

INSURANCE CONTRACTS AND SECURITIZATION

Neil A. Doherty
Harris Schlesinger

ABSTRACT

High correlations between risks can increase required insurer capital and/or reduce the availability of insurance. For such insurance lines, securitization is rapidly emerging as an alternative form of risk transfer. The ultimate success of securitization in replacing or complementing traditional insurance and reinsurance products depends on the ability of securitization to facilitate and/or be facilitated by insurance contracts. The authors consider how insured losses might be decomposed into separate components, one of which is a type of "systemic risk" that is highly correlated among insureds. Such a correlated component might conceivably be hedged directly by individuals but is more likely to be hedged by the insurer. The authors examine how insurance contracts may be designed to allow the insured a mechanism to retain all or part of the systemic component. Examples are provided that illustrate this methodology in several types of insurance markets subject to systemic risk.

INTRODUCTION

Insurance markets are undergoing a transformation as new risk management strategies are formulated and new financial instruments are created to supplement and/or replace traditional insurance/reinsurance products. This article takes a normative look at these emerging strategies and markets. In particular, the authors explain how an unbundling of the aggregate insured risk leads to a greater ability to lay off the unwanted parts of the risk. While this type of risk decomposition and risk management is becoming more fashionable to insurers, the authors show how an innovation in current insurance contracting can help facilitate improved risk retention on the part of insureds. In many cases, this innovation leads to results that mimic full access to financial risk management markets directly by the insureds.

An important driver of this innovation is an impaired ability of traditional insurance markets to cope with highly correlated risks. For example, reevaluations of catastrophe

Neil Doherty and Harris Schlesinger work at the University of Pennsylvania and the University of Alabama, respectively. The authors thank seminar participants at the Universities of Konstanz, Munich, and Toulouse as well as at Tulane University for helpful comments on an earlier version of this article. Comments from Henri Loubergé and from three anonymous reviewers were especially helpful.

exposures following events such as Hurricane Andrew and the Northridge earthquake suggest the plausibility of single catastrophes on the order of \$50–100 billion. Yet the total net worth of the entire property-liability insurance industry is only on the order of \$300 billion. While such losses are large enough to overwhelm the insurance industry, the appeal of securitizing this risk becomes apparent when one considers that losses of this magnitude are less than one standard deviation of the daily value traded in U.S. capital markets.¹

A similar demand for insurance substitutes arises in the market for general liability insurance. An implicit correlation arises from the changes in liability rules against which new claims will be resolved through new precedents and/or new legislation. Such rule changes have a common effect on whole groups of policyholders. For example, a legal precedent that extends common-law liability rules will apply to all subsequent suits in the same jurisdiction, unless overruled by a higher court. This judicial instability lies at the heart of the periodic availability crises that have been experienced in liability insurance markets. Yet in today's marketplace, insurers are currently expanding their liability insurance offerings to ever broader classes of insureds, especially in the area of professional liability. Moreover, these lines of insurance are becoming more specialized, such as in liability products for financial services, technology, and telecommunications (see Hofmann, 1996). Obviously, this specialization is partly a function of marketing strategy, but the results in this article imply that such specialization could also be an attempt to take advantage of the stronger correlations present within a more narrowly defined risk class. These narrower risk classes lend themselves more readily to the formation of new securitized products, based on an index of the correlated risk. The high correlation of the individual risk with the index mitigates the problem of basis risk.²

Another example of correlation occurs in the area of property insurance at replacement cost. Unanticipated changes in prices, as well as changes in fixed costs such as permit fees for rebuilding damaged real property, will likely affect all indemnity costs. Again, such changes in prices and costs are common to a group of insureds. Even if the absolute level of correlation for certain types of losses is relatively low, it may be fairly easy to factor out these highly correlated components of the losses and thus design products that improve insurer efficiency. An analogous situation exists in the market for health insurance.

Each case above has a systemic component of the collection of risks being pooled. When correlation exists between loss exposures, the optimal type of risk-sharing contract is one in which the risk can be decomposed into diversifiable and nondiversifiable parts, with the former fully insured and the latter shared with the insurer. This is the essence of mutual insurance.³ One can think of this sharing arrangement as one in which all individuals fully insure their idiosyncratic risk but receive a dividend from their insurer that is scaled to the aggregate experience of the insurance pool. Such

¹ See, for example, Cummins, Doherty, and Lo (2001) and Froot (2001).

² Basis risk need not be all bad. It may, for example, help to alleviate moral hazard problems. However, such issues go beyond the scope of this article. See Doherty (1997) for further discussion.

³ See, for example, Borch (1962), Marshall (1974), Smith and Stultzer (1990), and Dionne and Doherty (1993).

contracts are called *participating policies* and are fairly common for many types of insurance. The innovation in this article is that the authors allow for an endogenous level of participation, whereby the insured can choose a convex mixture of a fixed-premium contract and a participating policy. This *variable participation contract* allows the individual to selectively hedge both the diversifiable and nondiversifiable risk components.

The process of securitization entails the decomposition and repackaging of risk. Securitization can entail both the direct packaging of an individual insurer's loss liabilities for sale in the capital market or designing securitized products based on some economic index. Obviously, securitization is often a substitute for reinsurance in that it allows for insurers to transfer excess risk. However, securitization can offer an ability to carve out pieces of the risk, rather than treating the risk as a whole. In particular, the authors show how this ability to carve out the risk allows for an increase in consumer welfare.

Since securitization relies on the decomposition of risks, a knowledge of the mathematical structure of the loss correlations becomes important. Different mathematical structures may require different markets and different contracts to achieve efficient risk sharing. If losses can be decomposed into two additive components, one independent among insureds and the other highly correlated, then variable participation contracts can be replicated through the combination of a traditional nonparticipating insurance policy and a futures contract to hedge the systemic risk component. In this case, policyholders can, at least in theory, assemble the optimal hedge on their own account.⁴ If the two risk components are multiplicative, replicating the variable participation contract on the individual's own account may still be possible, but this will require insurer intermediation. As this article explains, efficient securitization requires that any systemic risk first be pooled through insurance. The optimal package is then a variable participation contract for the insured together with a futures contract for the insurer covering the aggregate policyholders' dividend risk. Thus, securitization will likely involve the sale of insurance derivatives to primary insurers, and these instruments will compete with, or complement, traditional reinsurance products. Given the relative size of capital markets as compared with that of reinsurance markets, securitization is clearly a necessary component. Indeed, reinsurers will also likely benefit by turning to capital markets themselves.

The focus herein is on how insurance contracts can be redesigned to improve efficiency. Although the authors do not consider the effects of this improved design on supply and demand within the insurance market *per se*, insurance companies that do not keep up with design improvements clearly stand to lose market share. Although the authors examine the effects of securitization in markets for which the systemic part of the decomposed risk (be it additive or multiplicative) is perfectly correlated, the

⁴ Although securitization may also affect markets through reductions in transaction costs (including agency costs), as compared to traditional insurance products, this is not a focus of the current article. Instead, the authors focus on the value of securitization aside from any effects upon transaction costs. Since transaction costs associated with insurance have often run on the order of 30 percent of premiums, the authors do not mean to imply that securitized products cannot have a large effect on cost efficiency. See, for example, Niehaus and Mann (1992) and Froot (2001).

authors also show how imperfect correlations introduce a type of basis risk. The authors conclude by considering several commonly known real-world examples of insurance markets that exhibit correlation, and the authors conjecture how the empirical pattern of securitization might develop.

ADDITIVE SYSTEMIC RISK

This theoretical analysis of optimal hedging uses only the assumptions of risk aversion, defined as an aversion to mean-preserving spreads, and of a preference for higher levels of wealth. In other words, consumers have a preference for second-degree stochastic dominance. The authors allow for risk aversion of both order 1 and order 2, as defined by Segal and Spivak (1990), which has the advantage of allowing for very generalized results. Results within particular frameworks, such as expected-utility theory, smooth nonlinear preference functionals (Machina, 1982), rank-dependent expected utility (Quiggin, 1982), and the dual theory (Yaari, 1987), are all obtainable as special cases.⁵

The securitization of insurance risk is at issue when risk can be decomposed into an idiosyncratic component and a systemic element that can be indexed. The authors consider two forms of this decomposition: additive and multiplicative. These forms carry different implications for the design of insurance derivatives, and they are useful in accessing well-known results on optimal hedging behavior. In this section, the authors consider the case in which the risk decomposition is additive. The authors examine the multiplicative case in the following section.

Let $(\Omega, \mathcal{F}, \mu)$ be a probability space, and consider the measurable functions $L: \Omega \rightarrow [t, T]$ and $\varepsilon: \Omega \rightarrow [-s, s]$, $t, T, s \in \mathcal{R}$. Let L and ε denote the random variables so defined. Consider an individual with initial wealth $W > 0$ that is subject to a loss of size $L + \varepsilon$; assume here that the scalars t, T , and s are chosen such that $0 \leq t - s \leq T + s \leq W$.⁶

An infinitely large population of consumers with identical loss distributions exists. Refer to this population as the risk pool, and assume that the idiosyncratic random loss components L_i are mutually independent from one another, so that L_i for the i th individual is independent from that of the j th individual, L_j . However, the second components of the loss decomposition are equal for all individuals, i.e., perfectly positively correlated, $\varepsilon_i = \varepsilon_j$. Assume that $E(\varepsilon) = 0$, where E denotes the expectation operator, and that ε and L are independent of one another for all individuals. Since ε is identical for all individuals, suppress its subscript. For example, if $\varepsilon = 100$, then losses are 100 higher for everyone. Although this additive case is less realistic than the multiplicative case that follows, it will be useful for establishing later results.

The authors will consider various scenarios for interpreting ε . A word of caution is in order here. The case in which $\varepsilon = 0$ need not be "expected" in any sense except as a long-run average. Indeed, the distribution of ε may be highly skewed. For example, in modeling catastrophes, one might think of ε as being slightly negative in almost every

⁵ See Machina (1995), Karni (1995), and Schlesinger (1997) for summaries of insurance results in these models.

⁶ The authors make this last assumption to avoid complications of modeling limited liability. The authors also wish to discourage thinking of ε as a loss amount itself. It is simply an adjustment to the long-run average loss that is experienced within a given year.

year. It might then be positive and large only on rare occasions. Hence, $\varepsilon = 0$ in no way should be interpreted as a “typical year,” or as a “forecast” for the year. One can view the catastrophe example as follows: Over the years, the average loss per individual insured is $E(L + \varepsilon) = EL$. However, the average loss is not stable over time, so that in many years it is less than EL , whereas in some years, and most notably catastrophe years, the average annual loss exceeds EL . For cases in which ε may represent an unexpected change in price levels, it may be more natural to assume that ε is slightly positive or negative with equal probability. Thus, this modeling of the “ ε risk” is meant to be rather general.

Optimal Risk Sharing With Consumer Access to Securitization

Suppose first that separate markets exist for hedging the risks L and ε . Assuming competitive markets, postulate an insurance market for L with actuarially fair pricing. In such a market, full insurance is purchased on L . This holds regardless of the treatment used for ε and regardless of whether risk aversion is of first or second order.⁷ To hedge the systemic risk component ε , postulate the existence of a futures market. Assuming a clientele of only speculators and individuals endowed with ε risk, such a futures market will exhibit normal backwardation, due to the natural hedging demand by insureds. Model this backwardation as replacing the random loss component ε with a fixed certain loss of $\gamma > 0$.⁸ What the authors envision here is a world in which the ε risk might not be fully diversifiable in global markets but can be reduced enough to allow insurability of the ε component. Cummins and Weiss (2000) label these two cases as “globally diversifiable” and “globally insurable,” respectively. Froot (2001) examines several reasons why γ might be lower in the capital markets as opposed to in traditional reinsurance markets.

Futures contracts are assumed to be fully divisible, and individuals choose the fraction of systemic risk ε that they wish to hedge on their own account. Because this article is using a static model and assuming perfect correlation of the ε risk, assume away problems associated with basis risk due to the timing of futures contracts or to the imperfect nature of the hedge instrument. Final wealth Y is thus given as

$$Y = W - EL - [b\gamma + (1 - b)\varepsilon], \quad (1)$$

where b denotes the fraction of ε hedged in the futures market, i.e., the hedge ratio. Since $\gamma > 0$ and $E(\varepsilon) = 0$, if preferences satisfy second-order risk aversion, then the

⁷ Doherty and Schlesinger (1983) prove this result for a model using differentiable expected utility. Since full coverage for any treatment of ε is optimal for all risk averters defined via expected utility, it follows from Zilcha and Chew (1990, Theorem 1) that such behavior is optimal for the broader class of risk-averse preferences examined here. The result still holds if the utility function is not differentiable everywhere, which follows from Segal and Spivak (1990).

⁸ If many pools of insureds exist, each with an ε that is independent of other groups', then γ possibly equals zero. The authors assume that a larger market does not exist to “pass off” the ε risk, so that $\gamma > 0$. The authors also do not consider a general equilibrium model, in which the existence of the types of contracts proposed in this article have an effect on market prices, including a type of feedback effect upon γ itself.

optimal hedge ratio equals $b^* < 1$. If risk aversion is of order 1, then $b^* < 1$, with the possibility of complete hedging, $b^* = 1$.⁹

Variable Participation Contracts

Since real-world futures markets are not likely to exist for hedging only a part of an individual's loss, the article examines an alternative contract available through the insurance market. In particular, the individual may buy insurance via a participating insurance contract. The innovation that is presented here is to make the degree of participation a choice variable of the individual. To this end, assume that insurers are willing to offer insurance with zero participation, with a market-determined premium loading factor of $\lambda > 0$. Since the total individual loss is $L + \varepsilon$ with ε identical (i.e., perfectly correlated) for all individuals, the market charges the risk premium λ for this ε component of the total loss. Assume here that, although the reinsurance market and capital market can absorb the ε risk, it cannot be fully diversified away. Since the L are all independent and identically distributed (i.i.d.), there is no additional amount of premium loading required due to the L ; rather, the L risks are fully diversifiable due to the assumed infinite number of independent risks. For simplicity, no other transaction costs exist in this model.

Consider now a fully participating policy with a premium equal to the *ex post* average indemnity paid by the insurer. Such a premium is (essentially) $\alpha(EL + \varepsilon)$, where the EL term is (essentially) guaranteed by the law of large numbers.¹⁰ Since the individual bears all of the ε risk in this case, the competitive-market insurance premium loading is zero. Rather than imposing zero or full participation, allow the individual to choose the degree of participation in the insurance market by setting the total premium P as follows:

$$P = \alpha[\beta(1 + \lambda)EL + (1 - \beta)(EL + \varepsilon)] = \alpha\{EL + [\beta\lambda EL + (1 - \beta)\varepsilon]\}, \quad (2)$$

where α denotes the proportion of loss indemnified by the insurer and where $\beta \in [0, 1]$ is a choice variable of the individual denoting the degree of participation, with $\beta = 1$ denoting a fixed premium and $\beta = 0$ noting full participation. The insurance policy

⁹ These conclusions follow easily along the lines suggested in note 7. Note that in Equation (1) only one source of uncertainty exists— ε . Note also that risk aversion is of order 1 if

$$\lim_{t \rightarrow 0^+} \pi'(tx) = 0,$$

where the limit is taken over positive values of t , x is a zero-mean random variable, $\pi(tx)$ is the risk premium such that $-\pi(tx) \sim tx$, and $\pi'(tx)$ denotes $\partial\pi(tx)/\partial t$ for $t > 0$. Risk aversion is of order 2 if $\pi'(0x) = 0$ but $\pi''(0x) \neq 0$. See Segal and Spivak (1990).

¹⁰ More realistically, one would need to be concerned with the timing of premium collections and indemnity payouts. However, the authors abstract from these nuances in the static model. The total premium as given above is random *ex ante*. The premium actually paid *ex post* is dependent on the realized value of ε . One can think of the individual paying an up-front premium of αEL . The individual is then assessed an extra premium of $\alpha\varepsilon$. In the case where $\varepsilon < 0$, this "assessment" is paid to the individual as a dividend. A negative modal value of ε would thus correspond to the payment of a dividend in most years under participating policies.

with a premium defined by Equation (2) is called a *variable participation contract*.¹¹ Final wealth Y is given by

$$Y = W - P - (1 - \alpha)(L + \varepsilon) = [W - \alpha EL - (1 - \alpha)L] - \varepsilon + \alpha\beta(\varepsilon - \lambda EL). \quad (3)$$

It is interesting to note in the decomposition in Equation (3) that, although the fixed insurance premium depends upon the loading factor λ , the λ only attaches itself to the ε risk in the decomposition. In other words, it is “as if” the individual is being offered fair insurance against the L risk, with a price for eliminating ε risk. This point is revisited in Equation (4) below.

Suppose that $\alpha \neq 0$, and suppose for the moment that one does not require $\beta \in [0, 1]$. Then note that the value of α in the last term in Equation (3) is irrelevant, since it can be “undone” by a choice of β . In particular, letting $\delta = \alpha\beta$, the choice variables in Equation (3) are effectively α and δ , so long as $\alpha \neq 0$.

For any fixed value of δ , the terms $-\varepsilon + \delta(\varepsilon - \lambda EL)$ are a random background risk, whereas the first three terms on the right-hand side of Equation (3), in brackets, represent the standard insurance choice problem with an actuarial fair premium. Thus, for any fixed value of δ , the optimal insurance level is $\alpha^* = 1$, i.e., full coverage.¹²

Now since $\alpha^* = 1$, the assumption of $\alpha \neq 0$ is redundant and the relaxation of the condition $\beta \in [0, 1]$ is irrelevant: The optimal β equals the optimal value of δ . Indeed, write Equation (3), under the assumption of full insurance coverage for any value of δ (i.e., any value of β), as

$$Y = W - EL - [\beta\lambda EL + (1 - \beta)\varepsilon]. \quad (4)$$

Compare Equation (4) with Equation (1) and assume that the market risk premium for ε risk would be the same, whether in a futures market or in an insurance market; i.e., assume that $\lambda EL = \gamma$. By using a variable participation contract, the mutual insurance market provides the exact same set of alternatives and same optimal solution (with $\alpha^* = 1$ and $\beta^* = b^*$) as obtains in two separate markets. Of course, if $\lambda EL \neq \gamma$, then the cheaper alternative is likely to be the one that prevails in the marketplace.

One might approximate the variable participation strategy by buying a fixed-premium contract and simply buying shares of the insurer’s stock, if it is a stock insurance company. However, stock prices include a broader view of company profitability and, in particular, a longer-term perspective. Thus, this strategy is likely to be dominated by one that develops a participation measure based on only the current aggregate L and ε development.

Market Structure With Additive Risk

The optimal variable participation contract entails full coverage, $\alpha^* = 1$, and $\beta = \beta^*$, which implies that some portion, β^* , of the systemic risk is transferred to external

¹¹ Of course, the individual may be able to self-construct an equivalent contract via the purchase of two separate contracts, one fixed-premium contract with coverage level $\beta\alpha$ and one fully participating contract with coverage level $1 - \beta\alpha$.

¹² This follows as in note 7.

investors. Two common mechanisms within the insurance industry are available for achieving this division of systemic risk. First, the insurer can be a stock insurer that issues a variable participation policy with portions $1 - \beta^*$ and β^* of the systemic risk allocated respectively to policyholders and to shareholders. Second, the insurer can be a mutual company that reinsures a portion β^* of its portfolio risk with independent stock reinsurers, thus passing the hedged systemic risk to the reinsurer's shareholders. Using securitization to handle the systemic risk would not affect the fundamental division of risk between policyholders and external investors (absent changes in transaction costs). Rather, securitization would achieve this division through a different set of contractual arrangements.

Although replication of the variable participation contract in separate markets also requires some insurance, securitization of the ε risk can be handled independently from the insurance contract. Whether the optimal contract is achieved via nonparticipating insurance with insureds hedging the ε risk directly in the futures market or via variable participation contracts with insurer-based securitization is likely to depend upon which method is more cost-effective once transaction costs are introduced. At least preliminary casual empirical evidence seems to support the latter. Moreover, a policy paying an indemnity based on L alone, rather than on $L + \varepsilon$, cannot adjust its claims until ε is learned *ex post*. Although $L + \varepsilon$ is observed immediately following a loss, one does not know the decomposition into L and ε until the end of the year, when the insurer has the data to observe $EL + \varepsilon$ for the year. A consumer is likely to prefer a current indemnity payment for the observed loss of $L + \varepsilon$, as presently exists, followed by an adjustment for ε (via a dividend or an assessment) at a later date. In theory, participation may require the insured to make additional large payments (assessments) after the policy period has expired, which are generally unacceptable to insurance regulators. More commonly, an insurer may "build in" most of the possible assessment as part of the insurance premium. This amount is then later refunded as a policy dividend (see note 9).

MULTIPLICATIVE SYSTEMIC RISK

Now assume that the insurable loss is of the form $(1 + \varepsilon)L$, where $E\varepsilon = 0$ and ε is independent of L_i for all i ; L_i are independent and identically distributed, and ε is identical (perfectly correlated) across all insureds.¹³ Thus, for example, if $\varepsilon = 0.05$, then everyone's realized loss is 5 percent higher than the long-run average. Final wealth Y in this case is given as

$$Y = W - (1 + \varepsilon)L = W - L - \varepsilon L. \quad (5)$$

Variable Participation Contracts

Maintaining the notation where $\beta = 1$ denotes a fixed premium and $\beta = 0$ denotes a fully participating policy, write the premium for coverage level α as

$$P = \alpha[\beta(1 + \lambda)EL + (1 - \beta)(EL)(1 + \varepsilon)] = \alpha EL\{1 + [\beta\lambda + (1 - \beta)\varepsilon]\}. \quad (6)$$

¹³ A sketch of the model in this section appears in Schlesinger (1999), who used it to examine losses from natural catastrophes.

Thus the consumer's wealth after the purchase of insurance is

$$Y = W - \alpha EL - (1 - \alpha)(1 + \varepsilon)L - \alpha EL[\beta\lambda + (1 - \beta)\varepsilon] \quad (7)$$

or equivalently

$$Y = (1 + \varepsilon)[W - \alpha EL - (1 - \alpha)L] - \varepsilon W - \alpha\beta EL[\lambda - \varepsilon]. \quad (8)$$

As in the previous section, suppose for now that $\alpha \neq 0$. Then note that the value of α in the last term in Equation (8) is irrelevant, since it can be "undone" by a choice of β . Once again letting $\delta = \alpha\beta$, the choice variables in Equation (8) are effectively α and δ , so long as $\alpha \neq 0$.

For any fixed value of δ , the last two terms in Equation (8) are a background risk, although the mean of this background risk is not zero. Note also that this background risk is linear in ε . The first term in Equation (8) has the same expected value for all choices of α . Let $H(\alpha)$ denote the random variable $W - \alpha EL - (1 - \alpha)L$, and let $F(\varepsilon)$ denote the last two terms in Equation (8). Thus Equation (8) can be written as $(1 + \varepsilon)H(\alpha) + F(\varepsilon)$. Because ε is statistically independent of L , it is statistically independent of H as well. Since $H(\alpha)$ has the same mean for all α , it follows from standard stochastic-dominance arguments that choosing $\alpha = 1$ will second-order-dominate every other choice of α for final wealth.¹⁴ Consequently, $\alpha^* = 1$ is optimal for any risk averter for a fixed level of δ . Now, since $\alpha^* = 1$, the assumption here that $\alpha \neq 0$ is redundant. From Equation (7), using $\alpha^* = 1$, $\delta^* = \beta^* < 1$, with $\beta^* = 1$ only in the case where preferences satisfy first-order risk aversion.

Optimal Risk Sharing With Consumer Access to Securitization

First establish that the individual cannot replicate the variable participation contract simply by purchase of a nonparticipating policy and a separate futures contract on ε . Since the individual's wealth prospect is $W - L - \varepsilon L$, a simple policy fully covering $L\varepsilon$ (i.e., replacing L with EL) must be supplemented with a futures hedge written not simply on ε but on $L\varepsilon$. In other words, the task facing the individual is to hedge a random number of units of the ε risk. One is left with a random optimal hedge ratio; i.e., the hedge ratio β ideally would need to be scaled according to the idiosyncratic and random L . In theory, financial markets could write contracts based jointly on the realizations of ε and L , but this would involve monitoring each individual's losses by the financial market, which is likely to be inefficient. However, it is possible for the insurer, who already tracks individual loss data, to intermediate here.

A market could relatively easily emerge in which contracts written on the ε risk are traded between individual investors and insurance companies, with the insurers also offering participating contracts. Note that each individual has a stochastically identical multiplicative term $L_i\varepsilon$, and that the L_i are all independent from one another and

¹⁴ Since all choices of α leave the mean of $H(\alpha)$ and thus of Y in Equation (8) unchanged, second-order stochastic dominance follows from Rothschild and Stiglitz (1970, Theorem 2), since one can write

$$(1 + \varepsilon)H(\alpha) + F(\varepsilon) = (1 + \varepsilon)H(1) + F(\varepsilon) + [H(\alpha) - H(1)](1 + \varepsilon),$$

where $E\{[(H(\alpha) - H(1))(1 + \varepsilon)] \mid (1 + \varepsilon)H(1) + F(\varepsilon)\} = E\{[(H(\alpha) - H(1))(1 + \varepsilon)] \mid \varepsilon\} = 0 \forall \varepsilon$.

from ε . It follows that there should be no risk-bearing cost in a competitive market (absent any transaction costs) for pooling the L risk.¹⁵ To see this, define $\mathcal{E} \equiv \sum L_i/n$ and note that $E(\mathcal{E}) = EL_i$. Now $VAR(\mathcal{E}) = (VARL_i)/n \rightarrow 0$ as $n \rightarrow \infty$, since the L_i are independent and identically distributed. It follows that, in the limit,

$$VAR(\mathcal{E} + \varepsilon\mathcal{E}) = (EL)^2VAR(\varepsilon).$$

Thus, for the insurer, \mathcal{E} can be approximated as a constant, and one can write $\mathcal{E} + \varepsilon\mathcal{E}$ in the approximate form $E(L) + \varepsilon E(L)$. In other words, each individual could pool his or her own εL_i and assume $\varepsilon E(L)$. That is, the individual swaps a random level of ε risk for a fixed level of ε risk. A competitive insurance market, in which participating policies are traded, could organize such pooling, for example. Under a pure mutual, with all idiosyncratic risk insured ($\alpha = 1$) and all systemic risk assumed by policyholders as dividends ($\beta = 0$), the individual's wealth would be

$$Y = W - EL - \varepsilon WL. \quad (9)$$

Letting ε' denote εEL and noting that ε' satisfies all of the requisite properties of ε for the case of additive risk components, it follows that the multiplicative risk component case is identical to the additive case, with ε' replacing ε . Thus, assuming a competitive insurance market and a futures market for ε that exhibits normal backwardation, the individual buys full insurance, $\alpha^* = 1$, together with an optimal futures market hedge, b^* , the optimal hedge ratio b for Equation (1), where here $\gamma = \lambda EL$ and $\varepsilon' = \varepsilon EL$ replaces ε . In other words, Equation (1) becomes

$$Y = W - EL - EL[b\lambda + (1 - b)\varepsilon]. \quad (10)$$

Note that Equation (10) is equivalent to Equation (7), with $\alpha^* = 1$. Thus, $b^* = \beta^*$ and once again the result is that two markets are equivalent to one insurance market with variable participating contracts. However, now the individual cannot readily access the futures market without help from the insurer (or some other financial intermediary) in pooling the random amount of ε risk, $L_i\varepsilon$.

Market Structure With Multiplicative Risk

To achieve the optimal contract ($\alpha^* = 1$, $\beta^* = b^*$) requires that the systemic risk of insureds be pooled, $\varepsilon' = \varepsilon EL$. This leads to two potential optimal contracting patterns. First, individuals can form a pure mutual insurance company, in which all systemic risk is passed back to policyholders in the form of dividends. Thus, each policyholder's wealth is $Y = W - EL - \varepsilon EL$, with the last term, εEL , being the dividend risk. Policyholders can then hedge a portion b^* of the dividend risk by trading on their own individual accounts. Second, the insurer can issue participating policies and can purchase a futures contract in the ratio β^* and pass the unhedged portion of the systemic risk ($1 - \beta^*$) back to the policyholder in the form of a dividend. With either structure, the individual's wealth prospect is as shown in Equation (10). However, the essential

¹⁵ Of course, as n gets larger, so does the variance of total losses. This would lead to a higher bankruptcy risk absent any increase in insurer capital. Assume $n \rightarrow \infty$ to circumvent this issue.

features of both contract structures are that the optimal contract can be assembled only if the systemic risk is pooled and insurers issue participating policies. Insurer intermediation is thus an essential component in securitizing the ε risk.

SECURITIZATION IN INSURANCE MARKETS: SOME EXAMPLES

Securitization is relatively new, and its full impact upon the business of insurance has yet to be fully determined. The authors believe current uses of securitization are still in the formative stages and have not yet fully self-developed within the marketplace. In this section, the three examples of insurance markets with correlated risks that were mentioned in the introduction are considered: (i) property insurance at replacement cost, (ii) liability insurance, and (iii) insurance for natural catastrophes. To the best of the authors' knowledge, only (iii) has been examined much at all in the literature. Each example is modeled using "multiplicative risk components." In each case, the authors examine how securitization and insurance contracting might work in an optimal setting.

Property Insurance at Replacement Cost

A futures market for any ε risk requires that it be clearly indexed. A very simple case is that in which policyholders are exposed to independent and identically distributed losses when measured in constant dollars, but in which a random (unexpected) inflation rate impacts all claims. The impact of inflation on each policyholder will depend on the size of each loss. This is clearly a case in which the systemic risk and the non-systemic risk are multiplicatively related. A random draw is taken from the inflation index, ε , and each constant dollar loss, L_i , is multiplied by the same revealed $(1 + \varepsilon)$. For example, if the realized value of ε is 0.05, then all loss claims cost 5 percent more than expected.

One way for insurers to handle the ε risk is to issue replacement-cost policies and to hedge the inflation risk themselves through options and/or futures markets. However, this approach may be costly to the insureds, who would prefer bearing some of the inflation risk themselves. At the other extreme, the policyholder could purchase an insurance contract that offered indemnification in constant dollars, such as one with an *ex ante* listing of insured values (a so-called valued policy, in the insurance world). However, the individual then would not have a fixed amount of ε risk to hedge in the futures market. If the insurer issues a participating replacement-cost policy, and insureds choose their own level of participation, $1 - \beta^*$, the insurer then can hedge the remaining portion β^* of the ε risk with some type of inflation-index derivative. Since the amounts of ε risk assumed by the insurer on individual policies are independent and are identically distributed except for a scaling factor due to differences in individual β_i , the insurer has (essentially) a fixed amount of ε risk to hedge.

Of course, in reality any type of inflation index will not affect all insured losses in the exactly same manner (i.e., the ε_i correlations will not be perfect), so that derivative securities on such an index will not be a perfect hedge. In the case of futures markets, the hedging strategy faces an added *basis risk*. Although such basis risk cannot be eliminated, it is possible that more restricted indexes will induce a more perfect correlation. For example, an inflation index on automobile parts would be a more efficient hedging instrument for automobile collision coverages than would a general consumer price index (CPI). Given their potential hedging purposes, one might

therefore expect to see derivative products arise that are based on more specialized price indexes. Unfortunately, options and futures written directly on price index have not yet seemed to have taken hold in real-world markets. A highly touted futures contract on the CPI, introduced in 1985, quickly failed, due to a lack of trading volume. However, other recent products, such as the U.S. Treasury's issue of inflation-indexed government bonds, may either prove to be viable hedging instruments on their own or at least increase the potential demand for inflation-index derivatives to the point where such derivatives are commercially viable.

Liability Insurance

The case of implicit correlation caused by changes in liability rules is more complex. Obviously, changes in liability rules can affect both the likelihood of verdicts and the average size of jury awards. Only the latter effect is considered here, although it is recognized that the former effect is likely to be just as important.¹⁶ Imagine that an index is taken of the changes in average liability awards (IALA), which corresponds to ε . If one assumes that the events that give rise to liability awards are stable (so that the number of awards is unaffected), one can discount individual awards by the IALA index to derive a "stable liability regime" (SLR) award. Accordingly, the SLR claims, i.e., the L_i , will be independent and identically distributed (or, more realistically, at least independent), although each L_i will not be observable. Only the awarded amount $(1 + \varepsilon)L_i$ is observed. Of course, if the ε are perfectly correlated between individuals, one can "back out" the value of L_i at the end of the year, when all of the award data are known. However, unlike a price index, the IALA index would only be determinable from the awards themselves (i.e., only from the indemnifiable claims), whereas, for example, a price index on auto parts would be able to use a broader set of market data unrelated to the insurable events.

Optimal contracting is achieved if the insurer covers the full loss, $(1 + \varepsilon)L_i$, but allows for a participating policy whereby the insured leaves the insurer with $1 - \beta^*$ of the ε risk. The participating dividend (which is negative, in the form of an assessment) is calculated after the end of the policy period, when the market has enough data to determine the IALA index. The insurer's share of the ε risk is once again (essentially) nonrandom and deterministic. Indeed, even though one cannot observe each L_i directly at the time the award is made, the insurer is assumed to know the value of EL_i *ex ante*. Therefore, the insurer can turn to the marketplace to hedge its aggregate ε risk.

Of course, the IALA is nothing more than a "wish-list item" for insurers at the moment, as is any market for derivatives on this index, should this index become a reality. However, another potential form of securitization that seems to be gaining some interest in the insurance world is the packaging of standardized risk units for direct sale to the public. A packaging of homogeneous risks, of a size sufficient to eliminate the average idiosyncratic risk component, would leave the purchaser(s) with only the ε risk. Since this risk would not be perfectly correlated with other market indexes (otherwise these

¹⁶ See Doherty (1991) for a more complete discussion of the many complicated effects involved. Also, see item "Insurance for Natural Catastrophes" ahead for a discussion on how one might model correlated likelihoods.

indexes could be used themselves to construct a hedge device), it should obtain a price in the marketplace, so long as the IALA, or some similar index, is verifiable.

Once again, the assumption that the ε_i are perfectly correlated is likely to be too strong to fit reality. If the correlation is high enough, this will only mean that the “packaged” liability losses, which contain (essentially) only ε risk in theory, will also contain some type of noise in the aggregate of the ε components, which can once again be treated as a type of basis risk. If the correlations are not high enough, the IALA will be too uninformative, and market-hedging strategies will not be effective. One way to make them more effective would be to define narrower classes of liability risks, ones within which the ε risks are highly correlated. Indeed, insurers currently seem to be in the process of developing liability policies within more and more specialized areas (see Hofmann, 1996). The reason is only partly demand driven, as existing products could easily be marketed to larger classes of insureds. But rather than just attracting new customers for existing products, insurers are continually developing newer product classes, especially in the area of professional liability. Modern databases can be readily fine-tuned to keep track of losses, making specialized versions of an IALA index feasible. Since one assumes here that the ε risks *across* the differing classes of liability risks are not perfectly correlated, insurers can diversify partly by taking separate positions in the various types of liability classes. With more narrowly defined classes, reinsurers become a potential source for diversifying the various types of ε risk, in addition to securitization for direct sale to the public.

Insurance for Natural Catastrophes

As a final example, consider homeowner’s insurance for people living in an area exposed to natural catastrophes, such as earthquakes in California or hurricanes in Florida. Indexes of catastrophe losses within such regions are now available. Catastrophe risk, such as hurricane or earthquake risk, commonly exposes those living within a fairly confined area to simultaneous losses. The standard model of variable participation contracts herein seems to fit this scenario fairly well. However, the catastrophic effects are modeled a bit differently here to exhibit how this model can be extended to handle frequency risk correlations. Although severities of losses also are likely to be affected, the authors focus only on the likelihood and assume that individuals possess loss severity distributions L_i that are independent and identically distributed and are conditional on suffering a catastrophic loss.¹⁷

In reality, policies typically cover a broad range of losses, so that the losses are separated off due to the potentially catastrophic peril. For example, although hurricane damage is covered under typical homeowner’s insurance, one can separate out the windstorm peril for the case of hurricanes, as was recently done for blocks of policies in the State of Florida. Let L_i denote the severity distribution conditional upon hurricane damage.

Correlation enters this model via the probability of damage. Let p_0 denote the long-term probability that a particular insured suffers a loss due to a catastrophic event.

¹⁷ In a related article, Schlesinger (1999) provided some details of modeling loss severity correlations. He also provided a few numerical examples for interested readers. An extension by Loubergé and Schlesinger (2000) examined the interaction of severity correlation and frequency correlation. The authors provide the frequency risk analysis herein as a building block for these models as well as for future research.

Assume that the group of insureds is defined such that the current-year probability of damage is identical for all insureds. One thus can view p_0 as the relative frequency of losses in an average year. Define δ_i to be a Bernoulli random variable taking on the value 1 with probability $\xi \equiv p_0(1 + \varepsilon)$ and taking on the value zero with probability $1 - \xi$. Maintain all previous assumptions about ε except that one now requires the support of ε to be restricted such that $0 < \xi < 1$. Assume that ξ is identical for all insureds.

One can view p_0 as representing the expected proportion of homeowners who suffer damage from the "catastrophic peril" in a typical year, and view $\xi \equiv p_0(1 + \varepsilon)$ as the random proportion of homeowners suffering damage in the current year. Note that ξ is the true probability that a randomly chosen homeowner suffered a loss in a given year, which is only known as a relative frequency *ex post*. This is *not* to be confused with Bayesian updating. The *ex ante* probability of loss is still viewed as p_0 . So, for example, suppose $p_0 = 0.10$ and that this year $\varepsilon = 0.50$. Then the insurer typically would see 10 percent of insured homes suffer a loss from the catastrophic peril. This is a long-term average frequency that includes both catastrophic and noncatastrophic years. However, the current year would see 50 percent more loss claims; i.e., a total of 15 percent of the insured properties would experience losses.

The δ_i are assumed to be independent and identically distributed within the insured group. Under these assumptions, the catastrophe risk held by each individual insured is given by $\delta_i L_i$. Although the individual $\delta_i L_i$ are all independent, the probability ξ is perfectly correlated among all individuals in the insured group. From the insurer's perspective, the *ex ante* risk per policyowner under full nonparticipating insurance is (essentially) $\xi EL \equiv p_0 EL(1 + \varepsilon)$. Since p_0 is a constant, this structure is identical to the multiplicative risk structure presented earlier in this article. Allowing for policy participation, the insurance premium is as given in Equation (6), with the exception that $p_0 EL$ replaces EL . The insured's wealth Y is thus

$$Y = W - \alpha p_0 EL - (1 - \alpha) \delta_i L_i - \alpha p_0 EL [\beta \lambda + (1 - \beta) \varepsilon]. \quad (11)$$

Note that Equation (11) is identical to Equation (7), with the random loss L replaced by $\delta_i L_i$, which in itself yields $E(\delta_i L_i) = p_0 EL$ in place of EL . As a result, the optimal variable participation contract is one for which $\alpha^* = 1$ and $\beta^* = b^*$, where b^* is the optimal hedge on $p_0 EL$ "units" of ε risk for an individual's personal account, if there were some type of market to hedge the ε risk directly. However, the problem here is trickier, since the ε risk is not monetary but probabilistic.

For the insurer, barring transaction costs, it really does not matter financially whether an ε of 0.50 represents the same number of losses, each 50 percent higher in magnitude, or an ε of 0.50 represents losses of the same average magnitude, but 50 percent more of them. At the individual level, rather than now having everyone with a 50 percent higher loss, 50 percent more of the individuals experience a loss. This makes it practically impossible to "undo" the ε effect at the individual level. In the case where the frequency is correlated, one can only undo the ε effect by randomly choosing 50 percent of households experiencing a loss and not paying them any indemnity. Although it might be possible to pay each household an indemnity that is 50 percent lower, this approach could cause problems for cases when $\varepsilon < 0$, when indemnities might need to exceed loss values. These types of problems need not be dealt with,

however, since the individual may simply purchase a variable participation contract and let the insurer be the one who turns to the use of securitized products.¹⁸

Once again, the use of such securitization depends upon ε being well defined. Although long-term probability projections for catastrophes still have not been perfected, much progress has been made. Certainly, the realized value of ε is calculable *ex post*, so long as one agrees on a projected value p_0 . The choices available to insurers for hedging the ε risk are growing rapidly. For example, an insurer can trade its ε risk for one type of catastrophe with the ε risk from a different insurer on a different type of catastrophe exposure. Such trades currently exist in the form of "CAT swaps." For example, an insurer in Florida may swap some of its hurricane risk in exchange for taking on some earthquake risk from a California insurer. Alternatively, an insurer can turn to the CBOT for options and futures on catastrophic loss indexes. Although trading volume to date at the CBOT in the area of insurance futures and options has been relatively light, insurers are just beginning to understand and recognize the potential hedging possibilities. Another possible course of action is for the insurer to issue CAT bonds, with a trigger set for the relevant catastrophe event. Of course, for hedging to be effective, the ε risk must be well defined and must be highly correlated within the insured group.¹⁹

One must be careful at this point to warn the reader that modeling the ε distribution accurately, along with an ability to observe the realized value of ε *ex post*, does not relieve us of the multitude of problems that may be associated with catastrophic risks. For example, in cases where the distribution of ε is heavily skewed, one may run into problems associated with extreme values.²⁰ This decomposition neither simplifies nor complicates the problems associated with estimating the overall loss distribution. This would be critical, for example, if one wished to develop a method for ascribing the appropriate insurance premia or for performing value-at-risk analyses. Indeed, the ε risk in this model needs to be determined by decomposing the estimated loss distribution.

The goal in this article is not to estimate the tail of the loss distribution but to point out how risk decomposition can aid in the design of better-performing insurance contracts. Such long tails in this model could lead to either extremely high values for the price of risk, γ , or even to the nonexistence of markets for the ε risk. Thus, the contracts

¹⁸ See Schlesinger (1999) for further discussion of the details of this case.

¹⁹ Recently, new catastrophe indexes from Property Claims Services (PCS) have been added to those already in use (which uses data from the Insurance Services Office) to extend the product line offered by the CBOT. Obviously, the CBOT agrees with this assessment, that securitization is still developing in the marketplace. Moreover, other exchanges are coming into existence. For example, the newly opened Bermuda Commodities Exchange (BCOE) trades options contracts on certain "atmospheric perils," such as tornadoes, hurricanes, and hailstorms, for specified regions within the United States. The contracts are based on a new index developed by a subsidiary of Guy Carpenter, and bidding takes place over the Internet.

²⁰ See Embrechts, Klüppelberg, and Mikosch (1997) for an excellent analysis of the multitude of problems involved in modeling long tails with extreme values. Another limitation in this model for dealing with extreme values is that the authors use the simplicity of proportional coverages, rather than the types of stop-loss contracts typically associated with such skewed distributions.

proposed in this model are best viewed as another tool for dealing with catastrophic risks, and not as any type of overall solution.

CONCLUSION

Since securitization is only now beginning to emerge as a viable risk-sharing technique for insurable losses, this model should be viewed as normative. Securitization may evolve through alternative means. If individuals must insure the noncorrelated (idiosyncratic) component of their loss exposure through a nonparticipating insurance contract, then securitized products may develop to allow the individual to directly hedge the correlated (systemic) risk component. Although this is theoretically simple to do in the case of additive risk components, markets for both of the above risk components have not yet shown any indication of evolving in real-world markets, as far as the authors know. For the case of multiplicative risk components, such a hedging strategy is unlikely, due to the random amount of systemic risk attributable to each individual. More likely, and what is seen developing to date, is that securitized products are used directly by insurers. These insurers can then offer "variable participation policies" to individual insureds and pass off any desired amount of the systemic risk, which is not assumed by the policyholders, in the capital market. One sees signs here of a burgeoning market for securitized products.

The efficacy of securitization depends crucially on the separability of the systemic and idiosyncratic risk components and on strong correlations within the systemic component. Recent designs of insurance products and insurance indexes to fit narrower bands of loss exposures with seemingly highly correlated systemic components indicate that the market is indeed developing in this way. This pattern is already appearing with insurers purchasing a growing number of different catastrophe options and futures at the CBOT. Moreover, the systemic component might itself be directly traded, such as is the case with CAT swaps. It also can be hedged via contingencies in debt and/or equity instruments. For instance, the insurer can issue CAT bonds, in which the hedge takes the form of a "forgiveness option" built into a debt issue made by the insurer, or it can issue "cataputs," which are essentially put options in which the insurer can sell new equity at a predetermined price.

Although the authors have not examined the issue in this article, the emergence of a market for securitized products must also obviously relate to transaction costs and/or contracting costs.²¹ Perhaps the most apparent saving lies in a reduction in the costs of financial distress to insurers. This is dramatically seen in catastrophe risk. A \$50 billion hurricane or earthquake loss in the United States represents about one quarter of the net worth of the entire domestic property-liability insurance industry. Such magnitudes of losses are not unfathomable. For instance, some estimates put the damage done during the 1995 Kobe earthquake in Japan at \$100 billion.²² Moreover, the events of September 11, 2001, have shown that man-made catastrophes can also be quite

²¹ This is the main focus of a recent article by Froot (2001).

²² Although the damages were high, the amount of damages insured was only approximately \$3 billion, according to data from Munich Reinsurance (see <http://www.munichre.de>). Munich Reinsurance also predicts that the maximum damage from a single California earthquake could be as high as \$150 billion.

severe, with losses of over \$100 billion and insurance losses and projected insurance payouts somewhere between \$35 and \$70 billion.²³

The prospect of such high losses creates significant costs in terms of both incentive conflicts between stakeholders and potential bankruptcy costs. Yet such losses are small compared with the \$20 trillion U.S. capital market. Moreover, since such losses exhibit close to zero correlation with most financial market indexes, the required rate of return for capital market investors should eventually move closer to the risk-free rate as markets develop and become more familiar to insurers and to investors. As insurers have increasingly more tools to decompose and hedge the components of their risks, some simple changes in primary insurance contracts, such as those proposed herein, will allow for more flexible risk management on the part of insureds. It will be interesting to see exactly how the market develops in the years to come. Hopefully, many of these conjectures will prove true. At the very least, the authors hope to stimulate thought within the academic and business worlds on how this topic might be approached.

REFERENCES

- Borch, K., 1962, Equilibrium in a Reinsurance Market, *Econometrica*, 30: 424-444.
- Cummins, D., N. Doherty, and A. Lo, 2001, Can Insurers Pay for "The Big One?" Measuring the Capacity of an Insurance Market to Respond to Catastrophic Losses, *Journal of Banking and Finance*, forthcoming.
- Cummins, D., and M. Weiss, 2000, The Global Market for Reinsurance: Consolidation, Capacity, and Efficiency, *Brookings-Wharton Papers on Financial Services*, R.E. Litan and A.M. Santomero, eds. (Washington, DC: The Brookings Institution), 159-222.
- Dionne, G., and N. A. Doherty, 1993, Insurance With Undiversifiable Risk, *Journal of Risk and Uncertainty*, 6: 187-203.
- Doherty, N. A., 1991, The Design of Insurance Contracts When Liability Rules Are Unstable, *Journal of Risk and Insurance*, 58: 227-246.
- Doherty, N. A., 1997, Financial Innovation in the Management of Catastrophe Risk, *Journal of Applied Corporate Finance*, Fall.
- Doherty, N. A., and H. Schlesinger, 1983, Optimal Insurance in Incomplete Markets, *Journal of Political Economy*, 91: 1045-1054.
- Embrechts, P., C. Klüppelberg, and T. Mikosch, 1997, *Modeling Extremal Events for Insurance and Finance* (Berlin: Springer-Verlag).
- Froot, K. A., 2001, The Market for Catastrophe Risk: A Clinical Examination, *Journal of Financial Economics*, 60: 529-571.
- Hofmann, M. A., 1996, Insurers Are Trying Harder to Develop Specialized Professional Liability Products, *Business Insurance*, 30: 13ff.
- Karni, E., 1995, Non-Expected Utility and the Robustness of the Classical Insurance Paradigm: Discussion, *Geneva Papers on Risk and Insurance Theory*, 20: 51-56.

²³ Numbers are based on projections by the New York State Insurance Department and the Insurance Industry, respectively, as reported in the *New York Post*, January 15, 2002.

- Loubergé, H., and H. Schlesinger, 2000, Optimal Catastrophe Insurance, Discussion paper no. 7, International Center for Financial Asset Management and Engineering (FAME).
- Marshall, J. M., 1974, Insurance Theory: Reserves Versus Mutuality, *Economic Inquiry*, 12: 476-492.
- Machina, M. J., 1982, "Expected Utility" Analysis Without the Independence Axiom, *Econometrica*, 50: 277-323.
- Machina, M. J., 1995, Non-Expected Utility and the Robustness of the Classical Insurance Paradigm, *Geneva Papers on Risk and Insurance Theory*, 20: 9-50.
- Niehaus, G., and S. Mann, 1992, The Trading of Underwriting Risk: An Analysis of Insurance Futures Contracts and Reinsurance, *Journal of Risk and Insurance*, 59: 601-627.
- Quiggin, J., 1982, A Theory of Anticipated Utility, *Journal of Economic Behavior and Organization*, 3: 323-343.
- Rothschild, M., and J. Stiglitz, 1970, Increasing Risk: I. A Definition, *Journal of Economic Theory*, 2: 225-243.
- Schlesinger, H., 1997, The Demand for Insurance Without the Expected-Utility Paradigm, *Journal of Risk and Insurance*, 64: 19-39.
- Schlesinger, H., 1999, Decomposing Catastrophic Risk, *Insurance: Mathematics and Economics*, 24: 95-101.
- Segal, U., and A. Spivak, 1990, First Order Versus Second Order Risk Aversion, *Journal of Economic Theory*, 51: 111-125.
- Smith, B., and M. Stultzer, 1990, Adverse Selection, Aggregate Uncertainty, and the Role for Mutual Insurance Contracts, *Journal of Business*, 63: 493-510.
- Yaari, M. E., 1987, The Dual Theory of Choice Under Risk, *Econometrica*, 55: 95-116.
- Zilcha, I., and S. H. Chew, 1990, Invariance of the Efficient Sets When the Expected Utility Hypothesis Is Relaxed, *Journal of Economic Behavior and Organization*, 13: 125-131.