The Journal of Risk and Insurance, 1997, Vol. 64, No. 4, 579–598.

Intertemporal Insurance 1

Bryan Ellickson
José Penalva-Zuasti

ABSTRACT

Assuming insurable events are generated by a marked point process, this article develops a framework in which insurance markets are dynamically complete in the sense of Kreps (1982). Insurance contracts can then be priced using the techniques of intertemporal finance: the equlibrium price of an insurance contract equals the discounted expected value of the payments which will be made in the event of an "accident" where expectation is taken with respect to the risk-neutral probability measure of Harrison and Kreps (1979). The difference between this expectation and the expectation under the true probability measure represents the equilibrium load.

INTRODUCTION

This article develops a framework that allows insurance contracts to be priced using the techniques of intertemporal finance. From Malinvaud (1972, 1973) to Cass, Chichilnisky, and Wu (1996), a significant part of the literature on insurance has been preoccupied with explaining how insurance markets manage to function despite the individualized nature of insurance contracts: if n individuals each face the prospect of having an accident or not, the state space is of size 2^n which seems to imply an impossibly large number of contingent contracts. As Malinvaud demonstrated, this apparent paradox can be resolved through an appeal to the law of large numbers. We suggest another approach to explaining why insurance markets work: introducing the possibility that accidents do not happen all at once, as in a static model, but gradually as part of a process that unfolds over time. This article develops the discrete-time version of the theory.

In our model, insurable events are represented as a marked point process, a stochastic process in which there is never more than one accident at any given

---

Bryan Ellickson is Professor of Economics and José Penalva-Zuasti is a doctoral student, both at the University of California, Los Angeles.

This research was supported by a grant from the UCLA Academic Senate. The authors thank David Levine, Joseph Ostroy, and William Zame, as well as participants at the UCLA Economics Department Theory Workshop, the 22nd Seminar of the European Group of Risk and Insurance Economists, and the Wharton Financial Institutions Center Conference on Risk Management in Insurance Firms for helpful comments on earlier versions of this article.

date. In this setting, the revelation over time of the true state of nature can be represented by a very simple information structure: an event tree with the same number of branches emerging from each node, one branch for each type of "accident" plus a branch corresponding to the possibility of no accident. As a consequence, market completeness in the dynamic sense of Kreps (1982) requires a number of contracts proportional to n, the number of consumers, rather than 2^n.

Establishing a framework in which dynamic spanning applies to insurance markets is the heart of this article; the rest simply applies standard finance in this specialized setting. As we shall see, finance provides some fresh insights into the nature of insurance. For example, even in a world of risk-averse consumers and investors, insurance contracts can be viewed as "actuarially fair" under the right probability measure: we show that the price of an insurance contract is the discounted expected value of future payments for the insured events where expectation is taken relative to the risk-neutral probability measure introduced by Harrison and Kreps (1979) in their synthesis of financial asset pricing theory. Insurance contracts provide a particularly natural interpretation of the Harrison-Kreps result: the risk-neutral probability measure incorporates society's aversion to risk and so builds into every "actuarially fair" insurance contract a load which compensates insurers for the risk they bear.

We begin with a formal description of a discrete-time marked point process and its application to insurance markets. Using a marked point process as the basic source of uncertainty, we define an intertemporal exchange economy on the event tree generated by the accident process and characterize the competitive equilibrium involving trade in Arrow-Debreu-Radner (ADR) date-event contingent commodities on this tree. The corresponding contingent prices allow us to price insurance contracts, which are typically *not* standard ADR contracts, as redundant securities. With this standard competitive equilibrium in place, we then "remove" the ADR contingent commodities, leaving only the insurance contracts to deal with risk. As we then show, insurance contracts alone yield dynamically complete markets in the sense of Kreps (1982), provided there is one insurance contract for each type of accident. We then illustrate the theory with a simple example yielding a natural interpretation of the Harrison-Kreps risk-neutral measure and a dramatic demonstration how trading a few insurance contracts can dynamically complete markets even though the state space is extremely large. Although the focus is on providing a bridge to the traditional insurance literature, our ultimate intent is to use intertemporal finance to develop a theory of how to insure events such as earthquakes or hurricanes for which the traditional theory seems less relevant. The conclusion describes some of the applications and extensions we have in mind.

THE MODEL

Our application of finance theory (as presented, e.g., in Dothan, 1990, or Huang and Litzenberger, 1988) to insurance markets hinges crucially on the assumption that markets are dynamically complete in the sense of Kreps (1982). Accidents are modeled as a discrete-time version of a marked point process with time intervals which should be interpreted as very short: short enough—say, a nanosecond—so that at most dates there is no "accident" and at no date is there more than one.[1] To simplify the discussion, we confine attention to a model of pure exchange with a single commodity available for consumption at each date.

Let $\boldsymbol{T} := \{0, 1, ..., T\}$ represent the time set.[2] An accident can happen at any date $t \geq 1$, but never more than one accident at any given date. If an accident happens at date t, it is assumed to occur and be known to all agents in the economy prior to trade or consumption at date t.

We assume that any accident occurring at date t can be classified into one of a finite number of accident types indexed by $\mathcal{K} := \{0, 1, ..., K\}$ with $k = 0$ signifying "no accident." All uncertainty in the economy is captured by the probability space $(\Omega, \mathcal{F}, P)$ with state space $\Omega := \mathcal{K}^T$, σ-algebra $\mathcal{F} = 2^\Omega$, and probability measure P. Throughout this article, we assume that $P(\omega) > 0$ for all $\omega \in \Omega$. For $k > 0$, let $N_k: T \times \Omega \rightarrow \mathbf{Z}_+$ represent the stochastic process which counts the number of accidents of type k: that is, $N_k(t,\omega)$ is the number of accidents of type k which have occurred up to date t. N_k is nondecreasing, $N_k(0,\omega) = 0$ for all $\omega \in \Omega$, at most one of the processes N_k jumps at any instant t and, when a jump occurs, it is always by exactly one unit. The vector-valued stochastic process defined by

$$N(t, \omega) := (N_1(t, \omega), ..., N_K(t, \omega))$$

is called a discrete time, K-variate counting process or, equivalently, a marked point process.

To illustrate these ideas more concretely, note that a sample point $\omega \in \Omega$ can be identified with a sequence $k_1, k_2, ..., k_T$ where $k_t \in \mathcal{K}$ denotes the type of accident occurring at time t. Thus, if $T = 10$ and $K = 5$, a sample point

$$\omega = 0300105300$$

corresponds to a realization in which there is an accident of type 1 at date five, accidents of type 3 at dates two and eight, an accident of type 5 at date seven, and no accident at any other date. N_3, the stochastic process which counts events of type 3, starts at $N_3(0, \omega) = 0$ at date zero, jumps to $N_3(2, \omega) = 1$ at date two and to $N_3(8, \omega) = 2$ at date eight while remaining constant otherwise.

In the usual way, the stochastic process N generates a filtration, a non-decreasing sequence of σ-algebras

$$\mathcal{F}_0 \subset \mathcal{F}_1 \subset \mathcal{F}_2 \subset \ldots \subset \mathcal{F}_T = \mathcal{F}$$

---

[1] A continuous-time formulation, which we are developing, is more elegant but also more complicated. For treatments of marked point processes see, for example, Brémaud (1981) or Last and Brandt (1995).

[2] We assume that $T < \infty$ in this section, but the analysis generalizes quite directly to the infinite horizon case. The example in the next section treats the case $T = \infty$.

with the property that $\mathcal{F}_t$ is the coarsest σ-algebra with respect to which the random variable $N(t): \Omega \to \mathbf{Z}_+$ is measurable. Let f_t denote the partition of Ω which generates the σ-algebra $\mathcal{F}_t$ with $f_0 = \{\Omega\}$ and $f_T = \{\{\omega\} \mid \omega \in \Omega\}$, corresponding to the absence of information at date t = 0 and complete information at date T. We will use a_t to denote a typical element of the partition f_t at date t.

Consumers in this economy are indexed by the finite set $I = \{1, ..., n\}$. For each consumer $i \in I$, a consumption process is a function $x_i : T \times \Omega \to \mathbf{R}$ and an endowment process a function $w_i : T \times \Omega \to \mathbf{R}$, each adapted to the filtration generated by the accident process N. Let L denote this vector space. Processes adapted to the accident filtration are constant on the sets $\{t\} \times a_t$ for each event $a_t \in f_t$.

Contingent commodities provide a convenient way to represent these consumption or endowment processes. The (t,a_t)-contingent commodity, representing one unit of consumption in event $a_t \in f_t$ at date t, is represented by the indicator function $1(t,a_t): T \times \Omega \to \mathbf{R}$ defined by[3]

$$1(t,a_t)(t',\omega') = \begin{cases} 1 & \text{if } t' = t \text{ and } \omega' \in a_t\ ; \\ 0 & \text{otherwise.} \end{cases}$$

Using these contingent commodities as a basis, a consumption process $x_i \in L$ for consumer $i \in I$ has the representation

$$x_i = \sum_{t=0}^{T} \sum_{a_t \in f_t} x_i(t,a_t)\, 1(t,a_t),$$

and an endowment process has the representation

$$w_i = \sum_{t=0}^{T} \sum_{a_t \in f_t} w_i(t,a_t)\, 1(t,a_t).$$

Letting

$$L_+ := \{x \in L \mid x(t,a_t) \geq 0\ \forall\, t \in T\ \&\ a_t \in f_t\}$$

represent the nonnegative orthant of L, assume that each consumer $i \in I$ has consumption set $X_i = L_+$, endowment $w_i \in L_+$, and preference relation $\succeq_i$ on $X_i \times X_i$ which is a complete preordering and strongly monotonic. Competitive prices are given by a linear functional $p: L \to \mathbf{R}$ with representation

[3] We write $1(t,a_t)$ rather than the more usual $1_{(t,at)}$ for typographical convenience.

$$p(x) = \pi \cdot x = \sum_{t=0}^{T} \sum_{a_t \in f_t} \pi(t, a_t)\, x(t, a_t),$$

where $\pi(t,a_t)$ is the price of a (t,a_t)-contingent commodity. Just as for consumption processes x, we can also view $\pi \in L_+$ as a stochastic process adapted to the accident filtration. A competitive equilibrium for this exchange economy consists of a feasible allocation x and a price system p such that x_i is in the demand set $\phi_i(p)$ of consumer i for every consumer $i \in I$.[4]

The ADR contingent commodities are, of course, purely theoretical constructs with little resemblance to actual insurance contracts. However, since markets are complete, any additional contracts which we choose to introduce are redundant assets and, consequently, they can be priced. Suppose we introduce a collection of assets which resemble actual insurance contracts, indexed by $j \in J$. Associated with each security $j \neq 0$ is a dividend process $d_j \in L_+$, where $d_j(t,a_t)$ represents the payout of insurance policy j at date-event (t,a_t), measured in units of the consumption good at (t,a_t). We assume that $d_j(0,\Omega) = 0$ for all $j \neq 0$. As for the price and consumption processes, each dividend process d_j can be viewed as a stochastic process adapted to the accident filtration, reflecting the fact that an insurance contract will pay off only on accidents that have already happened and not on those yet to come. In addition to these insurance policies, we also assume there exists a risk-free asset available at each date-event (t,a_t) which costs one unit of the consumption good at (t,a_t) and returns $d_0(t,a_{t+1}) = 1 + r(t,a_t)$ of the consumption good at each successor event $a_{t+1} \subset a_t$, $a_{t+1} \in f_{t+1}$. We will refer to $r(t,a_t)$ as the risk-free rate.

From now on, we adopt the price normalization[5] $\pi(0) = 1$ so that the competitive price functional has the representation

$$\pi \cdot x = x(0) + \sum_{t=1}^{T} \sum_{a_t \in f_t} \pi(t, a_t)\, x(t, a_t).$$

Given prices π, define for each event $a_t \in f_t$ and $a_{t+1} \in f_{t+1}$, where $a_{t+1} \subset a_t$, the *martingale conditional probability*

[4] The notation adopted here is from Ellickson (1993).

[5] Because the initial and terminal information partitions have a special structure,

$$f_0 = \{\Omega\} \quad \text{and} \quad f_T = \{\{\omega\} | \omega \in \Omega\},$$

it is convenient to abuse notation slightly, writing $\pi(0)$ and $x(0)$ in place of $\pi(0,\Omega)$ or $x(0,\Omega)$ and $\pi(T,\omega)$ or $x(T,\omega)$ in place of $\pi(T,\{\omega\})$ or $x(T,\{\omega\})$.

$$Q(a_{t+1}|a_t) := \frac{\pi(t+1, a_{t+1})}{\sum_{a'_{t+1} \subset a_t} \pi(t+1, a'_{t+1})}. \tag{1}$$

Assuming that prices of all insurance assets are ex dividend, we define for each $t < T$ the price process S_j for asset j according to the relation

$$\pi(t, a_t) S_j(t, a_t) = \sum_{s=t+1}^{T} \sum_{\substack{a_s \in f_s \\ a_s \subset a_t}} \pi(s, a_s)\, d_j(s, a_s),$$

the Arrow-Debreu valuation of the dividend stream following the date-event (t,a_t). At terminal nodes, $S_j(T,\omega) = 0$. Note that $S_j \in L_+$ so that security prices can also be viewed as non-negative stochastic processes adapted to the accident filtration. As with payouts, the security price $S_j(t,a_t)$ is measured in units of the (t,a_t)-consumption good.

As a simple consequence of the tree structure of the filtration,

$$\pi(t, a_t) S_j(t, a_t) = \sum_{a_{t+1} \subset a_t} \pi(t+1, a_{t+1})[S_j(t+1, a_{t+1}) + d_j(t+1, a_{t+1})], \tag{2}$$

where the sum is over events a_{t+1} belonging to the partition f_{t+1} and contained in a_t. In the special case of the risk-free asset, which costs one unit of the consumption good at (t,a_t) and pays $1 + r(t,a_t)$ at each of the immediate successor nodes, the above condition specializes to

$$\pi(t, a_t) = \sum_{a_{t+1} \subset a_t} \pi(t+1, a_{t+1})(1 + r(t, a_t)),$$

or, equivalently,

$$1 + r(t, a_t) = \frac{\pi(t, a_t)}{\sum_{a_{t+1} \subset a_t} \pi(t+1, a_{t+1})} \tag{3}$$

for all $t < T$. Using the definitions of the martingale conditional probability and the risk-free rate, equation (2) can now be written

$$S_j(t, a_t) = \frac{\sum_{a_{t+1} \subset a_t} [S_j(t+1, a_{t+1}) + d_j(t+1, a_{t+1})] Q(a_{t+1}|a_t)}{1 + r(t, a_t)}.$$

Letting

$$r(t) := \sum_{a_t \in f_t} r(t, a_t)\, 1(t, a_t),$$

denote the interest rate at date t and

$$S_j(t) := \sum_{a_t \in f_t} S_j(t, a_t)\, 1(t, a_t)$$

the price of the jth asset at date t, we have

$$S_j(t) = \frac{E_Q[S_j(t+1) + d_j(t+1) | \mathcal{F}_t]}{1 + r(t)}, \tag{4}$$

where $E_Q[\,\cdot\,|\,\mathcal{F}_t]$ denotes conditional expectation relative to the sigma-algebra $\mathcal{F}_t$ under the martingale measure Q.

Using the risk-free rate to discount insurance asset prices and their payouts, define

$$S_j^*(t) := \frac{S_j(t)}{\prod_{s=0}^{t-1}(1 + r(s))}$$

and

$$d_j^*(t) := \frac{d_j(t)}{\prod_{s=0}^{t-1}(1 + r(s))}.$$

Define the cumulative discounted dividend process for security j as

$$D_j^*(t) := \sum_{s=0}^{t} d_j^*(s).$$

Our first result asserts that insurance contracts are actuarially fair provided that expectations are computed according to the Harrison-Kreps risk-neutral (martingale) measure Q.

Theorem 1

For each insurance contract $j \in J$, $S_j^* + D_j^*$ is a martingale with respect to the measure Q.

Proof

From equation (4) it follows that

$$S_j^*(t)+D_j^*(t)=\frac{E_Q[S_j(t+1)+d_j(t+1)|\mathcal{F}_t]}{\prod_{s=0}^{t}(1+r(s))}+D_j^*(t)$$

$$=E_Q[S_j^*(t+1)+D_j^*(t)+d_j^*(t+1)|\mathcal{F}_t] \tag{5}$$

$$=E_Q[S_j^*(t+1)+D_j^*(t+1)|\mathcal{F}_t].$$

■

Our second result asserts that returns on insurance contracts satisfy a version of the capital asset pricing model. First, we decompose the rate of return of the jth insurance contract at date t+1 into a predictable and an innovation component,

$$\frac{S_j(t+1)+d_j(t+1)}{S_j(t)}-1=\mu_j(t+1)+\upsilon_j(t+1)-\upsilon_j(t),$$

where

$$\mu_j(t+1):=E_P\left[\frac{S_j(t+1)+d_j(t+1)}{S_j(t)}\middle|\mathcal{F}_t\right]-1$$

is the predictable component,

$$\upsilon_j(t+1):=\upsilon_j(t)+\frac{S_j(t+1)+d_j(t+1)}{S_j(t)}-E_P\left[\frac{S_j(t+1)+d_j(t+1)}{S_j(t)}\middle|\mathcal{F}_t\right]$$

is the innovation component, and $\upsilon_j(0)$ is an arbitrary constant. Define the likelihood ratio process *z* by

$$z(t):=E_P\left[\frac{Q}{P}\middle|\mathcal{F}_t\right]$$

for all $t \geq 1$. The following theorem establishes z as an aggregate risk factor for the economy, expressing the excess return of the jth insurance contract in terms of the conditional covariance between the innovation component of the return and the aggregate risk factor.[6]

Theorem 2

For each insurance contract $j \in J$ and each $t \geq 1$,

$$\mu_j(t+1) - r(t) = -\frac{1}{z(t)} covar_P[\upsilon_j(t+1), z(t+1)]. \tag{6}$$

Proof

From equation (4),

$$S_j(t) = E_Q\left[\frac{S_j(t+1)+d_j(t+1)}{1+r(t)} \middle| \mathcal{F}_t\right]$$

$$= \frac{1}{z(t)} E_P\left[\frac{S_j(t+1)+d_j(t+1)}{1+r(t)} z(t+1) \middle| \mathcal{F}_t\right]$$

$$= \frac{S_j(t)}{z(t)(1+r(t))} E_P\left[\frac{S_j(t+1)+d_j(t+1)}{S_j(t)} z(t+1) \middle| \mathcal{F}_t\right].$$

Therefore,

$$1+r(t) = \frac{1}{z(t)} E_P\left[\frac{S_j(t+1)+d_j(t+1)}{S_j(t)} z(t+1) \middle| \mathcal{F}_t\right]$$

$$= \frac{1}{z(t)} E_P\left[(1+\mu_j(t+1)+\upsilon_j(t+1)-\upsilon_j(t))z(t+1) \middle| \mathcal{F}_t\right]$$

$$= 1 + \frac{1}{z(t)} E_P\left[(\mu_j(t+1)+\upsilon_j(t+1)-\upsilon_j(t))z(t+1) \middle| \mathcal{F}_t\right]$$

where in the last step we use the fact that z is a P-martingale. Rearranging and using the fact that both z and υ_j are P-martingales,

[6] We follow closely the discussion leading up to Theorem 6.4 of Dothan (1990).

$$\mu_j(t+1) - r(t) = -\frac{1}{z(t)} E_P\left[(\upsilon_j(t+1) - \upsilon_j(t))(z(t+1) - z(t)) \,\middle|\, \mathcal{F}_t\right]$$

$$= -\frac{1}{z(t)} \mathrm{covar}_P\left[\upsilon_j(t+1), z(t+1)\right]$$

■

What we have shown is that, provided there is a complete set of ADR contingent contracts, standard insurance contracts can be priced. Assuming a complete set of ADR contingent contracts is clearly not very realistic. Fortunately, once insurance contracts have been introduced, the ADR contracts become redundant and therefore unnecessary. More precisely, our claim is that insurance contracts are capable of completing the market dynamically in the sense of Kreps (1982). The key to dynamic spanning is the index of the filtration which in an event tree context corresponds to the maximum number of branches leaving any node of the event tree: the number of accident types plus one in our model. As shown by Kreps (1982), the number of securities required for dynamic spanning is no greater than the index of the filtration. What this means for us is that, in addition to the risk-free asset, all that is required for dynamic spanning is one insurance contract covering each type of accident.

To illustrate, assume that $\mathcal{J} = \mathcal{K}$ where security $j = k \geq 1$ pays one unit of the consumption good at date t if and only if an accident of type k occurs at that date.[7] Let α represent an arbitrary event $a_{t-1} \in f_{t-1}$ and

$$B_t(\alpha) = \{a_t \subset a_{t-1} \mid a_t \in f_t\} = \{\beta_0, \beta_1, \ldots, \beta_K\}$$

the collection of its immediate successors. Let

$$\theta_i(t,\alpha) = (\theta_i^0(t,\alpha), \theta_i^1(t,\alpha), \ldots, \theta_i^K(t,\alpha))^T$$

represent the portfolio of securities purchased by consumer i at date-event $(t-1,\alpha)$ and

$$\theta_i(t+1,\beta) = (\theta_i^0(t+1,\beta), \theta_i^1(t+1,\beta), \ldots, \theta_i^K(t+1,\beta))^T$$

the portfolio acquired at date-event (t,β), $\beta \in B_t(\alpha)$, where T denotes transpose. Note that, by definition, the trading process θ_i is not only adapted to the filtration, but is also predictable: that is, for each t, $\theta_i(t)$ is measurable with respect to $\mathcal{F}_{t-1}$. Requiring predictability captures the economically natural restriction that a consumer must buy insurance prior to acquiring knowledge whether the insured event will occur. Finally, let

[7] In the terminology of Kreps (1982), these insurance contracts are long-lived securities.

$$\Delta_w x_i(t,\beta) := x_i(t,\beta) - w_i(t,\beta)$$

represent the net trade of consumer i at date-event (t,β). For each $\beta \in B_t(\alpha)$, budget balance in the spot market at date-event (t,β) requires

$$\Delta_w x_i(t,\beta) = \sum_{j\in J} \theta_i^j(t,\alpha)\left[S_j(t,\beta) + d_j(t,\beta)\right] - \sum_{j\in J} \theta_i^j(t+1,\beta)\, S_j(t,\beta),$$

which, letting

$$c_i(t,\alpha,\beta) := \Delta_w x_i(t,\beta) + \sum_{j\in J} \theta_i^j(t+1,\beta)\, S_j(t,\beta),$$

can be written

$$\sum_{j\in J} \theta_i^j(t,\alpha)\left(S_j(t,\beta) + d_j(t,\beta)\right) = c_i(t,\alpha,\beta). \tag{7}$$

For each date-event (t,α), define the K + 1 by K + 1 matrices

$$S(t,\alpha) = \begin{bmatrix} S_0(t,\beta_0) & S_1(t,\beta_0) & \cdots & S_K(t,\beta_0) \\ S_0(t,\beta_1) & S_1(t,\beta_1) & \cdots & S_K(t,\beta_1) \\ \vdots & \vdots & \ddots & \vdots \\ S_0(t,\beta_K) & S_1(t,\beta_K) & \cdots & S_K(t,\beta_K) \end{bmatrix}$$

and

$$D(t,\alpha) = \begin{bmatrix} d_0(t,\beta_0) & d_1(t,\beta_0) & \cdots & d_K(t,\beta_0) \\ d_0(t,\beta_1) & d_1(t,\beta_1) & \cdots & d_K(t,\beta_1) \\ \vdots & \vdots & \ddots & \vdots \\ d_0(t,\beta_K) & d_1(t,\beta_K) & \cdots & d_K(t,\beta_K) \end{bmatrix}.$$

The system of equations of type (7) at date-event (t,α) then takes the form

$$\left[S(t,\alpha) + D(t,\alpha)\right]\theta_i(t,\alpha) = c_i(t,\alpha), \tag{8}$$

where

$$c_i(t,\alpha) = (c_i(t,\alpha,\beta_0), c_i(t,\alpha,\beta_1), \ldots, c_i(t,\alpha,\beta_k))^T.$$

An ADR equilibrium allocation is said to be dynamically spanned by the set of securities J provided there is a portfolio $\theta_i(t,\alpha)$ which solves equation (8) for every date-event (t,α) and every consumer $i \in I$. If T is finite and the matrix $S(t,\alpha) + D(t,\alpha)$ is invertible for every date-event (t,α), then a set of dynamically spanning portfolio trades can be found by working backward from the terminal date T. In the next section, we illustrate a method of solution for an infinite horizon which exploits stationarity. In either case, the matrix $S(t,\alpha) + D(t,\alpha)$ must be invertible at each step, a condition which, as we shall see, holds quite naturally when there is a separate insurance contract for each type of accident.

ILLUSTRATING THE THEORY

This section illustrates our abstract theory of insurance pricing with an application to the type of risk situation which is the focus of the traditional insurance literature. In our model an accident can be any event—a flood, a car wreck, or an earthquake, for example—which affects the endowment or the preferences of one or more consumers. However, to facilitate comparison with the traditional insurance literature, here we consider a situation in which each accident affects only one individual and every consumer is affected by an accident in exactly the same way: for simplicity, we assume the consumer loses his entire endowment. Formally, let $\mathcal{K} = I \cup \{0\}$, where $k \geq 1$ represents an accident happening to individual $k = i \in I$ and $k = 0$ an accident happening to no one. Assume that consumer $i \in I$ has von-Neumann Morgenstern utility

$$u_i(x_i) = \ln(x_i(0)) + \sum_{t=1}^{T} \delta^t \sum_{a_t \in f_t} P(a_t) \ln(x_i(t,a_t)), \tag{9}$$

where $\delta \in [0,1)$ and $P(a_t)$ is the probability of event $a_t \in f_t$. Let

$$w(t,a_t) := \sum_{i \in I} w_i(t,a_t)$$

denote the aggregate endowment of the (t,a_t)-contingent commodity and

$$w(0) := \sum_{i \in I} w_i(0)$$

the aggregate endowment at date zero. Using the normalization $\pi(0) = 1$, the market-clearing price for the (t,a_t)-contingent commodity is given by

$$\pi(t,a_t) = \frac{\delta^t w(0)\, P(a_t)}{w(t,a_t)}. \tag{10}$$

For any $a_t \in f_t$ and $a_{t+1} \in f_{t+1}$, where $a_{t+1} \subset a_t$, assume that the conditional probability

$$P(a_{t+1}|a_t) = \begin{cases} \lambda_i \text{ if an accident happens to } i \in I \text{ at date } t+1, \text{ and} \\ \lambda_0 \text{ if no accident happens at date } t+1. \end{cases}$$

Assume also that the endowment of consumer i at date t will be

$$w_i(t, a_t) = \begin{cases} 0 \text{ if an accident happens to consumer i at date t, and} \\ Y \text{ otherwise,} \end{cases}$$

where $Y > 0$ is the same for all consumers. Thus, the hazard rate for an accident happening to consumer i remains constant over time and, when an accident happens to i, it destroys her entire endowment at the date on which the accident occurs.

Under these assumptions, equation (1) yields the following expressions for the martingale conditional probabilities:

No accident. If a_{t+1} is the event "no accident happens to any consumer at date t+1," then

$$Q(a_{t+1}|a_t) = \frac{\lambda_0}{\lambda_0 + \frac{n}{n-1}(1-\lambda_0)}.$$

Accident to consumer i. If a_{t+1} is the event "an accident happens to consumer i at date t+1," then

$$Q(a_{t+1}|a_t) = \frac{\lambda_i}{\frac{n-1}{n}\lambda_0 + (1-\lambda_0)}.$$

Thus, the insurance context provides a natural interpretation of the Harrison-Kreps "risk-neutral" probability measure Q: provided there is more than one consumer, Q increases the probability of an accident to each consumer and, consequently, lowers the probability of no accident. By inflating the risk involved, Q builds a "load" into each insurance contract compensating insurers for the risk they bear.

Specializing some more, suppose the hazard rate of no accident is $\lambda_0 = 1/4$, the number of consumers is even, and the population is split into two equal size groups: $I = I_1 \cup I_2$ with

$\lambda_i = 1/n$ for the *high-risk* consumers $i \in I_1$; and

$\lambda_i = 1/2n$ for the *low-risk* consumers $i \in I_2$.

For a_{t+1} the event "no accident at date t + 1,"

$$Q(a_{t+1}|a_t) = \frac{n-1}{4n-1}.$$

For a_{t+1} an event corresponding to an accident to some high-risk consumer,

$$Q(a_{t+1}|a_t) = \frac{4}{4n-1}.$$

And, finally, for a_{t+1} an event corresponding to an accident to some low-risk consumer,

$$Q(a_{t+1}|a_t) = \frac{2}{4n-1}.$$

As n, the number of consumers, approaches infinity, the martingale probability of no accident approaches one-fourth and the martingale probability of an accident to any particular consumer, whether high or low risk, approaches zero. In each of these cases,

$$\lim_{n \to \infty} \frac{Q(a_{t+1}|a_t)}{P(a_{t+1}|a_t)} = 1,$$

so that, consistent with the treatment in Malinvaud (1972, 1973), with a large number of consumers there is little difference between actuarial and martingale insurance pricing in this setting.

Specializing even more, assume that $n = 10$, $\delta = 12/13$, $Y = 390$, and $T = \infty$. Table 1 shows the equilibrium trades and net trades of the (t,a_t)-contingent commodity for a high-risk consumer (consumer 1) and a low-risk consumer (consumer 2) under three conditions: (a) there is no accident at date t (the accident type or "mark" $k(t) = 0$), (b) there is an accident to consumer i ($k(t) = i$), and (c) there is an accident to some consumer $i' \neq 1$ ($k(t) = i'$). Note that gross trades depend only on the "macro risk" in the economy, that is, whether an accident happens to some consumer or there is no accident, while net trades also depend on who has the accident.

Even though we allow at most one accident per "nanosecond," we face an explosion of contingent commodities: each node of the event tree is followed by 11 branches, so that there are 11 contingent commodities at date 1, 11^2 at date 2, 11^3 at date 3, and so forth. However, in our context, unlike that of Malinvaud, events unfold gradually over time, and, because time allows for dynamic spanning, the exponential explosion of states does not lead to an explosion of required insurance contracts. To illustrate, we consider how these Arrow-Debreu equilibria can be implemented under two insurance regimes, one offering short-term insurance to consumers and the other offering long-term contracts. Either regime cuts the number of securities needed for spanning dramatically: 10T short-term insurance contracts with the first regime or 10 insurance contracts of infinite duration plus a risk-free asset under the second regime doing the work of 11^T ADR contingent contracts.

Table 1
Equilibrium Trades and Net Trades

k(t)	$x_1(t,a_t)$	$x_2(t,a_t)$	$\Delta_w x_1(t,a_t)$	$\Delta_w x_2(t,a_t)$
0	380	400	−10	10
i	342	360	342	360
i'	342	360	−48	−30

Short-Term Insurance

We know that to achieve dynamic spanning, it is necessary to offer a separate insurance contract for each type of accident. In the first regime we consider, insurance contracts are short-term: one unit of insurance issued at date t on accidents of type k at date t+1 returns a payout at t+1 of

$$d_k(t+1, a_{t+1}) = \begin{cases} 1 \text{ if an accident of type k occurs at date t+1;} \\ 0 \text{ otherwise.} \end{cases}$$

Because security prices are ex dividend and this contract is short-term, it is worthless at date t + 1: $S_k(t+1) = 0$. Applying equation (4), we conclude that, for a high-risk consumer,

$$S_i(t, a_t) = \begin{cases} 4/39 \text{ if there was no accident at date t;} \\ 6/65 \text{ otherwise;} \end{cases}$$

while, for a low-risk consumer,

$$S_i(t, a_t) = \begin{cases} 2/39 \text{ if there was no accident at date t;} \\ 3/65 \text{ otherwise.} \end{cases}$$

To verify that we have dynamic spanning, equation (7) must be modified slightly to account for the fact that we are considering short-term insurance securities. At date t there are two securities of type k = i in existence: the "old" contracts issued at t − 1 and paying off at t and the "new" contracts issued at t and paying off at t + 1. Since contracts issued yesterday are worthless today (i.e., their ex dividend price is zero), we reserve $S_k(t,\alpha)$ to represent the price of contracts issued at date t and event $\alpha \in f_t$ on an accident of type k. With this modification, equation (7) becomes

$$\sum_{j=0}^{10} \theta_i^j(t,\alpha)\, d_j(t,\beta_k) = \Delta_w x_i(t,\beta_k) + \sum_{j=0}^{10} \theta_i^j(t+1,\beta_k)\, S_j(t,\beta_k) \qquad (11)$$

for each successor event $\beta_k \in B_t(\alpha)$.

Because we have assumed an infinite horizon in this example, we cannot solve for the trading portfolios by backward recursion. However, we might suspect that the portfolios are stationary, and that turns out to be the case. In equilibrium, neither the high-risk nor the low-risk consumer holds any of the risk-free asset. Each consumer buys $\theta_i^i(t,a_t) = 351$ units of insurance on accidents to herself and sells $\theta_i^k(t,a_t) = -39$ units of insurance to each consumer $k \neq i$. It is tedious but straightforward to verify that equation (11) is satisfied with this portfolio for each $\alpha \in f_t$ and $\beta_k \in B(\alpha)$ for all $t \geq 0$.

Long-Term Insurance

The short-term insurance contracts of the preceding section are, of course, simply Radner's variation on Arrow-Debreu contingent commodities: that is, contingent commodities traded at dates $t > 0$ rather than at the beginning of time. Although they do reduce the number of contingent contracts or securities required for spanning from 11^T to 10T plus a risk-free asset, such contracts seem only slightly more realistic than their ADR counterparts: issuing a new insurance contract every nanosecond puts a heavy burden on our tacit assumption of no transactions costs! Taking a step closer to reality, we now consider long-term insurance contracts. In the real world, insurance contracts cover all occurrences of some event over an extended period of time—say, a year. Here we assume for simplicity that the contracts last forever: specifically, an insurance policy on an accident of type k is a long-term obligation which, at a price $S_k(t,a_t)$ at date-event (t,a_t), returns one unit of the consumption good at each subsequent date-event at which an accident of type k occurs.

For the sake of symmetry, we also replace the risk-free asset with a bond which pays one unit of the consumption good at every date-event (t,a_t) from time one forward. For $a_t \in f_t$, let $\beta_k \in f_{t+1}$ be the immediate successor event in which an accident of type k occurs. From equation (4), the price of the bond at (t,a_t) is

$$S_0(t,a_t) = \frac{1 + \sum_{k=0}^{10} Q(\beta_k | a_t)\, S_0(t+1,\beta_k)}{1 + r(t,a_t)}.$$

Of the eleven successor events which could occur at t + 1, only two are distinct macro states in this simple economy: either an accident occurs or it does not. Let α^*, α^0 denote events in f_t in which an accident does or does not occur, respectively, and let β^*, $\beta^0 \in f_t$ denote the corresponding events in f_{t+1}. Applied to all $a_t \in f_t$, equation (4) reduces to two distinct equations:

$$S_0(t,\alpha^0) = \frac{1+(3/13)\,S_0(t+1,\beta^0)+(10/13)\,S_0(t+1,\beta^*)}{1+r(t,\alpha^0)},$$

and

$$S_0(t,\alpha^*) = \frac{1+(3/13)\,S_0(t+1,\beta^0)+(10/13)\,S_0(t+1,\beta^*)}{1+r(t,\alpha^*)}.$$

Assuming bond prices are stationary, so that

$$S_0(t,\alpha^0)=S_0(t,\beta^0) \text{ and } S_0(t,\alpha^*)=S_0(t,\beta^*),$$

we conclude that

$$S_0(t,a_t)=\begin{cases}13 & \text{if there is no accident at date t;}\\ 117/10 & \text{if there is an accident at date t.}\end{cases}$$

The prices of the insurance contracts are found in essentially the same way. Exploiting the fact that there are only two distinct macro states at any date and the hypothesis that the insurance contract price should be stationary, equation (4) implies that, for insurance on a high-risk consumer,

$$S_i(t,a_t)=\begin{cases}4/3 & \text{if there is no accident at date t;}\\ 6/5 & \text{if there is an accident at date t;}\end{cases}$$

while, for a low-risk consumer,

$$S_i(t,a_t)=\begin{cases}2/3 & \text{if there is no accident at date t;}\\ 3/5 & \text{if there is an accident at date t.}\end{cases}$$

Turning to the issue of dynamic spanning, equation (7) becomes

$$\sum_{j=0}^{10} \theta_i^j(t,\alpha)(S_j(t,\beta_k)+d_j(t,\beta_k)) = \Delta_w x_i(t,\beta_k) + \sum_{j=0}^{10} \theta_i^j(t+1,\beta_k) S_j(t,\beta_k) \quad (12)$$

for each successor event $\beta_k \in B_t(\alpha)$. Once again the hypothesis that stationary portfolios can achieve dynamic spanning turns out to be correct, but now, in contrast to the short-term insurance regime, low- and high-risk consumers hold somewhat different portfolios. Each high-risk consumer $i \in I_1$ sells $\theta_i^0(t,a_t) = -10$ units of the risk-free asset, buys $\theta_i^i(t,a_t) = 352$ units of insurance on accidents to herself, and sells $\theta_i^j(t,a_t) = -38$ units of insurance to each consumer $j \neq i$. Each low-risk consumer $i \in I_2$ buys $\theta_i^0(t,a_t) = 10$ units of the riskfree asset, buys $\theta_i^i(t,a_t) = 350$ units of insurance on accidents to herself, and sells $\theta_i^j(t,a_t) = -40$ units of insurance to each consumer $j \neq i$. It is easy to verify that these portfolios satisfy equation (12) for each $\alpha \in f_t$ and $\beta_k \in B(\alpha)$ for all $t \geq 0$ and that the security markets clear.

CONCLUSION

The example of the above section follows Malinvaud in assuming that an accident never affects more than one individual. What happens if, as in the case of an earthquake or a flood, many consumers are affected simultaneously? Malinvaud's framework leaves no scope for such a possibility, while the finance-based approach to insurance offered here covers such cases with no change in the theory. Suppose there are n consumers but only two types of accidents, so that $\mathcal{K} = \{0, 1, 2\}$. For the sake of interpretation, imagine a world consisting of two regions, either of which can experience an earthquake at date t. To apply the theory developed in this article, we must assume there is never more than one earthquake at a given date. Suppose 2n consumers are divided into subsets of equal size: I_1 who reside in region 1 and I_2 who reside in region 2. For each $i \in I := I_1 \cup I_2$, assume preferences are once again represented by the utility function (9), but the endowment of consumer $i \in I$ is now given by $w_i(t,a_t) = 0$ if an earthquake occurs in the region where she resides at date t and Y otherwise.

It is easy to see that the results we obtain are essentially those of the previous section in the special case of two consumers. In particular, the ADR prices, equilibrium net trades, and security prices are the same, and they do not change as n increases: insurance is priced competitively, but the law of large numbers is irrelevant.

As will be apparent to those familiar with the finance literature, the research reported here only begins to tap the potential for applying the tools of intertemporal finance to insurance markets. Insurance contracts are clearly more complex than the simple instruments described in this article, typically insuring a variety of types of accidents over varying periods of time with options to renew and the like. All such contracts are "redundant assets" in this setting and, as such, can be priced.

Insurance contracts also typically pay out in real rather than nominal terms, a distinction we have not addressed in our single-commodity version of the model but which clearly can be addressed in an extension of the model. Development of some version of a mutual fund theorem, in which consumers buy insurance on themselves and invest in a market portfolio of insurance contracts on others, is another obvious possibility.

As the theoretical discussion clearly shows, there is no reason to assume hazard rates are independent of the past history of the process and constant or that the effect of an accident is confined to the date at which it occurs. When one medical problem strikes, it may announce the increased chance of other problems arising. And an accident today may put a worker out of commission for months or years to come.

Finally, a major item on the agenda is to extend this model to continuous time. Much of the formalism of this article is aimed at making the transition from discrete to continuous time as effortless as possible.

References

Arrow, K., 1953, Le rôle des valeurs boursières pour la répartition la meilleure des risques, *Econometrie,* 11: 41–48 (Centre National de la Recherche Scientifique, Paris). [Reprinted as: The Role of Securities in the Optimal Allocation of Risk-Bearing, *Review of Economic Studies,* 31: 91–96 (1964).]

Brémaud, P., 1981, *Point Processes and Queues: Martingale Dynamics* (New York: Springer-Verlag).

Cass, C., G. Chichilnisky, and H.-M. Wu, 1996, Individual Risk and Mutual Insurance, *Econometrica,* 64: 333–341.

Dothan, M., 1990, *Prices in Financial Markets* (New York: Oxford University Press).

Duffie, D., 1988, *Security Markets: Stochastic Models* (Boston: Academic Press).

Duffie, D. and C. Huang, 1985, Implementing Arrow-Debreu Equilibria by Continuous Trading of Few Long-Lived Securities, *Econometrica,* 53: 1337–1356.

Ellickson, B., 1993, *Competitive Equilibrium: Theory and Applications* (Cambridge: Cambridge University Press).

Harrison, J. and D. Kreps, 1979, Martingales and Arbitrage in Multiperiod Securities Markets, *Journal of Economic Theory,* 20: 381–408.

Hirshleifer, J. and J. Riley, 1992, *The Analytics of Information and Uncertainty* (Cambridge: Cambridge University Press).

Huang, C. and R. Litzenberger, 1988, *Foundations for Financial Economics* (Amsterdam: Elsevier).

Kreps, D., 1982, Multiperiod Securities and the Efficient Allocation of Risk: A Comment on the Black-Scholes Option Pricing Model, in: J. McCall, ed., *The Economics of Uncertainty and Information* (Chicago: University of Chicago Press).

Laffont, J., 1989, *The Economics of Uncertainty and Information* (Cambridge, Mass.: MIT Press).

Last, G. and Andreas Brandt, 1995, *Marked Point Processes on the Real Line: The Dynamic Approach* (New York: Springer).

Malinvaud, E., 1972, The Allocation of Individual Risks in Large Markets, *Journal of Economic Theory,* 4: 312–328.

Malinvaud, E., 1973, Markets for an Exchange Economy with Individual Risks, *Econometrica,* 41: 383–410.

Radner, R., 1972, Existence of Equilibrium of Plans, Prices and Price Expectations in a Sequence of Markets, *Econometrica,* 40: 1365–1379.