

Lecture Notes
on
Information and the Strategic Timing of Investments

(Revised August 2009)

So far, our brief discussion of real options has confined itself to the decisions of a single firm. Although timing decisions are still involved (for example, what is the critical price of oil at which it is optimal to develop an undeveloped oil reserve?) the decisions are “non-strategic” in the sense that we have not worried about taking into account the actions of other firms. Now we will consider some of the strategic issues that arise when other firms may also be making investment decisions that can affect the ultimate market equilibrium and the profits to the first firm. There are two types of problems that can arise, and they have opposite effects on the timing of investment:

1. The market may be large enough to support only one or two firms, or for other reasons (e.g., strong network externalities) there may be a substantial first-mover advantage. This creates an incentive to invest early and preempt your competitors. Although evolving uncertainty over future market conditions will create an option value of waiting rather than investing now, the incentive to preempt may reduce this option value, or even eliminate it entirely.
2. Sometimes the investments of your competitors can yield important information. One example of this is exploration for new undeveloped offshore oil reserves. You and three or four other oil companies might all own leases that give you the right to explore on various different tracts in, say, the Gulf of Mexico. Is it worthwhile to spend a lot of money exploring for oil on your tracts? Because outcomes across tracts are correlated,

you could benefit by waiting to see whether the other companies that explore succeed or fail in finding oil. This creates an incentive to delay investment, even beyond the option value incentive that we have already examined.

1 Pharmaceutical R&D

A nice way to think about these issues is in the context of pharmaceutical R&D. Pharmaceutical R&D is usually directed at the development of new drugs for specific therapeutic categories, e.g., a new drug to reduce cholesterol, treat depression, or attack certain infectious diseases. Typically, several pharmaceutical companies will work on the development of a new drug, all based on the same biochemical mechanism.

When developing a new drug, a pharmaceutical company faces two basic types of uncertainty. First, there is uncertainty over the cost and time involved to complete the R&D, perform the necessary clinical testing, and ultimately gain FDA approval for marketing and selling the drug. Most new drugs, in fact, never even reach the final stage of FDA approval. Second, even if drug development and clinical testing are successful, there is often uncertainty over the size and value of the market. When Merck, for example, developed its protease inhibitor for the treatment of AIDS, it grossly underestimated the number of patients who would take such a drug (and have their insurance companies pay for it).

In addition, both types of strategic timing issues arise with respect to pharmaceutical R&D. First, there is usually a first-mover advantage, so that the first company to complete development and testing of the drug will get the largest share of the market, the second to enter will get a smaller share, and the third and fourth, smaller shares still. (For example, Lilly's Prozac was the first to market in the SSRI antidepressant therapeutic category, and had the largest market share; Pfizer's Zoloft was the second, with a smaller share; and Paxil, the third, had a smaller share still.) This creates an incentive for a pharmaceutical company to accelerate its investment in drug development.

Second, firms can learn from each other's experience about the difficulty of actually completing the R&D and testing, as well as the ultimate size of the market and hence the

value of a new drug. This creates an incentive to delay drug development or invest more slowly: Let the other guy find out how difficult the job is, and then we can decide whether it makes sense to go forward.

Thus there are counteracting incentives. On the one hand, there is an incentive to move quickly in order to preempt and be the first to market, but on the other hand, there is an incentive to wait and learn from the experience of other companies.

How can pharmaceutical companies balance these incentives, and determine the optimal timing of their R&D programs? First, they need to estimate the extent of first-mover advantage. First-mover advantage in pharmaceutical markets arises largely from brand-specific network externalities. Doctors will be more willing to prescribe, and patients more willing to take, a drug with an initially large market share because that share conveys positive information about “accepted practice,” as well as efficacy. Estimating the extent of these externalities is difficult, but recent statistical studies have begun to provide useful information. (Surprisingly, for many therapeutic categories, these externalities—and hence first-mover advantage—are not as large as one might have thought.)

Second, firms need to assess the nature and extent of uncertainty over the completion of R&D programs and over market size assuming a successful completion. This is something that pharmaceutical firms have had to do frequently, so although difficult, it is not at all alien to them. What is more difficult is assessing what they can learn from the activities of their competitors, and how to use that in making a strategic timing decision. We will address aspects of that issue below.

2 Learning from Others

I will assume here that the reader is already familiar with the basic ideas behind the “real options” approach to investment. (Those ideas are set forth in Chapter 2 of A. Dixit and R. Pindyck, *Investment Under Uncertainty*, which is in the 15.013 readings packet, and we will go over them in class.) Here we will consider what can happen if there are *two* firms considering an investment, and investing by one firm can yield information to the other.

Suppose that your firm and another firm are both considering producing digital widgets. The problem is that neither of you knows how much consumers are willing to pay for these widgets. Right now, you think they would pay \$50 per widget, or \$150, with equal probability. You will not be able to find out consumers' willingness to pay until you—or the other firm—actually makes an investment and begins producing the widgets.

Each firm can invest in one unit of capacity at a cost of \$800. Note that at this point there is no value of preemption—once you have made the investment, you can sell one widget per year, whether or not the other firm is also producing and selling widgets. Hence, for each firm, the net present value of investing now is:

$$\text{NPV}_i^{\text{NOW}} = -800 + \sum_{t=0}^{\infty} 100/(1+R)^t \quad (1)$$

Note that \$100 is the *expected* price that consumers would be willing to pay. We will assume that the discount rate, R , is ten percent, so that this NPV becomes $\text{NPV}_i^{\text{NOW}} = -800 + 1100 = \300 .

Suppose that Firm 2 is going to invest now. Should Firm 1 wait a year before deciding whether to invest? If it does wait, it will only invest if it learns that consumers are willing to pay \$150 per widget. Hence, assuming that Firm 2 will indeed invest now, the NPV for Firm 1 from waiting is:

$$\text{NPV}_1^{\text{WAIT}} = \frac{1}{2} \left[-\frac{800}{1.1} + \sum_{t=1}^{\infty} 150/(1.1)^t \right] = \$386 \quad (2)$$

Hence, in this situation it is better to wait.

The problem here is that Firm 2 is thinking the same thing, and would also like to wait for Firm 1. Suppose that as a result, neither firm invests now, and both wait a year, hoping (in vain) that the other firm will invest. If at the end of the year, both firms then go ahead and invest (without the benefit of any knowledge about the market), the NPV calculated as of today will be:

$$\text{NPV}_i^{\text{WAIT}} = 300/1.1 = \$273 \quad (3)$$

We thus have a gaming situation, the payoffs for which are shown below. Note that each firm would like the other one to invest first, but in all likelihood both firms will wait a year,

and in the end gain nothing in the way of new information. Assuming there is no longer any further possibility of waiting, they will then both invest, and be worse off than they would have been had they simply invested now. Also note that if the firms could collude, they would probably agree to toss a coin to see who will invest first. In that case, each firm would have an expected NPV of $\frac{1}{2}(300) + \frac{1}{2}(386) = \343 .

		<u>Firm 2</u>	
		Now	Wait
<u>Firm 1</u>	Now	300, 300	300, 386
	Wait	386, 300	273, 273

Assuming that collusion and agreement to toss a coin is not possible, how would you play this game if you were Firm 1? Note that you are in a *war of attrition*—each firm is hoping that the other will “blink” first. Hence, depending on your expectations regarding the other firm, it may be rational to invest now, or it may be rational to wait.

What you might do in this case is assign a probability to the outcome that Firm 2 will invest now. Suppose that probability is p . Then, your expected NPVs for investing now is:

$$\text{NPV}_1^{\text{NOW}} = 300, \tag{4}$$

and your expected NPV for waiting is:

$$\text{NPV}_1^{\text{WAIT}} = (p)(386) + (1 - p)(273) = 273 + 113p \tag{5}$$

Hence it is better to wait as long as $273 + 113p > 300$, or $p > .24$. The question then becomes, is it reasonable to think that the probability that Firm 2 will invest now is greater than .24? That is a difficult question to answer, because it depends not only on the characteristics of Firm 2’s managers, but also on what those managers are thinking about you and your characteristics.

Of course, there is no reason for this process to stop at the end of one year. Suppose that you and the other firm have both decided to wait a year. At the end of the year, both of you observe that (sadly) no one has invested. You are now in the very same situation. Should you wait another year, or go ahead and invest now? And, of course, if you both wait another year, the process will then repeat itself.

In this war of attrition, it is possible for a very long time to go by with neither firm investing. This is indeed likely, unless some other element is introduced to push the firms towards early investment. Take, for example, the case of undeveloped off-shore oil reserves. The payoffs from developing such reserves (by building platforms and drilling production wells) are very uncertain, because there is uncertainty over costs and over the quantity of oil that can be extracted. Hence there is an incentive for each firm to wait before developing a reserve. Each firm can learn by watching other nearby firms, with no fear of being “preempted,” because the world price of oil is independent of what these firms do. In this case the outside element that encourages early investment is a “relinquishment requirement” imposed by the government—the firms will lose their rights to the oil reserves if they do not begin development within a certain period of time.

Another element that can induce one or both firms to invest early is the fear of preemption. In our widget example above, we assumed that each firm’s demand is independent of what the other firm does. But suppose that by building enough capacity, each firm could gain a first-mover advantage that would reduce the potential returns from investing by the other firm. In this case, the incentive to preempt (and avoid being preempted) might outweigh the incentive to wait for information.

We saw how this might work in the context of pharmaceutical markets. In those markets there are significant network externalities associated with a therapeutic category of drug (i.e., a type of drug), and also with a *brand* of drug within a therapeutic category. However, recent research has suggested that the brand-level network externalities may be weak in many cases. (The strength of any network externalities at the level of a therapeutic category or the level of a specific brand is likely to depend on the particular type of drug.) If brand-specific network externalities were important, this would create an incentive for firms to

invest early and try to preempt. However, if brand-specific network externalities are small or nonexistent, it may be preferable to move slowly and be the second entrant in the market (as was Glaxo when it introduced Zantac). Hence, for strategic planning purposes, it is essential to evaluate the strength of these network externalities.

3 First-Mover Advantage and Investment Timing

Let us now put aside the possibility of learning from the investment activities of others, and simply examine the implications of first-mover advantage. We will use the same simple two-period framework that we used in the previous section. Now, however, we will assume that information is revealed by “nature,” and not by the activities of a competing firm.

Once again, suppose that *two* firms can invest. The first to market can sell one widget per year forever, at a price P . Currently, $P = \$100$; next year P will equal \$50 or \$150 with equal probability. The discount rate is 10 percent.

What about the second firm to market? Suppose that firm can also sell one widget per year forever, but only at a price βP , with $\beta < 1$. The cost of the investment for either firm is \$800. The payoff matrix for the two firms is shown below.

		<u>Firm 2</u>	
		Now	Wait
<u>Firm 1</u>	Now	$300, -800 + 1100\beta$ OR $-800 + 1100\beta, 300$	$300,$ $\max[0, -727 + 1500\beta]$
	Wait	$\max[0, -727 + 1500\beta],$ 300	$386, -727 + 1500\beta$ OR $-727 + 1500\beta, 386$

Note that if both firms invest now, the first to market receives an NPV of \$300, and the second to market receives an NPV of $-\$800 + \1100β . We don't know which firm will reach

the market first, so there are two possible outcomes for the (N_1, N_2) strategies, and likewise there are two possible outcomes for the (W_1, W_2) strategies.

If it were desirable to preempt, each firm could, in principle, spend additional amounts of money to complete the investment and get to the market first. As we have seen when we studied “winner-take-all” markets, this could lead to a destructive race, or war of attrition, where both firms lose money. To keep things as simple as possible, however, let us ignore this possibility. Suppose that the only investment allowed is the \$800 to build the widget factory, and if both firms invest at the same time, the one that gets to market first is determined randomly, e.g., by a coin toss. We can then write the payoff matrix in terms of *expected* NPVs. This is shown below.

		<u>Firm 2</u>	
		Now	Wait
<u>Firm 1</u>	Now	$-250 + 550\beta,$ $-250 + 550\beta$	$300,$ $\max[0, -727 + 1500\beta]$
	Wait	$\max[0, -727 + 1500\beta],$ 300	$-170.5 + 750\beta,$ $-170.5 + 750\beta$

There are several things that we can observe from this payoff matrix:

1. First, note that for any positive β , waiting by both firms dominates investing now by both firms.
2. Note that if $\beta > .31$, waiting by both firms yields a positive expected NPV. The reason is that both firms will end up investing only if the price of widgets goes up to \$150. If β is equal to, say, .40, the firm that arrives in the market second will have an ex post negative NPV, but the ex ante expected NPV is positive.

3. Suppose Firm 2 is going to invest now (and thus try to preempt). Might Firm 1 still want to wait? To answer this, just compare the expected payoffs from waiting versus investing now. Clearly, Firm 1 should wait if

$$-727 + 1500\beta \geq -250 + 550\beta$$

That is, the firm should wait if β is greater than or equal to .502.

What can we learn from this? First, if the gains from preemption are very large (e.g., if $\beta = 0$), neither firm will invest now or later, unless it has a means of preemption. The reason is that if both firms try to invest at the same time, the ex ante expected NPV is negative. If, on the other hand, the gains from preemption are moderate (e.g., $\beta = .6$), it is better to wait. The reason is that the value of waiting exceeds the gains from preemption. Of course, in this case *both* firms would want to wait. (For example, if $\beta = .6$, (W_1, W_2) yields an expected NPV of \$279.50 for Firm 2, while (W_1, N_2) yields \$173 for Firm 2.)

We can see from this that unless there is some action that one firm can take to preempt (and the other can't), *the value of early entry can be separated from the value of waiting and learning*. In other words, the game theoretic problem can be separated from the real options problem. Thus, a firm can begin by analyzing the pure value of waiting, ignoring preemption, and then given the resulting NPVs, analyze the gaming situation, taking preemption into account.

4 Informational Cascades

Learning from others sometimes leads to unfortunate results. Suppose you are considering investing in the stock of Ajax Corp., which is currently trading at \$20 per share. Ajax is a biotech company that is working on a radically new approach to the treatment of chronic boredom (a disease that sometimes afflicts students of economics). You find it difficult to evaluate the company's prospects, but \$20 seems like a reasonable share price. But now you see that the price is increasing — to \$21, \$22, then a jump to \$25 per share. In fact, some friends of yours have just bought in at \$25. Now the price reaches \$30. Other investors

must know something. Perhaps they consulted biochemists who can better evaluate the company's prospects. So you decide to buy the stock at \$30. You decided that there was positive information in the actions of other investors, and you acted accordingly.

Was buying the stock of Ajax at \$30 a rational (or "rationalizable") decision? It certainly might be. After all, it is reasonable to expect that other investors tried to value the company as best they could, and that their analyses might have been more thorough or better-informed than yours (even though they are not Sloan graduates). Thus the actions of other investors could well be informative and lead you to rationally adjust your own valuation of the company. In fact, we would expect that to be the case if the stock market is reasonably efficient. Likewise, the fact that other investors are buying up real estate in downtown Oshkosh could lead you to rationally conclude that the future of Oshkosh is quite bright indeed, so that you should also pick up some downtown real estate now, before it doubles or triples in price.

Note that in these examples, your investment decisions are based not on fundamental information that you have obtained (e.g., regarding the likelihood that Ajax's R&D will be successful or regarding real estate supply and demand conditions in Oshkosh), but rather on the investment decisions of others. And note that you are implicitly assuming that: (i) these investment decisions of others are based on fundamental information that *they* have obtained; *or* (ii) these investment decisions of others are based on the investment decisions of others still, which are based on fundamental information that *they* have obtained; *or* (iii) these investment decisions of others are based on the investment decisions of others still, which in turn are based on the investment decisions of still more others, which are based on fundamental information that *they* have obtained; *or* ... etc., etc. You get the idea. Maybe the "others" at the end of the chain based their investment decisions on weak information that was no more informative than the information you started with when you began thinking about Ajax and Oshkosh. In other words, your own investment decisions might be the result of an *informational cascade* — actions based on actions based on actions ..., etc., driven by very little fundamental information. You might think of this as simply a "bubble" or "herd behavior," but it is important to understand how it arises, and how it

can end (often unhappily for investors at the end of the chain).

4.1 A Simple Example

We want to distinguish between actions based on the arrival of fundamental information (which I will refer to as a “signal”) versus actions based on the actions of others, which might have some eventual link to the arrival of fundamental information. A simple example will help to clarify the difference, and illustrate how actions based on actions can lead to an informational cascade.¹

Let’s return to downtown Oshkosh. Suppose individuals are deciding in sequence whether to take an action — in this case, buy some real estate. If an individual acts, the ultimate payoff will be the fundamental value V of a piece of real estate, which is either 1 or -1 , and initially the probability of either payoff is $\frac{1}{2}$. (The payoff from not acting is zero.) Individuals receive *signals*, which can be either *High* (H) or *Low* (L). If in fact $V = 1$, the signal will be H with probability $p > \frac{1}{2}$ and L with probability $(1 - p) < \frac{1}{2}$, but if in fact $V = -1$, the signal will be H with probability $(1 - p) < \frac{1}{2}$ and L with probability $p > \frac{1}{2}$. In other words, a signal is informative, but it doesn’t eliminate all of the uncertainty. Note that if you receive one signal and it is H, your posterior probability that $V = 1$ becomes $p > \frac{1}{2}$, but if the signal is L, your posterior probability that $V = 1$ becomes $(1 - p) < \frac{1}{2}$.

Observable Signals. Now suppose that a new signal comes every week, and that all past signals are observed by all potential investors. In other words, information keeps accumulating, and is available to everyone. What will happen? As the number of signals increases, the uncertainty over the true fundamental value V is continually reduced, so that all investors will eventually settle on the correct choice — they will buy real estate if in fact $V = 1$ and they will not buy any if in fact $V = -1$. (If the signals are observed with noise, information will accumulate more slowly, but it will still eventually draw investors to the same correct decision.)

¹This example is drawn from S. Bikhchandani, D. Hirschleifer, and I. Welch, “Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades,” *Journal of Economic Perspectives*, Summer 1998, **12**, 151–170.

Observable Actions. Suppose instead that each individual receives one signal, and beyond your own signal, you can only observe the *actions* of other individuals. It is easy to see that this can lead to an informational cascade in which many people act (buy real estate) even though in fact $V = -1$, or alternatively many people do not act even though in fact $V = 1$.

Consider a sequence of risk-neutral individuals, labeled A, B, C, etc. We want to know what each individual will do given his or her own signal, and given the observed actions of his or her predecessors.

- Clearly A will act if his signal is H, but will not act if his signal is L. (His expected NPV is positive if the signal is H and negative if it is L.) And note that all other individuals can infer A's signal perfectly from his decision; if he acted he must have observed H, and if he didn't he must have observed L.
- Suppose A acted. Now what should B do? Clearly B should act if her signal is H.² However, if her signal is L, her posterior probability that $V = 1$ is $\frac{1}{2}$, so she is indifferent. We will assume that in this case she flips a coin to decide what to do.
- Individual C now faces three possible situations:
 1. A and B both acted. In this case, C will act, even if his signal is L. You can check that no matter what signal C receives, given that both A and B have acted, his NPV of acting is positive — even though B may well have flipped a coin after receiving a low signal.³
 2. A and B both rejected the investment. In this case, C will reject, no matter which signal he receives.

²Note that the probability that $V = 1$ given two signals of H is $p^2/(p^2+(1-p)^2)$. If $p = .6$, this probability is .69; if $p = .8$, this probability is .94.

³You can check that the probability that $V = 1$ given that A acted (which means his signal was H), B acted (which means that either her signal was H or it was L and she flipped a coin), and C's signal was L is $(p+1)/3$. Because $p > \frac{1}{2}$, this probability is greater than $\frac{1}{2}$, so that C's NPV of acting is positive.

3. A acted and B did not, or vice versa. In this case, C will act only if his signal is H.

- We will focus on the first situation: A and B both acted, and thus C acts even if his signal is L. In this case, what will D do? She will act, no matter what signal she receives. Likewise, E, F, etc., will all act.

We now have an informational cascade. It is quite possible that A, and only A, will have received a signal of H, and that all subsequent individuals received a signal of L. If D, E, etc. could have observed that B and C received signals of L, they would not have acted. But all they can observe is the *actions* of their predecessors. Everyone is acting rationally (expected NPVs are positive) even though no new information is being produced.

So, the price of real estate in downtown Oshkosh keeps going higher and higher. You observe that all 15 real estate investors who visited Oshkosh bought some property, so you — quite rationally — jump on the bandwagon and buy property yourself, pushing the price higher still.

Assuming that in fact $V = -1$, how does this process end? In the simple model described above, it only ends when we run out of investors. But in reality it probably ends when some new kind of signal becomes available to at least some investors. Perhaps some of the investors begin to notice that there are hardly any people living in Oshkosh, and that no tourist would ever want to visit the town. Those investors then start to sell their property. Other investors observe these actions, and — quite rationally — also sell. And the price of real estate in downtown Oshkosh plummets.

4.2 What to Do?

Our little story about real estate in Oshkosh shouldn't seem that far-fetched. There are plenty of real-world examples that followed this pattern: bank loans to Latin American countries during the late 1970s and early 1980s, the oil-based real estate boom in Texas and Oklahoma during the 1980s, the recent internet bubble, and perhaps the housing boom on the East and West Coasts that seems to be ending, just to name a few.

So what's the lesson from this story? Remember that the investors in Oshkosh who observed the actions of other investors were acting rationally — the NPVs were always positive. But what we didn't say much about was the *riskiness* of their investments. (We didn't need to worry about that because I assumed that the individuals were all risk-neutral.) Compare once again what happens with observable signals versus observable actions. With observable signals, uncertainty is reduced as more and more signals arrive, i.e., as more fundamental information accumulates. With observable actions, on the other hand, the informational cascade leads to no reduction in uncertainty. Each individual's NPV remains positive, but the uncertainty confronting individual C is the same as for D, for E, and so on. Thus the fact that 15 people bought real estate in Oshkosh is no more informative (and should be no more comforting) than the fact that 3 people did.

Does that mean you (the 16th investor) shouldn't buy real estate in Oshkosh? Not at all. But you should understand just how much (or how little) information your decision is based on. And you should understand how rational decisions based on the actions of others can involve much more risk than decisions based on the accumulation of fundamental information.