■ *Avinash Dixit is John J. F. Sherrerd '52 University Professor of Economics, Princeton University, Princeton, New Jersey.*

An example that combines upside and downside aspects is the very slow response of U.S. imports to the exchange rate. From 1980 to the end of 1984, the real value of the U.S. dollar increased by about 50 percent. The competitive advantage of foreign firms in the U.S. market rose dramatically. But import volume began its persistent rise only at the start of 1983: a lag far longer than the year or 18 months previously believed to be typical. In the first quarter of 1985 the dollar started to fall, and by the end of 1987 was almost back to its 1978 level. But import volume did not decrease for another two years; if anything it rose a little (Krugman and Baldwin, 1987, Figures 1 and 2). Once established in the U.S. market, foreign firms were very slow to scale down or shut down their export operations when the exchange rate moved unfavorably.

Some recent developments in the theory of investment under uncertainty have offered an interesting new explanation of these phenomena. This new approach suggests that textbook pictures of the dynamics of a competitive industry need substantial redrawing. More generally, it says that a great deal of inertia is optimal when dynamic decisions are being made in an uncertain environment. It builds on an interesting analogy between real investments and options in financial markets. The main merit of this approach is that it brings many disparate phenomena into a common framework. Most intriguingly, it even sheds new light on some non-economic matters. In this article I shall give a brief outline of the new view, and discuss several of its applications.

Timing of Investment and the Value of Waiting

Three features are common to most investment decisions, and they combine to yield effects like those in the examples above. First, almost as a matter of definition, an investment entails some sunk cost, an expenditure that cannot be recouped if the action is reversed at a later date. Second, the economic environment has ongoing uncertainty, and information arrives gradually. Finally, an investment opportunity does not generally disappear if not taken immediately; the decision is not only whether to invest, but also when to invest. The qualitative implication is easily stated. When these three conditions are present, waiting has positive value. In the evolving environment, time brings more information about the future prospects of the project. As long as the opportunity to invest remains available, a later decision can be a better one. And because there are sunk costs, it does not always pay to take a less perfect action now and change it later.

to only $13,500 (figures from the *Statistical Abstract of the United States*, 1990). These being national averages, there must have been many farming families who were earning much less than the opportunity cost of their own labor.

Of course, the value of waiting must be set against the sacrifice of current profit. If current conditions become sufficiently favorable, one should eventually take the action that is optimal according to the current calculation, and not wait any longer. But the "trigger" level of currently expected profit that makes it optimal to proceed exceeds the Marshallian normal return. Similarly, waiting has value when contemplating disinvestment. The Marshallian criterion of failing to cover variable cost should not trigger abandonment; the correct point is a critical negative level of operating profit.

This view of investment under uncertainty can be summarized as "a theory of optimal inertia," or "a benevolent tyranny of the status quo." It says that firms that refuse to invest even when the currently available rates of return are far in excess of the cost of capital may be optimally waiting to be surer that this state of affairs is not transitory. Likewise, farmers who carry large losses may be rationally keeping their operation alive on the chance that the future may be brighter.

The verbal argument above is purely qualitative; it says that waiting has a positive value, but not whether this value is typically large enough to have a significant impact on investment and disinvestment decisions. In the subsequent sections I shall show in some illustrative calculations that the effect can be very large indeed, and therefore merits serious attention.

The Example of a Discrete Investment Project

My first illustrative example is the simplest, namely a single discrete investment project. Suppose the project can be launched by incurring a sunk cost K, and once launched, lasts forever. Let R denote its flow of net operating revenues per unit time.

This is where the uncertainty comes in. Future revenues are only imperfectly predictable from the current observation. The probability distribution of future net revenues is determined by the present, but the actual path remains uncertain. This probabilistic law of evolution of R can take many forms, but a particularly simple specification proves insightful as well as realistic for many applications. We suppose that each period, R can either increase or decrease by a fixed percentage. The probabilities of increase and decrease need not be equal, so there can be a positive or a negative trend to R. In other words, R follows a random walk, whose steps are of equal proportions, that is, they form a geometric series. If the time period for each step of R is very short, then the distribution of the logarithm of R_t at a future time t, given the initial R_0 at time 0, is approximately normal. Then, R is said to follow a proportional or geometric Brownian motion.

Many economic time-series—exchange rates, prices of natural resources, prices of common stocks, and others—can to a reasonable first approximation be described as geometric random walks or Brownian motions. That makes the assumption particularly natural for this illustrative example.[3] Thus, the discrete project might be an oil well, whose future revenues are random as the price of oil fluctuates. Or it may be a manufacturing plant whose output is exported, so the future revenues fluctuate with the exchange rate. Purely for expository simplicity, I shall suppose that the trend rate of growth of R is zero. This does not affect the qualitative results, and I shall mention how a non-zero trend affects the quantitative ones.[4]

The Effect of Waiting

Suppose the aim is to maximize the expected (in the statistical sense of the mean or probability-weighted average) present value of profits. Let future revenues be discounted at a positive rate $\rho > 0$, the opportunity cost of riskless capital specified exogenously. Then, given a current level R of revenues, the expected present value of the discounted future stream of revenues is R/ρ. Observe that by focussing on the expected value of profits, I am making an implicit assumption that the investor is risk-neutral. The purpose of this assumption is to show that the value of waiting has nothing to do with risk-aversion. It is rather an intertemporal trade-off of present risk v. future risk.[5]

The textbook or Marshallian criterion would be to invest when the project has positive expected net worth (present value net of the sunk cost K), that is, when $R/\rho > K$. The borderline level M of the current revenue flow that would make one indifferent between investing and not investing is given by

$$M = \rho K. \tag{1}$$

The textbook recommends investment when the current revenue flow exceeds M; I shall call M the "Marshallian investment trigger."

But this criterion comes from thinking that the choice is between acting right now to get $R/\rho - K$, and not investing at all, which gets 0. What happens if the true menu of choices is wider, and waiting for a while and then reassessing the decision is also possible? Now at the Marshallian trigger, waiting is better than either investing right away or not investing at all. To see this, consider a particular alternative strategy: Wait for a fixed interval of time, and observe the value of R, say R_1, at its end. If $R_1 > M$ invest at once, otherwise

[3]The qualitative results are valid much more generally. What we need is "positive persistence" in R: a higher value today should shift the distribution of future values to the right. An additional Appendix, available by writing to the author, explains this point. Most investment problems will have this feature. It might fail when uncertainty is due to shocks to intertemporal preferences: a higher demand today then signals lower demand in the future.
[4]The appendixes develop the analysis with a general trend μ. See also McDonald and Siegel (1986), and Pindyck (1988).
[5]In fact, the case of a risk-averse investor can be treated using similar techniques and yields similar results. We need only modify ρ to take into account the project's systematic risk (beta); see Pindyck (1991).

never invest. (Of course, the alternative strategy is not itself optimal, but by showing it does better, we prove that the Marshallian criterion is not optimal when waiting is possible.) If the return at the end of the fixed waiting time exceeds the Marshallian trigger ($R_1 > M$), then the net worth of the investment must be positive at that time, and remains positive when discounted back to the starting time. If the expected return is less than the Marshallian trigger ($R_1 < M$), the net worth is zero because we do not invest. The probability-weighted average of a positive number and zero is of course positive. Therefore the proposed alternative strategy does better than either investing right away or not investing at all, each of which yields zero when the current revenue is exactly at the Marshallian trigger. By continuity, waiting remains better than investing for initial values of R slightly in excess of M.

The point is that waiting for a certain amount of time enables an investor to avoid the downside risk in revenues over that interval, while realizing the upside potential. This selective reduction in risk over time generates a positive value of waiting. On the other side, the cost of waiting is the sacrifice of the profit flow over the period of waiting. Therefore, if the current net revenue flow reaches a sufficiently high level, it won't pay to wait any longer. There is still a critical or trigger level, say H, such that investment is optimal when the current revenue exceeds it. But this H is larger than the Marshallian level M.

We can make the argument more precise, and explore what parameters determine the size of the difference between the optimal decision to wait and the Marshallian criterion for investment. As a first step, we see how the net worth of a project might be changed by a strategy of waiting until the expected revenue exceeds an exogenously given investment trigger H. This will furnish the tools for explaining how the investment trigger H itself should be optimally chosen.[6] Figure 1 illustrates the calculation.

The upward-sloping straight line labelled $i_1 i_2$ in Figure 1 represents the value to be received from investing immediately; that is, $R/\rho - K$. If the return R is zero, then the project would lose K. Otherwise, the value of this function increases with slope $1/\rho$ as the return R increases.

Now consider how the expected return from this project changes if the rule is applied that investment will occur only if the expected return R exceeds a trigger H. If the trigger is surpassed, then the investment project takes place, and the return is given by the thickly drawn portion of the line $i_1 i_2$ above the point h, where $R = H$. If the expected return is equal to the trigger, then the firm will be indifferent between waiting and investing immediately. If the expected return is less than the trigger, $R < H$, the rule tells us to wait. But there is a positive probability that at some future time R will climb above H and generate a positive net worth. Of course we rationally anticipate this

[6]This whole procedure is very rough and heuristic, and is adopted for ease of exposition. Readers who wish to see more rigorous arguments that the optimal policy takes this "trigger level" form, and fuller explanations of the subsequent mathematics, should read the additional Appendix, available by writing to the author, and the references cited there.

Figure 1
Values of Waiting and Investing

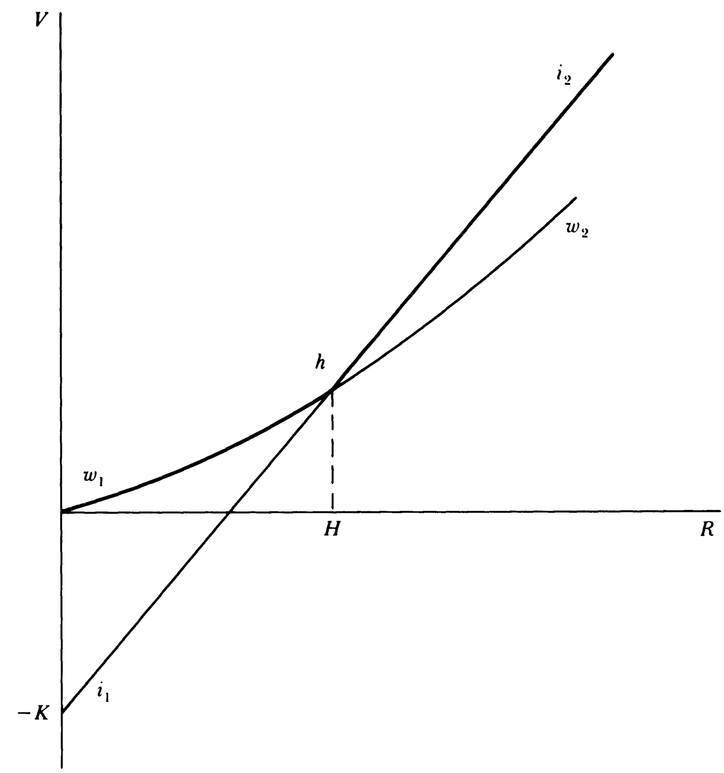

possibility, so the net worth is positive even now. The value is merely the value of waiting, or that of the opportunity or "option" to invest at some future time.[7]

Calculation of this option value needs some mathematical reasoning, which I relegate to the Appendix. Here is some intuition for its general form. The option value should approach zero if the current R is very low, because then the event of R climbing to H is unlikely except in the far future, and the discounted present value of that is quite small. Successively higher current values of R should raise the value of waiting increasingly rapidly. For R close to H but just below it, the probability of reaching H in the very near future approaches one, and the option value approaches the net worth of a live project at H. The result is shown as the convex curve labelled $w_1 w_2$ in Figure 1, starting at the origin and meeting the straight line $i_1 i_2$ at the point h. Only the

[7]I have assumed without stating so explicitly that the opportunity to invest is owned by a single firm or individual. If it is freely available to any of the usual infinity of potential entrants waiting in the wings, it cannot have a positive value. See the sections "Extensions and Qualifications" and "Competitive Industry Dynamics" later in this article.

thickly drawn portion w_1h to the left of H gives the value of waiting; beyond h investment takes place and the value of waiting is irrelevant.

The overall value of the opportunity to invest is then given by the thick curve w_1h and the thick line hi_2 taken together. The algebra of the Appendix gives the functional forms of these curves as

$$V(R) = \begin{cases} BR^{\beta} & \text{if } R \le H \\ R/\rho - K & \text{if } R \ge H. \end{cases} \tag{2}$$

The upper formula is the value of waiting (the convex curve w_1w_2 of Figure 1). The expression involves two constants, B and β, whose meaning will be explained soon. For now, just note that B is positive and β exceeds unity. The lower expression is the value of investing (the straight line i_1i_2 of Figure 1). The thickly drawn portions correspond to the value of waiting or investing in its range of validity, and the light portions show the continuation of the separate parts into the irrelevant regions. At H there is indifference, so the two expressions are equal.

There are two new terms in the first expression, showing the value of waiting, that require explanation. The power β depends on the discount rate ρ, and on the volatility of the revenue, which is measured by the variance σ^2 of the logarithm of R per unit time. The Appendix shows that

$$\beta = \frac{1}{2}\left[1 + \sqrt{1 + \frac{8\rho}{\sigma^2}}\,\right] > 1. \tag{3}$$

B is a multiplicative constant. It is determined by the condition that the two expressions for net worth $V(R)$ must be equal when R equals H. Therefore $BH^{\beta} = H/\rho - K$. Or to rephrase in a formulation which will be useful presently,

$$H/\rho = K + BH^{\beta}. \tag{4}$$

The intuitive meaning of B will be explored more in the next section.

The Optimal Policy

In the previous example, the investment trigger H was exogenously given. Now consider how the trigger should optimally be chosen. If the trigger value H is increased slightly above its value in Figure 1, that shifts the junction point h between the thickly drawn curve and line to the right. This can only be accomplished by raising the whole curve w_1w_2 representing the value of waiting. In equation (2), this corresponds to raising B in the upper formula.

To maximize value, such increase should be pushed as far as possible, that is, until the graph of the value of waiting—the curved line given by BR^{β}—becomes tangential to that of the straight-line return of investing immediately: $R/\beta - K$. Thus, the choice of an optimal trigger H is defined by the

Figure 2
Optimal Investment Policy

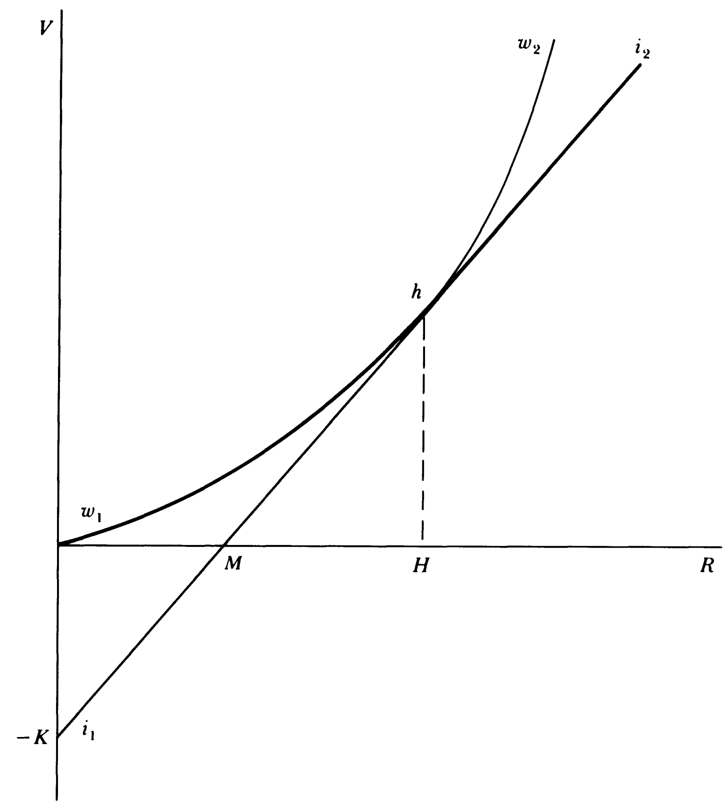

requirement that the graphs of the two formulas of the expression (2) should meet tangentially at H. This is called the "smooth pasting" condition.

Figure 2 shows the optimum H. The corresponding $V(R)$ function is drawn thicker, with the convex curve w_1h of the value of waiting to the left of H, and the straight line hi_2 of the net worth of the project to the right of that point. The Marshallian trigger M is where the value of investing just becomes positive, that is, where the straight line i_1i_2 crosses the horizontal axis. The optimum trigger H is obviously to the right of this.

The observant reader will have noted that the curve BR^{β} lies above the line $R/\rho - K$ to the right of H, and wondered if this means that investment is optimal only at the point H, and waiting again the preferred policy for higher values of R. The answer is no. The point is that the expression BR^{β} ceases to have a valid interpretation as the value of waiting when $R > H$. Otherwise it would create a pure speculative bubble; the value of waiting would be high because the prospect of reaching an even higher R would offer an even higher value of waiting, with no actual investment ever in sight. In the same way, increasing B even farther to lift the curve BR^β clean above the line $R/\rho - K$ is not a meaningful policy.

A sharper intuition into the relationship between the Marshallian and the optimal triggers for investment can be gleaned with some algebra. The smooth pasting condition equates the slopes of the value of waiting and the value of investing at the optimal trigger H. Therefore we differentiate each formula in (2) with respect to R, evaluate the derivatives at H, and equate the two expressions. This gives

$$\beta B H^{\beta-1} = 1/\rho. \tag{5}$$

Then use equations (4) and (5) to solve for H and eliminate B. We find that the optimal H is given by

$$H = \frac{\beta}{\beta - 1}\rho K. \tag{6}$$

Remember the Marshallian investment trigger M was to invest when $M = \rho K$. Therefore we have a very simple relation between the Marshallian and the optimal triggers: the latter is $\beta/(\beta - 1)$ times the former.

We can express the optimal trigger in a way that parallels the Marshallian formula even more closely. Define a new discount or hurdle rate ρ' that incorporates a correction for the value of waiting. Once this correction has been made, one can proceed in the Marshallian way; the project is worth undertaking when its net worth calculated using the corrected discount rate becomes positive. For this, we need $H = \rho' K$, or

$$\rho' = \frac{\beta}{\beta - 1}\rho. \tag{7}$$

This formulation will be applied in the next section to develop some estimates of the potential difference between the optimal and Marshallian trigger.

The reader should also be able to make greater intuitive sense of equation (4) at this point. Instead of correcting the discount rate for the value of waiting, we could correct the cost of investment. Immediate action has an opportunity cost, namely loss of the option to wait. This is valued at BR^β, and we must add it to the actual cost of investment K to get the full cost of immediate action. Then such action is justified when the benefit R/ρ exceeds this full cost. As (4) shows, this happens when the current revenue R reaches the trigger H.

Readers who happen to be familiar with elementary concepts of financial options—from theoretical study or their own practical investment experience—can sharpen their intuition by exploiting an analogy with financial options.[8]

[8]Other interested readers can find the basic concepts explained in the Symposium on Arbitrage in the Fall 1987 issue of this journal.

The opportunity to make a real investment is akin to an American call option—a right but not an obligation to buy a stock at a preset price called the strike price or exercise price. For the real investment project I am considering, the exercise price is the sunk cost K of the project. If the option is exercised, the firm acquires ownership of a stock that pays a dividend stream of expected present value R/ρ. The net worth, $R/\rho - K$, is called the "intrinsic value" of the option. But exercising the option at the instant its intrinsic value becomes positive is not optimal, because the option also has a value of waiting, called the "holding premium" or "time value." One should wait until the holding premium falls to zero. The "smooth pasting" condition that helps determine the optimal point of exercise has long been known in the theory of financial options. In fact, option pricing theory can claim credit for developing this condition that is now standard in the general theory of control of Brownian motion (Merton, 1973, fn. 60).[9]

The Importance of Option Values

Is the difference between the Marshallian trigger M and the optimal trigger H, or equivalently, the difference between the conventional discount rate ρ and the modified discount rate ρ', quantitatively so large that economists should alter our orthodox views on investment and rewrite our textbooks? Of course, the answer depends on the parameters. If the uncertainty is low, there can be only little value in waiting. If the uncertainty is high, on the other hand, setting a high trigger before taking action may avoid some very bad outcomes. If the discount rate is low, the future is valued relatively more and options that help avoid bad future outcomes become more valuable. Here are some sample calculations to show that for plausible parameter values the effect can be very large indeed.

For export projects whose revenues fluctuate with exchange rates, a coefficient of variation of 10 percent over one year fits the recent experience of exchange rate volatility (Frankel and Meese, 1987). If the project is an oil well or a copper mine, a much higher figure of 25 to 40 percent per year is closer to the experience of fluctuations in the prices of these resources (Brennan and Schwartz, 1985). Therefore, let us use a value in this range, say $\sigma = 0.2$, as a base case. Suppose the discount rate is 5 percent per year. Then we find $\beta = 2.15$, and the multiple $\beta/(\beta - 1)$ equals 1.86. Thus current revenues have to rise to nearly double the level that ensures a positive net worth before waiting ceases to be optimal. Using the alternative method of adjusting the discount rate, we find $\rho' = 9.3$ percent, which is quite a big correction to $\rho =$ 5 percent.

[9]For experts in financial economics, I should clarify that the project of my example is an option with an infinite expiry date; the finite-horizon case is treated by McDonald and Siegel (1986). Also, the stock (project) pays dividend (revenue flow); that is why exercise before the expiry date can be optimal.

For a more general sense of how the underlying parameters affect β and ρ', note from the definition of β given earlier in equation (3) that a lower discount rate ρ or a higher standard deviation σ of the revenues yield a lower β. In turn, a smaller β means a larger factor $\beta/(\beta - 1)$, and therefore the longer it is optimal to wait.

It is intuitively evident that when the future is less heavily discounted, the value of waiting for more information goes up. As an example, if in the above calculation we reduce ρ to 2 percent, which is closer to historic riskless real rates of interest, then β drops to 1.62, and the multiple $\beta/(\beta - 1)$ rises to 2.61. It is equally intuitive that greater uncertainty means a higher value of waiting. If in the numerical example we raise σ to 0.4 (while keeping ρ at 5 percent), then $\beta = 1.43$, $\rho' = 16.6$ percent and H is 3.32 times M.

Two limiting cases are worth mention. If the future is very heavily discounted (ρ large) or very certain (σ small), then β goes to infinity and $\beta/(\beta - 1)$ goes to 1. Option values become unimportant in this limit and the Marshallian criterion applies. In the opposite extreme, as ρ goes to 0 or σ goes to infinity, β goes to 1 and $\beta/(\beta - 1)$ goes to infinity; Marshallian analysis becomes totally misleading.[10]

To sum up, even when the cost of capital is as low as 5 percent per year, the value of waiting can quite easily lead to adjusted hurdle rates of 10 to 15 percent. Summers' (1987) finding of median hurdle rates of 15 percent is no longer a puzzle.

Extensions and Qualifications

The above example of a discrete investment project was deliberately oversimplified to highlight the value of waiting. In practical applications, of course, various complications and countervailing considerations must be recognized. Here I shall briefly outline some important matters of this kind.

The example can readily be generalized in many respects, and the essential lesson of the importance of the value of waiting survives unscathed. We can allow the scale of the initial investment to be a matter of choice, introduce some

[10]I have set the trend growth rate of revenues at zero in the above analysis, but this is a good place to mention its twofold effect. On the one hand, a faster expected rate of future revenue growth makes investment more attractive. If the trend rate of growth is μ, then the discounted present value of future revenues starting at R is $R/(\rho - \mu)$, so the Marshallian trigger is $M = (\rho - \mu)K$. On the other, the consequences of a given difference in the current revenue level become magnified as time goes on. Therefore the value of avoiding a given amount of downside risk increases, and with it the value of waiting. The adjusted discount rate ρ' is explicitly derived in the Appendix. The adjusted discount rate is given by $\rho' - \mu = [\beta/(\beta - 1)](\rho - \mu)$, and the optimal trigger is $H = (\rho' - \mu)K$. Now μ also affects β, and an increase in μ lowers β; this is the waiting effect. Numerical calculations show that the waiting effect generally wins. For example, if we keep the basic values $\rho = 5$ percent and $\sigma = 0.2$, but raise the trend growth rate from zero to 2 percent per year, then β falls from 2.15 to 1.58, and ρ' rises from 9.3 to 13.6 percent per year, and H is 3.9 times M.

cost to varying this scale, and also allow a choice of the level of operation at each instant by varying labor or other inputs. If the net revenues can sometimes become negative, we can allow temporary suspension or abandonment. This last extension raises some interesting new issues, and is considered in the next section.

The assumption that net revenues follow a Brownian motion embodies a restriction: uncertainty is roughly symmetric around the trend. In practice, distributions of future outcomes are sometimes quite lopsided. Therefore, we should know what they do to investment decisions.

Bernanke (1983) found the answer and called it the *bad news principle*: "of possible future outcomes, only the unfavorable ones have a bearing on the current propensity to undertake a given project." In other words, the downside risk is the primary force governing optimal investment decisions when waiting is possible. To be more precise, the total upside probability matters, but not the shape of the distribution of revenues to the right of the optimal trigger. The technical proof of this is in an additional Appendix available from the author. But the intuition is not difficult. Remember that when we decide whether to proceed or to wait a little more, what is at stake is not the uncertainty per se, but how it will resolve in the next small amount of time; that is, the tradeoff between current and future risks. Most of the upside potential remains whether action occurs right away, or after a small delay. The possibility of a downturn, and the ability to avoid an action that could thereby prove to be a mistake, is what makes waiting valuable. That is why the downside risk matters most when deciding whether to wait.

Let us turn to some other issues that were left out of the model. In the simple example, waiting had a positive value because it allowed further observations of the revenue fluctuations. More generally, the point is that the passage of time reveals more information. In reality there are other forces that also bear on the issue of whether to wait.

First, there may be a race to seize a scarce opportunity. In the simple example, the opportunity to invest was assigned to just one firm. But if it is available to any of several firms, then waiting is no longer feasible. The option to wait will expire because some competitor will seize the opportunity. Then some firm will invest as soon as the expected present value crosses zero, and the Marshallian trigger will be valid.

But when there are several firms, a more interesting scenario is that more than one firm can invest. When they do so, industry supply increases and price falls along the demand curve. This, or the expectation of such price fall, places a limit on the equilibrium investment. In other words, we have a competitive industry in a dynamic environment. This is the natural setting in which to explore the validity of the Marshallian story of entry at long run average cost and exit at short run variable cost. I consider this in the later section titled "Competitive Industry Dynamics." The outcome of the correct melding of each firm's choice of waiting or investing into a dynamic equilibrium process turns out to be quite far from the Marshallian picture.

Second, there are strategic situations where making the first move has a commitment value. The role of investment in altering the outcome of a Cournot oligopoly is well known; see Dixit (1980), and a richer dynamic version in Fudenberg and Tirole (1983). In practice, strategic considerations may call for early investment at the same time that information aspects suggest waiting; the optimal choice then has to balance the two.

Finally, when there are several firms, their information may differ. Then each firm has to consider how it can infer other firms' information from their actions, and in turn how its information may leak to others through its actions. For example, suppose each firm independently evaluates the prospects of a project, and the evaluation is subject to error. If one firm observes that no other firm has invested, it infers that their evaluations were insufficiently favorable, and adjusts its own evaluation downward. When all firms do this, they may all decide to wait. Conversely, once one firm invests, others conclude that its evaluation must have been very strongly favorable, and adjust their own judgments upward. Therefore the first firm may quickly be followed by others, resulting in a bunching of investment. For discussions of such matters, see Stiglitz (1989) and Leahy (1990, Chapter 3).

Abandonment and Hysteresis

In the simple example of the previous section, the net revenue flow from the project was always positive; therefore, there was no reason to suspend or abandon a project once it was launched. In reality, we see firms suffering operating losses. To capture this, let R now be *gross* revenue, and introduce a flow cost C of operation. Let R follow a geometric Brownian motion. For simplicity of exposition, suppose R has zero trend, and that C is constant. This does not alter the qualitative results; the former assumption is relaxed in the Appendix.

If temporary suspension of operation is possible, this will be done whenever R falls below C. This looks like the textbook Marshallian theory: disinvestment should take place when operating losses are being made. But suspension is not disinvestment. Most typically, if a firm ceases operation, it cannot restart at will without incurring some further cost. It is as if the machinery rusts when unused. To highlight this feature, I shall suppose that rusting is total and immediate. Then, suspension is the same as outright abandonment. If one ever wants to restart in the future, the whole sunk cost K must be incurred over again.[11]

[11]For an analysis of the case where temporary suspension is possible, see McDonald and Siegel (1985).

The possibility of waiting now influences the decision to abandon, just as it affected the decision to invest. The gross revenue R has to fall some way below the operating cost C before abandonment becomes optimal. The intuition is similar to that for investment. The investor is willing to tolerate some operating loss to keep alive the option of future profitable operation should R turn upward. Only when the current loss exceeds the value of the option does it pay to abandon. Let L be the critical low value of revenue that just triggers abandonment when option value is taken into account properly. Then our intuitive reasoning says that L must be less than C.

The expected present value of operating profits when the current revenue flow is R equals $(R - C)/\rho$. The Marshallian investment criterion would tell the firm to go ahead when this exceeds K, that is, when current revenue exceeds the trigger level $M = C + \rho K$. The right-hand side is just Marshall's long run cost, being the sum of the variable cost and the interest on the sunk cost. When the option value of waiting to invest is recognized, the trigger H is higher. Thus we have the chain of inequalities

$$L < C < C + \rho K < H. \tag{8}$$

In the earlier simple example, we could explicitly show the solution for the investment trigger H. But here, L and H are determined by a more complicated system of nonlinear equations that does not permit a closed-form solution. The derivation of this system is sketched in the Appendix; more details are in Dixit (1989b). Here I shall merely mention some numerical results.

First, let us estimate the quantitative significance of the option values. Let $\rho = 0.05$ and $\sigma = 0.2$ as in the base case of the previous section. Choose units of account so that $C = 1$, and suppose that $K = 2$. Then the normal return to capital or the interest on sunk costs is $\rho K = 0.1$. The long run average cost, or the Marshallian investment trigger, is $M = 1.1$. With these values we find $L = 0.72$ and $H = 1.62$. At the truly optimal entry trigger H, the operating profit is 0.62, which is more than six times the normal return to capital. At the exit trigger L, losses equal to nearly a third of variable costs are being sustained. Once again the departure from Marshallian theory is very dramatic for quite plausible values of the parameters. Dixit (1989b) considers a wide range of parameter values and finds similar results.

It is important to recognize that the triggers L and H are jointly determined by all the parameters of the problem. An increase in the sunk cost K will obviously raise the investment trigger H. But it will also lower the abandonment trigger L; the project will be continued through periods of greater losses for the option of keeping alive the larger sunk stake. Conversely, if abandonment is costly—for example, severance payments to workers or the cost of restoring the site of a mine—then the entry trigger is higher; firms are more cautious in undertaking a venture they may have to abandon later at a cost.

Optimal Inertia

Sunk costs alone will produce a zone of inaction between the two Marshallian triggers of the variable cost C and the total cost $M = C + \rho K$. If the current revenue flow is between these levels, then the optimal policy is to maintain the status quo. The project is not launched, but if already active, is not canceled.

But with uncertainty, the zone of inaction that takes option values into account is wider, expanding to between the triggers L and H. The numerical calculations show how big the gap can be. In the example just above, the Marshallian range of inertia extends from 1 to 1.1, while the optimal range goes from 0.72 to 1.62, which is quite a dramatic difference. Other plausible values of the parameters ρ, σ and so on have equally substantial effects. Dynamic economic choices should exhibit much greater inertia when there is uncertainty.

This has potentially important implications for macroeconomics. Small nominal or real frictions can produce even larger rigidities than those suggested by models that ignore evolving information, for example the "menu cost" models of Mankiw (1985) and Akerlof and Yellen (1985) or the "portfolio" model of Greenwald and Stiglitz (1989).

The implications for labor markets are also potentially dramatic. Tangible costs of hiring and firing workers are significant in almost all occupations, and quite large in some countries. If wages are sticky, then the response of employment to output demand fluctuations will be slower when employers, recognizing the option value of the status quo, hoard labor in downturns, and are slow to hire in upturns. Alternatively, very large wage fluctuations will be needed to maintain classical full employment.

If wages are sticky and employment responds slowly, the marginal product of labor may go quite some way above the wage without any hiring taking place, and below it without any firing (Bentolila and Bertola, 1990). Contrary to conventional theory, the wage in any occupation is not constantly equated to opportunity cost of labor. Economists usually dismiss the popular concern about a "loss of jobs" by invoking just such an equation: the person out of a job only ceases to earn in this occupation just about what he or she could have earned elsewhere. The view presented here suggests that the popular concern may have more justification.

Hysteresis

Picture a particular path of the stochastic evolution of net revenues through time. Let the numerical values be as above. Suppose the initial R equals 1, and it starts to rise. It crosses the Marshallian trigger of 1.1, but no investment takes place. Finally it rises above 1.62, and the project is launched. Then the revenue starts to fall, and comes back all the way down to 1. But this does not justify abandonment. The driving force behind the investment decision, namely the currently observed revenue, has been restored to its initial level. But its meandering along the way has left its mark, namely an active project where there was none before.

Similar effects have long been known in physics and other sciences. The closest for our purpose comes from electromagnetism. Take an iron bar and 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 left some lasting effect (the magnetized bar). This phenomenon is called hysteresis, and by analogy the failure of investment decisions to reverse themselves when the underlying causes are fully reversed can be called economic hysteresis.

If some electric current is passed through the wire in the opposite direction, the residual magnetism will be lost. With a strong enough opposing current, magnetism will be induced in the reverse direction. Similarly, if our project's current revenue falls even more, it will eventually be abandoned. Then a subsequent rise back to 1 in revenues will not restore the project; there is hysteresis in the reverse direction, too.

Sunk costs alone can cause hysteresis in textbook Marshallian analysis, as R moves in and out of the Marshallian zone of inaction between variable and total costs. But if such fluctuations are occurring, it behooves us to let the firm have rational expectations about its stochastic environment. When that is done, the uncertainty magnifies the effect quite dramatically; very large changes in R in the opposite direction are needed to reverse the effects of a temporary move in either direction.

In this light, the slowness of the U.S. imports to respond to the dollar appreciation of the early 1980s, and the even greater slowness to improve despite the subsequent fall back to the 1980 level, become quite understandable and even intuitive. Krugman (1989, Chapter 2) and Dixit (1989a) discuss this case in greater detail.

U.S. v.s. Japan

Observers of America's relative decline in manufacturing—for example, Dertouzas et al. (1990, pp. 61–65)—attribute part of the problem to the dominance of short-term thinking among U.S. managers, which causes them to apply high "hurdle" rates when considering investment decisions. The explanations given for this short-term emphasis include fear of hostile takeovers, high mobility of managers, and various kinds of uncertainty including that in the government's taxation, regulation, and trade policies. Our analysis of the effect of option values on investment suggests that uncertainty is even more important than previously realized. It can explain much or even all of the gap between typical hurdle rates and the cost of capital. The high rates might actually be optimal responses to uncertainty.

But the explanation is inadequate as it stands. The option value effect raises the optimal entry point, but by the same token it lowers the optimal exit point. If American firms are more hesitant to invest or enter new ventures because of uncertainty, they should be more ready to ride out bad periods. But the same observers find exactly the opposite tendency. In many sectors including color TVs, VCRs, and semiconductors, American firms have abandoned the field after short periods of losses, while Japanese firms hang in there (Dertouzas et al., 1990, especially Industry Studies C and F).

A common explanation of Japanese firms' willingness to absorb losses is the lifetime employment system. This makes labor a quasi-fixed factor, and reduces the variable component of cost. Conventional theory points out that lower variable costs mean that revenues must fall farther to cause abandonment. But by the same token, these larger sunk costs of Japanese firms should make them more reluctant investors! Reality is the opposite; they are particularly aggressive investors.

A better explanation may be found in Bernanke's bad news principle. Suppose the uncertainty facing Japanese firms is more lopsided; they are protected from the downside risk because the government supports them in various ways, including cartelization to avoid destructive competition in recessions. Then the value of waiting to invest, which is governed mainly by the downside risk, is quite small, and they invest more aggressively. For disinvestment, of course, the argument turns around and becomes the good news principle. The option value of keeping the operation alive is governed primarily by the upside potential, which is relatively more important for Japanese firms, and induces them to ride out bad periods that would drive American firms into dissolution.

This analysis is not intended to be comprehensive, but it does allow some conditional statements on policy. If the aim is to induce quicker new investment or entry of firms, it is especially important to reduce the downside risk. If the aim is to prevent disinvestment or exit, it is especially important to improve the upside potential.

Competitive Industry Dynamics

The analysis so far has dealt with a single project or a single firm, taking its revenues as exogenously determined. Now consider what happens to an industry populated by many active price-taking firms, and identical potential entrants. Each takes the price as evolving exogenously over time, albeit with some uncertainty. But the actions of all of them in turn determine the price path. What will be the overall equilibrium of this process?

Consider a simple structure that serves to bring out the essential points. The source of the uncertainty must now be something exogenous to all the firms; I shall suppose this to be a demand shock. Specifically, suppose an inverse demand curve for the industry, expressing price P as a function of quantity Q and the shock to demand X; assume this takes the algebraic form

Figure 3
Competitive Industry Dynamics

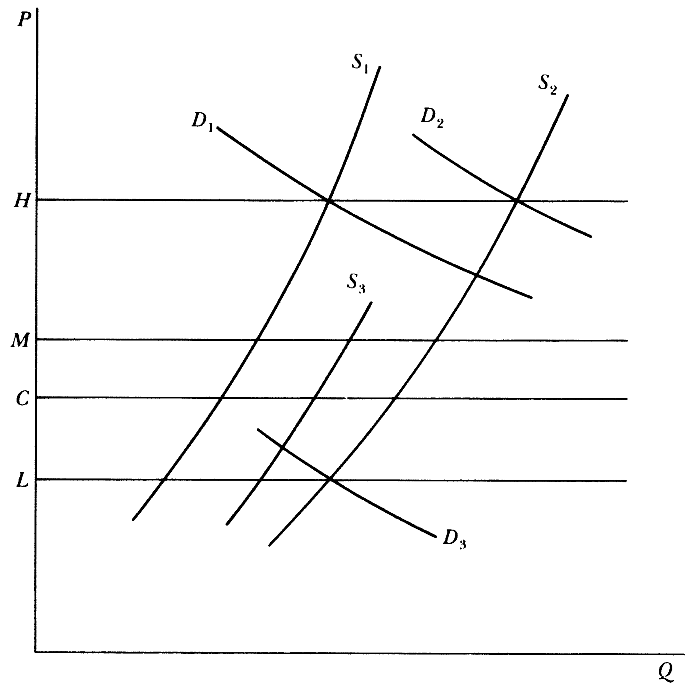

$P = XD(Q)$. Each firm has a very Marshallian technology. It becomes active by making an initial sunk investment of K. While active, it has a standard rising short run marginal cost curve that becomes its supply curve. Temporary suspension of operations is again assumed away. The industry supply curve at an instant is found by the usual horizontal summation of the supply curves of all active firms. Write C for the minimum short run average variable cost and M for the minimum long run average cost.

Figure 3 enables us to trace out the dynamics of the industry. Suppose the firms that are originally active generate the industry supply curve S_1. Fluctuations in the demand shock variable X will induce movements along S_1, leading to fluctuations in the price P. Suppose new entry is triggered when the price rises to the critical level H. Along S_1, this happens when demand rises to the position D_1. When some new firms enter, the supply curve shifts to the right; let S_2 be its new position. The price falls along the demand curve D_1. Thereafter, new demand shocks will cause movements along S_2, until either the critical high price H is reached again, triggering more entry, or a critical low price L is reached, triggering some exit. In the figure, the former occurs if the demand curve rises to the position D_2, and the latter if the demand curve falls to the position D_3. For the whole range of demand curves between these extremes, the number of firms stays unchanged and the price and quantity fluctuate with the shocks to demand. If demand hits the upper limit, further entry shifts the supply curve to the right of S_2; if it hits the lower limit, the supply curve shifts to the left to a position like S_3. The process then goes on.[12]

Consider the picture from the perspective of a potential firm. It knows that new entry will prevent the price from ever rising above H, exit will stop it from falling below L, and within this range the price will fluctuate as demand shocks evolve. Given this price process, it must choose its own entry and exit strategies, which will consist of an entry trigger H' and an exit trigger L'. For the industry's equilibrium with rational expectations and identical firms, the entry and exit triggers chosen by each firm should be the same as the ceiling and floor on the price process that each assumes in arriving at its optimal decision. That is, $H' = H$ and $L' = L$.

Determining such an equilibrium requires some mathematics (Lippman and Rumelt, 1985; Edleson and Osband, 1989; Leahy, 1990; Dixit, 1991b). But one important general property is easy to see: the entry trigger will allow supernormal profit, and the exit trigger will allow some operating losses. The reason is evident when we consider the net worth of a firm contemplating entry or exit. Suppose, contrary to the above assertion, that the equilibrium ceiling price coincides with the Marshallian long run average cost M. Now, each firm knows that new entry will prevent the price from ever rising above this level, but adverse demand shocks can drive the price below M from time to time. Then each firm expects that the operating profit will never be more than normal, but can be less at times. Therefore the expected average return to capital must be below normal, and net worth must be negative. Firms will not enter under such circumstances. The investment trigger H must exceed the long run average cost, to give firms the prospect of periods of supernormal returns mixed with periods of subnormal ones. The equilibrium level of H must be such as to ensure exactly normal average return, or zero net worth, for a potential investor.

Similarly, if the equilibrium floor price were the Marshallian average variable cost C, then firms at this point would see non-negative operating profit at all future dates, and positive at some dates. This does not call for exit. The equilibrium trigger for exit L must be sufficiently below C to make the consequences of staying in, namely operating losses for some periods and profits for other periods, average out to exactly zero net worth, and make each firm indifferent between staying and leaving.

The ranking of L, C, M and H in the industry's equilibrium is the same as the earlier ranking (8) when there was just one monopoly firm. Therefore we

[12]Actually there is a further subtlety when the demand shocks follow Brownian motion. Time being continuously variable for this process, a little entry quickly drops the price slightly below H. From there, the probability of demand rising to take the price to H again in a short interval is quite high. That induces some further entry, and so on. In other words, once the price hits the entry trigger, it is quite likely to keep on bouncing close to this level for a while, with gradual entry and increase in quantity.

should explore the connection between the above reasoning based on zero net worth, and the earlier analysis based on the value of waiting. When the entry of new firms places a ceiling at H on the price, this cuts off each firm's upside profit potential. That reduces the value of investing and waiting alike. In equilibrium, the reductions are such that the pure waiting value for each firm is zero, as it should be when there is a potential infinity of identical entrants and no firm has any scarce privilege. Likewise, the floor L cuts off the downside risk, and raises the value of exiting and staying alike. In fact, when the demand shock follows a geometric Brownian motion, the changes in the values of investing and waiting are equal, and each firm's entry and exit triggers L and H are exactly the same as they would have been if it were the only firm in the market, and therefore not subject to the ceilings and floors that result from other firms' entry and exit (Leahy, 1991).

Therefore we can take the previous numerical calculations for individual firms and apply them to the industry equilibrium. We see that the no-entry-no-exit range of prices in a competitive industry is likely to be quite wide. Remember that in our base case, the Marshallian range of inaction extended from 1 to 1.1, and the one that accounts for option values went from 0.72 to 1.62. Therefore we are likely to see significant periods of supernormal profits with no new entry, and of operating losses without exit, in the course of a competitive industry's equilibrium evolution.

This picture calls for a very fundamental rethinking, particularly in the matter of regulation and other industry policies. It is most important to regard the equilibrium of an industry as an organic process over time. Drawing inferences from snapshots at particular instants can be seriously misleading.

Suppose we observe such an industry at an instant when the price is between the Marshallian long run average cost M and the entry trigger H. We see established firms making supernormal profits, but no new entry taking place. Given our training in conventional microeconomics or industrial organization theory, we suspect the presence of monopoly power or entry barriers. We might be inclined to suggest antitrust action. But we would be wrong; the process viewed as a whole is fully competitive, and long run expected returns are normal.

Likewise, if the price is below the minimum average variable cost, that need not signal predatory dumping by the firms that are making the losses. They may merely and rationally be riding out the bad period to keep their sunk capital alive.

Only by observing the evolution of the industry for a long time can we spot genuine departures from the competitive norm. Basing policies on snapshots can result in harm despite the policy-maker's best intentions. In the picture above, temporary large profits are merely due to the swings of demand in a competitive industry that permits only normal profit as a long run average. If the government tries to reduce these supernormal returns using antitrust action or price-ceilings, this merely depresses new entry to suboptimal levels.

The resultant reduction of supply can actually raise the long run average price; see Dixit (1991c). Similarly, if the government pursues policies to support firms in bad periods, firms will anticipate this, leading to additional entry, which will aggravate the losses when bad times arrive.

Non-Economic Applications

Many personal, social, and political decisions that are not narrowly economic share the features of investment: they are costly to reverse, they must be made in an uncertain environment, and their timing is a matter of choice. Therefore option values and optimal inertia are significant for them, too. Quantification is a lot harder, but useful qualitative lessons can be drawn. I share the economist's usual temptation to indulge in amateur sociology, and will offer two quick examples.

The first is a dramatic example of the value of waiting. Hamermesh and Soss (1974) offer a Beckerian model of suicide. This involves comparing the expected utility of the rest of one's life and a suitable standard of zero. Far from being an excessively rational model, this is not rational enough, because it leaves out the option value of staying alive. We have seen that exit decisions are governed by Bernanke's good news principle: the prospect of an upturn is the primary determinant of the option value. As Micawber said in *David Copperfield*, "something will be turning up."

The second example shows how other considerations may offset the value of waiting. Once upon a time in New York City, there lived an Assistant Professor of Finance. He and his "spouse-equivalent" had separate rent-controlled apartments. Their relationship progressed to a point when the woman suggested that they should keep one of the apartments and give up the other. He explained to her the importance of keeping options alive: it was unlikely that they would split up, but given a positive probability, and so on. She took this very badly and ended the relationship.

Financial economists who hear this story say that it just proves how right the man was about option values. But the economics of information offers a more convincing explanation. The man misunderstood the situation. This was not a decision problem under uncertainty, but a signalling game. The woman was unsure how highly he valued her, and it was precisely his willingness to undertake the costly irreversible action of giving up the apartment that had value as a signal.[13] The man overlooked this, tried to sit on the fence, and fell flat on his face.

[13]Barry Nalebuff suggested a more fully game-theoretic resolution: the man knew what game was being played, and *meant* to send the signal that the woman correctly interpreted. But in my judgment, if he had wanted to convey such a signal, he could have done so in many other and more unambiguous ways.

Further Reading

For readers whose appetite is whetted by this sampling of new ideas on investment under uncertainty and their applications, the next stop should be Pindyck (1991). He gives a full survey of the literature, to which readers can then turn for even more details. Pindyck also develops the analogy with financial options particularly thoroughly.

Readers sufficiently intrigued to attempt their own research in this area must make a little investment in techniques. I hope they have learned the basic lesson of the story, and will be duly cautious in making this decision. For those who decide to go ahead, Dixit (1991a) provides a relatively gentle introduction aimed at economists, and gives references to rigorous but harder mathematical treatises.

■ *I am very grateful to the editors—Carl Shapiro, Joseph Stiglitz, and Timothy Taylor—for their excellent suggestions to improve the exposition. I also thank Ben Bernanke, Alan Blinder, Gene Grossman, James Hines, Eric Rasmusen, Lars Svensson, Timothy Taylor, and Christopher Williams for comments on earlier drafts, and the National Science Foundation for financial support under Grant No. SES-8803300.*

Appendix
Value of Waiting

Here I derive the expression for the value of waiting. I shall allow the net revenue R to have a non-zero trend growth rate μ. The proportional variance per unit time is σ^2. Consider the opportunity to invest as an asset that is held for a small interval of time dt while R stays below H. It pays no dividend, but its value $V(R)$ changes with R, so it may make a capital gain or loss. The change dR in R over this interval is random, with mean and variance

$$\mathrm{E}[dR] = \mu R\, dt, \qquad \mathrm{Var}[dR] = \sigma^2 R^2\, dt.$$

Then

$$\mathrm{E}[dR^2] = (\mathrm{E}[dR])^2 + \mathrm{Var}[dR] = \mu^2 R^2\, dt^2 + \sigma^2 R^2\, dt.$$

Therefore the expected capital gain is

$$\begin{aligned}\mathrm{E}[dV] &= V'(R)\mathrm{E}[dR] + \tfrac{1}{2}V''(R)\mathrm{E}[dR^2] \\ &= V'(R)\mu R\, dt + \tfrac{1}{2}V''(R)\left[\mu^2 R^2\, dt^2 + \sigma^2 R^2\, dt\right].\end{aligned}$$

In equilibrium, this should equal the normal return $\rho V dt$. Writing this equation, dividing by dt, and letting dt go to zero, we get the differential equation

$$\tfrac{1}{2}\sigma^2 R^2 V''(R) + \mu R V'(R) - \rho V(R) = 0. \tag{A.1}$$

This is a simple equation of the Cauchy–Euler type. Try a solution of the form $V(R) = R^x$. By substituting in (A.1), x must satisfy the associated quadratic equation

$$\tfrac{1}{2}\sigma^2 x(x-1) + \mu x - \rho = 0. \tag{A.2}$$

The left-hand side of (A.2) is negative at $x = 0$ when $\rho > 0$. It is also negative when $x = 1$ provided $\rho > \mu$, which we assume to ensure convergence of the expected discounted present value of the revenues. Then one root of (A.2) is negative, and the other exceeds 1. Call them α and β respectively. Then the general solution of (A.1) is

$$V(R) = AR^\alpha + BR^\beta, \tag{A.3}$$

where A and B are constants to be determined.

The value of waiting should go to zero as R goes to zero. Since α is negative, we see from (A.3) that we must have $A = 0$. That leaves $V(R) = BR^\beta$, which is just the expression in equation (2) of the text.

In the text I assumed $\mu = 0$. Then (A.2) becomes

$$x(x-1) = 2\rho/\sigma^2, \quad \text{or} \quad \left(x - \tfrac{1}{2}\right)^2 = \left[1 + 8\rho/\sigma^2\right]/4.$$

This immediately gives equation (4) of the text for β.

Adjusted Discount Rates

When $\mu \neq 0$, the future revenue stream is expected to grow at rate μ and is discounted back at rate ρ, so its expected present value is $R/(\rho - \mu)$. The net worth of investing is $R/(\rho - \mu) - K$. The Marshallian trigger is $M = (\rho - \mu)K$. Otherwise the calculation proceeds as in the text, and the optimal trigger H is still $\beta/(\beta - 1)$ times M. Therefore the adjusted discount rate ρ' is defined by

$$\rho' - \mu = \frac{\beta}{\beta - 1}(\rho - \mu).$$

Investment and Abandonment

Here it proves convenient to label the value of waiting and the value of a live project separately as functions of R; call them $V_0(R)$ and $V_1(R)$ separately.

The value of waiting still satisfies the same equality of capital gain and normal return, leading to

$$V_0(R) = B_0 R^{\beta}, \tag{A.4}$$

where the constant B_0 is to be determined. The value of an active project now includes a value of the option to abandon. To find it, follow the same steps as above, but note that the asset now pays a dividend, namely the revenue flow R. The normal return $\rho V_1(R)\,dt$ should now equal the sum of the dividend $R\,dt$ and the expected capital gain $\mathrm{E}[dV_1]$. This leads to the differential equation

$$\tfrac{1}{2}\sigma^2 R^2 V_1''(R) + \mu R V_1'(R) - \rho V_1(R) + R = 0. \tag{A.5}$$

The general solution now includes a term corresponding to a particular solution of the non-homogeneous equation. Trying a linear form kR, we find that k must equal $1/(\rho - \mu)$. Therefore

$$V_1(R) = R/(\rho - \mu) + A_1 R^{\alpha} + B_1 R^{\beta}.$$

The first term is just the expected present value of revenues; the rest is the value of the option to abandon. This option is very far from being exercised if R goes to infinity, so the option value should go to zero there. For that we need $B_1 = 0$, leaving us with

$$V_1(R) = R/(\rho - \mu) + A_1 R^{\alpha}. \tag{A.6}$$

At the investment trigger H the increase in value upon investing should equal the cost of investing:

$$V_1(H) - V_0(H) = K, \tag{A.7}$$

and the two value functions should meet tangentially, leading to the "smooth pasting condition"

$$V_1'(H) - V_0(H) = 0. \tag{A.8}$$

Similarly, at the abandonment trigger L, we have

$$V_1(L) - V_0(L) = 0; \tag{A.9}$$

if there is a direct cost of abandonment such as severance payment, then minus that will appear on the right-hand side. There is also the smooth pasting condition

$$V_1'(L) - V_0'(L) = 0. \tag{A.10}$$

Substituting the functional forms (A.4) and (A.6) into equations (A.7)–(A.10), we get four equations that must be solved simultaneously for the two constants A, B and the two triggers H, L. An explicit analytical solution is not possible, but some properties of the solution can be obtained by analytical methods, see Dixit (1989b). Numerical solution is quite easy using simultaneous nonlinear equation solving routines such as the one available in GAUSS.

References

Akerlof, George, and Janet Yellen, "A Near-Rational Model of the Business Cycle with Wage and Price Inertia," *Quarterly Journal of Economics*, Supplement, 1985, *100*, 823–838.

Bentolila, Samuel, and Giuseppe Bertola, "Firing Costs and Labor Demand: How Bad is Eurosclerosis?", *Review of Economic Studies*, July 1990, *57*, 381–402.

Bernanke, Ben S., "Irreversibility, Uncertainty, and Cyclical Investment," *Quarterly Journal of Economics*, February 1983, *98*, 85–106.

Brennan, Michael J., and Eduardo S. Schwartz, "Evaluating Natural Resource Investments," *Journal of Business*, April 1985, *58*, 135–157.

Dertouzas, Michael, Richard K. Lester, and Robert M. Solow, *Made in America*, New York: Harper Paperback Edition, 1990.

Dixit, Avinash, "The Role of Investment in Entry Deterrence," *Economic Journal*, March 1980, *90*, 95–106.

Dixit, Avinash, "Hysteresis, Import Penetration, and Exchange Rate Passthrough," *Quarterly Journal of Economics*, May 1989 (a), *104*, 205–228.

Dixit, Avinash, "Entry and Exit Decisions Under Uncertainty," *Journal of Political Economy*, June 1989 (b), *97*, 620–638.

Dixit, Avinash, "The Art of Smooth Pasting," working paper, Princeton University, 1991 (a).

Dixit, Avinash, "Irreversible Investment and Competition Under Uncertainty," working paper, Princeton University, 1991 (b).

Dixit, Avinash, "Irreversible Investment with Price Ceilings," *Journal of Political Economy*, June 1991(c), *99*, 541–557.

Edleson, Michael, and Kent Osband, "Competitive Markets with Irreversible Investment," working paper, Rand Corporation, 1989.

Frankel, Jeffrey, and Kenneth Meese, "Are Exchange Rates Too Variable?" In Fischer, Stanley, ed., *NBER Macroeconomics Annual*, Volume II. Cambridge: MIT Press, 1987, 117–153.

Fudenberg, Drew, and Jean Tirole, "Capital as Commitment: Strategic Investment to Deter Mobility," *Journal of Economic Theory*, December 1983, *31*, 227–250.

Greenwald, Bruce, and Joseph E. Stiglitz, "Toward a Theory of Rigidities," *American Economic Review*, Papers and Proceedings, 1989, *79*, 364–369.

Hamermesh, Daniel S., and Neal M. Soss, "An Economic Theory of Suicide," *Journal of Political Economy*, January–February 1974, *82*, 83–90.

Krugman, Paul, *Exchange Rate Instability*, Cambridge: MIT Press, 1989.

Krugman, Paul, and Richard Baldwin, "The Persistence of the U.S. Trade Deficit," *Brookings Papers on Economic Activity*, 1987, No. 1, 1–44.

Leahy, John, *Aggregation and Adjustment Costs*, Doctoral Dissertation, Princeton University, 1990.

Leahy, John, "The Optimality of Myopic Behavior in a Competitive Model of Entry and Exit," working paper, Harvard University, 1991.

Lippman, Stephen A., and R. P. Rumelt, "Demand Uncertainty and Investment In Industry-Specific Capital," working paper, UCLA Graduate School of Management, 1985.

Mankiw, N. Gregory, "Small Menu Costs and Large Business Cycles: A Macroeconomic Model," *Quarterly Journal of Economics*, May 1985, *100*, 529–539.

McDonald, Robert, and Daniel Siegel, "Investment and the Valuation of Firms when There Is an Option to Shut Down," *International Economic Review*, June 1985, *26*, 331–349.

McDonald, Robert, and Daniel Siegel, "The Value of Waiting to Invest," *Quarterly Journal of Economics*, November 1986, *101*, 707–727.

Merton, Robert C., "The Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, Spring 1973, *4*, 141–183.

Pindyck, Robert S., "Irreversible Investment, Capacity Choice, and the Value of the Firm," *American Economic Review*, December 1988, *78*, 969–985.

Pindyck, Robert S., "Irreversibility, Uncertainty, and Investment," *Journal of Economic Literature*, September 1991, *26*:3, 1110–1148.

Stiglitz, Joseph E., "Incentives, Information and Organizational Design," *Empirica—Austrian Economic Papers*, 1989, *16*:1, 3–29.

Summers, Lawrence H., "Investment Incentives and the Discounting of Depreciation Allowances," in *The Effects of Taxation on Capital Accumulation*, Martin Feldstein, ed., Chicago: University of Chicago Press, 1987.