



The Ecology of Risk Taking

FRANÇOIS DEGEORGE

University of Lugano, via Giuseppe Buffi 13, CH-6904 Lugano, Switzerland; HEC School of Management (on leave), Paris; Centre for Economic Policy Research, London

BOAZ MOSELLE

Ofgem, London, UK

RICHARD ZECKHAUSER*

richard_zeckhauser@harvard.edu

Kennedy School of Government, Harvard University, Cambridge, MA 02138, USA

Abstract

We analyze the risk level chosen by agents who have private information regarding their quality. We show that even risk-neutral agents will choose risk strategically to enhance their reputation in the market, and that such choices will be influenced by the mix of other agents' types. Assuming that the market has no strong prior about whether the agents are good or bad, good agents will choose low levels of risk, and bad agents high levels. Empirical evidence is gathered on 2462 firms over 24 years. The results support the model: agents of higher quality have less variable performance.

Keywords: risk, incomplete information

JEL Classification: D81, D82

Ecology studies the relationships between organisms and their environment. One major finding is that the numbers of different types of organisms in a particular ecosystem can dramatically affect how the members of a particular type behave. Types may be different species, or sub-groups within a species that are defined by characteristics such as age, sex, or dominance rank. For example, the behaviors of predators and prey, of hosts and parasites, of herbivores and plants, tend to shift when the number of either rises or falls. Similarly, within a species, the males' mating strategies—examples run from coho salmon and walrus to squirrels and bighorn sheep—may depend on their dominance status and the mix of types in the population at that point in time.¹

This paper looks at the ecology of risk taking, studying humans or organizations with differing prospects. An agent's choice of risk is normally thought of as an individualistic activity. In the real world, however, the risk level an agent selects often depends on the choices of other agents, both of his and other types (species in the ecological metaphor). If the payoffs received from a lottery were the sole component of final

*To whom correspondence should be addressed.

utilities, there would be no reason to consider the choices of others. However, one's reputation may yield a second component of utility. When types cannot be observed directly, the payoffs from one's chosen lottery—besides their intrinsic value—serve as signals to others about one's type, e.g., one's productivity or quality. Bayesian analysis processes these signals to yield a probability distribution on types, i.e., one's reputation. Thus, a college student gets grades, a direct payoff, but also a signal that helps to define his reputation and post-graduation salary. Firms secure the direct payoff of earnings; but earnings also signal firm quality, which in turn determines the reward to reputation, namely the stock price. Professors develop a publication record, to the same purposes.

In all the instances we consider, there are parties on the other side of the market, e.g., employers, tenure committees and investors, who do the observing. We call them buyers. Buyers calibrate signals to decide what offer, if any, to make to an agent. To simplify, we assume competitive conditions in the market, apart from the information asymmetries we study. This supports our assumption that reputations are rewarded with their expected marginal product. Here ecology comes into play: That is, the distribution of agent types and the lotteries they choose become relevant, because the reputation reward flowing from a specific payoff will depend on what the market, using Bayesian analysis, perceives to be the distribution of agent types given a signal. Each agent knows this, and also anticipates the risk-choice strategies of others. Our results are developed for a Bayesian world where agents act strategically, and buyers draw appropriate inferences.

An agent's type will determine the alternative lotteries (risk choices) available to him. For any level of risk, good agents get a higher return on average than bad agents. An agent faces a tradeoff when maximizing his total expected utility. A sacrifice in his direct payoff may be beneficial if it affects the information flow to the market about his type in a way that enhances the expected return to his reputation. Much the same behavior goes on in traditional ecosystems when the types are different species of animals. Humans in the wild try to make themselves known as such, since most animals know or sense their superior capabilities. Many animals, by contrast, misrepresent themselves; e.g., a threatened blowfish, at some resource expense, enhances its perceived size, hoping to look like a bigger, fiercer species.

The following intuition underlies our analysis. Agents with rosy prospects have an incentive to engage in low-risk activities, which are unlikely to upset market priors. In contrast, agents with bleak prospects may prefer riskier ventures in which the performance signal is noisy and may usefully—from the agent's viewpoint—reverse market priors. The applications of this general lesson apply in a great range of contexts, from labor markets to college admissions, from junior professors building their publication record for a tenure position to suitors seeking a marriage partner.

The alert reader will have noticed that we talk about rosy and bleak, which refers to the overall environment, not good and bad, which identifies the agent's type. The reputational reward to a particular direct payoff will depend on the whole distribution of types in the population, which alongside an agent's type affects how rosy or bleak is a situation. If the distribution is very favorable, even relatively poor agents should choose low risk levels,

since they are likely to get strong rewards anyway. Conversely, in a weak population, good agents may have to gamble lest they merely be construed as bad agents who got lucky. For any agent, the reward function will depend on his type and the distribution of types.

To motivate our analysis, we shall often talk in terms of good and bad firms securing earnings, though the reader should easily extrapolate to his or her area of interest, e.g., discussing student grades or professorial standing instead of firm earnings. For our firms example, the buyers on the other side of the market are investors. In our empirical analysis, we focus on the risk choices of firms. This empirical focus offers two advantages: (1) there are relatively unambiguous measures of the signals firms generate, and (2) there are vast amounts of data on firm performance. In the case of firms, we expect the proportions of good and bad quality to be roughly equal in the population. In this case, our theoretical model predicts that good firms choose low levels of risk, and bad firms high levels. Our empirical results support this prediction.

At least four literatures are relevant to the current analysis, those on: contests and superstars, signaling, adverse selection, and statistical discrimination.² Beyond this, this framework presents a new consideration for the challenge of *Choosing the Right Pond* in which to operate.^{3,4}

The paper proceeds as follows. Section 1 introduces our model. First, it presents a discrete example to foster intuition. It then presents the continuous case. Section 2 develops results for the continuous case. Section 3 provides empirical tests. Section 4 concludes.

1. The model

We start, as do all signaling theories, from the premise that agents differ in unobservable ways, and that agents know more than buyers do about their true qualities. Agents are risk neutral. At the outset agents learn their type, which can be good or bad. The proportion of good agents, p , is public knowledge. Agents then make a risk choice, which determines the distribution of payoffs. We assume that the market does not observe this risk choice.⁵ Chance determines a payoff. Using Bayes' rule, the market then assesses the likelihood that the agent is good or bad. Agents are then put up for sale. In the competitive market, their sale price is their discounted expected value, which depends on their reputation.⁶ In effect, this is a three-player game involving good agents, bad agents, and the market.

Discussion is facilitated if we focus on the example of firms. Agents are firm managers. The market's buyers are investors. The direct payoff is a level of earnings. Reputation could be the stock price after earnings are announced. The agent/manager has a short time horizon because the whole firm might be sold at the end of the period (the case illustrated here). Alternatively, he might merely expect to leave and cash his options.

Discrete example. Half the firms are good; half are bad (that is, $p = 0.5$). The two possible strategies are safe and risky. The distribution of period payoffs is given by the following table:

Probabilities of payoffs given player and strategy								
Period payoff	Safe				Risky			
	1	2	3	4	1	2	3	4
Good	0	0	1	0	0	0.19	0.6	0.21
Bad	0	1	0	0	0.21	0.6	0.19	0

Expected payoffs are given by the following table:

	Safe	Risky
Good	3	3.02
Bad	2	1.98

It is slightly actuarially preferable for a good type to play risky (expected period payoff is 3.02 rather than 3) and for a bad type to play safe (expected period payoff is 2 rather than 1.98). Given our assumption that the agent is sold at the end of the period, once the sale is completed the new owners will have good agents play risky and bad agents play safe, in perpetuity. We ask what each type will choose in the first period. The problem for the market, and the source of any efficiency loss, is that agents may sacrifice some first-period payoff in order to increase their expected sale price. Assume the discount rate is 0.05. Evaluated at the beginning for the first period, the expected sale price is

$$\frac{1}{0.05} \{3.02E[\text{Prob}(\text{good} \mid \text{first period payoff})] + 2E[\text{Prob}(\text{bad} \mid \text{first period payoff})]\}.$$

Assume the market believes that in the first period, the good type plays safe and the bad type plays risky. We will now show that these market beliefs are sustainable in equilibrium. The posterior probabilities that the agent is good are:

$$\begin{aligned} \text{Prob}(\text{good} \mid \text{first period payoff is 1}) &= 0, \\ \text{Prob}(\text{good} \mid \text{first period payoff is 2}) &= 0, \\ \text{Prob}(\text{good} \mid \text{first period payoff is 3}) &= \frac{p}{p+0.19(1-p)} = \frac{0.5}{0.5+0.19(1-0.5)} = 0.84, \text{ and} \\ \text{Prob}(\text{good} \mid \text{first period payoff is 4}) &= 1. \end{aligned}$$

Consider the possible outcomes and posterior market beliefs for a good firm playing safe:

Outcomes, posteriors and expected payoffs for good type if he plays safe

First-period payoff	Probability of first-period payoff	Market's posterior that agent is good given first-period payoff	Probability of first-period payoff times market posterior
4	0	1	0.000
3	1	0.840	0.840
2	0	0	0.000
1	0	0	0.000

Thus for a good firm playing safe, the expected probability that it will be perceived as good after the first-period payoff is 0.84. Its overall expected payoff of playing safe is:

$$3 + \frac{1}{0.05}(0.84 \times 3.02 + 0.16 \times 2) = 60.14.$$

Consider now the possible outcomes and posterior market beliefs for a good firm playing risky:

Outcomes, posteriors and expected payoffs for good type if he plays risky

First-period payoff	Probability of first-period payoff	Market's posterior that agent is good given first-period payoff	Probability of first-period payoff times market posterior
4	0.21	1	0.21
3	0.6	0.840	0.504
2	0.19	0	0.000
1	0	0	0.000

Thus for a good type playing risky, the expected probability that it will be perceived as good after the first-period payoff is 0.714. Its overall expected payoff from playing risky is:

$$3.02 + \frac{1}{0.05}(0.714 \times 3.02 + 0.286 \times 2) = 57.59.$$

Similar calculations show that the overall expected payoff for a bad type playing risky is 45.24, compared to 42 if he plays safe. Hence, both types have the incentive to adhere to market beliefs, namely that good types play safe and bad types play risky.

Payoffs given market beliefs: Good plays safe and bad plays risky

	Agent actually plays safe	Agent actually plays risky
Good receives	60.14*	57.59
Bad receives	42	45.24*

*Preferred strategy.

The intuition supporting the equilibrium is that playing risky gives a bad type a reasonable chance to be thought good. By contrast, playing safe considerably cuts down a good type's chance of being considered bad. Further calculations show that the equilibrium is unique: for the parameters chosen, no other set of market beliefs is consistent with the good and bad types' optimal choices given these market beliefs.⁸

Matters shift if the proportion of good types is much smaller. If p is 10 percent, for example, the only equilibrium has both types play risky. With an unfavorable prior, a good type gambles to have a greater chance of distinguishing himself from bad types.⁹ We shall see echoes of these outcomes in the substantially more general continuous model. That is, we will find that only bad types gamble when priors are relatively favorable, but both do when they are unfavorable.¹⁰

The model for the continuous case. Firms invest in projects, e.g., undertake R&D or hire employees. Each project contributes, positively or negatively, to both the mean and variance of the firm's portfolio of undertakings. Agents select a portfolio for their firm, i.e., they choose its degree of risk. It is positionally convenient to talk in terms of the firm making the choice.

The risk choice is assumed to be continuous; i.e., there is sufficient project divisibility.¹¹ Formally, each firm chooses from a one-dimensional family of random variables indexed by its variance, V , where $\underline{V} < V < \bar{V}$. The random variable \tilde{x}_V represents first-period earnings. It is distributed normally with variance V and mean $\mu(V, \theta) = \mu(V) + \theta$, where θ indexes the type of the firm. We will refer to the $\theta = 0$ type firm as bad and the $\theta = \Delta$ type as good, where $\Delta > 0$. The prior probability that a firm is of the good type, denoted by p , is common knowledge, as is the mean-variance schedule $\mu(V)$.

Each firm chooses a point on the mean-variance schedule given by $\mu(V)$ and for given V a firm of type θ has returns

$$\tilde{x} \sim N(\mu(V) + \theta, V).$$

Figure 1 illustrates the case when $\mu(V)$ is single-peaked and concave, i.e., there is an increasing marginal cost of adding variance.¹² In fact, the mere existence of an interior maximum is sufficient for most of our theoretical results—single-peakedness and concavity are not required. The existence of an interior maximum is no more than the usual assumption that

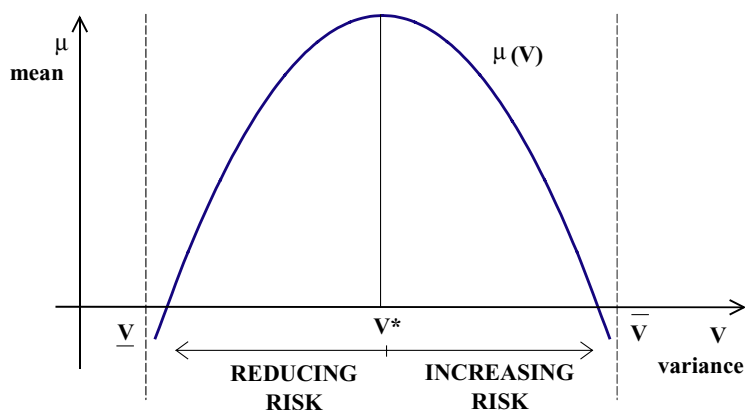


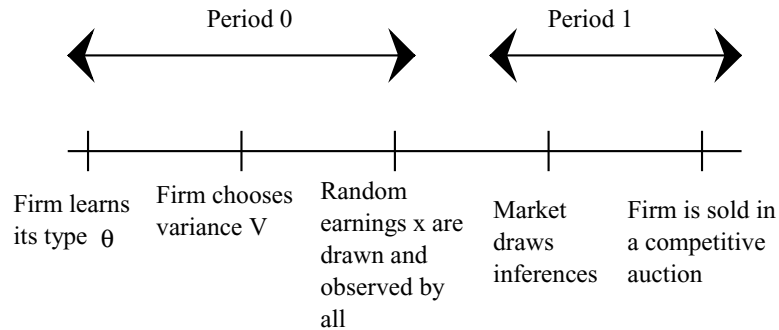
Figure 1. Mean variance schedule, $\mu(V)$, for earnings.

the firm has available to it no more than a finite set of investment opportunities yielding positive NPV. A natural way to think about the shape of $\mu(V)$ is to imagine that the manager allocates investment funds over a finite set of potential projects. A very simple special case occurs when the projects are finite in size and independent, with some having positive expected returns, others negative, and the manager has no binding budget constraint. It is trivial to see in this case that $\mu(V)$ will be generically single-peaked and concave.^{13,14} Our computer-based simulations are carried out with a function that is smooth and concave, although our theoretical results do not assume concavity. It is convenient (though by no means essential) to assume, as in Figure 1, that $\mu(V) \rightarrow -\infty$ as $V \rightarrow \underline{V}$ or $V \rightarrow \bar{V}$, and we maintain this assumption throughout.

We are implicitly ruling out leverage as a way to increase both variance and expected return. That is because leverage, being readily observable, would accomplish nothing from an informational point of view: the informativeness of the company's real earnings would be preserved, and the signaling mechanism we analyze in this paper would be unaffected.

Firms and investors are assumed to be risk-neutral, so that in a full-information setting firms would choose the mean-maximizing level of variance V^* . We investigate how asymmetric information leads to departures from this optimum. As illustrated in Figure 1, choosing a variance less than V^* represents risk-reducing behavior; choosing a variance greater than V^* increases risk.

At the beginning of period 0, a firm learns its type, then chooses its desired variance V . Its random return x is then drawn according to the equation above, and observed by all, and the period ends. At the beginning of the period 1, market participants draw inferences, and the firm is sold in a competitive auction to a new and permanent owner-manager.¹⁵



It will clearly be optimal for the new owner to choose the level of variance V^* that maximizes expected returns (we assume risk-neutrality throughout), so the expected return to owning a firm of type θ from next period onward is

$$\sum_{t=1}^{\infty} \frac{\mu(V^*) + \theta}{(1+r)^t} = \frac{\mu(V^*) + \theta}{r},$$

where r is the discount rate. Thus, the expected present value (period return plus expected price) of a firm with returns x this period is

$$x + \frac{\mu(V^*)}{r} + \frac{1}{r} E[\theta \mid \text{all available information}],$$

where buyers use Bayesian analysis to compute the expectation term.

Firms choose V to maximize their expected total payoff, which is given by their present value, namely

$$\Pi = \max_V E \left[\tilde{x} + \frac{\mu(V^*)}{r} + \frac{1}{r} E[\theta \mid \text{all available information}] \mid \text{choice of } V \right].$$

If θ were observable, this problem would be trivial; we would have

$$E[\theta \mid \text{all available information}] = \theta,$$

and firms would always set $V = V^*$. θ is private information, so the firm is faced with a trade-off in maximizing its expected total payoff: clearly, setting $V = V^*$ maximizes the expected current return; however, deviating from V^* will change the information that flows to the market, and hence may increase the expected sale price.

2. Results for the continuous case

All proofs are gathered in the Appendix.

Claim 0. There is a region $[V_{\min}, V_{\max}]$ containing V^* , such that all strategies outside this region are strictly dominated.

An obvious corollary is that the shape of $\mu(V)$ outside the region $[V_{\min}, V_{\max}]$ does not affect the outcome of the game. In particular, where our proofs rely on single-peakedness, this need only apply within the region $[V_{\min}, V_{\max}]$.¹⁶

It is helpful to begin with the intuition explaining the equilibrium we obtain. A firm will maximize a weighted average of its current period expected return and its “reputation,” that is, the market’s expected posterior probability that it is good. When considering changing its level of risk to a level below or above the full-information maximizing level V^* , it faces a tradeoff. It will lose on average in current earnings; on the other hand, perhaps deviating from V^* is beneficial if it brings about the kind of performance that makes it more likely that investors think highly of it and therefore pay a higher price for it.

This tradeoff is best understood by looking at Figure 2, which plots the posterior probability that the firm is good as a function of realized current earnings, for various values of p , the market’s prior probability that the firm is good. For a normal density function, the curve is S-shaped. The more favorable the market’s prior—that is, the higher is p —the more to the left the curve lies. Here μ_b and μ_g represent the expected earnings corresponding to a choice of risk of V^* , for a bad and good firm respectively. Any deviation from V^* shifts expected current earnings somewhat to the left of μ_b and μ_g .

Consider the case $p = 0.5$ —that is, there are as many good firms as bad firms in the population. In this case, μ_g is in the concave region of the curve, suggesting that a good firm

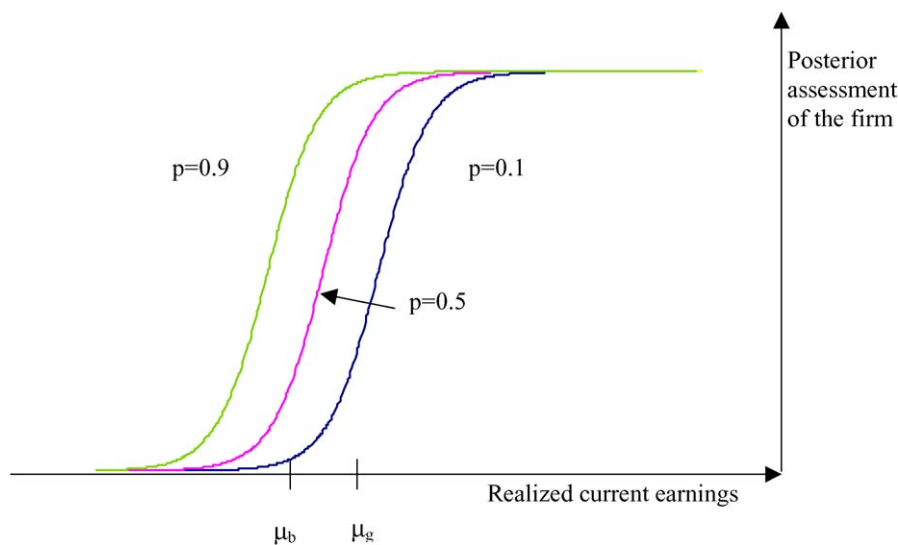


Figure 2. Posterior probability that a firm is good, as a function of realized current earnings, for various values of p , the prior probability that a firm is good.

should choose a level of risk smaller than V^* in order to maximize the market's expected posterior assessment of the good firm. As it reduces risk, the good firm's mean μ_g shifts somewhat to the left. Not only does it sacrifice some current expected earnings; it also moves to a less concave region of the curve, which makes risk reduction less attractive. This process limits the benefits to risk reduction.

By contrast, for $p = 0.5$, μ_b is in a convex region of the curve, suggesting that a bad firm should choose a level of risk higher than V^* in order to maximize the market's expected posterior assessment of its quality.

Our discussion so far has been heuristic. We now proceed to establish our claims rigorously. Since V is not observed, investors can only form conjectures about variance choice by observing current returns. We will denote the market's conjectures about the choice of variance by bad and good firms by W_b and W_g , and the actual choices of variance by V_b and V_g . Each firm makes its choice of variance in response to these market conjectures. This gives rise to a "best response" mapping

$$(W_b, W_g) \mapsto (V_b^*(W_b, W_g), V_g^*(W_b, W_g)).$$

The relevant equilibrium concept is that of rational expectations. Thus in order for market conjectures W_b, W_g , and firm choices V_b, V_g , to constitute an equilibrium, the two following conditions must be true:

- taking the market conjectures W_b, W_g as given, each type of firm is maximizing expected profits, and
- market conjectures are correct ($W_b = V_b$ and $W_g = V_g$).

Equivalently, market equilibria correspond to *fixed points* of the best response mapping above.¹⁷ As is typically the case, explicit solution of this fixed-point problem proves analytically intractable, and we resort to numerical methods.

Equilibrium variance choices and Bayes' rule. Each firm chooses V to maximize its expected total payoff

$$\Pi = E \left[\tilde{x} + \frac{\mu(V^*)}{r} + \frac{1}{r} E[\theta \mid \text{all available info}] \mid \text{choice of } V \right].$$

With variance unobservable, this is equivalent to

$$\max_V r \cdot \mu(V) + E[\theta \mid \text{market beliefs, choice of } V],$$

or equivalently,

$$\max_V r \cdot \mu(V) + \int 0 \cdot \text{Prob}(\theta = 0 \mid x) + \Delta \cdot \text{Prob}(\theta = \Delta \mid x) dF(x, \theta, V).$$

The probabilities on θ are derived from Bayes' Rule, using the market conjectures and the common knowledge prior:

$$\text{Prob}(\theta = \Delta | x) = \frac{pf(x, \Delta, W_g)}{pf(x, \Delta, W_g) + (1 - p)f(x, 0, W_b)},$$

and similarly for $\theta = 0$. Taking first order conditions for this problem for each type of firm gives the best response map just described. We have not shown that this is a contraction mapping, but in practice the usual iterative approach has provided a solution.¹⁸

We can also use this approach to prove the existence of a mixed-strategy equilibrium:

Proposition 1. *A market equilibrium always exists in mixed strategies.*

It is important to note that our proof does *not* assume the monotonic likelihood ratio property. That property would not be satisfied if the two firms selected different levels of variance: an extreme outcome at either end makes the high variance firm more likely.

In order to apply numerical methods, we choose a particular functional form for the mean-variance schedule:

$$\mu(V) = -\frac{1}{V} - \frac{1}{1 - V}.$$

The general shape of this mean-variance schedule is a broad-based inverted U. The full-information optimal choice of variance for both types is $V^* = 0.5$.¹⁹

Equilibrium for $p = 0.5$. As the heuristic discussion above suggests, when $p = 0.5$ good firms choose a lower level of risk than in the full-information optimum, whereas bad firms choose a higher level. For example, for $r = 0.11$ and $\Delta = 3.0$, the equilibrium choices of variance are $V_g = 0.47$ and $V_b = 0.53$. Bad firms have an incentive to increase the risk in their performance, even at a cost, to raise the chances that they will be mistaken for good firms. Good firms reduce the risk in their earnings to make current returns a more revealing signal. In equilibrium, the good firm selects a below-optimum variance and the bad firm selects an above-optimum variance.

An increase in variance for the bad firm has four effects; one is a direct effect on current returns, and the other three affect the firm's reputation. First, raising variance above V^* is costly; current expected earnings are reduced. Second, because of the loss in expected current return, the increase in the variance diminishes the overlap between the good and the bad firms' distributions in the region of likely outcomes, which decreases the posterior probability that the firm is good. Third, for somewhat less likely events, the gain in the variance will increase the overlap with the good firm's distribution, which will help the bad firm's posterior expectation. Fourth, for extremely unlikely events, an increase in variance by the bad firm hurts its posterior expectation, because the likelihood ratio approaches infinity as x tends to infinity; that is, a very high return will reveal risk-increasing behavior.

The effects of the good firm's lowering its variance are analogous. First, if reducing risk is costly, cutting the variance directly sacrifices expected earnings. Second, the loss in current expected earnings shifts the good firm's distribution toward that of the bad firm, which will

hurt the good firm's posterior expectation. Third, and more positively, cutting its variance enables the good firm to diminish the overlap of the left tail of its distribution with the bad firm's distribution. Fourth, for very unlikely events, the reduction of variance will bolster the good firm's posterior expectation.

The third effect appears to dominate for both types. Because of it, incomplete information introduces a payoff asymmetry. High draws help a bad firm's reputation more than low draws hurt it. The situation is reversed for a good firm. The welfare implications are clear. Bad firms could always just maximize for themselves, ignoring the good firms. Since bad firms instead increase risk, they must be better off by doing so. Since the aggregate payoff cannot be expanded beyond its level at (V^*, V^*) , and signal-jamming is costly, there must be a transfer from good to bad firms. Moreover, asymmetric information also leads to an efficiency loss: both types choose lower mean earnings than they would under full information.

Two important parameters help determine the degree of risk reduction and increase we observe: Δ , the gap between the mean earnings schedules of the good and bad types, and r , the discount rate.

Our results are summarized in Figure 3. When good and bad firms are similar, that is, Δ is small, the distributions for the two types are close to each other. Consequently, the second effect is large for both types, and the departure from V^* is relatively small. As Δ increases, however, the third effect gains importance; by reducing (increasing) risk, good (bad) types are able to diminish (enlarge) the overlap with the other type's distribution. Thus, good firms reduce risk more and bad firms increase them more.

Above a certain threshold of Δ , when the good firms achieve a commanding lead, this phenomenon is reversed. The struggle by the bad firms to stay up is no longer worth it. By looking at current earnings, investors can readily distinguish between good and bad firms. It

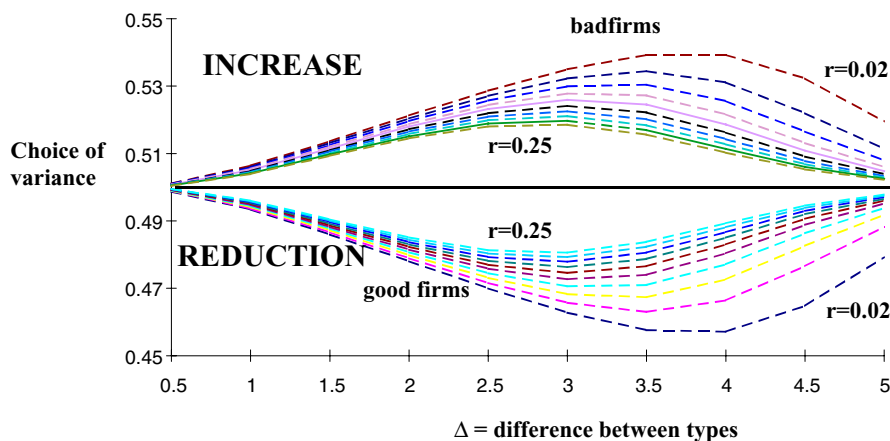


Figure 3. Risk reduction and increase as a function of the difference between types (Δ) and the discount rate r (p is fixed at 0.5).

becomes less likely that increasing risk levels will significantly alter investors' perceptions about bad firms. As a consequence, bad firms diminish their risk, and good firms are able to trim risk less.

As current earnings become more important, that is, as r increases, signal manipulation diminishes and ultimately vanishes. The effect of r on risk choices is straightforward: the more costly either activity becomes (as r increases), the less departure from V^* occurs in equilibrium. The first effect, a reduction in current expected earnings, further dampens firms' incentives to reduce or increase risk.

Equilibrium for $p \neq 0.5$. When $p = 0.5$, the prediction of our model accords well with the "good types reduce risk, bad types increase it" intuition. As p approaches 0 or 1, the results become more complicated. For a sufficiently high value of Δ , the equilibrium is similar to the $p = 0.5$ case. However, when $p < 0.5$, for a low enough value of Δ both types of firms increase risk. Conversely, when $p > 0.5$, both types reduce risk (see Figure 4).

The intuition is as follows. When the prior probability that a firm is good is high (p is high), and a bad firm is not very different from a good firm (Δ is low), a bad firm gains little by increasing risk; chances are that it will be mistaken for a good firm anyway, so there is no point in rocking the boat. For both types, low draws, which suggest that the firm is bad, are very costly, so both types reduce risk. By contrast, when the market's prior probability is unfavorable (p is low), prudent behavior will tend to preserve the status quo assessments. This is unfortunate for the good firm, which would be almost certainly labeled as bad. Thus the good firm might as well take its chances and increase risk. A favorable outcome for the gamble would enable the good firm to distinguish itself, since such an outcome is so unlikely for a bad firm. Figure 5 illustrates equilibrium variance choices for $p = 0.1$.²⁰

In sum, when substantial proportions of both good and bad firms inhabit the population, and their capabilities differ significantly (Δ is large), good firms will reduce risk while bad

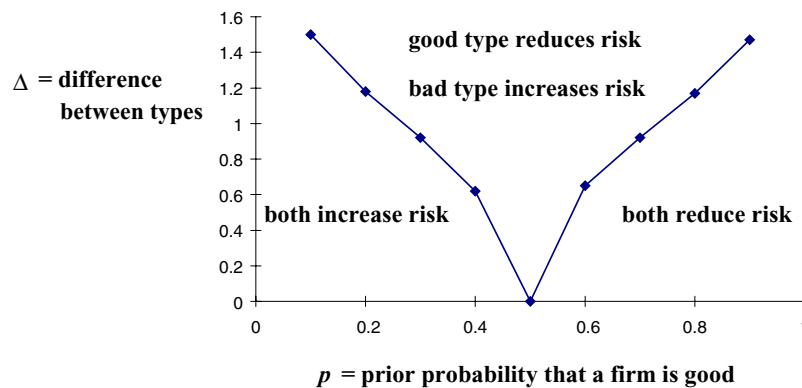


Figure 4. Threshold value of Δ (the difference between types) as a function of p (the prior probability that a firm is good).

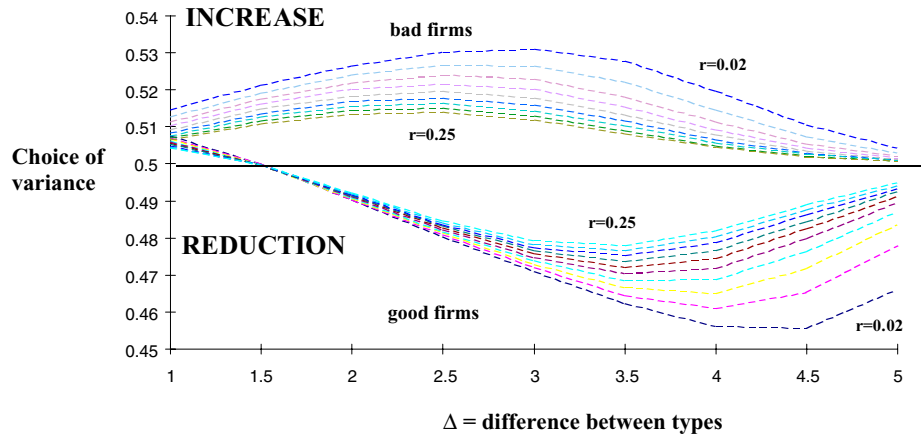


Figure 5. Risk reduction and increase as a function of the difference between types (Δ) and the discount rate r (p is fixed at 0.1).

firms will gamble. Only when the market knows that virtually all firms are bad, will the good firms gamble alongside (or even more than) the bad.

3. Empirical evidence: Firm quality and performance variability

We conduct our empirical analysis on the earnings of major firms. Our model makes a precise prediction: If there are reasonably balanced numbers of good and bad firms, good firms should take fewer risks, and have less variable earnings, than bad. That prediction is the focus of our empirical analysis.

Measuring performance variability and firm quality. To determine the variability of firm performance we use four accounting measures of performance: operating income after depreciation, income before extraordinary items, net income, and earnings per share. For each firm and each performance measure, we regress performance on year. We define variability of performance as the mean square error of this regression divided by the mean performance over the period.²¹

To measure firm quality we used two approaches. One is simply to use the mean value of our accounting performance measures. Another is to proxy firm quality with the firm's average Tobin's q over the period. The two approaches yielded similar results. We present only our results using Tobin's q because they are less subject to confounding interpretations than are those using mean values of accounting measures.

Sample and construction of empirical variables. We gathered data on all firms in COMPUSTAT having more than \$10 million in total assets in 1975, for the 1975–1998 period. (This sample stops short of the Internet-bubble period, and its accompanying accounting irregularities.) This yielded a sample of 3079 firms. For each firm-year we calculate Tobin's q as (market value of equity + book value of debt)/(book value of assets).²²

Table 1. Firm quality and performance variability. Spearman rank correlations between average Tobin's q over the sample period and the mean square error of a regression of performance on year, within 1-digit SIC codes.

1-Digit SIC code and industry description	Performance measure			
	Operating income after depreciation	Income before extraordinary items	Net income	Earnings per share
1. Mining and construction	-0.0756 (125)	-0.1338 (97)	-0.1010 (93)	0.0400 (111)
2. Manufacturing	-0.3404 (472)	-0.3596 (432)	-0.3594 (432)	-0.4137 (430)
3. Manufacturing	-0.0443 (756)	-0.1292 (666)	-0.1954 (663)	-0.1682 (687)
4. Transportation, communication and public utilities	0.0221 (298)	-0.0214 (284)	-0.0772 (278)	-0.0287 (281)
5. Wholesale and retail trade	-0.2598 (271)	-0.2362 (224)	-0.2746 (229)	-0.3218 (240)

Our sample consists of all firms in COMPUSTAT with more than \$10 million in total assets in 1975, over the 1975–1998 period. We calculate Tobin's q as (market value of equity + book value of debt)/(book value of assets). For any accounting measure of performance (operating income after depreciation, income before extraordinary items, net income, or earnings per share) we define the performance variability of a firm as the mean square error of a regression of the performance measure on year, divided by the mean performance over the period. We restrict attention to firms for which we have at least ten years of data, and discard firms with a negative mean performance over the period. (The number of observations in each cell is indicated in parentheses.)

We use one- and two-digit SIC codes as industry controls. To make sure that our interpretation of our results is correct, we must make sure that SIC codes adequately control for industry effects (Tobin's q is clearly industry-specific). Within each SIC code we computed the average correlation among firms of the change in Tobin's q . Then for each SIC code we generated 1000 random control groups drawn from our sample, each with the same number of firms as in the SIC code. In each control group we computed the average correlation (among firms in the control group) of the change in Tobin's q . We counted the percentage of times that the average correlation in a control group exceeded the correlation in the SIC code. We think of this percentage as a bootstrapped p -value, measuring the probability that a higher correlation in the SIC code is due by chance. The lower these p -values, the more confident we are that using SIC codes as industry controls makes sense.

Our results support our use of SIC codes as industry controls. Our sample has firms in 71 2-digit SIC codes. Tobin's q can be most appropriately measured for firms in 2-digit SIC codes 10–59 (mining and construction; manufacturing; transportation, communication, and public utilities; and wholesale and retail trade).²³ Our sample has firms in 46 of these industries. For 31 of them, our bootstrapped p -value is equal to or less than 0.05 (for 17 of them, it is exactly equal to zero, meaning that in none of the 1000 random control groups was the average correlation higher than the average correlation in the 2-digit SIC code).

Table 2. Aggregate t -statistics for the Spearman rank correlations between firm quality and performance variability, given various categories and controls.

Category/control	Performance measure			
	Operating income after depreciation	Income before extraordinary items	Net income	Earnings per share
1-digit SIC code	-6.23	-7.38	-8.79	-8.58
2-digit SIC code	-7.98	-7.54	-9.27	-9.26
Size decile 1975 (total assets)	-9.80	-8.98	-10.13	-10.39
Size decile 1975 (market value)	-5.22	-6.32	-7.81	-7.29
Size decile 1975 (book value of equity)	-8.91	-8.68	-9.98	-10.14
Size quintile 1975 (total assets) within 1-digit SIC code	-9.26	-9.49	-10.71	-11.27

Our sample consists of all firms in COMPUSTAT with more than \$10 million in total assets in 1975, over the 1975–1998 period. We calculate Tobin's q as (market value of equity + book value of debt)/(book value of assets). For any accounting measure of performance (operating income after depreciation, income before extraordinary items, net income, or earnings per share) we define the performance variability of a firm as the mean square error of a regression of the performance measure on year, divided by the mean performance over the period. We restrict attention to firms for which we have at least ten years of data, and discard firms with a negative mean performance over the period. (The number of observations in each cell is indicated in parentheses.) The t -statistics refer to the Spearman rank correlations between average Tobin's q over the sample period and the mean square error of a regression of performance on year.

Interestingly, the results are much less clear for firms in 2-digit SIC codes equal to 60 or more (financial sector, services, and public administration), codes in which Tobin's q cannot be measured properly. Our sample has firms in 25 of those industries. Our bootstrapped p -value is equal to or less than 0.05 for only 6 of them (it is exactly equal to zero for none of them).

The results are even more striking for 1-digit SIC codes. In each of the 1-digit SIC codes 1-5, our bootstrapped p -values are zero. For 1-digit SIC codes 6 or higher, they are much higher. We conclude from this exercise that, for industries in which Tobin's q can be properly measured (1-digit SIC codes 1-5), SIC codes provide a suitable way of controlling for industry effects. In the empirical analyses that follow, we restrict attention to firms in our sample with 1-digit SIC codes 1-5 (2462 firms).

The risk levels of good and bad firms. Our maintained assumption is that firm performance persists throughout our sample period—that is, we indeed have good and bad firms. We

checked this persistence assumption by dividing our sample period in half (1975–1986 and 1987–1998), and computing each firm’s mean Tobin’s q in each period. We then computed the Spearman rank correlation between mean Tobin’s q in the first period and mean Tobin’s q in the second period. The correlations are strongly positive and statistically significant, suggesting strong persistence, and supporting our assumption that there are “good” and “bad” firms.²⁴ We now proceed to examine the key empirical prediction of our model: firms of higher quality have earnings that reflect lower risk choices. To test this prediction we compute the Spearman rank correlation between performance variability and firm quality, controlling for industry. Table 1 shows that for each of the four accounting measures, the correlation between firm quality and performance variability (risk) is almost always negative. We checked the statistical significance of our results, controlling for industry, using one- and two-digit SIC codes, size (using 1975 numbers for accounting assets, market value of equity, and book value of equity), and simultaneously controlling for industry (one-digit SIC codes) and size. We report our findings in Table 2. In all cases,²⁵ the aggregate t -statistics are very significantly negative.²⁶

4. Conclusion

We analyzed the levels of risk good and bad agents select when they know their quality but buyers do not. If risk choices cannot be observed, as surely many cannot, and if the market is fairly uncertain about whether agents are good or bad, then good types will reduce risks and bad types will increase them. Good types reduce risk to avoid being wrongly identified; bad types increase risk in the hope of being wrongly identified. If the market’s prior beliefs are highly unfavorable, both types gamble, hoping to alter these beliefs. If the prior beliefs are highly favorable, both types reduce risk, hoping to avoid refuting the market’s beliefs.

Our brief empirical study reveals that higher-quality firms appear to have less variable performance, precisely in accord with our model, assuming that the market’s prior beliefs are not extreme.

Wherever agents have some but not full control over their own destinies, the choice of lotteries is implicitly a subject of discussion. Thus, a young person choosing a career path must choose a lottery much the way a firm must choose among investment alternatives. This analysis shows that much more than a tradeoff between risk and return is involved. Other factors aside, agents will wish to increase or reduce the variability of their immediate payoffs to enable them to look more or less like others. Like animals struggling to survive and reproduce, their risk-taking behaviors will depend not only on their own type but on the mix of types in their ecosystem.

Appendix

Proof of Claim 0: By choosing variance V^* , a firm of type θ guarantees itself an expected payoff of at least

$$\theta + \mu(V^*) + \frac{\mu(V^*)}{r}$$

This lower bound is the expected payoff if the choice of V^* leads the market to assume that the firm is of the bad type with probability 1.

The expected payoff to a firm of type θ that chooses some other variance V is at most

$$\theta + \mu(V) + \frac{\mu(V^*) + \Delta}{r}$$

This upper bound is the expected payoff if the market concludes that the firm is of the good type with probability 1.

Any V for which the latter payoff (the upper bound) is strictly less than the former payoff (the lower bound) is strictly dominated by V^* . To be explicit, choosing any V for which

$$\theta + \mu(V) + \frac{\mu(V^*) + \Delta}{r} < \theta + \mu(V^*) + \frac{\mu(V^*)}{r}$$

is a strictly dominated strategy. Simple algebra shows that this is equivalent to

$$\mu(V) < \mu(V^*) - \frac{\Delta}{r}$$

so the V_{\min} and V_{\max} mentioned above can be defined as smallest and the largest solutions to

$$\mu(V) = \mu(V^*) - \frac{\Delta}{r}. \quad \square$$

Proof of Proposition 1:

1. We assume that $\mu(V) \rightarrow -\infty$ as $V \rightarrow \underline{V}$ or $V \rightarrow \bar{V}$, hence there exists a closed (and therefore compact) interval $C = [\underline{V}, \bar{V}] \subset (\underline{V}, \bar{V})$ such that whatever the market believes, any optimal choice of V lies in C .
2. Consequently, we can without loss of generality define a “mixed strategy” for the firm to be a pair of Borel measures on C , one for each firm type. We write $B[C]$ for the set of Borel measures on C , endowed with the topology of weak convergence of measures.
3. We further define a “mixed conjecture” to be market beliefs represented by a pair of elements of $B[C]$.
4. Define a correspondence $\phi : B[C] \times B[C] \rightarrow B[C] \times B[C]$ to be the mixed strategy version of the “best response” correspondence: given $(\tau_1, \tau_2) \in B[C] \times B[C]$, we interpret (τ_1, τ_2) as “market conjectures” as to the strategies of good and bad types. We then define a posterior function $\xi(\tau_1, \tau_2, x)$ to be the Bayesian-updated posterior probability that the market assigns to the firm being of the good type, after observing the outcome x . For later purposes, the key observation (easy to check) is that this posterior function $\xi(\tau_1, \tau_2, x)$ is continuous in (τ_1, τ_2) with respect to the given topology on $B[C] \times B[C]$ (the product of the weak convergence topologies).

The best response correspondence is therefore given by:

$$\begin{aligned}\phi_1(\tau_1, \tau_2) &= \{\sigma \in B[C] \mid \text{supp}(\sigma) \subseteq \arg \max(\text{objective function of good type})\} \\ \phi_2(\tau_1, \tau_2) &= \{\sigma \in B[C] \mid \text{supp}(\sigma) \subseteq \arg \max(\text{objective function of bad type})\}\end{aligned}$$

5. Note that a fixed point of ϕ represents a mixed-strategy equilibrium of the game.
6. Finally, apply Glicksberg's theorem to prove the existence of such a fixed point:²⁷ Let K be a nonempty compact convex subset of a locally convex Hausdorff space, and let the correspondence $F : K \rightarrow K$ have closed graph and nonempty convex values. Then the set of fixed points of F is compact and nonempty.²⁸
7. The space $B[C] \times B[C]$ (with the product topology from the topology of weak-convergence) satisfies the requirements on K .
8. It therefore only remains to show that the correspondence ϕ has closed graph and nonempty convex values. Non-emptiness follows from the existence of maxima for both types' objective functions, which is immediate from the continuity of those functions over the compact set C . Convex-valuedness is immediate from the definition. The closed graph property is proved by using Berge's theorem ("theorem of the maximum") and the continuity of the posterior function $\xi(\tau_1, \tau_2, x)$.

Acknowledgments

We thank Fondation HEC and the Center For International Affairs at Harvard for financial support; the Harvard Business School for hospitality; and Miriam Avins, Sandeep Baliga, William Bossert, Darrell Duffie, Bernard Dumas, Kate Emans, Edith Ginglinger, Gary Gorton, Ulrich Hege, Jayendu Patel, Didier Pène, Kip Viscusi and a referee for helpful comments and discussions. Seminar participants at Essec, HEC, London Business School, the London School of Economics, the Stockholm School of Economics and Université Paris-Dauphine provided helpful comments. Degeorge was at HEC Paris while this study was completed.

Notes

1. For example, dominant Rocky Mountain bighorn rams tend and defend ewes as their primary mating strategy. Subordinate rams, by contrast, "fight dominants for temporary copulatory access (lasting seconds) to defended ewes," a behavior called coursing. "Coursing brawls may cause death. . . [it is an] alternative mating tactic that is high-gain and high-risk" (Hogg and Forbes, 1997, p. 33).
2. Where superstar situations pertain, even very talented agents may find it worthwhile to gamble, since productivity depends on one's relative ranking. This was illustrated in the 2002 World Cup, on numerous occasions. The players were hoping that their teams would win, but also that they would do something spectacular, get noticed, and move to a better team. In the semi-final between England and Brazil, Ronaldinho—one of Brazil's top players, but not the superstar Ronaldo—took a direct shot on goal with a free kick from 35 yards out. It would have almost certainly been better to pass to a player near the goal. However, Ronaldinho scored an impossible goal. His reward is likely to be a move from his decent French team to one of the rich Spanish, Italian or British teams that pay much better. Had he missed, his French job was not at risk. See O'Keefe,

- Viscusi, and Zeckhauser (1984) on contests, Rosen (1981) on superstars, and Riley (2001) for an excellent survey on signaling.
3. See Frank (1986). Frank is primarily concerned with the structure of individual preferences, and the likelihood that relative standing matters. By contrast, our focus is on the rewards others give you after inferring your type (capability). In both cases, the mix of types in the population matters.
 4. Breeden and Viswanathan (1998) has some resemblance with our work. However, their focus is on financial hedging. Our focus is on risk-taking involving real activities. Moreover, our theoretical setup is more general (we allow for a continuum of actions and payoffs), and we also conduct empirical work.
 5. Many actions with strong variance implications are likely to be unobservable—for example, a firm's R&D decisions. DeGeorge, Moselle, and Zeckhauser (2002) analyze the equilibrium in the observable risk choice case. Consider students facing two versions of an aptitude test. It is common knowledge that the first version of the test measures aptitude precisely, while the second is less precise. In this case the market, in its assessment of a student's ability, will take into account both the test result and the information conveyed by the student's choice of which version to take—precise or imprecise. Good students will prefer the precise version of the test. Bad students might be tempted to choose the imprecise version—but in so doing, they would automatically reveal themselves to be bad. The authors show that under such conditions all agents, regardless of type, have strong incentives to choose identical, low—often below the mean maximizing—levels of risk.
 6. If the payoff were to a worker, the payoff might be a long-term employment contract.
 7. This posterior probability emerges because no matter what the bad type plays, he can never reach a first-period payoff of 4.
 8. This equilibrium is sustainable for $0.18 < p < 1$.
 9. If bad types had a slightly worse distribution of outcomes from playing risky, the equilibrium for this low value of p would have good types play risky and bad types play safe.
 10. If there are but two strategy choices, it is possible to have a low- p equilibrium where good types play risky and bad types play safe.
 11. If the projects were indivisible, discrete methods rather than calculus would be required to produce some of our results. Since these are real projects, there is an upper bound to how much firms can invest in them (finite project size), and a lower bound (no, or at least limited, ability to short the projects).
 12. We have drawn $\mu(V)$ symmetric around V^* , but symmetry is not crucial, as we show below.
 13. In other cases $\mu(V)$ may be less tractable. For example, if the manager's budget constraint is binding at the peak then there may be discontinuities, non-concavity and even non-monotonicity in the $V > V^*$ part of the curve. However, even in this more difficult case, standard portfolio theory implies concavity for $V < V^*$ (Huang and Litzenberger, 1988, pp. 63–80).
 14. In the first best case the firm invests 100% in every project with positive expected return. This gives the peak of the $\mu(V)$ curve. If the manager wishes to increase variance, he will invest in the negative return project with the lowest absolute value of the mean-to-variance ratio. Once this opportunity to add variance is exhausted (by the finite size assumption), the manager identifies the second lowest mean-to-variance ratio project and the process continues. To meet a precise variance requirement may entail some reordering of projects.
 15. The assumption that the new owner is a permanent one enables us to abstract away from a range of dynamic issues. Multiple periods of earnings would give investors some chance to infer variance from an earnings series, but the basic principles would still apply. The analysis becomes exceedingly complex, even for two periods, since each firm is engaged in a contingent, dynamic signaling strategy, assuming it can change what it does in the second period depending on its earnings in the first.
 16. For example, suppose that V^* is “much better” than any other local maximum, in the sense that $\mu(V^*) > \mu(V) + \Delta/r$ for any other local peak V (so that any other local maximum is a dominated strategy). Then there is no need to require $\mu(V)$ to be single-peaked globally, since it will automatically be single-peaked within the region $[V_{\min}, V_{\max}]$. By definition, any V in that region satisfies $\mu(V^*) < \mu(V) + \Delta/r$ and therefore by the assumption that V^* is “much better” it cannot be a local peak unless $V = V^*$.
 17. Our previous assertion that the symmetry of $\mu(V)$ is not essential is best understood in the light of this best response mapping. The best response to (W_b, W_g) does not depend on the value of $\mu(V)$ at or in the neighborhood of the “mirror points” $(V^* + (V^* - W_b), V^* + (V^* - W_g))$ any more than on its value at or around other points. The use of a symmetric function for $\mu(V)$ is helpful in that it simplifies computation, and removes one degree of freedom for achieving results.

18. That is, start with a guess (W_b, W_g) to “seed” the process, and set $V_b(0) = W_b, V_g(0) = W_g$, then apply the best response map iteratively: $V_b(n+1) = V_b^*(V_b(n), V_g(n)), V_g(n+1) = V_g^*(V_b(n), V_g(n))$. If this process converges, then by construction the point it converges to is indeed a fixed point of the mapping as required.
19. This functional form implies a negative mean. Our results would obviously be unaffected by the addition of a constant.
20. In terms of Figure 2, when $p = 0.1$, both μ_b and μ_g are in a convex region of the curve, inducing risk-increasing behavior.
21. We restrict attention to firms for which we have at least ten years of data. Because of the scaling we use, we discard firms with a negative mean performance over the period.
22. The literature on Tobin’s q (e.g. Chung and Pruitt, 1994) shows this measure to be very close to the one using the Lindenberg and Ross (1981) procedure.
23. See for instance Servaes (1991).
24. The results are confirmed when we control for industry or size. Detailed results are available upon request.
25. In the interest of space we do not report these results in detail. They are available upon request.
26. The aggregate t -statistics were computed for a null-hypothesis assuming independence in observations across industries. On rank correlation methods, see e.g. Mosteller and Rourke (1973) or Kendall and Gibbons (1990).
27. For an application of Glicksberg’s theorem to a space of Borel measures, see Milgrom and Weber (1985).
28. Aliprantis and Border (1994).

References

- Aliprantis, Charalambos D. and Kim C. Border. (1994). *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. New York: Springer-Verlag.
- Breeden, D. and S. Viswanathan. (1998). “Why do Firms Hedge? An Asymmetric Information Model,” mimeo, Duke University.
- Chung, Kee H. and Stephen Pruitt. (1994). “A Simple Approximation of Tobin’s q ,” *Financial Management* 23, 70–74.
- DeGeorge, François, Boaz Moselle, and Richard J. Zeckhauser. (2002). “The Ecology of Risk Taking,” Working Paper #03-053, Division of Research, Harvard Business School.
- Frank, Robert H. (1986). *Choosing the Right Pond: Human Behavior and the Quest for Status*. New York: Oxford University Press.
- Hogg, J.T. and S.H. Forbes. (1997). “Mating in Bighorn Sheep: Frequent Male Reproduction via a High-Risk ‘Unconventional’ Tactic,” *Behavioral Ecology and Sociobiology* 41, 33–48.
- Huang, Chi-fu and Robert H. Litzenberger. (1988). *Foundations for Financial Economics*. New York: North-Holland.
- Kendall, Maurice and Jean D. Gibbons. (1990). *Rank Correlation Methods*. London: Edward Arnold.
- Lindenberg, Eric B. and Stephen A. Ross. (1981). “Tobin’s q Ratio and Industrial Organization,” *Journal of Business* 54, 1–32.
- Milgrom, Paul and Robert J. Weber. (1985). “Distributional Strategies for Games with Incomplete Information,” *Mathematics of Operations Research* 10, 619–632.
- Mosteller, Frederick and Robert E.K. Rourke. (1973). *Sturdy Statistics*. Reading: Addison-Wesley.
- O’Keefe, Mary, W. Kip Viscusi, and Richard J. Zeckhauser. (1984). “Economic Contests: Comparative Reward Schemes,” *Journal of Labor Economics* 2, 27–55.
- Riley, John G. (2001). “Silver Signals: Twenty-Five Years of Screening and Signaling,” *Journal of Economic Literature* 39, 432–478.
- Rosen, Sherwin. (1981). “The Economics of Superstars,” *American Economic Review* 71, 845–858.
- Servaes, Henri. (1991). “Tobin’s q and the Gains from Takeovers,” *The Journal of Finance* 46, 409–419.