"Option Learning as a Reason for Firms to Be Averse to Idiosyncratic Risk"

June 26, 2019

Eric Rasmusen

Abstract

The standard reasons why widely held corporations might be averse to idiosyncratic risk as well as systematic risk are based on the principal-agent problem, bankruptcy costs, external finance, and tax convexity. This paper offers a different reason: idiosyncratic risk make learning the quality of business decisions more difficult. It is well known that risk can actually increase the value of investment projects because of option value. We must distinguish, however, between risk over the expected value of profits ("value risk") and risk over the volatility of cash flows ("cash-flow noise"). Value risk is good, because an unprofitable policy can be abandoned. Cash-flow noise is bad, because it makes learning when to abandon a decision more difficult. This distinction is different from Knightian risk versus uncertainty, is unrelated to ambiguity aversion, and matters even if there is no principal-agent problem.

Rasmusen: Dan R. and Catherine M. Dalton Professor, Department of Business Economics and Public Policy, Kelley School of Business, Indiana University. 1309 E. 10th Street, Bloomington, Indiana, 47405-1701. (812) 855-9219. erasmuse@indiana.edu, http://rasmusen.org.

This paper: http://rasmusen.org/papers/risk_aversion.pdf.

Keywords: idiosyncratic risk, risk aversion, business risk, real options, learning, capital budgeting, risk management, value risk, cash-flow noise.

JEL Codes: D21, D25, D81, G32, L21.

I would like to thank Philip Bromiley, Brett Green, David Hirshleifer, and Robert Pindyck for helpful comments, and Aaron Kolb for exceptionally helpful input.

Aimed at AER-Insights. 5427 words as of June 21 or so. Abstract is 77 words.

1. Introduction

Risk is commonly divided into two categories: systematic and idiosyncratic. A company's systematic risk is the part of its risk correlated with the returns an investor could earn by holding other assets in the economy. Its idiosyncratic risk is the part which is not. Under the CAPM and other theories of asset pricing, a publicly traded corporation should be averse to systematic but not to idiosyncratic risk, because its shareholders can diversify away idiosyncratic risk by holding a variety of assets. Thus, firms should not take idiosyncratic risk into consideration.

Yet it seems that firms do avoid idiosyncratic risk in both investment decisions (Panousi & Papanikolaou, [2012]) and operating decisions ((Amit & Wernerfelt [1990], Chatterjee, Lubatkin & Schulze [1999]). They manage enterprise risk (Bromiley, McShane, Nair & Rustambeko [2015]), hedge against movements in commodity prices and currencies (Stulz [1996], Smith [2008]), smooth accounting earnings (Burgstahler & Dichev [1997]), and buy insurance (Harrington, Niehaus & Risko [2002]). There may even be a "diversification discount": investors prefer firms with diversified lines of business (Villalonga [2004], though this is unclear: see Laevena & Levine [2007]).

The most common explanation is that to incentivize managers the firm must link pay to firm-specific performance. As a result, managers will wish to reduce risk regardless of shareholder desires, and even if the shareholders can control the managers they will need to design complex contracts and pay them extra for bearing risk (Smith [2008] Panousi & Papanikolaou, [2012], Glover & Levine [2017]). A second explanation is bankruptcy costs. Bankruptcy arises even from idiosyncratic risk and creates not just direct costs in reorganization but indirect costs because of anticipatory moves by lenders, suppliers, and customers (Stulz [1996], Kuerstein & Linde 2017]). A third explanation is that external finance is more costly than retained earnings as a source of capital. If retained earnings are variable, the firm will sometimes be short of capital and sometimes have more than it has use for

(Froot, Scharfstein & Stein [1993]). A final explanation is convex taxation, which makes it advantageous to smooth profits (Nance, Smith & Smithson [1993]).

This article explores a different explanation: risk's effect on the quality of business decisions. If the firm encounters more noise in evaluating whether an activity is profitable or not, it will make worse decisions, so it choose avoid risky activities. Option value has been the topic of much study, some focussed on adjustment costs and competitive conditions (Hartman [1972], Abel [1983], Kim & Kung [2017]), some on valuation (Dixit & Pindyck [1994], Trigeorgis [1996], Lambrecht [2017]), some on management implications (Trigeorgis & Reuer [2017]). The central idea is that option value increases with volatility. If an activity might have either high or low returns, the firm has the option to abandon it if the return is low, whereas the upside potential is unlimited. This, however, suggests that corporations should actually be risk-loving, not risk-averse.

On the other hand, volatility makes learning more difficult, and idiosyncratic risk is even worse than systematic risk, since it cannot be deduced from the returns of the stock market. Two forces are at work: the advantage of a more valuable option, and the disadvantage of not knowing what to do with it. Thus, let us distinguish between "value risk" and "cash-flow noise". Value risk increases with divergence in possible expected values. Cash-flow noise increases with the variability of cash flows increases, for given expected value. Value risk has positive option value because the firm has the option to discontinue the policy if its learns its expected value is low. Cash-flow noise hurts option value because it makes it more difficult for the firm to tell when it should exercise that option. This paper will discuss the difference using examples and a simple Brownian-motion model.

Value Risk versus Cash-flow Noise

There are two reasons one might say, "This project is risky." The first is that one doesn't know whether the project is profitable. The second is that even though its profitability is known, the cash flow varies over time. Profitability is akin to Aristotle's "substance", a data-generating process, with cash flows as the data (Aristotle's "accidents"). We will call dispersion in possible expected values "value risk", and dispersion in cash flows around the expected value as "cash-flow noise". More formally, let activity A have its cash flow over time of X^A_{it}, $t = 1, \dots T$ be given by stochastic process X^A_i, $i = 1, \dots N^A$ with probability F^A_i from distribution F^A. Using the standard partial ordering of riskiness in Rothschild & Stiglitz (1970), we will say that activity A has greater *value risk* than B if any agent with strictly concave utility would prefer $\{E_t(X^B_i), F^B_i\}$, $i = 1, \dots N^B$ to $\{E_t(X^A_i), F^A_i\}$, $i = 1, \dots N^A$. We will say that activity A has greater *cash flow noise* than B if any agent with strictly concave utility would prefer $\{X^B_i - E_t(X^B_i), F^B_i\}$, $i = 1, \dots N^B$ to $\{X^A_i - E_t(X^A_i), F^A_i\}$, $i = 1, \dots N^A$.

Confusing the two kinds of risk is easy because we model them the same way. Compare project X with a 70% chance of yielding a steady cash flow of +1 and a 30% chance of 0, and project Y with a 70% chance of yielding a cash flow of +1 and a 30% chance of 0 in any particular year. They each have an annual expected value of .7, but project X has just value risk and project Y has just cash-flow noise.

Example 1: Simple or Complex Product?
A firm is considering which of several new brands of soap to introduce for a three-year product cycle (after which the brand will be replaced by something else). The firm is either "well-suited" or "ill-suited" with respect to this particular product, with probability .5 of each, and does not know its ability in advance. Market conditions will be "normal" in some years and "difficult" in others. Product A is the safest, with no risk whatsoever. Product B's value is unknown, but at least its cash flow will not depend on market conditions. Product C is the most complicated, yielding profit in normal years, but losses in difficult years. Assume zero time discounting.

Product A yields a yearly cash flow of 0, regardless of market conditions or whether the firm is well-suited.

Product B yields a yearly cash flow of +1 if the firm is well-suited and -1 if ill-suited, regardless of market conditions. Product B's expected cash flow is 0.

Product C yields a cash flow of +2 in a normal year, which has 70% probability, regardless of ability. If the firm is well-suited, it yields a cash flow of -4/3 in a difficult year, for an expected yearly cash flow of +1. If the firm is ill-suited, it yields a cash flow of -8 in a difficult year, for an expected flow of -1. Product C's overall expected cash flow is also 0.

If the firm chooses Product A, it has no decisions to make, and the product's lifetime value is 0. If the firm chooses Product B, it will immediately discover if it is ill-suited because the cash flow will be -1 the first year, and it will cancel the product. The expected three-year profit consists of a 50% chance of -1+0+0 and a 50% chance of +1+1+1, an expected profit of +1, even though if the firm refused to cancel the expected value would be 0, as with Product A.

If the firm decides to choose Product C, the best strategy is to keep selling until a year in which the cash flow is -8, which would make plain that the expected cash flow is -1, and then to cancel. If the firm is ill-suited, its probability of selling for one year would then be .3, for two years, .21, and for three years, .49 (which includes probability .343 of never discovering the truth), for an expected profit of -2.19. If the firm is well-suited, its expected three-year profit is +3. Overall, expected profit is .405.[1] Product C is better than Product A, because it has option value, but worse than B, because learning C's value is more difficult.[2]

[1] $.405 = .5[.3(-8) + .21(2-8) + .147(4-8) + .343(6)] + .5[3]$.

[2] Whether the firm can switch products in the middle of the three-year product life cycle does not matter. Even with switching, if product C is chosen and the firm then discovers it was ill-suited the first year, C would be discontinued, but B would be no more profitable for the ill-suited firm.

In Example 1, value risk is desirable. Product B's uncertainty makes it superior to Product A— but only because the firm both (a) learns, and (b) has the option to cancel the product. Product C has even more uncertainty than B, but the wrong kind: cash-flow noise. That uncertainty does not affect option value directly, because the firm has the same option to cancel an unprofitable product. Indirectly, however, cash-flow noise makes learning more expensive: it costs .5(-1) to learn that product B is unprofitable, compared to .5(-2.19) for product C.

This phenomenon is not ambiguity aversion, the Ellsberg Paradox's idea that even if two lotteries have the same risk, if one of them is a "lottery of lotteries" people will avoid it (Ellsberg [1961], Ju & Miao [2012]). Product C is indeed a lottery of lotteries, a 50-50 gamble between two different 70-30 lotteries rather than a 50-50 gamble between two integers like product B. As with the Ellsberg Paradox, easier decisionmaking is preferred. An executive viewing an MBA's powerpoint proposal to undertake product C might well worry about the cognitive cost of calculating the value of complex lotteries, the likelihood of making a mistake, and the danger of being fooled, but here what is driving his distaste for C is that it requires more expensive experimentation. This will be apparent in the next section, where we will construct a model in which the two alternatives are equally simple, with cash flow distributions that differ only in a single variance parameter.

3. A Continuous-Time Model

Stochastic calculus has often been used to look at adoption and abandonment decisions (see Pindyck [1993], Dixit & Pindyck 1994, p. 345], Bernardo & Chowdhry [2002], Decamps , Mariotti & Villeneuve [2006], Kwon & Lippman [2011]). The model below makes no claim whatsoever to technical novelty; indeed, it is a special case of Lippman & Ryan (2003), despite its different look.[3]

Consider a single project whose type $\theta \in \{H, L\}$ is initially unknown, drawn from a prior distribution with $\Pr(\theta = H) = p_0 \in (0, 1)$. The project delivers a stochastic cash flow X_t according to a continuous-time diffusion process:

$$dX_t = \mu(\theta)dt + \sigma dW_t, \tag{1}$$

where W_t is a standard Brownian motion process for positive mean $\mu(\theta)$ and standard deviation σ. Let the possible means of the process be $\mu_H = m$ and $\mu_L = -m$. The firm discounts cash flows at rate r. Its only decision, which it makes every instant, is whether to continue the project or abandon it.

Here the parameter m represents the amount of value risk and σ the amount of cash-flow noise. As time passes, the firm continuously updates its posterior belief p_t about $m(\theta)$ using Bayes's rule. Rather than analyzing the cash flow path, X_t, it is most convenient (see section 2.3 of Daley & Green [2012] or section A1 of Kolb [2019]) to analyze the belief path, p_t, converting to log-likelihood beliefs, Z_t that update linearly and vary on $[-\infty, +\infty]$ as p_t varies on [0,1]:

$$Z_t \equiv \ln \frac{p_t}{1 - p_t} \tag{2}$$

Using filtration results standard in the literature (e.g., section 3 of Bolton & Harris [1999]), the law of motion for the posterior belief conditional on the type being $\theta \in \{H, L\}$ is

$$Z_t^H = \frac{\phi^2}{2}dt + \phi dW_t \tag{3}$$

$$Z_t^L = -\frac{\phi^2}{2}dt + \phi dW_t, \tag{4}$$

where $\phi \equiv \frac{\mu_H - \mu_L}{\sigma} = \frac{2m}{\sigma}$.

[3]Lippman & Ryan (2003) uses m_H and $-m_L$ rather than m and $-m$. Imposing symmetry makes my propositions simpler and stronger (because less general), and allows value risk to be parameterized by m.

The firm's optimal strategy is to abandon the project whenever the belief Z_t falls to some lower threshold z^*, as stated in the following proposition.

Proposition 1. *The firm's optimal abandonment threshold is*

$$z^* = \ln\left[\frac{m^2 + r\sigma^2 - m\sqrt{m^2 + 2r\sigma^2}}{r\sigma^2}\right] < 0, \tag{5}$$

and its value function is

$$V(z) = \begin{cases} \frac{e^z-1}{e^z+1}\frac{m}{r} + \frac{1-e^{z^*}}{e^z+1}e^{m(z-z^*)}\frac{m}{r} & \text{if } z \geq z^* \\ 0 & \text{otherwise,} \end{cases} \tag{6}$$

where $m = \frac{1}{2} - \sqrt{\frac{1}{4} + \frac{r\sigma^2}{2m^2}} < 0.$

Proof. By Ito's Lemma, since the expected flow return is $Ep_\theta = pm + (1-p)(-m) = \frac{e^z-1}{e^z+1}m$, the value function prior to abandonment satisfies

$$rV(z) = \frac{e^z-1}{e^z+1}m + \frac{e^z-1}{e^z+1}\frac{2m^2}{\sigma^2}V'(z) + \frac{2m^2}{\sigma^2}V''(z). \tag{7}$$

This differential equation has solutions of the form $V(z) = \frac{m}{r}\frac{e^z-1}{e^z+1} + C_1\frac{e^{mz}}{e^z+1} + C_2\frac{e^{(1-m)z}}{e^z+1}$, where m is defined in the proposition. We know that the value is $\frac{m}{r}$ if the project type is high, which is belief $p_t = 1$ (and thus $Z_t = ln\frac{p_t}{1-p_t} \to \infty$), so it must be true that $\lim_{z\to+\infty} V(z) = \frac{m}{r}$ and we can conclude that $C_2 = 0$. Pinning down z^* and $C_1 = \frac{m}{r}\left(1 - e^{z^*}\right)e^{-mz^*}$ then follows from the standard value matching and smooth pasting conditions $V(z^*) = 0$ and $V'(z^*) = 0$ (Dixit [1993]). Simplifying the second term yields (6). □

Figure 1:
Cash Flows for High and Low Volatility

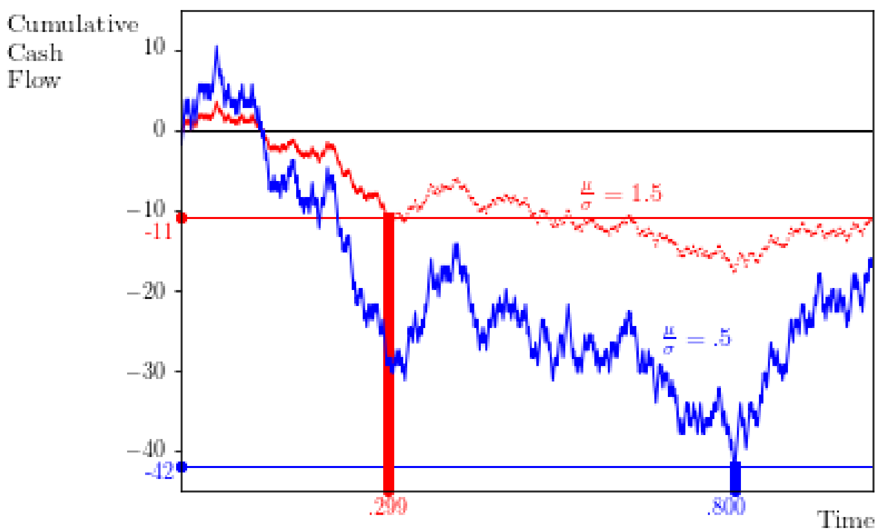

Figure 2:
Beliefs for High and Low Volatility

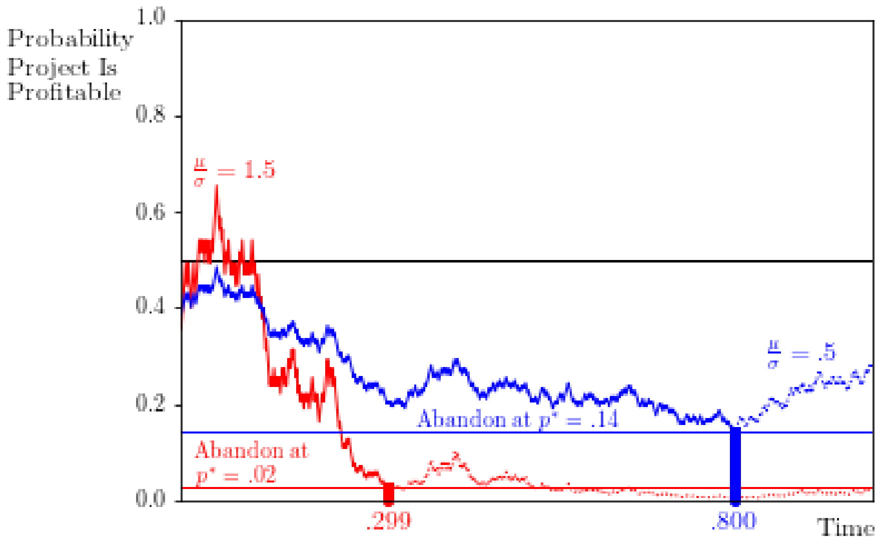

Figures 1 and 2 show examples of paths over time when $r = .12$, and the belief starts at 40% probability the project is profitable but it is actually unprofitable ($\mu(\theta) = -15$).[4] The red paths are for $\sigma = 10$ and the blue paths for $\sigma = 30$, using the same discrete approximation realization of Weiner process dW_t. The optimal thresholds $p^*(\sigma)$ are .02 and .14, which are below .50 because of option value. With low volatility, the belief that the project is high must fall all the way to 2% before abandonment, but the information is so good that this happens quickly, in period 299, after losses have only accumulated to 11. With high volatility, the belief only needs to fall to 35%, but the information is poor enough that this happens in period 800, after losses have accumulated to 42.

Returning to the optimal value function, equation (6), note that it is composed of two terms. The first term represents the expected value of running the project forever, $(2p_t - 1)\frac{m}{r}$, since $\frac{e^{z_t}-1}{e^{z_t}+1} = 2p_t - 1$. That term is positive if $p_1 > 0$ and negative if $p_t < 0$. The second term represents the gain from the ability to abandon the project after poor returns. It is always positive, because $z^* < 0$ so $1 - e^{z^*} > 0$. If $z_t = z^*$, then $\frac{1-e^{z^*}}{e^{z^*}+1} = 2p^* - 1$ and $e^{m(z_t - z^*)} = 1$, so the second term's magnitude is exactly equal to the first term's. The value to running the project another instant is $rV(z_t)$, so when $z_t = z^*$, the positive learning flow from continuing is exactly cancelled by the negative cash flow. This, of course, is what we used in solving the $V(z)$ function, because for any lower z_t the option value would not be worth the cash flow loss. The next proposition shows how value risk and cash-flow noise affect the tradeoffs.

Proposition 2. *As cash-flow noise σ rises, the abandonment threshold v^* increases and the project's value falls. As value risk m rises, the abandonment threshold decreases and the project's value rises.*

Proof. In what follows, we use $C_1 = \frac{m}{r}\left(1 - e^{z^*}\right)e^{-mz^*}$ from Proposition 1's proof so we can write the positive part of the value function as, for

[4] The Python 3 code is available at http:\rasmusen.org/papers/risk_aversion.py.

$z > z^*$,

$$V^+(z) = \frac{m}{r}\frac{e^z - 1}{e^z + 1} + C_1 \frac{e^{mz}}{e^z + 1} \tag{8}$$

We will define $F(z)$ to be the function which extends $V^+(z)$ over the entire real line. From Proposition 1's proof, $V(z^*) = 0$ and $V'(z^*) = 0$, and $V(z) > 0$ for large enough z, so since $F(z)(e^z + 1)$ is an exponential polynomial with at most two sign changes it follows that $F(z)$ is nonnegative with a single minimum at $z = z^*$.

First, consider the effect of an increase from σ to $\tilde{\sigma}$ and let tilde denote the solution components for $\tilde{\sigma}$. We claim that $F(z)$ and $\tilde{F}(z)$ can intersect at most once, since $\Delta(z) := (e^z + 1)[\tilde{F}(z) - F(z)] = \tilde{C}_1 e^{\tilde{m}z} - C_1 e^{mz}$, where the expression in brackets is an exponential polynomial with exactly one sign change. Now $\frac{\partial z^*}{\partial \sigma} = \frac{2m}{\sigma\sqrt{m^2+2r\sigma^2}} > 0$, so $\tilde{z}^* > z^*$. The minimum of $\tilde{F}(z)$, but not of $F(z)$ is at $\tilde{z}^*$, so $\tilde{F}(\tilde{z}^*) = 0, F(\tilde{z}^*) > 0$, and $\Delta(z^*) < 0$. Since $\Delta(z^*) > 0$ by the same kind of argument, there exists a root $\hat{z} \in (z^*, \tilde{z^*})$ of $\Delta(z)$ where $\Delta(\hat{z}) = 0$ and $\tilde{F}$ intersects F from above. Since that is the only sign change $\Delta(z)$ can have, it follows that that $F(z) > \tilde{F}(z)$ for $z > \hat{z} < \tilde{z}^*$. Recall that $V(z)$ is identical to $F(z)$ except for equalling zero on $z < z^*$. Hence, $\tilde{V}(z) = 0$ for all $z < \hat{z} < \tilde{z}^*$, so $V(z) > \tilde{V}(z)$ for all $z > z^*$ (and $V(z) = \tilde{V}(z) = 0$ for $z \leq z^*$).

Next, consider the effect of an increase from m to $\tilde{m}$, and use tilde to denote solution components for $\tilde{m}$. We have $\Delta(z) := (e^z+1)[\tilde{F}(z) - F(z)] = \frac{\tilde{m}-m}{r}(e^z - 1) + \tilde{C}_1 e^{\tilde{m}z} - C_1 e^{mz}$. Now $\frac{\partial z^*}{\partial m} = -\frac{2}{\sqrt{m^2+2r\sigma^2}}$, so $\tilde{z}^* < z^*$, and thus Δ has a root in $(\tilde{z}^*, z^*)$ where Δ crosses 0 from below. We claim that any root above $\tilde{z}^*$ must cross from below, and thus there is only one root above $\tilde{z}^*$. Note that if z is a root, then $\tilde{F}(z) = F(z)$ so $\Delta'(z) = \Delta'(z) - \tilde{m}\Delta(z) = \frac{\tilde{m}-m}{r}[e^z - \tilde{m}(e^z - 1)] + C_1 e^{mz}(\tilde{m} - m)$. It is easy to verify that $\tilde{z}^* = \ln \frac{-\tilde{m}}{1-\tilde{m}}$, so the first term is positive for all $z > \tilde{z}^*$. The second term is always positive because m increases in m, so the claim holds. It follows that for all $z \geq z^*$, we have $\Delta(z) = \tilde{F}(z) - F(z) > 0$. Moreover, for $z \leq \tilde{z}^*$, $\tilde{V}(z) = V(z) = 0$ and for $z \in (\tilde{z}^*, z^*)$, $\tilde{V}(z) > 0 = V(z^*)$. We conclude that $\tilde{V}(z) \geq V(z)$ for all $z \in \mathbb{R}$, with strict inequality for all $z > \tilde{z}^*$. $\square$

The belief about project type changes faster if m is big and σ is small. This is equivalent to faster learning as news comes in from the cash flows. A bigger σ means bigger movements in cash flows, a more variable cash flow process, but when converted to belief space it means smaller movements in beliefs, because a given change in cash flow is less communicative of whether the state is high or low. As the cash-flow noise σ increases, the signal-to-noise ratio of the cash flow signal decreases, and the firm learns more slowly about the project's type. This reduces the option value to continuing the project to try to learn more, the second term in the value function. When the option value falls, that makes the firm worse off and causes it to abandon projects sooner. Thus, the firm should be averse to cash-flow noise because it makes decisionmaking harder, whether or not its managers and shareholders are risk-averse in their utility functions.

On the other hand, when value risk increases the spread between the expected returns $-m$ and m of low and high projects, the firm benefits twofold. First, the upside potential m/r from running a good project forever increases, while the downside potential is still bounded below at zero, as the firm can always abandon it. This is the standard result that increasing variance increases an option's value. This first benefit would apply even if the firm's learning was instantaneous, as with Product A in Example 1. The second benefit is that the signal-to-noise ratio increases and the firm can more quickly learn about the project's type. The firm can then respond by with more aggressive experimentation. Its abandonment threshold is lower because it is willing to continue the project even under a pessimistic belief because quicker learning means quicker abandonment if the project type is indeed low.

We have described the model as one of project choice, but it could apply to any of a business's activities. In this setting of normally distributed risk, we can also interpret a higher value of σ as the result of adding normally distributed zero-mean noise to an existing activity, since adding an independent zero mean normal variable to an existing normal variable increases the variance while leaving normality and the mean unchanged. In particular, one could view accounting systems as attempts to reduce cash-flow noise. In practice, standard business earnings are not cash flows, which are higher-variance measures of profitability, but flows of accruals and depreciation, which require careful definition intended to balance current economic profits against variability across time that makes determining profitability more difficult. Thus, we can view accounting imperfections as a source of risk in "cash flows". Deficiencies in accounting quality would often merely be idiosyncratic risk, but it is risk that hurts the quality of investment decisions.

Knightian Uncertainty, and the Principal-Agent Problem

Splitting risk into two types brings to mind Knight's distinction between risk and uncertainty (Knight [1921]). What Knight was getting at remains controversial (see Leroy & Singell [1987]), but one interpretation is that Knightian risk is unpredictability with known probabilities in repeated situations, the kind of risk that is insurable, whereas Knightian uncertainty is unpredictability in one-time situations where the probabilities are more subjective. In Example 1, the cash-flow noise was about whether the market was normal or difficult, something objective that would seem to be Knightian risk. The value risk was that the firm might be either well-suited or ill-suited, something "unknowable" even though the firm will have subjective probabilities, and would thus be Knightian uncertainty. Similarly, in the continuous-time model, the cash flows were Knightian risk, while the expected value was either high or low, but unknown: Knightian uncertainty. Example 2 will illustrate that this is not a reliable equivalence.

Example 2: Hedging Foreign Exchange Risk

An American firm has set up a London office for a three-year period. The new office is a "good" or "bad" idea with equal probability. The office buys and sells in both dollars and pounds, so it is unclear how the exchange rate affects its profitability. If the office is a good idea, though, the profit will be +1 if the exchange rate moves beneficially, which has probability .70, and 0 if it moves harmfully, an expected annual profit of +.7. If the office is a bad idea, the profits will be 0 or -1, an expected profit of -.3. The firm's optimal strategy is to open the office, and shut it down if and only if its profit is seen to be -1 in some year. If the idea is a good one, the expected profit will then be 2.1. If the idea is a bad one, the office will have 30% probability of closing after one year, 21% after two years, 14.7% after three years due to losses in the third year, and 34.3% of operating profitably for three years and closing down only because the three-year period is over. Its expected profit if the idea is bad is -.657. Overall, balancing the probability the idea is good with that it is bad, the expected profit is .7215.[5]

Now suppose the firm instructs the office to hedge its foreign exchange risk. If the office is a good idea and hedging is costless, its profit each year is +.7, while if it is a bad idea, its profit each year is -.3. The firm will know to shut down the first year if its profit is -.3, so the expected profit over the three years is .5(-.3) + .5 (.7+.7+.7) = .90. This is higher than .7215, so the hedging helps, and the firm would be willing to pay as much as .1785 in transactions costs to be able to hedge.

As in Example 2 and the continuous-time model, the exchange rate risk is insurable and objective, so it would seem that again cash-flow noise corresponds to Knightian risk. Whether the London office is a good idea or not appears to be subjective and uninsurable, so value risk would correspond to Knightian uncertainty. Knightian uncertainty would thus be good for profit, and Knightian risk would be bad.

But Example 2 can be interpreted to give the opposite results. Suppose the firm has extensive experience setting up foreign offices, and

[5] Another strategy is to close down the London office the first year if -1 is observed, or the second year if 0 is observed in the first year and -1 or 0 in the second. That would have expected profit of -.51 if the idea is bad, of -.51 and 1.547 if it is good, an average of .5185.

has found that in exactly 50% of them the idea is a good one. Then the value risk is Knightian risk. Moreover, while exchange rate risk is insurable, the big problem is that the firm doesn't know enough about the inner workings of the office to know whether it benefits from a strong dollar or a weak one. Thus, the cash-flow noise is Knightian uncertainty. The questions of whether there is opportunity for learning, value risk, and whether learning will be difficult, cash-flow noise, are separate from whether the probabilities are subjective or objective.

For our last example, we will return to the principal-agent problem. We have emphasized in our discussion that option learning is an alternative to managerial risk aversion as an explanation for corporate aversion to idiosyncratic risk. It is also, however, a supplement to the principal-agent as an explanation, a supplement that applies even if the agent is risk-neutral. That noise worsens a principal-agent problem is not a new idea. Holmstrom's classic 1979 "sufficient statistic" paper notes how the informativeness of a signal is hurt by additional risk and the negative welfare effect shows up naturally in many models, even if, it seems, there has been little or no comment on the implication for company decisions. (One exception is D. Hirshleifer & Suh [1992], in which a firm avoids high-variance projects because of the difficulty of incentivizing the manager.) Our next example will be of the problem of incentivizing a risk-neutral manager, to show that option learning does not depend on anyone in the model having concave preferences.

Example 3: The Principal-Agent Problem. Management is considering one of two protocols for managing customers. Under Protocol X, the worker dealing with a customer reads from a script. If the worker exerts high effort, the customer is satisfied, a value of +10 for the firm. If he exerts low effort, the customer is dissatisfied, a value of -10. Under Protocol Y, the worker is trained to try to adapt to the particular customer. If the worker exerts high effort, the customer is highly satisfied with probability .5, a value of +40 for the firm and dissatisfied with probability .5, a value of -20, which comes to an expected value of +10. If he exerts low effort, the customer is satisfied with probability

1/6 and dissatisfied with probability 5/6, which comes to an expected value of -10. The workers are risk-neutral, with the utility function $w - e$ for wage w and effort $e = 0$ or 3, and their expected utility must equal at least 1 to take the job. They cannot be paid a negative wage.

It is easy to see the optimal wage for Protocol X. If the company pays the worker 4 if the customer is satisfied and 0 otherwise, the worker would have utility of 1 (= 4-3) from high effort and 0 from low effort and so would choose high effort. With the contract $(w_{sat} = 4, w_{diss} = 0)$, the firm ends up obtaining +10 from the customer and paying 4 to the worker.

Protocol Y takes a bit more consideration. The participation constraint is $E(w - e|high\ effort) = .5w_{sat} + .5w_{diss} - 3 \geq 1$. The incentive compatibility constraint is $E(w - e|high\ effort) \geq E(w - e|low\ effort)$, which is $.5w_{sat} + .5w_{diss} - 3 \geq (1/6)w_{sat} + (5/6)w_{diss} - 0$. Because of the nonnegative-wage constraint, the biggest gap between w_{sat} and w_{diss} is if $w_{diss} = 0$, so the optimal contract will choose that. The participation constraint then requires $.5w_{sat} + .5(0) - 3 \geq 1$: the worker is willing to take the job if the contract is $(w_{sat} = 8, w_{diss} = 0)$. But the incentive compatibility constraint tells us that $.5w_{sat} + .5(0) - 3 \geq (1/6)w_{sat} + (5/6)(0) - 0$: the lowest w_{sat} that will induce high effort is $w_{sat} = 9$. As a result, if it uses Protocol Y the firm will wish to use an "efficiency wage", a contract that gives the worker more than his reservation utility because the firm needs to be generous to induce appropriate effort (Yellen [1984], Rasmusen [2007 ch. 8]). With the contract $(w_{sat} = 9, w_{diss} = 0)$, the firm would end up obtaining an expected value of +10 from the customer and paying an expected wage of 4.5 to the worker.

What Example 3 illustrates is that even if both employer and employees are risk-neutral, risk in cash flows can create contracting costs. The basic idea in principal-agent models is that if agents are risk-averse then when output is noisy the optimal contract balances the incentivizing effect of incentive pay against the extra risk that requires a higher expected wage. In Example 3, the agent needs no compensation for bearing risk since he is risk-neutral, but noisy output requires divergence in pay to incentivize effort. When pay is bounded below by zero, the firm needs to increase the upper bound and thus the expected value, of pay.

Concluding Remarks

What's wrong with risk? Our usual answer is concavity of utility: sometimes you have too much, sometimes too little. We focus so much on utility that we neglect another answer: Risk makes it hard to figure out what to do. When the author's parents retired to a farm, his mother hung up a placque on the front porch saying: "Simplify". Executives live a frenzied existence. Much of their job is to solve unexpected problems. That does not leave enough time to evaluate people and projects. They want simplicity, and a steady bottom-line, whether positive or negative, simplifies decisionmaking.

References

Abel, Andrew B. (1983) “Optimal Investment under Uncertainty,” *American Economic Review*, 73: 228– 233.

Amit, Raphael and Birger Wernerfelt (1990) “Why Do Firms Reduce Business Risk?” *The Academy of Management Journal*, 33: 520–533.

Bernardo, Antonio E. and Bhagwan Chowdhry (2002) “Resources, Real Options, and Corporate Strategy,” *Journal of Financial Economics*, 63: 211–234.

Bolton, Patrick and Christopher Harris (1999) “Strategic Experimentation,” *Econometrica*, 67: 349–374.

Bromiley, Philip, Michael McShane, Anil Nair, and Elzotbek Rustambeko (2015) “Enterprise Risk Management: Review, Critique, and Research Directions,” *Long Range Planning*, 48: 265–276.

Burgstahler, David and Ilya Dichev (1997) “Earnings Management to Avoid Earnings Decreases and Losses,” *Journal of Accounting and Economics*, 24: 99–126.

Chatterjee, Sayan, Michael H. Lubatkin, and William S. Schulze (1999) “Toward a Strategic Theory of Risk Premium: Moving beyond CAPM,” *The Academy of Management Review*, 24: 556–567.

Daley, Brendan & Brett Green (2012) “Waiting for News in the Market for Lemons,” *Econometrica*, 80: 14331504.

Decamps, John-Paul, Thomas Mariotti, and Stephane Villeneuve (2006) “Irreversible Investment in Alternative Projects,” *Economic Theory*, 28: 425–448.

Dixit, Avinash K. and Robert Pindyck (1994) *Investment under Uncertainty*, Princeton University Press, Princeton, New Jersey.

Dixit, Avinash K. (1993) *The Art of Smooth Pasting*, vol. 2, Taylor & Francis.

Ellsberg, Daniel (1961) “Risk, Ambiguity, and the Savage Axioms,” *The Quarterly Journal of Economics*, 75: 643–669.

Froot, Kenneth A., David S. Scharfstein and Jeremy C. Stein (1993) “Risk Management: Coordinating Corporate Investment and Financing Policies,” *The Journal of Finance*, 48: 1630–1658

Glover, Brent and Oliver Levine (2017) “Idiosyncratic Risk and the Manager,” *The Journal of Financial Economics*, 126: 320–341.

Harrington, Scott E., Greg Niehaus, and Kenneth J. Risko (2002) “Enterprise Risk Management: The Case of United Grain Growers,” *Journal of Applied Corporate Finance*, 14: 71–81.

Hartman, Richard 1972) “The Effects of Price and Cost Uncertainty on Investment,” *The Journal of Economic Theory*, 5: 258–266.

Hirshleifer, David and Yoon Suh (1992) “Risk, Managerial Rffort, and Project Choice,” *Journal of Financial Intermediation,* 2: 308–345.

Holmstrom, Bengt (1979) “Moral Hazard and Observability,” *The Bell Journal of Economics,* 10: 74–91.

Ju, Nenjiu and Jianjun Miao (2012) “Ambiguity, Learning, and Asset Returns,” *Econometrica*, 80: 559–591.

Kim, Hyunseob and Howard Kung (2017) “The Asset Redeployability Channel: How Uncertainty Affects Corporate Investment,” *The Review of Financial Studies*, 30: 245–280.

Knight, Frank L. (1921) *Risk, Uncertainty, and Profit*, Boston: Houghton Mifflin Company.

Kolb, Aaron M. (2019) “Strategic Real Options,” *The Journal of Economic Theory*, forthcoming. 1.

Kuersten, Wolfgang and Rainer Linde (2011) “Corporate Hedging versus Risk-Shifting in Financially Constrained Fiirms: The Time-Horizon Matters!” *The Journal of Corporate Finance*, 17: 502–525.

Kwon, H. Dharma and Steven A. Lippman (2011) “Acquisition of Project-Specific Assets with Bayesian Updating,” *Operations Research*, 59: 1119–1130.

Laevena, Luc and Ross Levine (2007) “Is There a Diversification Discount in Financial Conglomerates?” *The Journal of Financial Economics*, 85: 331367.

Lambrecht, Bart M. (2017) “Real Options in Finance,” *The Journal of Banking and Finance*, 81: 166–171.

LeRoy, Stephen F. and Larry D. Singell (1987) “Knight on Risk and Uncertainty,” *The Journal of Political Economy*, 95: 394–406.

Lippmann, Steven and Reade Ryan (2003) “Optimal Exit from a Project with Noisy Returns,” *Probability in the Engineering and Informational Sciences*, 17: 435–458.

Nance, Deanna R., Clifford W. Smith, and Charles W. Smithson (1993) “On the Determinants of Corporate Hedging,” *The Journal of Finance*, 48: 267–284.

Panousi, Vasia and Dimitris Papanikolaou (2012) “Investment, Idiosyncratic Risk, and Ownership,” *The Journal of Finance*, 67: 1113–1148.

Pindyck, Robert S. (1993) “Investments of Uncertain Cost,” *The Journal of Financial Economics*, 34: 53-76.

Rasmusen, Eric (2004) *Games and Information: An Introduction to Game Theory*, 4th edition, Oxford: Blackwell Publishing.

Rothschild, Michael and Joseph E. Stiglitz (1970) “Increasing Risk I: A Definition,” *The Journal of Economic Theory*, 2: 225–243.

Smith, Clifford W. (1995) "Corporate Risk Management: Theory and Practice," *The Journal of Derivatives*, 2: 21–30.

Smith, Clifford W. (2008) "Managing Corporate Risk" *Handbook of Corporate Finance: Empirical Corporate Finance*, Ed by Espen B. Eckbo, vol. 2, ch. 18: 201–213.

Stulz, Rene M. (1996) "Rethinking Risk Management," *The Journal of Applied Corporate Finance*, 9: 8–25.

Trigeorgis, Lenos (1996) *Real Options: Managerial Flexibility and Strategy in Resource Allocation*, Cambridge, The MIT Press.

Trigeorgis, Lenos and Jeffrey Reuer (2017) "Real Options Theory in Strategic Management," *Strategic Management Journal*, 38: 42–63.

Villalonga, Belen (2004) "Diversification Discount or Premium? New Evidence from the Business Information Tracking Series," *The Journal of Finance*, 59: 479-506.

Yellen, Janet L. (1984) "Efficiency Wage Models of Unemployment," *The American Economic Review: Papers and Proceedings*, 74: 200–205.