

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Options, inaction, and uncertainty

This is a pre print version of the following article:

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1513052> since 2015-08-25T09:44:20Z

Published version:

DOI:10.1111/j.1467-9485.2009.00517.x

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Options, Inaction, and Uncertainty

Giuseppe Bertola*
(Università di Torino)

December 2009

Abstract

I outline and discuss economic insights afforded by the elegant and powerful continuous-time techniques that economists discovered some 20 years ago. I do so in the context of a simple action-timing problem that does *not* exploit those techniques, in order to illustrate how the role of uncertainty in shaping the optimality of adoption depends on substantive modelling details. I then outline the clearer implications popular and tractable continuous-time continuous processes, noting that they may fail to represent key features of some real-life business investment problems, but provide a suitable toolkit for treatment not only of financial phenomena but also of real option problems, such as those encountered in dynamic consumption problems, where small and unpredictable random changes are realistic.

1 Introduction

In 2009, I was kindly asked to be a keynote speaker at a nice conference on "Real Options," and the *Journal of Economic Literature* asked me to review an excellent recent textbook on *The Economics of Inaction* (Stokey, 2009). These pleasant invitations brought me back in time to the late 1980s, when I was trying to make progress on a Ph.D. dissertation, and other economists and I discovered something that seemed very important: continuous-time optimization techniques developed in the engineering literature make it possible to obtain elegant, concise, and rigorous quantitative solutions to infrequent action problems.

The problems themselves were not new. Problems exist in nature, and it had been discovered a long time before then that inaction is realistic and important. Among others, Hicks (1950) emphasized the stop-and-go character

*This paper is based on the author's keynote lecture at the CESifo Summer Institute Workshop "Operating uncertainty using real options," Venice, 8 July 2009. I am grateful to Michael Funke and to the participants for their attention and comments.

of business cycles bound by capacity ceilings and disinvestment floors, and Arrow (1968) showed that periods of capital shortage and of excess capacity are optimal when investment programs face predictable fluctuations and installed capital cannot be put to alternative uses. These and other authors' interest in the optimality of inaction and the infrequency of investment was heightened by the contrast between such results and the prevailing approach to economic analysis, where the balance of marginal costs and benefits determine the intensity of investment at all times. From that perspective, it was somewhat amusing to find that the optimal attitude, even in a rapidly changing environment, would often be one of "doing nothing" (Bar Ilan and Blinder, 1988). By the 1980s, however, it had become clear that inaction is often qualitatively realistic, and lets inertia and expectations shape interestingly unusual dynamics. But quantitative and empirical analysis still was so much easier framed in terms of continuous control.

When trying explicitly to model uncertainty and rational expectations, researchers found it difficult to step out of the linear-quadratic framework of Sargent's (1979) textbook and of its convenient certainty-equivalent separation of estimation and control in dynamic systems. Back then, it was not easy for economists to become familiar with continuous-time formal concepts and analytical tools. The short textbook by Malliaris and Brock (1981) offered a somewhat cursory technical treatment of continuous time stochastic processes. When discussing economic optimization, however, it was limited to applications in that environment of investment problems where standard convexity assumptions implied internal solutions, with continuous investment. It was interesting to see that uncertainty could be tacked on to such problems, and to learn that optimal portfolio and consumption choices could similarly be modelled in an uncertain continuous-time environment. A formal approach to the economics of optimal inaction, however, remained elusive.

Twenty or so years ago, the news was that at least some optimal inaction problems became tractable as economists learned how to apply technical tools developed in other disciplines. When trying to make progress along these lines, we had to look in unfamiliar library shelves for applied mathematics and statistics textbooks that were very rigorous and very dry, because their authors made no effort to motivate assumptions and discuss results in terms of the economics we knew and of the problems we wanted to solve.

Like the much more concise continuous-time optimization toolbox in Bagliano and Bertola (2007, Section 2.8 and exercise 24), Stokey (2009) allows its readers to appreciate economic applications of controlled Brownian motion and their relationship to more familiar dynamic optimization problems. This makes it much easier for students and applied economists to appreciate and use techniques that used to require them to seek obscure mathematical treatises in unfamiliar sections of their universities' libraries.

Of course, however, it would be wrong to use Brownian processes and optimal control techniques just because they have become easier. In what follows, I recall insights, problems, and solutions encountered when economists began to applied Brownian processes to the economics of inaction, and discuss briefly how and when they are still useful.

2 Insights

Inaction may well be optimal when action entails discrete changes in a decision-maker's costs and future outlook. Obviously, a wait-and-see attitude affords some benefits if future developments may lead one to regret having acted: hence, when choosing whether to act, the option to wait has to be weighed against the value of action.

This intuitive point applies to the discrete action/inaction choice, i.e. to the timing of decisions, the same principles that in fact apply to any dynamic optimization problem. Regardless of whether the alternative to action is complete inaction, or different actions, or the intensity of action, it is always essential for optimization in dynamic environment to consider the impact of current actions for future situations, and for the actions that will be called for by future developments.

2.1 The linear-quadratic case

The future is generally uncertain, of course, and should be discounted. In a very special case, uncertainty about the implications of future developments for future optimal actions happens to be irrelevant to current decisions. Dispersion of future realizations around their expected value is irrelevant to current decisions in problems that optimize quadratic benefit and cost functions, because the first-order conditions that pin down the optimal intensity of action are linear, and expectations of linear functions coincide with the functions evaluated at the expected value of their argument.

This implies a "certainty equivalence" property for optimal solutions. In general, decision makers know that they will be surprised by future developments, and that surprises may be pleasant or unpleasant. In linear-quadratic problems, surprises have symmetric implications for future conditions and decisions. Since the implications of current choices for future regrets are balanced, there is no need to worry about them, and choices can be based on point expectations of future developments.

For example, a firm subject to increasing marginal adjustment costs of investment will wish to have anticipated an unexpected increase in capital's marginal productivity, accumulating capital in slower and less costly fashion. But this does not change its ex ante choice in a linear-quadratic problem. Just because the increase is unexpected, it might just as well have turned

out to be a decrease, leading the firm to regret not to have planned its investment so as to meet the future with a lower capital stock.

The case where marginal costs are linear and the implications of possible positive and negative surprises offset each other perfectly is, of course, very special. Any deviation from linearity of the relationships between exogenous forcing variables and endogenous optimal choices implies that decisions should not be based on point expectations of exogenous processes. By Jensen's inequality, in fact, expectations of nonlinear functions are different from the functions evaluated at the expectation. Thus, the expectations relevant to optimal decisions depend on the dispersions as well as on the expectation of future variables. (as originally pointed out by Hartman, 1972, and discussed in detail in Bagliano and Bertola, 2007, Section 2.4).

Linear-quadratic decision criteria that neatly separate forecasts from control activities can only approximate real-life optimal problems. Approximations can of course be very good, but can also be very bad, and linearized decision rule are so bad an approximation as to be ill-defined when the costs of action not only are not quadratic, but are not differentiable, preventing an approximation interpretation for quadratic functional forms.

2.2 First-order adjustment costs

When acting is discretely rather than infinitesimally more costly than not acting, then optimal decisions are not only quantitatively, but also qualitatively different from the behavior implied by linear-quadratic problems. When more intense action implies a smooth cost increase above the zero cost of no action, acting at all times and modulating the amount of adjustment is optimal. Decision makers faced by adjustment costs that are first-order larger for action than for inaction should instead choose inaction for discrete, well-defined periods, and find it optimal to bunch actions in time and at specific points in the optimization problem's state space.

Certainty equivalence cannot hold, as it never does outside of the special linear-quadratic case. Uncertainty also matters because during periods of inaction some of it is resolved. Intuitively, when the future is more uncertain a wait-and-see attitude is more appealing, because there is more to see. In a highly uncertain environment, where exogenous developments may be so drastic as to imply very deep regret, actions entailing unrecoverable costs are less attractive than in a steadier environment, where the future is unlikely to differ much from the present. In a very random environment, where fast-arriving news are likely to quickly lead the optimizer to regret taking action, hesitation may indeed be highly justified.

These are important qualitative insights, based on the simple intuitive argument that dynamic optimization aims at avoiding regret. Before concluding that in a more uncertain environment we should expect to see more waiting and less action, however, it is important to inspect the mechanics

of the relevant problems in more detail, paying attention to the sources and character of uncertainty, and bearing in mind the implications of discounting and expectations for the character of optimal solutions.

3 An illustrative problem

In 1990, I had the privilege to teach investment theory within an advanced Macroeconomic course, which was for historical reasons dubbed "Money" in the graduate sequence. Researchers at the time were enthusiastic about the new (to Economics) optimal control techniques. The students and I spent a substantial amount of time and effort studying those techniques and comparing their implications to those of more traditional approaches. The exam I assigned, however, asked students to solve and discuss a problem that was not expressed in continuous time, and did not feature the continuous processes—on which more in the next section—borrowed from the financial and engineering literatures.

3.1 A Princeton graduate exam

The problem was rather more ordinary and day-by-day than the vexed question of what determines business investment: as reported in Box 1, the exam question asked students to characterize how should one choose when to pick a mushroom that grows over time and, if left in the ground, may disappear with some known probability.

Box 1: When should a mushroom be picked?

This question concerns investment: part a) is a simple exercise, part b) asks you to write a related essay. Answer both.

a) A wild mushroom grows at rate g ,

$$M_t = M_0 e^{gt}, \quad g > 0$$

You go by the mushroom on your way to work, every D time units. The utility you get from picking the mushroom is linear in its size, and the utility loss from going out of your way to pick it is a constant K . You discount future utils exponentially at rate r . There is a chance that, if you do not pick the mushroom, somebody else will: let this be a Poisson event with arrival rate p .

Consider two picking strategies:

(1) Pick the mushroom if $M_t > K$.

(2) Pick the mushroom if $M_t - K > E_t e^{-rD} [M_{t+D} - K]$

Discuss informally the intuitive appeal of the two strategies.

In particular, consider the limiting case when $D = 0$, i.e. you monitor the mushroom in continuous time. Let M_0 be small, so that you do not pick

the mushroom at time zero. As of time zero, what is the value of strategy (1)? Of strategy (2)? Can the latter be smaller than the former?

Discuss the role of the parameters g , r , p in determining the difference between the two strategies, "the value of waiting".

b) The simple model in Part a) tries to capture certain aspects of realistic investment decisions. Discuss how available models of firms' investment in physical capital deal with those insights. What difficulties does a theory of investment stressing these aspects encounter? (include in your discussion issues of aggregation, empirical implementation).

The first point highlighted by the exercise is that waiting can be optimal if time elapsed matters. This is the case, even without uncertainty, if the future is discounted and the relationship between costs and benefits of action changes over time.

In the exercise, the constant cost of action is to be compared to a benefit (the size of the mushroom) that grows over time at a constant rate. Clearly, the benefits of picking the mushroom exceeds the cost of that action as soon as $M_t > K$, so the exercise's strategy (1) has positive value. Since we are told that the size of the mushroom is growing at a positive rate, however, the benefits would be even larger if we wait.

To see the implications of the passage of time, which are of course crucial in any dynamic situation and particularly clear in an optimal timing problem, consider the case where $p = 0$. There is no uncertainty, and the mushroom's size will somewhat unrealistically grow forever towards infinity. Supposing that the growth rate g is lower than the discount rate r ensures that the problem is well defined (otherwise, the mushroom has infinite present value at all times, and it would be hard to decide when to consume it): the optimal picking time maximizes the (finite) present value of the problem.

As long as the mushroom is not eaten, of course, no actual benefit accrues to the decision maker. Postponing the decision delays the cost of action and reduces its discounted value, and also allows the mushroom to grow larger. Strategy (2) is, quite intuitively, better than strategy (1). It is a good idea to compare, in discrete time, the benefits of immediate picking to the discounted benefits of delayed picking, and chooses action only when the latter fall short of the former. As long as the inequality in (2) is reversed, it is optimal to wait.

The exam question asked students to characterize not only the role of discounting and growth, but also that of uncertainty, represented in the example by the risk that somebody else would pick the mushroom. Formal derivations are, as usual, easier and cleaner in continuous time, not least because instead of involving inequalities the solution can be expressed neatly

in terms of a precise "picking" time T^* that solves

$$\max_x E [e^{-rx} (M_0 e^{gx} - K)],$$

the problem of maximizing the discounted, expected size of the mushroom net of picking costs, $M_0 e^{gT^*} - K$.

3.2 Growth and discounting

Instead of analyzing fully the example proposed in the exam problem, it seems better here to proceed quickly towards a more explicit economic problem, referring to the mushroom-picking situation only when it may help strengthen intuition. As in the seminal paper by McDonald and Siegel (1986), suppose a firm owns an opportunity to invest that, once taken, will yield a cash flow. Unlike McDonald and Siegel, however, we rule out cash-flow uncertainty and suppose that, like the exercise's mushroom, the cash flow of an adopted project grows steadily, at a rate denoted ϑ in what follows. We also suppose, in analogy to the exercise's picking cost K , that adopting the project entails paying a cost C .

Since the cost C is not paid unless the firm chooses to do so, letting time pass implies not only that the project's cash flow will start from a higher value if $\vartheta > 0$ but also that, if the opportunity cost of funds is $r > 0$, delaying investment yields financial benefits. If the project is adopted at time T , cash flow will begin to accrue starting at the level $\Pi_T = \Pi_0 e^{\vartheta T}$ and continuing to grow at rate ϑ .

Hence, as of time T the adopted project is worth

$$\int_0^\infty \Pi_0 e^{\vartheta x} e^{-rx} dx = \frac{\Pi_T}{r - \vartheta},$$

provided of course that $\vartheta > r$ to ensure convergence of the integral.

Adoption is profitable whenever the value of the project exceeds the adoption cost, i.e., whenever $\Pi_T / (r - \vartheta) > C$. The condition $\Pi_T > (r - \vartheta)C$, however, identifies the optimal adoption rule only if the project is a "take it or leave it" opportunity, presented to the firm suddenly. In fact, if the firm owns the project and can choose when to adopt it, waiting to invest only until such a time means accepting that the project's cash flow will equal its cost, yielding exactly zero net value.

Obviously, this is better than a negative net value: as long as adoption (or picking a mushroom) is optional, it would be silly voluntarily to experience an avoidable loss. Just as obviously, however, if the project's net value reaches zero from below then it will become positive. Since delaying adoption makes it possible to lift its value above zero, waiting is clearly a good idea. But, just as clearly, it cannot be optimal to wait forever to pick a mushroom, since that would mean never eating it.

To determine for how long should adoption be delayed, consider that waiting has costs and benefits. As long as the project is not adopted, there will be no cash flow. Just like a person who holds out on the mushroom-picking option, the firm foregoes "dividends" when holding an unexercised option to adopt. But waiting also generates capital gains on that option. As long as these offset the lost cash flow, it will be optimal to delay payment of C , and adopt the project when its cash flow's present discounted value has grown larger.

In the absence of uncertainty, in fact, it is easy to see that the value of the investment opportunity as a function of the adoption time T ,

$$V_0(T) = e^{-rT} \left(\frac{\Pi_0 e^{\vartheta T}}{r - \vartheta} - C \right), \quad (1)$$

is the product of two terms that depend on T in opposite ways. The first term, the discount factor, decreases with T if $r > 0$ (as discounting decreases the value of future cash flows). The second term increases with T if $\vartheta > 0$ but, as long as $C > 0$, its increase is slower than exponential. Hence, as shown in Figure 1, the value of the investment opportunity increases only up to a point, and its peak as a function of the adoption time T identifies the optimal strategy.

With positive discounting, it can be a good thing to postpone payment of the investment cost C . Formally, the optimal T^* solves the problem

$$\max_T e^{-rT} \left(\frac{\Pi_0 e^{\vartheta T}}{r - \vartheta} - C \right),$$

and the first order condition that sets to zero the value's derivative identifies as the optimal adoption strategy the rule

$$\Pi_{T^*} = rC.$$

If $\vartheta > 0$, the cash flow that triggers adoption of projects with an option to wait is larger than its $(r - \vartheta)C$ counterpart of take-it-or-leave-it investment opportunities. If $\vartheta < 0$, as might be the case if mushrooms shrink or deteriorate when left unpicked, it is smaller. In both cases, the rule is the same, and a familiar one from Jorgenson's (1963) textbook investment theory. Since by waiting the firm forsakes the cash flow Π_T , but also saves the opportunity cost rC , it will be optimal to wait until flow benefits and flow costs become equal.

3.3 Roles of uncertainty

Uncertainty, in the problem we are discussing, is described by Poisson process, a modelling device that was also part of the macroeconomist's toolkit at the time when unemployment, for example, was beginning to be studied in terms

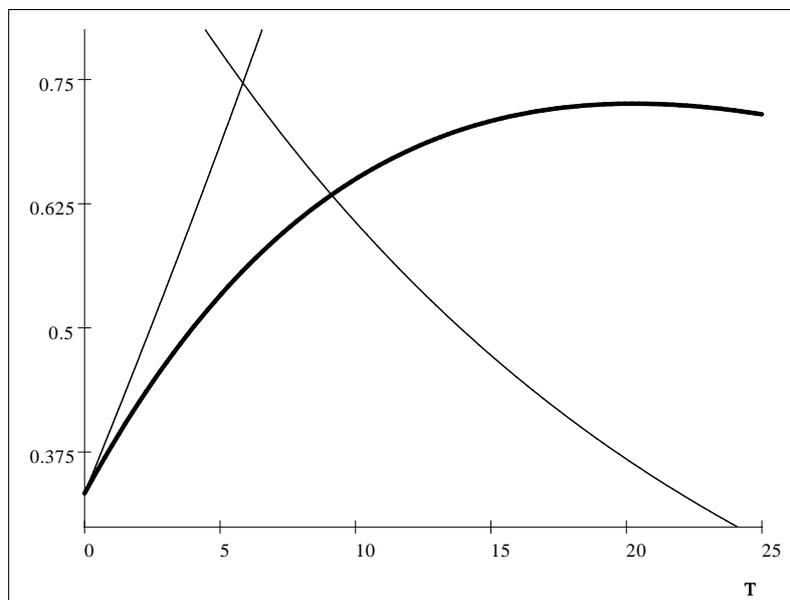


Figure 1: The solid line plots $e^{-rT} \left(\frac{\Pi_0 e^{\vartheta T}}{r - \vartheta} - C \right)$, the discounted value of the project adopted at time T . The discount factor e^{-rT} is plotted by the decreasing line, and the net value of the project as of the adoption time is plotted by the increasing line. Other parameters: $\Pi_0=0.1$, $C=3$, $r=0.05$, $\vartheta=0.02$.

of job-finding and/or job-losing probabilities, along the lines of the Shapiro and Stiglitz (1984) or Pissarides (1987) models which were also included in the readings of courses like the one I was teaching.

Just like a mushroom picked by others, so a business project may become unavailable if the firm waits. A Poisson process lets the probability of this event be $1 - e^{-\delta\Delta t}$ over a period of length Δt . This process is memoryless: since the probability of the event is constant, the future always look the same as time goes by. As far as a mushroom picker or a business investor is concerned, however, the future does look different if the action is taken because, by assumption, this rules out all uncertainty.

This admittedly special feature of the problem we are considering illustrates a more general insight. The passage of time influences the costs and benefits of waiting not only because of discounting and growth, along the lines discussed above, but also because it produces information. At any point in time, the decision maker knows whether an event that bears on costs and benefits has yet occurred, while the future remains uncertain. This bears on his or her incentives to take action: knowing that new information can arrive makes it more or less worth to wait, inasmuch as it affects the probability that acting or failing to act will lead to future regrets.

Whether these considerations imply that higher uncertainty will foster action or inaction depends on the structure of the specific problem. In the one under consideration, choosing a time T for action will engender "regret" if, by then, the project or mushroom will turn out to have disappeared, so that earlier action will look preferable. If T is the time of adoption, the choice will yield a discounted value $e^{-rT} \left(\frac{\Pi_0 e^{\vartheta T}}{r - \vartheta} - C \right)$ with probability $e^{-\delta T}$, and zero otherwise. Following the same steps above, the optimal time T^* solves

$$\max_T e^{-(r+\delta)T} \left(\frac{\Pi_0 e^{\vartheta T}}{r - \vartheta} - C \right),$$

and the first order condition implies that

$$\Pi_{T^*(\delta)} = \frac{(r + \delta)(r - \vartheta)}{r + \delta - \vartheta} C.$$

Quite intuitively, when $\vartheta > 0$ this is again larger than the take-it-or-leave-it trigger $(r - \vartheta)C$. And it is also intuitive to find that the larger disappearance probability implied by a larger δ implies that the project is adopted at a lower cash flow level.¹ A smaller—albeit growing—mushroom is worth picking when not doing so may allow others to pick it instead.

A positive δ decreases the optimal adoption or picking size and, for a given growth rate of the unexploited opportunity's payoff, also implies less

¹Formally,

$$\frac{d}{d\delta} \frac{(r + \delta)(r - \vartheta)}{r + \delta - \vartheta} = \frac{-\vartheta}{r + \delta - \vartheta} \frac{r - \vartheta}{r + \delta - \vartheta} < 0$$

for $\vartheta < r$.

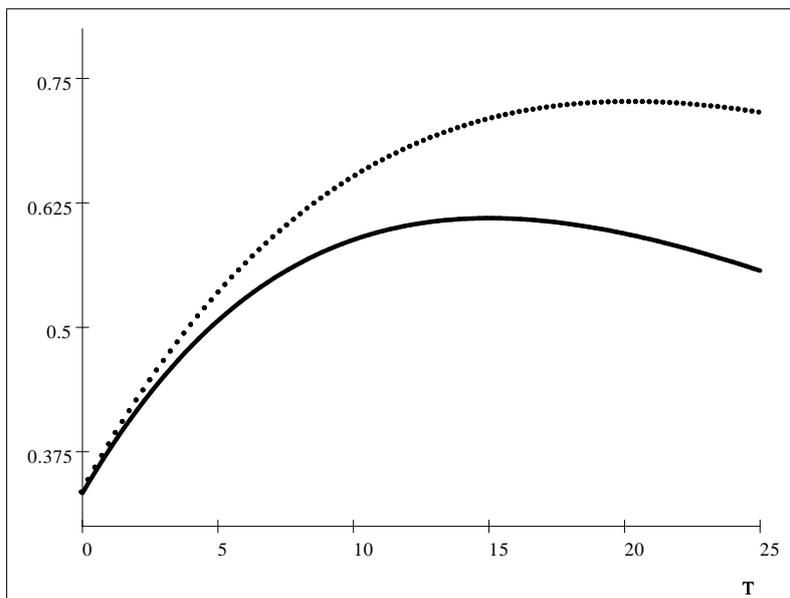


Figure 2: The solid line plots $e^{-(r+\delta)T} \left(\frac{\Pi_0 e^{\vartheta T}}{r-\vartheta} - C \right)$, the expected discounted value of a project that will be adopted at time T but may have disappeared by then with probability intensity $\delta > 0$. The dotted line plots the same formula setting $\delta = 0$. Other parameters as in Figure 1.

waiting: as shown in Figure 2, the expected present value of the opportunity peaks sooner, as a function of waiting time T , when δ is larger than zero.²

Action is accelerated by uncertainty in this problem, for the good reason that further inaction may be prolonged forever by disappearance of the option to act. The quicker action implied by more uncertainty decreases the probability of disappearance, at the same time as more uncertainty makes disappearance more likely within any spell of time. Thus, it is natural to wonder whether optimal adoption of the project is more or less likely when δ is higher.

As shown in Figure 3, the optimal probability $\exp(T^*(\delta))$ of adoption turns out in this problem to be ambiguously related to the probability intensity of the opportunity's disappearance. The decision maker waits less as δ increases above zero, and waiting becomes less likely to pay off. Initially, however, the optimal strategy also ends up adopting less—picking not only smaller, but also fewer mushrooms.

This illustrates simply a more general insight: uncertainty does delay action but, just because optimal behavior is motivated by the possible dis-

²Formally, as long as $\vartheta < r$ and the problem is well defined, $\Pi_0 e^{\vartheta T^*} = \frac{(r+\delta)(r-\vartheta)}{r+\delta-\vartheta} C$, and the optimal adoption time $T^*(\delta) = \frac{1}{\vartheta} \ln \left(\frac{(r+\delta)(r-\vartheta)}{r+\delta-\vartheta} \frac{C}{\Pi_0} \right)$ is decreasing in δ .

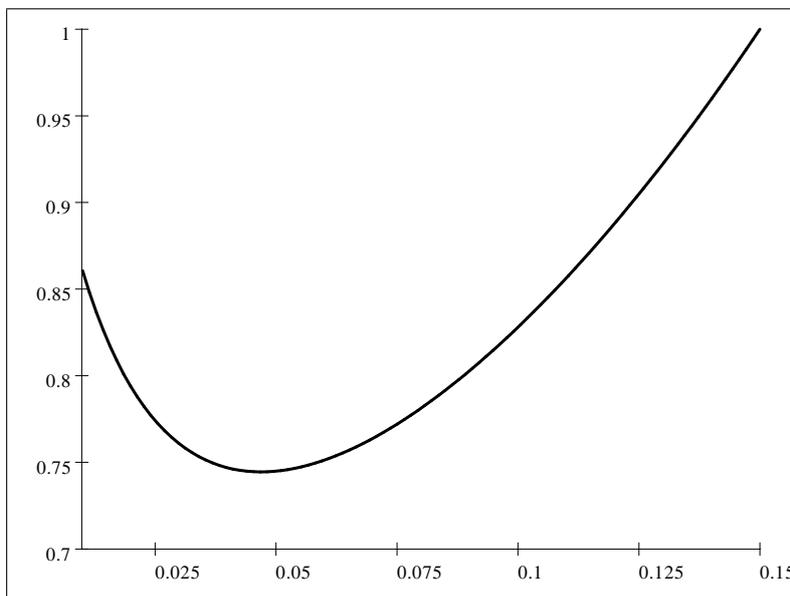


Figure 3: Plot of $e^{-\delta \left(\frac{1}{\vartheta} \ln \left((r+\delta) \frac{C}{r+\delta-\vartheta} \frac{r-\vartheta}{\Pi_0} \right) \right)}$ as a function of δ . Other parameter values as in the previous figures.

appearance of the option, it may or may not do so as strongly as to make successful adoption more likely. More uncertainty does not necessarily imply more (or less) inaction. In this problem, adoption is speeded up sufficiently to imply a higher probability of action when the probability intensity of the opportunity's disappearance increases above a critical level, that depends on the problem's parameters.³

This type of problem has room for analysis of further interactions between the optimal time and optimal size effects of uncertainty. In the setting we are considering, a higher likelihood of disappearance implies not only more uncertainty, but also less advantageous investment opportunities. And if a higher likelihood of disappearance accelerates adoption, then it also reduces the optimal size of the adopted project (or picked mushroom) but, of course, the latter is a direct consequence of the former, as the growth rate

³The derivative with respect to δ of the probability's exponent,

$$\begin{aligned} \frac{d}{d\delta} \left[-\frac{\delta}{\vartheta} T^*(\delta) \right] &= -\frac{1}{\vartheta} \ln \left((r+\delta) \frac{C}{r+\delta-\vartheta} \frac{r-\vartheta}{\Pi_0} \right) \\ &\quad + \delta \frac{C}{\Pi_0} \frac{r-\vartheta}{(r+\delta-\vartheta)^2} \left((r+\delta) \frac{C}{r+\delta-\vartheta} \frac{r-\vartheta}{\Pi_0} \right)^{-1} \end{aligned}$$

has different sign depending on the value of δ and other parameters, notably inasmuch as the latter determine the distance between the initial and optimal values of the adopted project, or picked mushroom.

ϑ is given less time to work.

To isolate the role of uncertainty, we may compare the solution of problems which differ not only in terms of disappearance probabilities, but also in terms of growth rates in the period leading up to disappearance or adoption. Specifically, consider the implications of letting the growth rate ϑ depend linearly on the disappearance intensity δ according to

$$\vartheta = g + \delta.$$

This assumption serves the purpose of isolating the implications of uncertainty from those of the expected discounted value of the investment opportunity. For any given adoption time T , in fact, the probability $e^{-\delta T}$ that the project will not have disappeared falls at the same rate at which the cash flow $\Pi_0 e^{\vartheta T}$ grows. Hence, the passage of time does not directly influence the value of waiting through the expected value of the project.

In the investment project interpretation (if not in the mushroom-eating example), optimal adoption timing depends on the growth rate of cash flows not only before adoption, but also after adoption, because this determines the proportionality factor of the project's value to the cash flow as of the adoption time. Hence, further insight into the relationship between the problem's parameters and the "waiting" characterization of the optimal solution is gained by considering how adoption may affect the decision maker's outlook on further development. To this end, let the project's cash flow grows at rate h once adopted, a parameter which may differ from ϑ and/or from g .

The optimal timing choice maximizes the expected discounted value

$$e^{-(r+\delta)T} \left(\frac{\Pi_0 e^{\vartheta T}}{r-g} - C \right) = e^{-(r+\delta)T} \left(\frac{\Pi_0 e^{(g+\delta)T}}{r-h} - C \right),$$

which is well defined as long as $r > g$ and $r > h$. The first order condition for choice of T can, as before, be simplified and rearranged to read

$$(r-g) \frac{\Pi_T}{r-h} = (r+\delta)C.$$

Hence, waiting should optimally stop at such a time T^* that

$$\Pi_{T^*} = \Pi_0 e^{(g+\delta)T^*} = (r+\delta) \frac{r-h}{r-g} C.$$

For given h and g , the project should now be adopted when its cash flow has grown to a level that is an increasing function of the disappearance intensity δ . The decision maker's outlook on growth expectations and fear of regrets still favors waiting for a larger-size project (or mushroom). In the current specification of the problem, however, not only the the likelihood of adoption

but also the optimal picking time are ambiguously related to the probability intensity of disappearance.⁴

In summary, the decision maker's inclination to act and the likelihood of action are determined in this problem by interactions (between the speed of growth, discounting, and uncertainty) that, while intuitive, are intricate enough to yield overall ambiguous relationships between the relevant parameters and the solution's characterization.

4 Business investment under uncertainty

The bright graduate students who took the exam reported in Box 1 were amused by the mushroom example, and managed to perform at least some derivations like those outlined above. They did not do so well on the essay portion of the exam, and I blamed myself for failing to express the question as clearly as I should have to make them realize that the mushroom problem considered in the Section 3 has room for many different forms of uncertainty, and so does its translation in terms of a business investment opportunity, where the project's cash flow could also (or instead) be random after adoption. One could model such events in terms of the Poisson process, along the lines of the derivation above. The "real options" literature, however, has focused on continuous-time Brownian processes, and it is possible and important to understand exactly how continuous-time techniques can help make more general insights precise.

This section sketches technical issues encountered in this context, aiming to highlight the approach's properties in terms of tractability and realism. In applications to business investment, the counterpart of the simple problem above is the marginal contribution of the capital stock to a firm's operating profits, a function $F(K_t, Z_t)$ of the installed capital stock K_t and of other, exogenous processed indexed by Z_t . If capital can be adjusted without cost, or if its optimal stock grows steadily without uncertainty, the optimal investment rule satisfies the standard Jorgenson (1963) condition: as any student of Economics textbooks knows, the marginal revenue product of capital should at all times be equal to the user cost of capital, which in the absence of depreciation equals in turn the required rate of return r times

⁴Formally, the derivative of T^* with respect to δ is

$$\frac{d}{d\delta} \left(\frac{1}{g+\delta} \ln \left((r+\delta) \frac{r-h-C}{r-g\Pi_0} \right) \right) = \frac{1}{g+\delta} \left(\frac{1}{r+\delta} - \frac{1}{g+\delta} \ln((r+\delta)) \right),$$

which becomes negative for large values of δ . The probability of adoption is $e^{-\frac{\delta}{g+\delta} \ln \left((r+\delta) \frac{r-h-C}{r-g\Pi_0} \right)}$, and

$$\frac{d}{d\delta} \left(-\frac{\delta}{g+\delta} \ln \left((r+\delta) \frac{r-h-C}{r-g\Pi_0} \right) \right) = -\frac{g}{(g+\delta)^2} \ln((r+\delta)) - \frac{\delta}{g+\delta} \frac{1}{r+\delta}$$

can be positive or negative depending on parameter values.

the purchase price C of a unit of capital:

$$\frac{\partial F(K_t^*, Z_t)}{\partial K} = rC.$$

When instead investment is not perfectly reversible, however, each unit of capital is in fact a project, like that considered in Section 3: if (all, or part of) the investment cost C cannot be recovered, adoption has to take into account the possible arrival of bad news, which may lead the firm to regret having invested.

An elegant solution is available when, for fixed K_t ,

$$\Pi_t \equiv \frac{\partial F(K_t, Z_t)}{\partial K}$$

follows a *geometric Brownian motion*: its exponential growth rate is expected to be a constant θ_Π , but is continuously disturbed by independent shocks with variance σ_Π^2 per unit of time. Formally,

$$\frac{d\Pi_t}{\Pi_t} = \theta_\Pi dt + \sigma_\Pi dW_t,$$

where dW_t is the change, over the small interval of time dt , of a *standard Brownian motion* or Wiener process, $\{W(t)\}$. This process can be defined by its probabilistic properties: for $t \geq t_0$, probability statements about $W(t)$ as of time zero can be made viewing $W(t)$ as a normally distributed random variable, with mean $W(t_0)$ and variance $t - t_0$: realizations of $W(t)$ are quite concentrated for small values of t , while more and more probability is attached to values far from zero for larger and larger values of t .

This stochastic processes is simply defined as a random walk on infinitesimal time and state steps, a simple and powerful way to model uncertainty: nature takes infinitely frequent random steps, and probabilistic assessments of events far into the future are more uncertain inasmuch as the number of random steps taken in the meantime is larger.⁵ Large changes become unlikely as smaller and smaller time intervals are considered, because the process is continuous. Since variation over time is always normally distributed over any finite interval of time and as the normal distribution assigns positive probability to any finite interval of the real line, however, large variations have positive probability on short (but finite) intervals of time.

Normality of the process's increments follows naturally from averaging infinitely many independent random variables, and is useful in applications because it is preserved by linear transformations, but has peculiar implications. As increments have to be independent over any interval, the direction

⁵In fact, a good way to get acquainted with continuous-time stochastic processes is building them up from discrete-time, discrete-state space processes, as in Dixit (1991).

of movement must be random at arbitrarily close points: hence, independence over time of continuous variation has the somewhat awkward implication that derivatives with respect to time cannot be well defined. A random differentiable function of time would move in locally predictable ways, but Brownian processes move up and down so frequently that, when crossing a given level, they visit it so often as to lead to continuous accumulation of the "local time" spent there, and have *infinite variation*: the absolute value of the increments over infinitesimally small subdivisions of an arbitrarily short time interval is infinite.

Mathematical manipulation of such processes requires analytical tools that, while conceptually equivalent to Leibnitz's and Newton's calculus, were only developed in the middle of the 20th century by the Japanese mathematician T.Itô and others. An introduction to the relevant tools may be found in Bagliano and Bertola (2007), and an exhaustive treatment in Stokey (2009) and its references. Here, it is perhaps only useful to mention the main difference between "stochastic" and classic calculus.

If a differentiable function of time $z(t)$ has first derivative $z'(t) = dz(t)/dt = \dot{z}$, and function $f(\cdot)$ has first derivative $f'(z) = df(z)/dz$, then $df(z) = f'(z) \dot{z} dt$. When a twice differentiable function's argument is an Itô process with increments

$$dZ(t) = \mu_Z dt + \sigma_Z dW(t),$$

conversely, the "chain rule" takes not only the first but also the second derivative into account:

$$\begin{aligned} df(Z(t)) &= f'(Z(t)) [\mu_Z dt + \sigma_Z dW(t)] + \frac{1}{2} f''(Z(t)) (\mu_Z dt + \sigma_Z dW(t))^2 \\ &= f'(Z(t)) [\mu_Z dt + \sigma_Z dW(t)] + \frac{\sigma_Z^2}{2} f''(Z(t)) dt, \end{aligned} \quad (2)$$

where the second line follows from the fact that while $(dt)^2 = 0$ in the infinitesimal limit, the limit $(dW(t))^2$ of the Wiener process's variance is of the same order as dt .

In this context, as in more general problems, the role of uncertainty is twofold.

First, the marginal revenue product flow may be a nonlinear function of exogenous random variables, and through Jensen's inequality uncertainty may affect the expectation of future cash flows and the appeal of investment (as in Hartman, 1972, and unlike in the linear-quadratic settings briefly reviewed above). For example, suppose that capital's marginal contribution to the firm's cash flow has constant elasticity to the installed stock and to a summary indicator Z_t of the firm's demand and cost conditions,

$$\frac{\partial F(K_t, Z_t)}{\partial K} = K_t^\alpha Z_t^\beta;$$

a constant-elasticity specification, which is linear in growth rates, preserves the (log)normality of the process's distribution, and would make it possible to accommodate tractably depreciation at a constant rate of the installed capital stock.

Then, for fixed K_t the dynamics of $\Pi_t = K_t^\alpha Z_t^\beta$ are given by Itô's rule (2) as

$$d\Pi_t = K_t^\alpha Z_t^{\beta-1} \beta dZ_t + K_t^\alpha Z_t^{\beta-2} \beta (\beta - 1) \frac{\sigma_Z^2 Z_t^2}{2} dt,$$

and if $\beta \neq 1$ the expected growth rate $\theta_\Pi = \mu_Z + \frac{\sigma_Z^2}{2} \beta (\beta - 1)$ of the marginal revenue product flow depends on the variability of the exogenous Z_t process that (nonlinearly) determines it.

Second, and more importantly, uncertainty also influences the behavior of marginal revenues *around* their expectation. When investment is irreversible—or, more generally, when marginal adjustment costs are nonlinear and asymmetric—uncertainty increases the likelihood of both positive and negative developments, and the probability distribution the firm's regrets of its actions.

While the optimal investment policy again rules out arbitrage opportunities, on the basis of present-value considerations, the optimal solution cannot be expressed in term of an optimal adoption time T^* , as in the mushroom problem of Section 3, because at every time exogenous sources of uncertainty may have determined many different paths. To pin down the solution, it is necessary to express it in terms of the level of the exogenous processes, which (to the extent that the future is unpredictable, and along with the process's parameters) also summarizes the probability distribution of further developments.

That approach is also valid in the simpler setting where uncertainty is absent. Replacing 0 with a generic t in (1), the value at time t of a plan to adopt at time T a project (or pick a mushroom) promising cash flows that in the meantime grow at rate ϑ ,

$$V_t(T) = e^{-r(T-t)} \left(\Pi_t e^{\vartheta(T-t)} - C \right), \quad (3)$$

grows steadily at the rate of time discount r . Writing its value as a function $v(\Pi_t; \Pi^*)$ of the current level Π_t and of the "picking" size Π^* , which are one-to-one related to t and T^* in this problem, the chain rule in the absence of uncertainty implies

$$\frac{dv(\Pi_t; \Pi^*)}{dt} = \frac{\partial v(\Pi_t; \Pi^*)}{\partial \Pi_t} \vartheta,$$

and the value of the opportunity to adopt satisfies an equation in the form

$$rv(\Pi_t; \Pi^*) = \frac{\partial v(\Pi_t; \Pi^*)}{\partial \Pi_t} \vartheta.$$

Among the solutions of this differential equation, one corresponds to the best adoption strategy in the time-dependent expression (3), and satisfies a first order condition with respect to the optimal picking size $\Pi^* = \Pi_0 e^{\rho T^*} = rC$.

In the case where the growth of Π_t is disturbed by the increments of a Wiener process, this solution procedure also results in a differential equation that has to be satisfied by the value $v(\Pi_t; \Pi^*)$ of the unexploited opportunity to invest when a critical value Π^* is reached. That differential equation involves the second derivative of $v(\cdot)$ with respect to Π_t , as expected "capital gains" have to be computed on the basis of the stochastic chain rule (2), and again its solution has to satisfy a first-order condition with respect to the chosen adoption trigger Π^* .

Since readers may refer to Stokey (2009) or Bagliano and Bertola (2007) for details of the solution procedure, suffice it to say here that the optimal adoption rule when investment is completely irreversible takes the form

$$\Pi^* = \frac{x}{x-1}(r - \theta_{\Pi})C, \quad (4)$$

where x is the positive solution to

$$r = x\theta_{\Pi} + (x-1)x\frac{\sigma_{\Pi}^2}{2}.$$

Rearranging this quadratic equation to read

$$\frac{x}{x-1} = \left(r + \frac{1}{2}x\sigma_{\Pi}^2 \right) / (r - \theta_{\Pi}),$$

the optimal rule (4) implies that it is optimal to invest so as to prevent Π_t from exceeding

$$\Pi^* = \left(r + \frac{1}{2}\sigma_{\Pi}^2 x \right) C > rC.$$

Hence, irreversibility inserts a positive wedge between the marginal revenue and user cost of capital at times of positive investment. Intuitively, as the firm knows that it will be impossible to reduce the capital stock in response to future negative developments, it aims at avoid the regret implied by excessive capacity in such instances by restraining investment in good times. Recognizing that the marginal revenue product process considered takes as given one of the many possible levels of installed capital K_t , irreversibility of investment implies future decisions to install capital may only increase its stock and—under decreasing returns—reduce its marginal revenue product.

It is precisely the expectation of future excess capacity (and low marginal revenue products) that makes the firm reluctant to invest. In present expected value terms, in fact, capital's marginal revenue product fluctuates around the same user-cost level that would determine it in the absence of adjustment costs. And since fluctuations are wider in the presence of more

uncertainty, a larger σ_{Π}^2 implies a larger "irreversibility premium" in this model. Higher uncertainty also imply longer delays in investment from any given initial condition. The expected first passage time at Π^* of a geometric Brownian motion starting at Π_t is $\ln(\Pi^*/\Pi_t) / (\theta_{\Pi} - \sigma_{\Pi}^2/2)$, and increases in σ_{Π}^2 for given Π^* . Since Π^* is in turn an increasing function of σ_{Π}^2 , more uncertainty implies that, from any given starting point, investment will take longer to occur and, on average, will be less frequent in the long run.

5 Concluding comments

Stochastic calculus has proved to be extremely useful in Economics, and the irreversible investment model in a Brownian framework elegantly highlights the implications of uncertainty of a specific form. Its simple and elegant structure makes it possible to derive equally simple and elegant results, but simplicity comes at a price, as the interesting qualitatively ambiguous relationships discussed in Section 3 are replaced by the unambiguous association between higher uncertainty and less likely, more delayed action derived in Section 4 and its references. Thus, the stark assumptions and implications of models set in a Brownian motion framework may hide the qualitatively ambiguous interactions highlighted by the growing-but-possibly-disappearing mushroom problem of Section 3, where what matters for the shape and character of optimal waiting behavior is not only the intensity, but also the specific character of uncertainty (in particular, such exogenous elements of irreversibility as are implied by sudden disappearance of investment opportunities).

Continuous-time techniques nowadays are just one of several tools available for the study of inaction's causes and consequences. Qualitative insights into interactions between uncertainty, adjustment costs, and inaction may be easier to grasp in simpler technical frameworks, such as that of Section 3 above, or the discrete environment discussed in Bagliano and Bertola (2004, Chapter 3). The elegant formulas of continuous-time models should be used wisely, so that their tractability does not come at the expense of realism and substantive insights. Independent, unpredictable, continuous variation around expected trends is realistic for the price of financial assets, which in efficient markets should not offer safe arbitrage opportunities, and react only to news. It may also be a good approximation in "real" applications framed in financial contexts. For example, if it is possible to borrow and lend without constraints at a given rate, nondurable consumption variation should not be predictable (Hall, 1978). This may justify using Brownian techniques to characterize the role of consumption uncertainty in shaping the frequency and size of discrete consumer durable purchases (Bertola, Guiso, and Pistaferri, 2005), a setting where, of course, uncertainty affects optimal choices also through more standard risk and precaution effects.

Another important qualification of Brownian models relates to their usefulness for representing and interpreting real-life data that are not only unavoidably collected in discrete time, but typically vary much more often, and more smoothly, than would be implied by observation of infrequently controlled continuous processes. This can be explained by the fact that observations are aggregated, across investors and also across decisions (as each investor chooses to act along different types of capital), so as to average and inaction. The resulting inertia in the aggregate series' response to shocks is qualitatively similar to the "marginalist" continuous relationships that would be implied by linear-quadratic approximations: the extent of such similarity depends on many details of the underlying problems, and generally needs to be studied numerically, in discrete settings that, unfortunately but unavoidably, forsake the elegant simplicity of continuous time derivations.

References

- Arrow, Kenneth J. (1968) "Optimal Capital Policy with Irreversible Investment" pp. 1-19 in J.Wolfe (ed) *Value, Capital, and Growth*, Chicago: Aldine Publishing Company.
- Bagliano, Fabio, and Giuseppe Bertola (2004), *Models for Dynamic Macroeconomics*, Oxford: Oxford University Press.
- Bertola, Giuseppe (1998) "Irreversible Investment (1989 unpublished draft)" *Research in Economics*, 52, 3-37.
- Bertola, Giuseppe, Luigi Guiso, and Luigi Pistaferri (2005) "Uncertainty and Consumer Durables Adjustment" *Review of Economic Studies*, 72:4, 973-1007.
- Bar-Ilan, Avner, and Alan S. Blinder (1988) "Consumer Durables and the Optimality of Usually Doing Nothing" NBER working paper 2488.
- Dixit, Avinash K. (1991) "A Simplified Exposition of the Theory of Optimal Control of Brownian Motion" *Journal of Economic Dynamics and Control*, 15, 657-673
- Dixit, Avinash K., and Robert S.Pindyck (1994), *Investment Under Uncertainty*, Princeton: Princeton University Press.
- Hall, Robert E. (1978) "Stochastic Implications of the Permanent Income Hypothesis. Theory and Evidence," *Journal of Political Economy*, 96, 971-987.
- Hartman, Richard (1972) "The Effect of Price and Cost Uncertainty on Investment," *Journal of Economic Theory*, 5, 258-266.
- Hicks, John (1950) *A Contribution to the Theory of the Trade Cycle*, Oxford: Clarendon Press.
- Jorgenson, Dale W. (1963) "Capital Theory and Investment Behavior," *American Economic Review (Papers and Proceedings)*, 53, 247-259.
- Malliari, A.G. and William A. Brock (1981) *Stochastic Methods in Economics and Finance*, Amsterdam: North-Holland.
- McDonald, Robert, and Daniel Siegel (1986) "The Value of Waiting to Invest" *Quarterly Journal of Economics* 101:4, 707-728.
- Pissarides, Christopher (1987) "Search, Wage Bargains and Cycles" *Review of Economic Studies* 54 473-83.
- Sargent, Thomas J. (1979) *Macroeconomic Theory*, Orlando: Academic Press, Inc.
- Shapiro, Carl, and Joseph E. Stiglitz (1984) "Equilibrium Unemployment as a Worker Discipline Device" *American Economic Review* 74:3, 433-444.
- Stokey, Nancy L. (2009), *The Economics of Inaction – Stochastic control models with fixed costs*, Princeton University Press, 2009.