

The definitive article is available on the publisher's website:
<http://link.springer.com/article/10.1023%2FA%3A1021952927149>



The Geneva Papers on Risk and Insurance Theory, 27: 153–165, 2002
© 2003 The Geneva Association

The Pricing of Event Risks with Parameter Uncertainty

KENNETH A. FROOT

Harvard Business School, Soldiers Field, Morgan Hall 391, Boston, MA 02163, USA

STEVEN E. POSNER

Goldman Sachs & Co., 85 Broad Street, New York, NY 10004, USA

Received March 8, 2000; Revised September 26, 2002

Abstract

Financial instruments whose payoffs are linked to exogenous events, such as the occurrence of a natural catastrophe or an unusual weather pattern depend crucially on actuarial models for determining event (e.g., default) probabilities. In many instances, investors appear to receive premiums far in excess of these modeled actuarial probabilities, even for event risks that are uncorrelated with returns on other financial assets. Some have attributed these larger spreads to uncertainty in the probabilities generated by the models. We provide a simple model of such 'parameter uncertainty' and demonstrate how it affects rational investors' demand for event risk exposures. We show that while parameter uncertainty does indeed affect bond spreads, it does not tend to increase spreads by much. Indeed, the spread increases due to parameter uncertainty in our numerical examples are on the order of only 1–2 basis points. Moreover, in many instances, including those that have the most sensible correlation settings, parameter uncertainty tends to *decrease* the size of bond spreads. We therefore argue that parameter uncertainty does not appear to be a satisfactory explanation for high event-risk returns.

Key words: parameter uncertainty, event risk, catastrophe risk, asset pricing

JEL Classification No.: G22

1. Introduction

Bonds and other instruments with payments linked to explicit events are growing in popularity among investors and intermediaries. Examples include catastrophe bonds and other insurance-linked securities, weather-related and credit-risk derivatives. Most Wall Street institutions have committed resources to analyze these products for the purposes of both issuing and trading them. Of particular interest is the pricing of catastrophe ("cat") event and weather risks. These risks are perhaps unique to financial economists because their actuarial properties can be objectively modeled with considerable sophistication. They are unique to practitioners because their yields have fallen dramatically over the past few years as traditional reinsurance markets have become more contestable. The appropriate level for these yields going forward is a matter of considerable debate among practitioners and academics.

Cat and weather risks tend to be uncorrelated with the returns on diversified financial portfolios.¹ In addition, each risk is individually small relative to the risks in financial markets. In most equilibrium models this would imply that event risks should yield the risk free rates plus an unbiased actuarial estimate of expected loss. This statement arises from well-known concepts in financial theory. However, the theory is applicable only in markets that have reached an equilibrium so that risks are widely shared—a state which some markets most certainly have not reached.

But many observers argue that event risks have not been, and should not be, priced at actuarial levels, even in equilibrium. Many practitioners expect the risks to provide a yield considerably greater than the actuarially expected loss. There is little agreement, however, on the determinants of this premium spread. For cat-linked instruments, the premium is most commonly determined as a fixed constant times the volatility (or variance) of loss. More sophisticated users add to the premium based on the skewness of the loss profile. Others argue that the expected loss *conditional* on an event, effectively explains cat-linked yields (see Morton Lane [1998], for example).

Yet another group argue that the uncertainty associated with actuarial probabilities is an important explanatory factor for high event yields. Our paper addresses this argument. It begins by asking what the pricing of event bonds ought to be in an equilibrium where actuarial probabilities are uncertain. We first show that uncertainty in actuarial probabilities does not affect the pricing of the event instruments as long as the actuarial probabilities are unbiased and are uncorrelated with the event itself. Even if the uncertainty about probabilities is correlated with the event outcomes—a relatively unlikely state of affairs—we find that the uncertainty about event probabilities is insufficient to explain the level of cat-event spreads. If the uncertainty about event probabilities is correlated across a number of different events, there is a detectable, though generally small, upward influence on spread. As a result, we argue that parameter uncertainty about the probability of events is unlikely to explain why event spreads are high.

Of course, we can not rule out the hypothesis that event-risk spreads are high not because of uncertainty per se, but because participants believe that the true average probabilities of event occurrence are higher than those put forth by third-party modelers. In other words, the market believes that the typically-measured event probabilities are downward biased. While this hypothesis is possible, it strikes us as unlikely to be true: it would be extraordinary to claim that unsophisticated “gut” views of event probabilities are able to avoid the biases that plague sophisticated models that use historical data, simulation and other techniques. Thus, detectable bias is unlikely to be the reason that event-risk yields are so far in excess of their associated event probabilities.

Why then have yields exceeded actuarially-fair levels? For catastrophe event risks in particular, we argue that risk pricing is currently determined by reinsurers. This is inefficient, since reinsurers are equally as informed as modeling-firms yet they hold large risk positions (positions that are not diversified, yet benchmarked against absolute returns rather than against the performance of aggregated event risks). As the market learns that it can profitably underbid reinsurers for these risks, future spreads are likely to decline.

2. A simple model of event risk pricing

In this section we lay out a simple framework for pricing event risk, one that allows us to take parameter uncertainty explicitly into account. Our approach is to assume that there is a value-maximizing investor (or institution) that can choose among risky assets including event risk. We consider how that investor values event risk under alternative assumptions about the distribution of that risk and the uncertainty associated with the distribution's parameters. Our portfolio approach allows us to understand how a firm should either:

- Price an event instrument in the marketplace, either as a buyer or a seller;
- Decide (if and) how much of an offered event instrument should be added to an existing portfolio.

Of course, the holder of a large portfolio of instruments must view any alterations or additions to the portfolio in the larger context of the entire portfolio. We focus on how event allocations affect the higher moments of that portfolio. We allow the investor to choose among a portfolio of risky and riskless assets, in addition to any exposure to event risk. This allows the investor to allocate assets optimally at all times, and yet is flexible enough to apply to investors who traditionally have had no direct event exposure, as well as to others (e.g., reinsurers), whose portfolios contain concentrated allocations to event risk.

Based on this analysis, we derive a series of optimal event-risk demand curves. These show the optimal allocation for alternative risk levels and given event yields (or, reciprocally, the minimum yield required for any given size allocation to event risk).

2.1. Basic assumptions and notation

To keep things simple, our setup makes several basic assumptions. First, we assume the investor can choose among three assets: a risky market portfolio; a riskfree asset; and catastrophe exposure. We denote the end-of-year value of the market asset by M , where M is assumed to be lognormally distributed:

$$\ln M \sim \text{Normal}(\mu, \sigma),$$

as we normalize the initial value of the market to unity. The market's expected value is therefore

$$E[M] = e^{\mu + \sigma^2/2} \equiv 1 + r + E(M_e),$$

where r is the annualized risk free rate and

$$M_e = M - (1 + r)$$

is the annual excess arithmetic return. This market portfolio represents a preexisting, well-diversified portfolio held by the investor.

The event risk is assumed to be held through a simple binary bond. One dollar invested in this bond has a one year excess return of C_e . This takes the value s (the promised spread) with probability $1-P$ or $s-1$ (a cat loss) with probability P . The expected excess return is therefore given by, $E[C_e] = s - P$. The notation can be simplified by defining a binary random variable, $B(P)$, which is equal to 0 with probability $1-P$ and 1 with probability P , such that:

$$B(P) = s - C_e.$$

Second, we allow the investor to consider an allocation of wealth among the market portfolio, the risk free asset, and the event exposure. The portfolio has a stochastic excess return of:

$$R_e = w_1 M_e + w_2 C_e, \quad (1)$$

where the weight allocated to the riskfree asset is given by $1 - w_1 - w_2$. Below, we will make alternative assumptions about the distributions of M and C risks.

Third, we need some sort of penalty for bearing risk. One way to think of this at a completely general level is that the investor wishes to maximize expected return, holding a function of all the other moments of the portfolio return distribution to be a constant:

$$\max_{w_1, w_2} E[R_e] \quad \text{subject to} \quad f(E[R_e^2], E[R_e^3], E[R_e^4], \dots) = \lambda \quad (2)$$

where λ is a constant. A sufficient condition to solve Eq. (2) is to maximize return subject to holding each individual moment constant:

$$\max_{w_1, w_2} E[R_e] \quad \text{subject to} \quad E[R_e^2] = \lambda_2, E[R_e^3] = \lambda_3, E[R_e^4] = \lambda_4, \dots, \quad (3)$$

where each λ_i is a constant. The constraints in Eqs. (2) and (3) are very general. However, in some instances we will want to provide explicit solutions and/or simulate the results. In those instances, one can simply assume that only the first two moments matter, so that we can maximize the portfolio's Sharpe ratio, i.e., the ratio of expected excess return divided by its standard deviation.

3. Equilibrium expected excess returns on independent event risks are zero

Our first, and very general, point concerns the pricing of event risk in equilibrium. We think of this as the reward in excess of the riskfree rate that event risk would pay if all investors held it equally in proportion to their wealth. We consider individual event risk exposures that are relatively small in comparison with total financial wealth, so that the fraction of overall wealth devoted to these exposures is near zero. In addition, we assume for the moment that individual event exposures are independent of other risks. Under such idealized circumstances, the equilibrium expected return on event risk is the riskfree rate. In other words, the promised yields on event exposures would equal the expected loss rate.

This conclusion follows from the fact that the value of the objective function in Eqs. (2) and (3) remains unchanged if an investor adds a small amount of event risk at her reservation return. If the investor has no preexisting allocation to event risk and event risk is independent of other financial risks, then this reservation expected return must be the risk free rate. A proof of this statement proceeds as follows.

First, if in equilibrium small changes in w_2 are to have no effect on the moments in Eqs. (2) and (3), then the excess return on event risk must be zero. To see this, note that the n 'th moment of the portfolio, is given by $E[(w_1 M_e + w_2 C_e)^n]$. When the event weight w_2 changes, the change in the portfolio's n th moment is $nE[(w_1 M_e + w_2 C_e)^{n-1} C_e]$. Given that C_e is independent of other portfolio returns, and that w_2 is small in equilibrium (i.e., $w_2 = 0$), this expression simplifies to $nE[(w_1 M_e)^{n-1}]E[C_e]$, which is zero only if the expected excess return on event risks is zero, $E[C_e] = 0$.² Because this is true for all n , none of the moments change when the excess return on event risk is zero. And if there is no changes in any of the moments in Eqs. (2) and (3), there can be no change in the objective function. Thus, regardless of the shape of the event risk distribution, its independence and small size imply that $E[C_e] = 0$ is the equilibrium expected excess return on event risk.

4. Parameter uncertainty in an equilibrium model for multiple event risks

4.1. Independent events and independent parameter uncertainty

The second general point we can make concerns the consequences of parameter uncertainty in event distributions.³ To gain intuition, recall that for a simple binary event bond, the one year excess return, C_e , is either s with probability $1-P$ or $s-1$ with probability P , summarized by $C_e = s - B(P_e)$. We now introduce parameter uncertainty, in that the observed probability, p , is itself a random variable modeled as $p = P + \varepsilon$, where ε is an arbitrary zero mean distribution (of appropriate support to be sensible). The mean of ε is zero so that p is an unbiased estimator of P .

It is easy to see that the moments of the event bond excess return are independent of the distribution of ε . The n 'th moment of the event bond excess return is unrelated to the distribution of ε , provided that its mean is zero:

$$E[C_e^n] = E\{E[C_e^n | p]\} = E\{(1-p)s^n + p(s-1)^n\} = (1-P)s^n + P(s-1)^n \quad (4)$$

This simple logic—together with the assumption that event risk is independent of market risk—shows that the variability in the estimation of P is irrelevant for the pricing of binary event bonds. That conclusion follows because, from Eq. (4), variability in P has no effect on any of the higher-order portfolio moments in Eqs. (2) and (3).

It is important to note that this result holds for *any* allocation to event risk (i.e., independently of w_2). To see this, note that each moment of the portfolio can be written as a set of terms in the separate, lower-order moments of market and event risk:

$$E[(w_1 M_e + w_2 C_e)^n] = \sum_{i=0}^n a_{i,n} E[(w_1 M_e)^{n-i}] E[(w_2 C_e)^i], \quad (5)$$

where $a_{i,n} = \binom{n}{i}$. Since all portfolio moments can be expressed as functions of event moments, and since parameter uncertainty has no effect on event moments, parameter uncertainty has no effect on the pricing of independent binary event risks, regardless of the importance of these risks in the portfolio.

If there are multiple event risks, it is easy to show that the above logic will apply if each event is independent and has independent parameter uncertainty. To see this, suppose that we split the event return, C_e , into a portfolio of two event instruments with returns, $C_{e,1}$ and $C_{e,2}$. $C_{e,i}$, $i = 1, 2$ takes the value s_i with probability $1 - p_i$ or worth $s_i - 1$ with probability p_i . As before, we can write this as

$$C_{e,i} = s_i - B_i(p_i), i = 1, 2, \quad (6)$$

where $B_i(p_i)$ is a binary random variable, equal to 1 with probability p_i and 0 with probability $1 - p_i$.

We assume, as before, that the true probabilities are not known and that we observe p_i , where $p_i = P_i + \varepsilon_i$, $i = 1, 2$, where ε_1 and ε_2 have zero mean and identical higher-order moments. For simplicity, we assume that $P_1 = P_2 = P$. For the moment, we treat ε_1 as independent of ε_2 (e.g., parameter uncertainties are unrelated) and $C_{e,1}$ as independent of $C_{e,2}$ (e.g., hurricane and earthquake). Given this, the last expectation term of Eq. (5) can be written as:

$$\begin{aligned} E[(w_2 C_e)^n] &= g^n \sum_{i=0}^n a_{i,n} (-1)^i \sum_{j=0}^i a_{j,i} \hat{w}_{21}^{i-j} \hat{w}_{22}^j E[(B_1(p_1))^{i-j} (B_2(p_2))^j] \\ &= g^n E[B_1(p_1) B_2(p_2)] \sum_{i=0}^n a_{i,n} (-1)^i \sum_{j=0}^i a_{j,i} \hat{w}_{21}^{i-j} \hat{w}_{22}^j, \end{aligned} \quad (7)$$

where w_{21} and w_{22} are weights given to the two event risks, $g = w_{21}s_1 + w_{22}s_2$ and $\hat{w}_{2i} = w_{2i}/g$. Equation (7) shows that the expected product of the binary event outcomes, $E\{B_1 B_2\}$, is the critical term. It is easy to show that this term is given by:

$$\begin{aligned} E\{B_1 B_2\} &= E\{E[B_1 B_2 | p_1, p_2]\} = E\{E[B_1 | p_1, p_2] E[B_2 | p_1, p_2]\} \\ &= E[p_1 p_2] = P^2 \end{aligned} \quad (8)$$

This expression says that if the outcomes and parameter uncertainties are independent, the moments of the event portion of the portfolio can be written as functions of the true probabilities, but not as a function of the uncertainty of these probabilities. As a result, parameter uncertainty again can have no impact on the objective functions in Eqs. (2) and (3). This is true regardless of the weight of the event risks in the portfolio (i.e., w_{21} and w_{22} need not be small). Of course, an investor will demand a higher spread to hold a large fraction of her portfolio in a single event risk. But this additional charge is not affected by uncertainty concerning the distribution of that risk.

4.2. Correlated events and correlated parameter uncertainty

Next, we relax the assumption that the event occurrences and parameter uncertainties are all independent of one another. In doing so, we want to allow for two types of correlation. The first correlation is between the two event outcomes. An example would be catastrophic hurricane losses in two adjacent Florida zip codes. For such events, the probability of a hurricane to one zip code is related to whether damage occurred in the adjacent area. Mathematically, this means the covariance of event outcomes, *for given ex ante probabilities*, is nonzero.

The second correlation is that between the two parameter uncertainties. For example, the same model might be used to evaluate the two Florida zip codes. Any modeling error in the first zip code is likely to spill over into the second.

With these two types of correlation, the expected product of event outcomes becomes:

$$\begin{aligned}
 E\{B_1 B_2\} &= E\{E[B_1 B_2 | p_1, p_2]\} \\
 &= E\{\text{cov}[B_1, B_2 | p_1, p_2] + E[B_1 | p_1, p_2]E[B_2 | p_1, p_2]\} \\
 &= E\{\text{cov}[B_1, B_2 | p_1, p_2]\} + E\{p_1 p_2\} \\
 &= E\{\text{cov}[B_1, B_2 | p_1, p_2]\} + \text{cov}[p_1, p_2] + P^2
 \end{aligned} \tag{9}$$

Subtracting P^2 from both sides, yields:

$$\text{cov}[B_1, B_2] = E\{\text{cov}[B_1, B_2 | p_1, p_2]\} + \text{cov}[p_1, p_2]. \tag{10}$$

By using the assumption that the correlation coefficient, ρ , is a fixed parameter independent of the p 's, the first term on the right-hand side of Eq. (10) can be expressed as:

$$\text{cov}[B_1, B_2 | p_1, p_2] = \rho \sqrt{p_1(1-p_1)p_2(1-p_2)} \tag{11}$$

Hence we have that

$$\text{cov}[B_1, B_2] = \rho E \sqrt{p_1(1-p_1)p_2(1-p_2)} + \text{cov}[p_1, p_2] \tag{12}$$

The first term on the right-hand side of Eq. (12) can be approximated to second order (see appendix for a proof), yielding:

$$\text{cov}\{B_1, B_2\} = \rho P(1-P) - \rho \sigma_p^2 - \rho(\sigma_p^2 - \sigma_{p_1, p_2}) \frac{(1-2P)^2}{4P(1-P)} + \sigma_{12} \tag{13}$$

where

$$\begin{aligned}
 P &= E[p_1] = E[p_2] \\
 \sigma_{12} &= \text{cov}[p_1, p_2] \\
 \sigma_p^2 &= \sigma_{p_1}^2 = \sigma_{p_2}^2 = \text{var}[p]
 \end{aligned}$$

where ρ is the *conditional* correlation of the outcomes of the two bonds—conditional on the probabilities being known or independent (so that $\text{cov}[p_1, p_2] = 0$).

Since we have for simplicity set as equal the variances of the parameter uncertainties, we can define the correlation of the parameter uncertainties, u , to be

$$\sigma_{12} = u\sigma_p^2.$$

Equation (13) can then be rewritten as

$$\text{cov}\{B_1, B_2\} = \rho P(1 - P) + \{u - \rho\}\sigma_p^2 - \rho\sigma_p^2(1 - u)\frac{(1 - 2P)^2}{4P(1 - P)}$$

By recognizing that $\sigma_B^2 = P(1 - P)$ we can then define the *unconditional* correlation between the bonds—the correlation in payoffs allowing for correlated parameter uncertainties—as:

$$\rho_u = \frac{\text{cov}\{B_1, B_2\}}{\sigma_B^2} = \rho + \{u - \rho\}\frac{\sigma_p^2}{\sigma_B^2} - \rho\frac{\sigma_p^2}{\sigma_B^2}(1 - u)\frac{(1 - 2P)^2}{4\sigma_B^2} \quad (14)$$

Note that for the special case of where the two bonds are identical (which would imply that $\rho = u = 1$), we have that $\rho_u = \rho$, i.e., that parameter uncertainty is irrelevant. This is what we would expect given the result for a single event bond above: that regardless of the allocation to that bond, parameter uncertainty is irrelevant for pricing.

Simplifying, we then have that

$$\begin{aligned} \rho_u &= \rho + \frac{\sigma_p^2[u(1 + a\rho) - \rho(1 + a)]}{(1 - P)P} \\ &= \rho + \frac{\sigma_p^2[u(4P(1 - P) + \rho(1 - 2P)^2) - \rho]}{4(1 - P)^2 P^2}, \end{aligned} \quad (15)$$

where $a = \frac{(1 - 2P)^2}{4P(1 - P)}$ and noting that $1 + a = \frac{1}{4P(1 - P)}$.

This result has some intuitive properties. First, the correlation of the outcomes can either increase or decrease with degree of parameter uncertainty (given by the variance of p , σ_p^2). A sufficient condition for greater parameter uncertainty to decrease the correlation of outcomes is that $\rho > u$, i.e., that the conditional correlation of the events is higher than the correlation of the parameter uncertainty. In this case, the parameter uncertainty behaves like noise, reducing the correlation of the outcomes.

Second, the greater is the correlation of the parameter uncertainties, the greater is the correlation of the observed outcomes relative to the conditional correlation of the events.

Third, in the presence of parameter uncertainty (i.e., $\sigma_p^2 > 0$) a given increase in the conditional correlation of the events increases by a smaller amount the correlation of observed outcomes.

Some special cases are interesting and worth noting:

$$\begin{aligned}\rho = 0 &\Rightarrow \rho_u = \frac{u\sigma_p^2}{\sigma_B^2} \\ u = 0 &\Rightarrow \rho_u = \rho \left(1 - \frac{\sigma_p^2}{4\sigma_B^4} \right) \leq \rho \\ u = 1 &\Rightarrow \rho_u = \rho + (1 - \rho) \frac{\sigma_p^2}{\sigma_B^2} \geq \rho\end{aligned}$$

First, if there is conditionally zero correlation between the events then unconditional correlation will have the same sign as the correlation between the parameter uncertainties (u).

Second, if $u = 0$ then the unconditional correlation is always smaller than the conditional correlation.

Third, if there is perfect correlation between the parameter uncertainties ($u = 1$) then the unconditional correlation is always greater than the conditional correlation.

5. Numerical examples

We next consider the impact that parameter uncertainty has on spreads. Consider a holder of one unit of a cat bond C_1 . This cat bond pays a premium with spread of s_1 over riskfree rates. The probability of this bond being triggered is assumed to be P (in expectation).

This investor is now considering the acquisition of x units (assumed to be positive, i.e., that he goes long the new bond) of cat bond C_2 . This bond also has trigger probability P in expectation. The criterion by which the investor chooses a portfolio allocation is that of Sharpe-ratio invariance. That is,

$$\frac{s_1 - P}{\sqrt{P(1 - P)}} \leq \frac{s_1 - P + x(s_2 - P)}{\sqrt{(P(1 - P))(1 + x^2 + 2x\rho)}} \quad (16)$$

For now, we use ρ to denote correlation in the traditional sense—this will be refined momentarily.

The inequality in Eq. (16) implies that

$$(s_1 - P) \frac{\sqrt{1 + x^2 + 2x\rho} - 1}{x} + P \leq s_2 \quad (17)$$

The left-hand side of Eq. (17) gives the minimum spread on cat bond 2 required for a given allocation level x as a function of correlation. For sufficiently large negative correlations the threshold may be negative, i.e., the investor would accept a negative spread. This is intuitive since with sufficient negative correlation, the second bond is a natural hedge for the first bond so the investor will in effect pay to hedge the first bond.

Now we use this setup to demonstrate the pricing impact of parameter uncertainty. We compute the minimum spread required for the second bond under two cases: first where ρ represents the basic conditional correlation and where ρ represents the unconditional correlation, ρ_u . In the example, we use Eq. (15) and assume that $P = \sigma_p = 1.00\%$, $x = 1$ (so that the portfolio is initially entirely invested in the first event bond), and $s_1 = 4.00\%$ (which is a realistic approximation of the spreads on 1% event bonds). We vary below both ρ and u (the correlation of the parameter uncertainty). The table displays the difference between the minimum spread required for the second cat bond using the unconditional correlation and that using the conditional correlation. In other words, the table shows how much larger the second spread must be in order to compensate for the presence of parameter uncertainty.

There are several interesting observations to make about how parameter uncertainty affects spreads. First, take the case in which the basic events are independent (e.g., hurricane and earthquake), so that that $\rho = 0$. The presence of parameter uncertainty can still affect spreads, but only by inducing a correlation in the unconditional correlation of the outcomes. If, for example, the correlation between the p 's is zero, there is no effect on spreads, because investors recognize that the parameter uncertainty in one bond has no implications for the outcomes of the other bond. However, if the correlation between the p 's is positive, investors recognize that the probabilities of the two independent events tend to be overstated (or understated) at the same time. With positive correlation of the p 's, the conditional probability of an earthquake occurrence is rationally higher, given the occurrence of a hurricane. Thus, the correlation of the p 's actually induces a positive correlation in the bond outcomes, though the increase in spreads is very small. Table 1 shows that the spread increases by only as much as 2 basis points with 1 percent true event probabilities.

Table 2 shows the impact on the unconditional correlation of the outcomes in this example. In the case with $\rho = 0$ and $u > 0$, Eq. (15) shows that the correlation between outcomes becomes $\rho_u = \frac{u\sigma_p^2}{(1-P)P} > 0$. Note, however, that as Table 2 shows, this induced correlation is very small, reaching a maximum of only 1 percent when $u = 1$. The intuition is from Eq. (15), which shows that what matters is the ratio of the variance of the parameter uncertainty to the variance of the event outcome. This ratio becomes zero for small event

Table 1. Additional event bond spread induced by parameter uncertainty.

		(correlation between p 's) u (%)				
		0	25	50	75	100
ρ	-100	2.14	1.88	1.57	1.19	0.60
	-75	0.70	0.55	0.40	0.25	0.07
	-50	0.36	0.28	0.21	0.13	0.05
	-25	0.15	0.12	0.09	0.06	0.03
	0	0.00	0.01	0.01	0.02	0.02
	25	-0.12	-0.09	-0.05	-0.02	0.01
	50	-0.23	-0.17	-0.11	-0.05	0.01
	75	-0.32	-0.23	-0.15	-0.07	0.00
	100	-0.40	-0.29	-0.19	-0.10	0.00

Table 2. Correlation of the outcomes of two event bonds in the presence of parameter uncertainty.

		u (correlation between p 's) (%)					
		0	25	50	75	100	
ρ	-100	-74	-80	-86	-92	-98	
	-75	-56	-60	-65	-69	-73	
	-50	-37	-40	-43	-46	-48	
	-25	-19	-20	-21	-22	-24	
	0	0	0	1	1	1	
	25	19	20	22	24	26	
		50	41	44	47	51	
		75	56	61	66	70	75
		100	74	81	87	94	100

probabilities, given our assumption that the distribution of the parameter uncertainty is more symmetric than are binary event outcomes around their means.

In the case where the correlation between the events is positive, $\rho > 0$, but there is uncorrelated parameter uncertainty, $u = 0$, the uncertainty actually *lowers* the spread on the second bond by about 40 basis points compared to what it would have been without parameter uncertainty (see Table 1). The reason is clear from Table 2, which shows that the correlation of outcomes is lower than ρ unless u is sufficiently large. Essentially, the presence of parameter uncertainty introduces noise into what would otherwise be perfectly correlated outcomes. Unless this noise is also highly correlated, the result is a decline in the correlation of outcomes. Of course, if the parameter uncertainty is very highly correlated (i.e., 100%) it can induce a slight increase in the observed correlation, again only on the order of 1%. For example, for $\rho = 0.50$ and $u = 1.00$, the observed correlation becomes slightly higher, at $\rho_u = 0.51$.

In fact, the only instance where the spread of the second bond is increased substantially by parameter uncertainty is the case in which the conditional correlation, ρ , is *negative* (e.g., earthquake this year and next year). The result is that the spread on the second bond must be higher because investors cannot be sure to receive as much diversification as they would hope to if the probabilities were known. Thus, the spreads become higher as does the correlation of the outcomes. This effect can be large, as in the case where $\rho = -100\%$ and $u = 0$; the spread on the second bond increases by 214 basis points over that of the first bond and the correlation between the bonds increases from -100% to -74% .

The more plausible case, is where the conditional events are uncorrelated or somewhat positively correlated and the parameter uncertainty is more highly positive correlated. This is likely to be the case, since model runs probably tend to have similar biases from one exposure measurement to the next. In that case, the basic effect of parameter uncertainty is not to raise the spread on the second bond, but actually to *lower* it. Thus, in this sense, parameter uncertainty not likely to provide a satisfactory explanation for high event-risk spreads.

These numerical examples are, of course, special to all of the assumptions that we have made. Nevertheless, the qualitative results are fairly robust. We have assumed that the initial portfolio is comprised completely of a single, preexisting event exposure. Portfolios that contain less parameter uncertainty (say, because they include exposures other than event risk) will be less sensitive to the parameter uncertainty. Thus, the impact on the second bond's spreads will be even lower than shown in Table 1.

6. Summary

Parameter uncertainty is often alleged to be an important reason why rational investors require high compensation for bearing uncorrelated event risk. We have shown that while parameter uncertainty can affect bond spreads, it does not tend to increase spreads by much. Indeed, the spread increases due to parameter uncertainty that we have seen in our numerical examples are on the order of only 1–2 basis points, which is two orders of magnitude smaller than that observed in the market. Indeed, in many instances, including those that have the most sensible correlation settings, parameter uncertainty tends to *decrease* the size of bond spreads.

Appendix

We prove that

$$\begin{aligned} & E\{\text{cov}[B_1, B_2 | p_1, p_2]\} \\ &= \rho P(1-P) \left(1 - \frac{\text{var}[p]}{P(1-P)} - (\text{var}[p] - \text{cov}(p_1, p_2)) \frac{(1-2P)^2}{4P^2(1-P)^2} \right) \end{aligned}$$

The main tool we use is a second-order Taylor series approximation of $\sqrt{1+x} \approx 1 + \frac{x}{2} - \frac{x^2}{8}$:

$$\begin{aligned} & E\{\text{cov}[B_1, B_2 | p_1, p_2]\} = \rho E\sqrt{p_1(1-p_1)p_2(1-p_2)} \\ & E\sqrt{p_1(1-p_1)p_2(1-p_2)} = E\sqrt{(P+\varepsilon_1-(P+\varepsilon_1)^2)(P+\varepsilon_2-(P+\varepsilon_2)^2)} \\ &= P(1-P)E\sqrt{\left(1 + \frac{\varepsilon_1(1-2P)-\varepsilon_1^2}{P(1-P)}\right)\left(1 + \frac{\varepsilon_2(1-2P)-\varepsilon_2^2}{P(1-P)}\right)} \\ &= P(1-P)E\sqrt{1 + \frac{\varepsilon_1(1-2P)-\varepsilon_1^2}{P(1-P)} + \frac{\varepsilon_2(1-2P)-\varepsilon_2^2}{P(1-P)} + \frac{\varepsilon_1(1-2P)-\varepsilon_1^2}{P(1-P)} \frac{\varepsilon_2(1-2P)-\varepsilon_2^2}{P(1-P)}} \\ &\approx P(1-P)E\left\{1 + \frac{\varepsilon_1(1-2P)-\varepsilon_1^2}{2P(1-P)} + \frac{\varepsilon_2(1-2P)-\varepsilon_2^2}{2P(1-P)} + \frac{\varepsilon_1\varepsilon_2(1-2P)^2}{2P^2(1-P)^2} - \frac{(1-2P)^2}{8P^2(1-P)^2}(\varepsilon_1+\varepsilon_2)^2\right\} \\ &= P(1-P)\left\{1 - \frac{\sigma_1^2}{2P(1-P)} - \frac{\sigma_2^2}{2P(1-P)} + \frac{\sigma_{12}(1-2P)^2}{2P^2(1-P)^2} - \frac{(1-2P)^2}{8P^2(1-P)^2}(\sigma_1^2 + \sigma_2^2 + 2\sigma_{12})\right\} \\ &= P(1-P) - \sigma_p^2 + \frac{\sigma_{12}(1-2P)^2}{2P(1-P)} - \frac{(1-2P)^2}{4P(1-P)}(\sigma_p^2 + \sigma_{12}) \\ &= P(1-P) - \sigma_p^2 + \frac{(1-2P)^2}{4P(1-P)}(\sigma_{12} - \sigma_p^2) \end{aligned}$$

Acknowledgments

This research was supported by Marsh & McLennan Securities Corporation, where Posner was formerly employed, and by Harvard Business School. The authors would like to thank Ernie Asaff, Chris McGhee, and Roberto Mendoza for helpful discussions. Responsibility for any errors and omissions and for all views lies solely with the authors.

Notes

1. See Froot et al. [1995].
2. Given that the expected excess return on the market portfolio is positive, an efficient portfolio would always allocate a positive amount to it.
3. Major [1999] discusses the complex and somewhat paradoxical interactions between uncertainty and bias.

References

- FROOT, K., MURPHY, B., STERN, A., and USHER, S. [1995]: "The Emerging Asset Class: Insurance Risk," Special Report from Guy Carpenter and Company, Inc.
- LANE, M.N. and MOVCHAN, O.Y. [1999]: "Risk Cubes or Price Risk and Ratings," *The Journal of Risk Finance*, 1, 71–86.
- MAJOR, J. [1999]: "Taking Uncertainty into Account: Bias Issues Arising from Parameter Uncertainty in Risk Models," *CAS Forum*, Casualty Actuarial Society, 153–195.