The *I-r* hump: irreversible investment under uncertainty

By Colin Rose

Theoretical Research Institute, Bellevue Hill, Sydney, NSW 2023, Australia;
e-mail: colin@tri.org.au

It is well known that if investment is irreversible and uncertain, there exists a benefit to waiting. When such benefits are taken into account, the relationship between interest rates and investment may be quite complex. In particular, when net revenues follow a Gaussian random walk, model investment tends to zero at both high and low interest rates. That is, investment is a hump-shaped function of r.

1. Introduction

Recent developments in the theory of irreversible investment under uncertainty can explain the existence of inertia in a remarkable variety of scenarios, both economic and purely social. In particular, standard Marshallian investment rules have been shown to be sub-optimal because they ignore the benefit of waiting in an uncertain environment. Dixit and Pindyck (1994) provide seminal coverage. Given this setting, Section 2 constructs a simple expository framework to help analyse the relationship between interest rates and aggregate investment. It does so by comparing intertemporal expected net present values in a discrete-time model in which a firm has a two-period window of opportunity. This approach is of some interest in its own right on two counts. First, it extends the analysis to the Gaussian random walk.[1] Second, tractability is enhanced: the costs and benefits of waiting can easily be identified, and there is no need for stochastic/differential calculus. Section 3 then shows that the effect of interest rates on aggregate investment can be counter intuitive. In particular, model investment tends to zero at both high and low interest rates. That is, investment is a hump-shaped function of the interest rate. Model assumptions are then relaxed, and policy implications are discussed. A Coda and Appendix follow.

[1] While a geometric Brownian motion for net revenues allows firms to make unboundedly large gains, it assumes that firms never make operating losses. This is sometimes modified by subtracting a fixed cost per period, so that losses have a precisely known lower bound. This is still unrealistic, since in reality firms do make unexpected losses. By contrast, a Gaussian random walk resolves this problem by allowing net revenues to become unboundedly large or small, with infinitesimal probability.

2. Model

There are three central elements to such analyses. First, the investment must involve an irreversible expenditure, the size of which is K. This expenditure yields an infinite sequence of net revenues $\langle R_t \rangle$. Second, there must be some uncertainty—here net revenues are stochastic, and follow a discrete time random walk

$$R_{t+1} = R_t + \varepsilon_{t+1} \tag{1}$$

where $\varepsilon \sim N(0, \sigma_\varepsilon^2)$ (Gaussian White Noise).

Third, the firm can delay its investment. In particular, we consider the case where the firm is able to delay investment by a single period of time, of arbitrary length. That is, the firm owns a two-period window of opportunity within which it can invest. The decision to invest is thus an intertemporal one: either invest today, or wait one period until tomorrow and then decide again given the new information set. Let I_0 denote the expected net present value of investment today; let I_1 denote the expected net present value of investment tomorrow, given that we only invest tomorrow if it is viable to do so, and given tomorrow's information set; finally let I_1^E denote I_1 conditional on today's information set. To emphasise that the results have nothing to do with risk aversion, it is assumed that firms are risk neutral. Then, the optimising firm should

$$\begin{aligned} &\text{Invest today} \quad \textit{iff}\ I_0 > I_1^E \\ &\text{Wait 1 period}\ \textit{iff}\ I_1^E > I_0 \end{aligned} \tag{2}$$

By contrast, the Marshallian risk neutral firm invests today if the project has a positive net present value. That is, it will invest if $I_0 > 0$. Let future revenues be discounted at a positive rate r, the opportunity cost of riskless capital, set exogenously by policy. Then I_0 and I_1 may be expressed as

$$I_0 = \frac{R_t}{r} - K \qquad\qquad I_1 = \frac{1}{1+r}\begin{cases} \dfrac{R_{t+1}}{r} - K & \text{if } R_{t+1} \geq rK \\ 0 & \text{if } R_{t+1} < rK \end{cases} \tag{3}$$

R_{t+1} is not known at time t. However, the distribution of R_{t+1} is known, and is given by equation (1) as $R_{t+1} \sim N(R_t, \sigma_\varepsilon^2)$ with pdf $\phi(R_{t+1})$ and distribution function $\Phi(R_{t+1})$. Hence, we can calculate the expectation of I_1 conditional on information at time t, which yields I_1^E. Note that I_1 is equivalent to the payoff of a call option, whereas I_1^E is equivalent to the value of this call option. As $R_t \to \infty$, the Prob $[R_{t+1} \geq rK] \to 1$. It then follows from (3) that as

$$R_t \to \infty,\ I_1^E \to \frac{1}{1+r}\mathrm{E}\left[\frac{R_{t+1}}{r} - K\right] = I_0/(1+r)$$

Similarly, as $R_t \to -\infty$, Prob $[R_{t+1} < rK] \to 1$ and $I_1^E \to 0$. We now have sufficient information to qualitatively plot both I_0 and I_1^E, as illustrated in Fig. 1.[2]

[2] The *Mathematica* package `Bounded.m` was used to generate diagrams. It is available in Rose (1993), and as item 0205-399 on MathSource <http://www.mathsource.com/>.

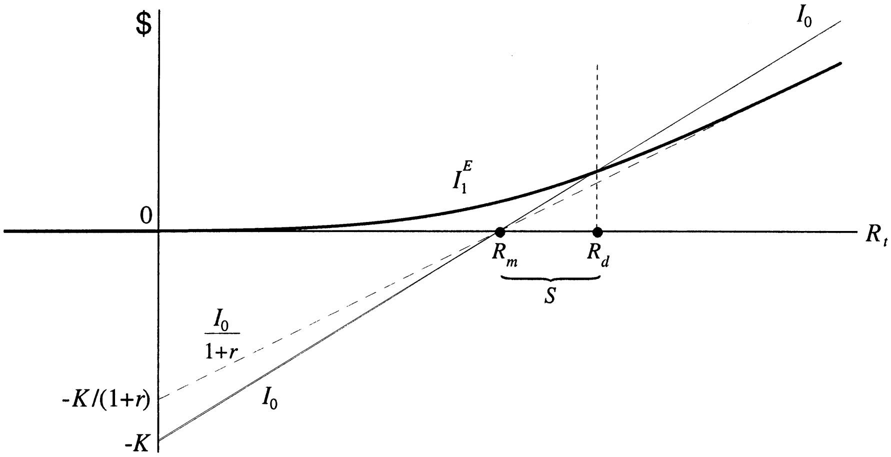

Fig. 1. The expected NPV of investing today and tomorrow

Let R_m denote the value of R_t at which $I_0 = 0$, and let R_d denote the value of R_t at which $I_1^E = I_0$. Then a Marshallian firm invests if $R_t > R_m = rK$, while an optimising firm invests if $R_t > R_d$. Since I_1^E is necessarily positive, it follows that I_0 must also be positive when $I_1^E = I_0$, and hence that

$$R_d > R_m \tag{4}$$

Then there must exist a non-empty set $S = \{R_t : R_m < R_t < R_d\}$ within which Marshallian investment rules are strictly sub-optimal: if $R_t \in S$, Marshallian rules state 'invest today', whereas optimality prescribes waiting one period and then evaluating the problem again given the new information set (as per (2)). The problem of whether or not to wait one period really amounts to evaluating the benefits and costs of waiting one period. Referring to Fig. 1, it seems quite natural to define

$$C = \text{(expected) Cost of Waiting} \quad = I_0 - \frac{I_0}{1+r}$$

$$B = \text{(expected) Benefit of Waiting} \quad = I_1^E - \frac{I_0}{1+r} > 0$$

More formally, if investment is to take place, waiting is costly because it delays the expected income stream by one time period. An investment today yields an expected income stream with NPV of I_0. If, however, we wait one period and then make the same investment, we expect to receive the identical income stream, but now delayed by one time period (the NPV of which is $I_0/(1+r)$). The benefit of waiting is that, in doing so, we can avoid potentially poor investments, at least in the light of newly received information. The appropriate investment rule is now

$$\begin{array}{l} \text{Invest today} \quad \textit{iff } C > B \\ \text{Wait 1 period } \textit{iff } B > C \end{array} \tag{5}$$

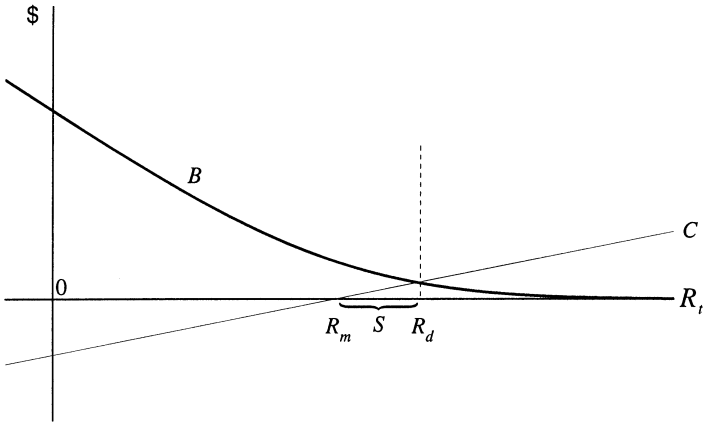

Fig. 2. The benefits and costs of waiting

Note that rule (5) is formally equivalent to rule (2), since $B = C$ if and only if $I_1^E = I_0$. Also note that the Marshallian risk neutral firm will always invest today if the cost of waiting is positive ($C > 0$ *iff* $I_0 > 0$). This is just the well-known result that Marshallian investment rules are suboptimal because they only consider the cost of waiting; all benefits are ignored. Figure 2 illustrates the benefits and costs of waiting. Note that when net revenues R_t are small or negative, the benefit from waiting is large, and the cost small (in fact negative) so that waiting is optimal. Contrariwise, when R_t is large, the benefits from waiting are small, and the costs large, such that waiting is not viable. As σ increases, the B curve shifts out and to the right, causing R_d to increase. By analogy with option pricing, these results are expected, for as is well known, the value of a call option is an increasing function of σ.

An explicit algebraic expression for I_1^E can be derived quite easily.[3] To do so, note that eq. (3) implies that

$$I_1^E = \frac{1}{1+r}\int_{rK}^{\infty}\left(\frac{R_{t+1}}{r} - K\right)\phi(R_{t+1})dR_{t+1} \tag{6}$$

As shown in the appendix, this integral may be expressed as

$$I_1^E = \frac{1}{1+r}\left\{I_0[1 - \Phi(rK)] + \phi(rK)\cdot\sigma^2/r\right\} \tag*{(6)*}$$

By providing an explicit form for I_1^E, eq. (6)* makes it possible to plot I_1^E with I_0 in real examples.

3. The hump-shaped relationship

Instead of but one investor, suppose that at any time t, there exists a pool of potential investors. For simplicity, firm i's decision to invest is assumed indepen-

[3] More broadly, dynamic programming will yield a general solution when the firm's window of opportunity exceeds two periods, but with consequential loss of elegance.

dent of firm j for all $i \neq j$: for instance, each investor i may have monopoly rights in geographic area i. As before, all firms can make an irreversible expenditure K which yields a stochastic sequence of net revenues $\langle R_t \rangle$. Equation (1) becomes

$$R^i_{t+1} = R^i_t + \varepsilon_{t+1}$$

where $\varepsilon \sim N(0, \sigma^2_\varepsilon)$ and where ε now denotes industry-level exogenous shocks. We introduce heterogeneity by allowing different firms to have different current revenues R^i_t, just as empirically, one observes a distribution of returns across an industry. In this vein, let the distribution of R^i_t across the pool of potential investors be given by $G(\cdot)$ defined over $[\theta_1, \theta_2]$ for $\theta_2 > \theta_1 \geq 0$.[4] We can now contrast two different worlds: a Marshallian world in which firms use Marshallian decision rules, and an optimising world in which firms take into account the benefit of waiting. Then, the proportion of firms $P \in [0, 1]$ investing at time t is given by

$$P_d = 1 - G(R_d) \text{ in an optimising world}$$

and by

$$P_m = 1 - G(R_m) \text{ in a Marshallian world.}$$

Several results then follow:

Proposition 1 If investment is irreversible and uncertain, and firms behave optimally, then the level of investment is always smaller than Marshallian rules suggest.

Proof $R_d > R_m$ (by (4)). Thus, $P_d < P_m$. □

Proposition 2 As $r \to \infty$ investment tends to zero in both worlds.

Proof As $r \to \infty$, $R_m = rK \to \infty$ and thus $P_m \to 0$. Hence $P_d \to 0$ (by Proposition 1). □

Proposition 3 As $r \to 0$ $\begin{cases} \text{(i) Marshallian investment tends to unity} \\ \text{(ii) Optimal investment tends to zero} \end{cases}$

Propositions 1 and 2 are well known. By contrast, Proposition 3 is somewhat surprising. Its Marshallian component (i) is of course trivial, since as $r \to 0$, $R_m \to 0$, and thus $P_m \to 1$. To derive the optimising component (ii), it helps to think in terms of the costs and benefits of waiting, and to evaluate the limit of these. It is easy to show that as $r \to 0$, $C \to R_t$. This is also intuitive: recall that waiting is costly because we delay the expected income stream by one period. If $r = 0$, the future is not discounted, so the cost of waiting one period is just the net revenue forgone in that period which is R_t. While the cost of waiting tends to R_t (finite) as $r \to 0$, the benefit of waiting tends to infinity (proof below). Hence, as $r \to 0$, it always pay to wait and thus investment P_d tends to zero.

[4] $G(\cdot)$ is defined with positive support because we constrain our interest to the potential set of entrants. Given $K > 0$, a firm is defined to be a potential entrant if there exists some combination of r, $\sigma > 0$ at which a firm with current revenues R^i_t would invest. Under model assumptions, no such combination exists for firms with negative current revenues.

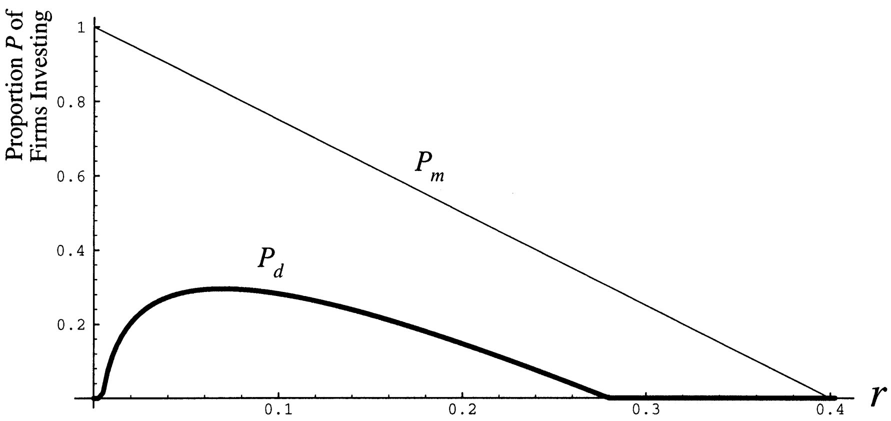

Fig. 3. Optimising investment and Marshallian investment

To prove $\lim_{r \to 0} B = \infty$

Recall $B = I_1^E - \dfrac{I_0}{1+r} = \dfrac{1}{1+r}\left[\dfrac{H}{r}\right]$ where $H = \sigma^2\phi(rK) - (R_t - rK)\Phi(rK)$ by (6)*. Hence, $\lim_{r \to 0} B = \infty$ if $\lim_{r \to 0} H > 0$ where $\lim_{r \to 0} H = \sigma^2\phi(0) - R_t\Phi(0)$. The Appendix shows that $\mathrm{E}[R_{t+1}|R_{t+1} < 0] = R_t - \sigma^2\dfrac{\phi(0)}{\Phi(0)} < 0$. Hence $\lim_{r \to 0} H > 0$. □

Since $R_m = rK$, it is clear that P_m is a decreasing function of r, as per standard analysis. By contrast, Propositions 2 and 3 combined suggest that P_d is a hump-shaped function of r. Figure 3 illustrates this surprising result. In this example, $K = 2{,}500$, $\sigma = 500$ and $G(\cdot)$ is uniform over $[0, 1{,}000]$. As per Proposition 1, the optimising investment curve always lies below the Marshallian curve.

3.1 Explaining the hump intuitively

The right-hand-side of the hump is of course intuitive: as $r \to \infty$, the present value of cash inflows tends to zero rendering projects unprofitable. The left-hand-side of the hump is less intuitive, but we do know that it arises formally because the benefit of waiting tends to infinity as the discount rate tends to zero (see proof above). Thus, to understand the hump intuitively, all we have to do is explain why $B \to \infty$ as $r \to 0$. In doing so, it helps to think in terms of the distribution of payoffs (denoted by y), instead of the distribution of net revenues. Let $y_t = R_t/r - K$ where R_t is known at time t. Similarly, let $y_{t+1} = R_{t+1}/r - K$ where $R_{t+1} \sim N(R_t, \sigma^2)$ and thus $y_{t+1} \sim N(y_t, \sigma^2/r^2)$. By waiting one period, we censor this distribution of future payoffs, reducing the downside risk while leaving the upside potential $(y_{t+1} > 0)$. To make this absolutely clear, we can re-express (3) as follows

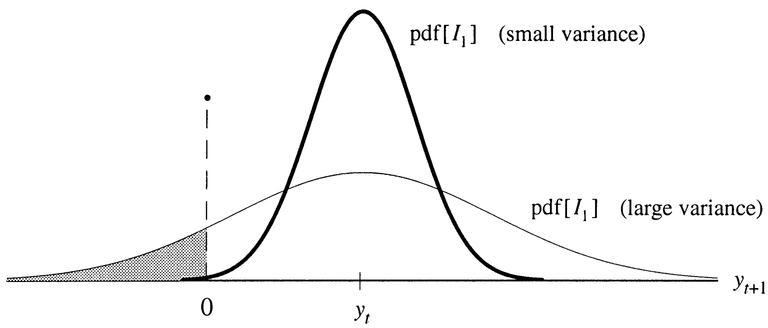

Fig. 4. Censoring becomes important as the variance increases

$$I_0 = y_t \qquad I_1 = \frac{1}{1+r}\begin{cases} y_{t+1} & \text{if } y_{t+1} \geq 0 \\ 0 & \text{if } y_{t+1} < 0 \end{cases}$$

As the variance of y_{t+1} increases, the benefit of censoring increases: more of the distribution is pushed into the tails, causing the upside potential to increase, but with minimal effect on the downside (see Fig. 4). It follows that as $r \rightarrow 0$, this benefit must tend to infinity, since $y_{t+1} \sim N(y_t, \sigma^2/r^2)$.

3.2 Policy implications

If investment is irreversible and uncertain, then the interest-rate effects of policy are complex. One can perhaps imagine embedding the investment hump into the basic IS-LM macro model to give a hump-shaped IS curve which is positively sloped at low rates of interest. In this perverse region, expansionary monetary policy lowers interest rates and output, while expansionary fiscal policy encourages economic activity both directly, and indirectly through higher interest rates.[5] The implications are perhaps reminiscent of the Keynesian liquidity trap, though rather more perverse. At the very least, by Proposition 1, monetary policy is always less effective than Marshallian rules suggest. By contrast, there exist other policy initiatives that are more effective; whereas monetary policy moves one around a given hump contour, policies that reduce uncertainty σ not only increase the magnitude of the hump, but shift the peak to the left, thus reducing the domain over which the effect of interest rates on investment is perverse.[6] Figure 5 illustrates. As per Fig. 3, $K = 2{,}500$ and $G(\cdot)$ is uniform over $[0, 1{,}000]$.

3.3 Generalising the results

We consider six extensions:

(i) The censoring argument above is quite general; it is not distribution specific.

(ii) For expository reasons, we assumed the firm can delay its investment by one period of time, of arbitrary length. That is, our firm had a simple 2-period window

[5] This assumes the LM curve is (locally) flatter than the IS curve.

[6] From p. 3, we know R_d is increasing in σ; hence, P_d decreases with σ.

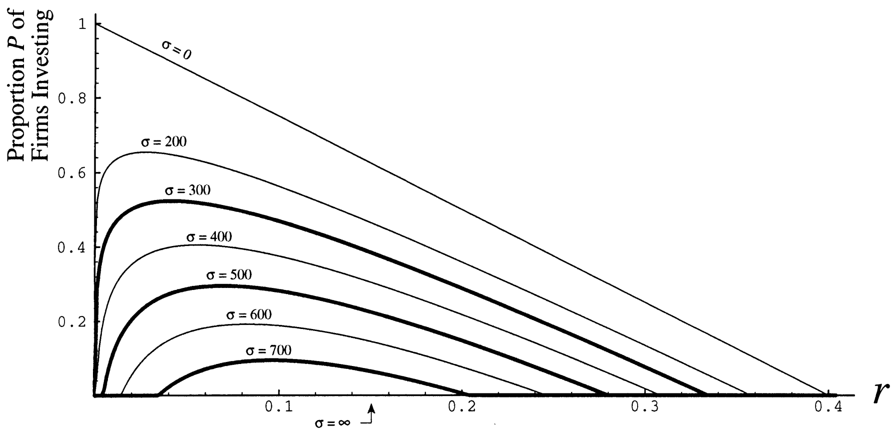

Fig. 5. Optimal investment at different levels of uncertainty

of opportunity. Proposition 3 then showed that as $r \to 0$, $P_d^{n=2} \to 0$ where n denotes the size of the window. In fact, the hump holds for all $n \geq 2$. The proof is easy (by induction):

To prove as $r \to 0$, $P_d^n \to 0 \,\forall\, n$; $n = \{2, 3, 4 \ldots\}$

By Proposition 3, as $r \to 0$, $P_d^{n=2} \to 0$. It is then sufficient to show that $P_d^{n+1} \leq P_d^n$. Our firm holds an option—an option to wait $(n-1)$ periods before committing itself. By virtue of the owner's right of exercise, the value of this option must be a non-decreasing function of n, just as the value of a call option is a non-decreasing function of the time to maturity. It follows that $R_d^{n+1} \geq R_d^n$ and $\therefore P_d^{n+1} \leq P_d^n$. □

(iii) This paper has shown that as $r \to 0$, $B \to \infty$. This is a sufficient condition for a hump to exist. However, it is not necessary for B to become infinite, and this is easy to show:

The most that firm i can lose by waiting is its current net revenue R_t^i, so $C_i < R_t^i \,\forall\, i$. Since the distribution of potential investors has support $R_t^i \in [\theta_1, \theta_2]$, it follows that $C_i < \theta_2 \,\forall\, i$. Thus, if $B > \theta_2$, there is zero investment. $\therefore B$ does not need to grow to infinity to achieve this result.

For instance, in Fig. 5, $\theta_2 = 1{,}000$. This important distinction between necessary and sufficient conditions establishes the framework for extensions (iv) and (v).

(iv) In the usual expository fashion, model investment (once made) was assumed to be infinitely lived. This generates the $1/r$ on the RHS of eq. (3), and allows one to neatly prove that $B > C$, as $r \to 0$. If investment is finitely lived, B and C are both finite, so one can no longer claim unambiguously that B is necessarily always larger than C. Nevertheless, as $r \to 0$, the same censoring principles are at work, striving to produce the same hump-shaped result, and provided the project is sufficiently long-lived, r will have sufficient leverage to generate the hump. Note that a project's life does not need to be particularly substantial for

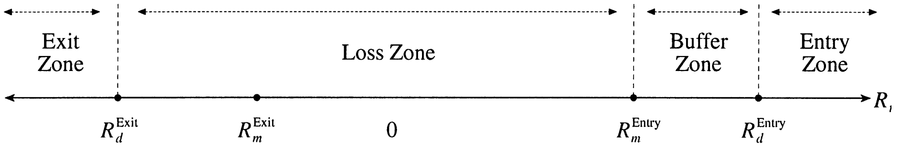

Fig. 6. Optimal entry and exit

it to have similar leverage (at the margin) to even an infinitely lived project. For instance, the NPV of an annuity of \$1 for m periods is

$$Z(m) = \sum_{t=1}^{m} (1+r)^{-t}$$

If $r = 0.1$, $Z(\infty) = \$10$ whereas say $Z(25) = \$9.077$: not dissimilar to the infinitely lived project.

(v) What if our firm not only has an option to enter, but also owns an option to exit which it can exercise by paying an exit fee $F > 0$ if events turn out badly? We then have two intertwined option pricing problems which need to be solved simultaneously. This is a non-trivial exercise that does not permit closed form solutions. Numerical methods are then required, and the end result is then example specific. Nevertheless, the essence of such a model is well known from the principles of irreversible investment: corresponding to the Marshallian entry (R_m^{Entry}) and exit ($R_m^{\text{Exit}} = -rF$) trigger points, there will now exist optimising entry ($R_d^{\text{Entry}} > R_m^{\text{Entry}}$) and exit ($R_d^{\text{Exit}} < R_m^{\text{Exit}}$) points that do take into account the benefit of waiting.[7] Figure 6 illustrates.

Is there then a residual benefit of waiting when firms have an option to exit? Yes! An option to exit only guards against outcomes in the Exit Zone (in which case we still make a loss of $K + F$). By contrast, the benefit of waiting strives to avoid all potentially poor investments: that is, it attempts to avoid both the Loss Zone in Fig. 6 and the Exit Zone (a potential saving of $K + F$).[8] Thus, if a firm has an option to exit, its worst case scenario is to lose $K + F$, so that an upper bound on the benefit of waiting one period is $(K + F)/(1 + r)$. Thus, in Fig. 2, as $R_t \to -\infty$, B will now tend to $(K + F)/(1 + r)$ rather than ∞, as Fig. 7 illustrates.

Two implications can be drawn: On the one hand, B and C will both be finite again, so one can not claim unambiguously that B is always larger than C as $r \to 0$, just like extension (iv). On the other hand, we know from extension (iii) that B does not need to be infinite to generate a hump, and therefore neither does F. For

[7] Note that R_d^{Exit} is always strictly negative. This is because even in the special case of costless exit $R_m^{\text{Exit}} = 0$, the firm will still optimally choose to incur some running losses rather than exiting (since $K > 0$).

[8] Note that an exit option does not necessarily prevent firms from making negative losses into perpetuity. After all, in order to exit, the firm pays an exit fee F which is itself tantamount to receiving a negative payoff of $R_t = -rF$ into perpetuity. It is not the duration of the loss that is critical, but its net present value.

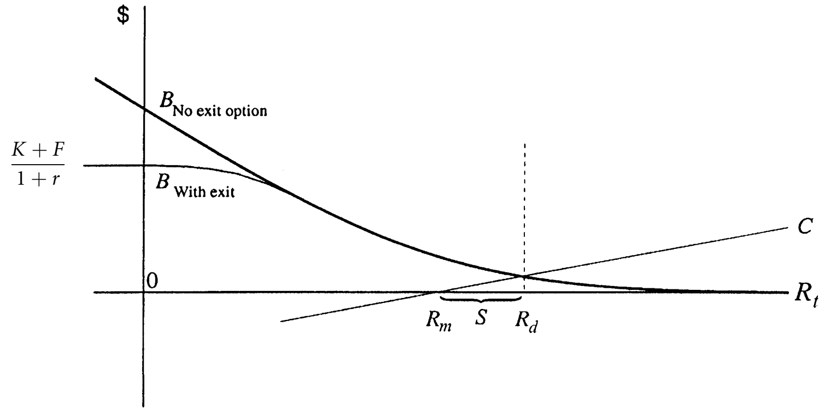

Fig. 7. The benefit and cost of waiting (with and without an exit option)

any given cost of exit $F > 0$, hump-shaped relationships will always exist, but will now depend on the example and in particular on σ, F, and the distribution of R_t^i across the industry. As F tends to 0, such examples may become special.

(vi) Risk aversion: In our world of risk-neutral agents, the optimising firm invested if $R_t > R_d$. If agents are now risk-averse, they will require an additional risk premium δ, so that investment only takes place if $R_t > R_d + \delta$. Consequently, there will be even less investment than in the risk-neutral case described above.

4. Coda

This paper adopted a discrete-time continuous state space model to simplify the analysis of irreversible investment under uncertainty. This model showed intuitively that waiting has benefits as well as costs. As the benefits are non-linear, the effects of interest rate policy on investment can be complex. In particular, the paper showed that model investment tends to zero at both high and low interest rates. The latter is surprising, and can be explained as follows. By waiting, we censor potentially bad pay-offs, whilst retaining the upside gain. As the interest rate becomes small, the variance of the payoff distribution becomes large, and thus so too must the benefit of waiting. By contrast, the cost of waiting grows more slowly, so that it eventually pays to wait. Model investment is then a hump-shaped function of the interest rate.

Adding an option to exit and/or shortening the length of the income stream reduces this variance and so may dilute the hump-effect. By contrast, the variance increases with the length of a firm's option to wait, which will exacerbate the hump. Further, the more risk averse the firm, the greater is the premium required to compensate for this variance, which will also accentuate the hump. It remains for the empirical significance of the hump to be tested. This serves as an intriguing topic for future research, especially in view of the important implications such work has for economic policy.

Acknowledgements

I thank Avinash Dixit and Jeff Sheen for helpful comments, two anonymous referees of this journal, seminar participants at LSE (Financial Markets Group), Michigan, Oxford (IES), the Reserve Bank of Australia and UCLA, and conference participants at the Australasian Econometric Society meetings.

References

Dixit, A. and Pindyck, R. (1994), *Investment under Uncertainty*, Princeton University Press, Princeton, NJ.

Rose, C. (1993), 'Bounded and Unbounded Stochastic Processes', in H. Varian (ed.), *Economic and Financial Modelling with Mathematica*, Springer-Verlag, New York, NY.

Appendix

Let $Z \sim N(0,1)$ with pdf $h(z)$ and distribution function $H(z)$. Let $X \sim N(\mu, \sigma^2)$ with pdf $\phi(x)$ and distribution function $\Phi(x)$. If $Z = \dfrac{X-\mu}{\sigma}$, one can show that

(i) $\phi(x) = \dfrac{h(z)}{\sigma}$

(ii) $\Phi(x) = H(z)$

(iii) by direct integration $\displaystyle\int_{\underline{z}}^{\bar{z}} zh(z)dz = h(\underline{z}) - h(\bar{z})$

By using the change of variable $x = \mu + \sigma z$, it follows that

$$\begin{aligned}
\int_{\underline{x}}^{\bar{x}} x\phi(x)dx &= \int_{\underline{z}}^{\bar{z}} (\mu + \sigma z)\frac{h(z)}{\sigma}(\sigma dz) && \text{applying (i), with } \underline{z} = \frac{\underline{x}-\mu}{\sigma}\cdots \\
&= \mu \int_{\underline{z}}^{\bar{z}} h(z)dz + \sigma \int_{\underline{z}}^{\bar{z}} zh(z)dz \\
&= \mu[H(\bar{z}) - H(\underline{z})] - \sigma[h(\bar{z}) - h(\underline{z})] && \text{applying (iii)} \\
&= \mu[\Phi(\bar{x}) - \Phi(\underline{x})] - \sigma^2[\phi(\bar{x}) - \phi(\underline{x})] && \text{applying (ii) and (i)}
\end{aligned}$$

Thus, if $R_{t+1} \sim N(R_t, \sigma^2)$ two results follow. First:

$$\int_{rK}^{\infty} R_{t+1}\phi(R_{t+1})dR_{t+1} = R_t[1 - \Phi(rK)] + \sigma^2\phi(rK) \qquad (\text{set } \underline{x} = rK, \bar{x} = \infty) \tag{A1}$$

The transition from (6) to (6)* is then straightforward. Second:

$$\mathrm{E}[R_{t+1}|R_{t+1} < 0] = \frac{1}{\Phi(0)}\int_{-\infty}^{0} R_{t+1}\phi(R_{t+1})dR_{t+1} = R_t - \sigma^2\frac{\phi(0)}{\Phi(0)} \quad (\text{set } \underline{x} = -\infty, \bar{x} = 0) \tag{A2}$$