

INSURANCE WITH FREQUENT TRADING

A Dynamic Analysis of Efficient Insurance Markets

By Jose S. Penalva Zuasti¹

*Department of Economic and Business, Universitat Pompeu Fabra
Ramon Trias Fargas, 25-27, Barcelona, 08005, SPAIN*

E-mail: jose.penalva@econ.upf.es

Key Words: risk sharing, insurance, complete markets, insurable risks.

This paper extends existing insurance results on the type of insurance contracts needed for insurance market efficiency to a dynamic setting. It introduces continuously open markets that allow for more efficient asset allocation. It also eliminates the role of preferences and endowments in the classification of risks, which is done primarily in terms of the actuarial properties of the underlying risk process. The paper further extends insurability to include correlated and catastrophic events. Under these very general conditions the paper defines a condition that determines whether a small number of standard insurance contracts (together with aggregate assets) suffice to complete markets or one needs to introduce such assets as mutual insurance.

1. INTRODUCTION

Insurance companies have an image of being very conservative and reluctant to be taken up by the enthusiasm of emerging financial markets and financial innovations. In the past, insurance contracts were sold for a relatively long period of time, prices readjusted little, and the risk from insurance contracts was dealt with either by reselling it as reinsurance

¹I would specially like to thank Bryan Ellickson for providing constant encouragement, comments and suggestions on the many iterations of this paper, William R. Zame for long hours of helpful discussion, comments and suggestions, and Juanjo Ganuza as well as James Costain, Arturo Kohatsu and Manuel Moreno. I also wish to thank audiences at the ESEM 1999 meeting in Santiago, Brown University, New York University, the William E. Simon Graduate School of Business, UC Los Angeles, UC Santa Cruz, Universitat Pompeu Fabra, Universidad Carlos III, and Universidad de Alicante for their helpful questions and comments. A prior version of this paper is available as UPF working paper #461. All errors are my own. I also wish to acknowledge financial support from Argentario via its "Convenio de Colaboración entre Argentario y la Universitat Pompeu Fabra", and the Generalitat de Catalunya (1999 SGR00157)

or absorbing and managing it using relatively stable capital reserves. As financial theory develops and financial markets are liberalized, more financial type approaches are being introduced into risk management - via the entry or influence of banks, and via a greater financial training among insurance professionals. This phenomenon is moving the industry towards more active risk management and requires a more dynamic theoretical treatment of insurance risks. We propose a framework which can be used to interpret the behaviour of insurance companies in the midst of financial market liberalization, and relate it to the problem of insurability and efficient risk sharing.

Risk is a major component of modern economic analysis. The efficient management of risk in an economy involves, for the most part, a reallocation of this risk: optimal risk sharing. Arrow and Lind(1970) argue that most of these risks are of an idiosyncratic nature and Arrow(1965) underlines the importance of understanding the scope of risk sharing and of understanding the restrictions imposed by insurability. This has been studied in an essentially static context and our paper extends the analysis by allowing dynamic asset trading. Our insights extend the initial extension of insurance to dynamic trading made by Ellickson and Penalva(1997). As we will show, dynamic trading allows the efficient management of correlated and catastrophic risks, but the number and type of assets needed will depend on how nicely does information about risks flow.

Insurance markets are considered to be the main vehicle for dealing with agent-specific risk, hence it becomes crucial to understand when such markets provide an efficient reallocation of risk. Malinvaud(1972) has argued that the key property is the possibility of applying the Law of Large Numbers. This is because private insurance contracts cannot deal with economy-wide scarcity in addition to the risk of an agent's endowment loss. To ensure ex-post market clearing there cannot be aggregate risk. Idiosyncratic risks (and hence insurable risks) are those that can be "socially removed by the operation of the law of large numbers" (Malinvaud, 1972).

Borch(1990) has taken a more pragmatic point of view. Given the possibilities provided by Lloyd's of London to ensure a wide range of risks, a large number of which would be classified in the uninsurable category, he concentrated on the issues of moral hazard and adverse selection as the main obstacles for a risk to be insurable. Every risk that does not include moral hazard or adverse selection is then insurable.

Cass, Chilchinski and Wu(CCW, 1996) consider an intermediate case: a finite number of agents and several types. All agents of the same type have identical preferences, identical ex-ante endowments and uncertainty over their final endowment is described by an exchangeable

random variable, which allows for aggregate risks. They show that efficient risk-sharing can be attained with a sufficiently low number of insurance contracts and Arrow securities. This extends the types of risks that can be dealt with efficiently but it introduces two additions to insurance markets: the need for “mutual insurance” (contracts written on individual and aggregate events), and the use of non-insurance related securities to deal with aggregate risk.

These advances in the theory of insurance have largely taken place independently of advances in finance theory. Ellickson and Penalva(1997) started to bridge this gap. They construct a model with discrete time consumption and no intermediate trading. They impose that at most one accident in the whole economy occurs between consumption dates and should be considered as a discrete approximation to a model with continuous time consumption and trading. Our model extends theirs by allowing intermediate trading and explicitly modelling the continuous revelation of information. We extend their results by providing a much more detailed analysis of risks and their effect on the type and number of asset traded to attain complete markets in equilibrium, more general definitions of insurance contracts, and explicit consideration of correlated risks and catastrophic risks.

In the finance literature, the role of financial markets for optimal diversification has been addressed under the heading of complete markets and effective completeness. Christensen et al(1999) and Zhou(1995) try to refine Duffie and Huang’s (1985) result on the minimal number of assets needed for complete markets in a dynamic context. They use Brownian motion and consumption at every trading date. We find both assumptions more appropriate for finance than insurance. They also characterize insurance as contracts that cover the idiosyncratic component of agents’ risks and are interested solely on the number of assets needed for effective completeness, without much comment on the insurance implications or interpretations. This gap between insurance and finance theory needs to be covered both for theoretical completeness and to be able to address the specific issues that arise in the insurance industry as financial markets liberalize and innovate.

Our approach extends the study of insurability to a dynamic setting, extending the results of Malinvaud and CCW and including them as special cases. This requires the definition of an appropriate general notion of risks and of the types of contracts used: private insurance, mutual insurance, and diversified portfolios. In this new setting the issue of whether risks are insurable or not continues to depend on whether insurance markets are effectively complete or not, and whether simple insurance contracts suffice, or one has to introduce more complex contracts such as mutual insurance as proposed by CCW. We depart from

the finance literature by not using Brownian motion or general Levy processes to describe the revelation of uncertainty, but rather prefer the use of a discrete marked point process, a generalization of Poisson processes. Also, the types of risks we are analyzing have a fundamental agent-specific component which does not arise in most finance models, and the contracts we look at are based on standard private insurance contracts and not on risk conditioned on the aggregate state of the world, as proposed by Zhou and Christensen et al.

Our analysis shows that insurability does not depend on there being a large number of identical agents, or on the preferences and endowments of those agents. The important characteristics for insurability are those of the risks involved, where risk is defined as in done in insurance practice: the characteristics of the potential loss and the distribution of the loss function (i.e., the properties of the stochastic loss process). We allow for risks to have aggregate effects as well as a very rich set of correlations and joint effects. We find a condition that determines whether such risks can be decentralized using standard insurance contracts or whether they require more sophisticated contracts such as mutual insurance (as proposed by CCW).

This paper is structured as follows: this introduction is followed by a description of the model and the main assumptions. We then consider an initial description of risk in the economy, and define insurance contracts and markets. Section four presents the first result describing how insurance markets could function with the initial description of risk, which is then generalized in subsequent sections. This generalization starts in section 5, which defines risk in a generic sense and introduces several ways in which this risk could be described (independent risk, exchangeable risk, Markov risk). Section 6 states the key condition that determines whether simple insurance contracts suffice or not, and the corresponding general results. Section 7 includes extensions and discussion and section 8 concludes. All proofs are relegated to the Appendix.

2. THE MODEL

Other authors studying optimal risk sharing in a dynamic context have used continuous consumption and trading models¹. I choose to separate consumption from trading dates both for theoretical and for practical reasons².

¹Aase(1993), Christensen et al(1999), Zhou(1995).

²For the theoretical part, see Hindy and Huang(1993) on the problems raised by continuous time consumption. In practice, consumption data is gathered at time intervals much wider than financial data. For more detailed discussion on these issues see Penalva[1997].

The separation of consumption and trading dates implies having a standard discrete time economic framework on which consumption dates are defined $\mathbf{T}; = \{0, 1, \dots, T\}$, and allow trading in between consumption dates, so we have a continuum of trading dates: $\mathcal{T} := [0, T]$. This approach (which extends that of the original asset pricing models of Harrison and Kreps(1979) and others - see Merton(1990) for further references) combines the benefits of existing continuous time and discrete time models in a single framework, while avoiding some of the problems with modeling time either as continuous or as discrete.

The fundamental uncertainty in the economy is described by $(\Omega, \mathcal{F}, P, \mathbf{F})$, which is defined by \tilde{N} , a Marked Point Process³(MPP), where $\tilde{N}(t) = (N_1(t), N_2(t), \dots, N_K(t))$, is the vector of counting functions generated from the marks, $N_k(t)$, $t \in \mathcal{T} := [0, T]$ (T can be either finite or infinite). Basically each element of the vector \tilde{N} , $N_k(t)$, keeps track of how many times something (whatever we choose k to represent: an earthquake, agent i 's house burning down, ...) has taken place⁴ up to date t .

Assumption 1. The marked point process, \tilde{N} admits a uniformly bounded, absolutely continuous intensity $\tilde{\lambda} = (\lambda_1, \lambda_2, \dots, \lambda_K)$ and for each j , N_j , admits the intensity λ_j . For all $t \in \mathcal{T}$, define

$$N(t) = \sum_{j=1}^K N_j(t).$$

Let Ω be the space of all possible paths of the MPP \tilde{N} . \mathbf{F} is the filtration generated by $\tilde{N}, \mathbf{F} := (\mathcal{F}_t)_{t \in \mathcal{T}}$, where⁵ $\mathcal{F}_t := \sigma(\tilde{N}(s), s \leq t)$, and $\mathcal{F} =$

³For more details on MPPs see Bremaud(1981) or Last and Brandt(1995). Essentially, they are a vector extension and great generalization of Poisson processes.

⁴To be more formal, $N_k(t)$ is a random variable which counts the number of jumps of type k that have taken place up to and including time t . Hence, $\tilde{N}(t)$ is a random vector which describes the total number of jumps of each type up to time t , and $(\tilde{N}(s))_{s \in \mathcal{T}}$ is a vector process. Note that although I have used the vector process $\tilde{N}(t)$ to describe the MPP, I could equally well have defined the MPP over sequences of random variables, (τ_n, Z_n) , where τ_n is the time of the n -th jump and $Z_n \in \{1, 2, \dots, K\}$ the mark of that jump. This implies that $N_k(t)$ can be written as

$$N_k(t) = \sum_{n=1}^{\infty} \mathbf{1}_{\tau_n \leq t} \mathbf{1}_{Z_n = k}.$$

where $\mathbf{1}$. is an indicator function. For the most part we will use the process description of the MPP, but at some points we will find it useful to revert to the jump time/mark specification.

⁵Note that $\sigma(x(s), s \leq t)$ denotes the σ -algebra generated by the process $x(t)$ up to and including time t . Also, I assume throughout the paper that for all $t \in \mathcal{T}$, all random variables, $x(t)$, are measurable with respect to the internal history of \tilde{N} , \mathcal{F}_t , unless explicitly stated otherwise.

$\forall t \in \mathcal{T} \mathcal{F}_t$. $\tilde{\lambda}$ is \mathbf{F} measurable and predictable. P is the probability measure on (Ω, \mathcal{F}) implied by $\lambda = (\lambda_1, \dots, \lambda_k)$.

Assumption 1 describes the kind of uncertainty facing the economy. The main things to note are:

- Uncertainty is exogenous and entirely described by \tilde{N} . This assumption is quite restrictive but is the necessary starting point. It precludes problems that involve moral hazard and adverse selection, but such issues are difficult for general equilibrium frameworks in general and are already recognized as a problem for insurability.
- By assuming that \tilde{N} admits a bounded intensity, I am giving the rate at which accidents can take place a uniform finite upper bound.
- A bounded intensity means that accidents are spread out over $[0, T]$. This ensures that with probability one, the number of jumps in any finite time interval will be finite⁶.
- Standard statistical methods are applicable. By assuming that each of the point processes associated to a mark admits its own absolutely continuous intensity we can define the intensity of the aggregate process⁷, N , $\lambda = \sum_{j=1}^K \lambda_j$, and can talk about the probability of a jump of type j conditional on past history (which is now well-defined).

I will use a standard perfectly divisible commodity space:

Assumption 2. The commodity space, L , is the space of non-negative absolutely bounded real-valued functions⁸ on $\Omega \times \mathbf{T}$, measurable on \mathcal{F}_t for all $t \in \mathbf{T} := \{0, 1, \dots, T\}$, the index set of consumption dates $L := L_\infty(\Omega)^{\mathbf{T}} = \{x \mid x : \Omega \times \mathbf{T} \rightarrow \mathbf{R}, \forall t \in \mathbf{T}, x(t) \in \mathcal{F}_t, \text{ bounded}\}$. The dual of L (the space of prices) is denoted $L^* := L_1(\Omega)^{\mathbf{T}}$ as we use the Mackey topology.

The set of agents will be finite and with standard time-separable, state-independent utility functions:

Assumption 3. There are $n < \infty$ agents indexed by $i \in I := \{1, 2, \dots, n\}$. Each agent is described by a consumption set, $X_i = L_+$, an endowment, $e_i \in L_+$, and Von Neuman-Morgenstern preferences of

⁶Uniform boundedness ensures λ implies a measure in case $T = \infty$.

⁷The intensity of a point process, it is intimately linked to the intuitive notion of the rate at which jumps take place and is a generalization of the statistical concept of a hazard rate (also referred to as a survival rate). See Brémaud(1981) chapter II for further details and the corresponding definitions and theorems: D7, T8, T15.

⁸Naturally, I identify functions as being equal up to sets of measure zero, P a.s.

the form

$$U_i(x) = E_P \left[\sum_{t \in \mathbf{T}} \beta_i^t u_i(x(t)) \right], \quad (1)$$

where $u_i(x)$ is a monotone increasing, concave real-valued function satisfying the standard Inada conditions. Denote the aggregate endowment by $e = \sum_{i \in I} e_i$. Assume $e > 0$ P -a.s.

Note that I am assuming agents have common priors. This isolates prices as a reflection of scarcity, rather than of differences of opinion (and also eliminates issues of asymmetric information, moral hazard, etc, as mentioned above). Also, it ensures that ‘insurance companies’ can calculate the risk, a condition for insurability one finds in insurance textbooks. The assumptions on endowments, consumption sets and preferences are quite standard.

This economy is completely described by $\mathcal{E} := (L, (U_i, e_i)_{i \in I})$.

3. INSURANCE IN THE DYNAMIC SETTING

Our economy in discrete time, \mathcal{E} , is a very general yet well-behaved economy. It has a Walrasian equilibrium in state-contingent commodities satisfying the two basic welfare theorems (Bewley(1972)). For any Walrasian equilibrium, the equilibrium price functional can be expressed as a function of the aggregate endowment, i.e. there is a representative agent; and, agents’ optimal consumption allocations are a function of the aggregate endowment (Huang(1987)). Furthermore, one can apply Duffie and Huang’s(1985, DH) result to show that the Walrasian equilibrium can be decentralized as a Radner equilibrium with as many long-lived risky assets as the martingale dimension of $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in \mathcal{T}}, P)$, which in our case is at most K , the dimension of the vector process \tilde{N} .

Ellickson and Penalva(1997) assumed that there were enough assets for markets to be complete. We want to use the structure provided by insurance markets and inquire deeper. In particular, we want to address the issue of effective completeness: does one really need K risky assets (plus a bond) or can one use fewer; and how does an agent trade in an insurance context: does an agent need to use all available assets or only a subset, and how does she use them.

3.1. RISK

In insurance circles we find: “risk is defined as uncertainty concerning the occurrence of a loss” [Rejda(1995)]. This is particularly important

in the context of actuarial science where a great deal of research has gone into analyzing the properties of risk in its two dimensions of (I) the probability of the occurrence of a loss, and (II) the magnitude of that loss when it occurs. Economic theory, on the other hand, has stayed away from this conception of risk and focused primarily on uncertainty, to which we owe the global scope of the economic theory of risk which applies to finance as well as insurance. On the other hand, we find ourselves drawn back to the insurance definition. As we will see, the number and types of assets traded in insurance markets depends on the classification of the risks in the economy using the insurance definition, and not preferences and endowments.

As the different types of risks can considerably cloud the discussion, we will start by assuming there is a single and simple type of risk in the economy and establish our results in that context. Then we will allow for a wider variety of risks. The type of risk in the economy is similar to that in classic insurance theory papers (Malinvaud, Arrow, Borch, ...), i.e. for every $i \in I$ the following holds:

Assumption 4.a. The endowment of agent $i \in I$, $e_i(t) \in \mathcal{F}_t$ $t \in \mathbf{T}$, is described by a fixed quantity, $w_i \in \mathbf{R}$, at each consumption date $t \in \mathbf{T}$, which is subject to potential losses all of which are:

- (1) Of the same magnitude, L ,
- (2) With arrival distribution driven by an independent Poisson process with parameter, λ .

To ensure that nobody's endowment becomes negative or that $e \leq 0$ (Assumption 3 holds), we impose the global feasibility conditions: $L \geq \min_i w_i$ and for at least one agent $w_i > L$, and, the Poisson process is truncated so that each agent can have at most one accident between t and $t + 1$.

In terms of our MPP, $\tilde{N}(t)$, this means that we have to track the accidents suffered by each agent, so we let $k = i$: the process $N_i(s)$ counts how many accidents agent i has had up to date $s \in \mathcal{T}$. Our restriction on losses means for any $t \in \mathbf{T}$, $N_i(t + 1) - N_i(t) \in \{0, 1\}$.

3.2. Insurance Markets

The first thing to note is that in our finite agent economy, this risk has both an agent-specific component, in that each lost unit of consumption is a unit of consumption taken away from some specific agent's endowment, and an aggregate component, in that it also represents one less unit of consumption available in the whole economy.

We need to define formally what it means for an asset to depend on what happens to an agent. We have defined agent i 's risk using the

process N_i hence define the filtration generated by this process: let $\mathbf{F}_i = ((\mathcal{F}_t^i)_{t \in \mathcal{T}})$, where $\mathcal{F}_t^i = \sigma(N_i(s), s \leq t)$. The filtration \mathbf{F}_i describes how information about the process N_i is made known over time.

DEFINITION 3.1. The set of possible insurance contracts on agent i is the set of real bundles⁹ which are measurable on the filtration \mathbf{F}_i at consumption dates, $t \in \mathbf{T}$.

Which means that an insurance contract on agent i pays at date t conditional on what has happened to agent i up to that date.

A key ingredient of modern financial markets is stock exchanges. We want to include them in the usual, highly stylized way, by allowing agents to trade assets continuously. Stock exchanges deal with claims on companies, in our case, insurance companies. There is no explicit mention of companies in the model but they can be easily incorporated as bundles of assets and liabilities. These would be held by agents in the economy and stock exchanges provide the forum in which to trade them. As we have efficient and frictionless financial markets, all of the firm's unique risk will be diversified away so that agents only hold fully diversified portfolios. These portfolios are modelled as contracts specifying payments conditional on aggregate events (i.e. the number of earthquakes with the relevant consumption dates they affect) and we refer to them as mutual funds (although they could be fully diversified portfolios of insurance company stocks, derivatives, reinsurance contracts, ...).

In order to define these securities formally, as we have done with insurance contracts, we need to determine the relevant processes describing the aggregate endowment, which in our simple case is the total number of accidents, $N(t)$. As done with insurance contracts, define the filtrations over trading dates: $\mathbf{F}_e = ((\mathcal{F}_t^e)_{t \in \mathcal{T}})$, where $\mathcal{F}_t^e = \sigma(N(s), s \leq t)$ is all the information known about the process N up to date t . Then:

DEFINITION 3.2. The set of mutual funds is the set of real consumption bundles which are measurable on the filtration \mathbf{F}_e .

We proposed that our model would include previous results as special cases. For that we need to introduce what CCW refer to as mutual insurance contracts, which require an additional definition. Construct the joint filtration $\mathbf{F}_{i \vee e}$, where $\mathcal{F}_s^{i \vee e} = \sigma(\mathcal{F}_s^i \vee \mathcal{F}_s^e)$ for $s \in \mathcal{T}$. This is the filtration generated by information on e_i and e .

⁹I am working with real, not nominal assets, although as long as markets are complete the difference is not important.

DEFINITION 3.3. The set of possible mutual insurance contracts on agent i is the set of real bundles which are measurable on the filtration \mathbf{F}_{iVe} .

It will be useful to have a risk-free asset to transfer consumption from one date to the next. For simplicity and without loss of generality I will assume that there exists a sequence of zero coupon bonds that are issued at every non-terminal consumption date $t \neq T$ and pay one unit of consumption at date $t + 1$ for sure. This sequence of bonds I will refer to as a riskless bond, and count as a single asset eventhough it clearly is not. Nevertheless, this is equivalent to the standard normalization of prices in asset pricing models and equivalent to a single additional asset (for example, see Merton(1990)).

As far as asset prices are concerned, we use the standard description used in financial markets, namely, for any asset described by dividend process $(d(t))_{t \in \mathbf{T}}$, its price is a semi-martingale on the space $(\Omega, \mathcal{F}, (\mathcal{F}_s)_{s \in \mathbf{T}}, P)$. Prices from the real economy will be extended to financial markets by extending Q to \mathcal{T} using a procedure similar to that used in Penalva(1997) (see Lemma A.3 in the Appendix).

4. EFFICIENT INSURANCE MARKETS

The question to be solved is: when is it possible to specify what type of concrete contracts will be efficient? In particular: when can consumers attain Pareto optimal consumption allocations in an economy where they can only trade in the type of restricted contracts described above (insurance contracts, mutual funds, and a bond)? How many of these assets will they need to do so? What will their informational needs be? And, what will their trading behaviour look like?

From the previous literature we can either look at DH, and say agents need to trade in K long-lived assets (in this case $K = n$), or following CCW, one would need to trade in “mutual insurance” contracts for the agent-specific risk and Arrow securities for the aggregate endowment. As CCW does not allow for dynamic trading, we could expect the mutual funds to replace the Arrow securities and that one would need fewer ones, but we cannot tell ex-ante what would replace mutual insurance.

As there are n private insurance contracts DH tell us that in general markets will be complete. Unfortunately, the ensuing Radner equilibrium would imply that agents are buying and selling other people’s private insurance contracts. It is as if each of us had to buy and sell insurance on the milkman, the pizza guy, a Wall Street broker, ... - not a very believable state of affairs! What we have found is that if one allows

trading in mutual funds, as defined above, such complex trading is not necessary. In particular, let Assumption 4.a. hold for all $i \in I$ together with the feasibility conditions on \tilde{N} , and refer to this economy as “ \mathcal{E} with a single risk”. Then:

THEOREM 4.1. *The economy \mathcal{E} with a single risk requires at most n insurance contracts, a single mutual fund and a riskless bond to decentralize the Pareto optimal allocation from any Walrasian equilibrium. Every agent, $i \in I$, needs at most to trade only in her own private insurance contract, the mutual fund and the bond.*

This means that agents do not need to buy and sell each other’s insurance contracts. Each agent purchases insurance on herself from the market, and trades in a mutual fund which in our case could be a diversified portfolio of shares in insurance companies. Insurance and the fund’s prices are such that one’s demand for insurance is exactly covered by the capital provided via the stock exchange to insurance companies. Hence markets clear, even if there is aggregate risk (as in CCW) but in contrast with previous results, the mutual fund is itself a redundant asset - naturally, as it is the sum of all the individual’s private insurance contracts.

Note two additional points: one, we have allowed agent heterogeneity in preferences and endowment (w_i). The restriction we have imposed is only on risk (the magnitude and probabilities associated with losses). Two, the type of assets used are insurance contracts of the type one is used to seeing in practice: if you lose x I will pay you y - they depend only on what happens to the agent and do not make statements about what happens in the economy as a whole.

5. GENERALIZED RISK AND INSURABILITY

In the introduction we have claimed that the above result implies that insurability extends beyond there being a large number of identical agents. It should be clear that we have already shown that in the sense: (a) the above result is true for any n , in particular, it is true for a small n , say 10 or 15 agents¹⁰; and, (b) we have allowed agent heterogeneity in terms of preferences, (β_i, u_i) , and endowments, w_i . Nevertheless, we can go still further - beyond identical and independent risks.

5.1. Modelling Risk

¹⁰Although in such cases the Walrasian, price-taking behaviour may not be considered generically valid but this an issue entirely outside the scope of this paper.

In the abstract model there is no specific notion of a risk, only uncertainty described by \tilde{N} . In our single risk case, the link was: a change in $N_i(s)$ at $s \in (t, t + 1]$ implies that agent i 's endowment at date $t + 1$ goes from w_i to $w_i - L$. In order to generalize the potential risks we want to consider we need to formally establish this relationship between the processes describing uncertainty, \tilde{N} , and risk - the consequences on economic variables. We proceed by first defining when the underlying uncertainty generates risk and how to classify risks.

The underlying uncertainty, $(N_k(t))_{k=1}^K$, is a vector process whose elements change by a fixed amount (in steps of one) every time there is a jump¹¹. If we are being very strict about interpretation, the jump cannot mean a direct loss of endowment as it occurs at times when endowments are not defined, yet they can be interpreted as announcements of future endowment losses. The marks of the underlying process describe what the content of the announcement is. Potentially, these announcements could be quite varied and could have very diverse consequences: it could mean a single agent loses a fixed amount at the next consumption date (as it did above), it could mean that ten different agents lose ten different amounts at the next consumption date, or that an agent loses 20 units this consumption period and 30 all periods thereafter, The first thing is to establish what the risk associated with announcements of a given type, i.e. given an announcement indexed by $k \in \{1, \dots, K\}$, what are its economic consequences.

Assume that uncertainty is non-anticipated by real economic factors (preferences, endowments and agents' information at date t are all measurable on \mathcal{F}_t). Given an economic process, x , i.e. a process defined at consumption dates, \mathbf{T} (such as agent i 's endowment), we want to determine the processes in \tilde{N} , which affect x . For this we need to define all events that are relevant for x , i.e. the filtration $\mathbf{F}_x = (\mathcal{F}_t^x)_{t \in \mathbf{T}}$, where $\mathcal{F}_t^x = \sigma(x(s), s = 0, \dots, t)$ - the careful reader would have seen the parallel to \mathbf{F}_i and \mathbf{F}_e above; the reason should become obvious as we proceed.

Given the MPP, \tilde{N} , consider the following thought experiment: fix a point in time $s \in \mathcal{T}$ and a process $N_k(t)$, $k \in \{1, \dots, K\}$. Consider deleting all the jumps of this process at this point in time, i.e. impose $\forall \omega \in \Omega \Delta N_k(s) = 0$. Then one can distinguish between the history up to time t along path ω , $\omega(t)$, and the same history without $\Delta N_k(s) = 1$, $\omega(t) - (s, k)$. Naturally, $\omega(t) = \omega(t) - (s, k)$ for all $t < s$ and for all $\omega \in \Omega$

¹¹In actuarial science and the general theory of MPPs one allows for the sizes of the jumps to belong to more general spaces than \mathbf{R}^K . In our context such generalization would cause substantial problems for market completeness (finding a unique equivalent martingale measure), which is a generally recognized problem of such models.

such that $\Delta N_k(s) = 0$ before the deletion operation. Then, $\omega(t) - (s, k)$ could define a (possibly different) history for \tilde{N} up to time t , and we will assume that $\omega(t) - (s, k)$ is always possible¹². Recall that a history is also an event in \mathcal{F} so that we can speak of probabilities conditional on that event in the usual way. Hence, define:

DEFINITION 5.1. The process $N_k(t)$, $k \in \{1, \dots, K\}$ is a *generator of risk for process x* iff there exists an $s \in \mathcal{T}$, $\omega(t)$ $t \in \mathbf{T}$, $n \in \mathbf{N}$ and $A \in \mathcal{F}_{t+n}^x$ such that $P(A|\omega(t)) \neq P(A|\omega(t) - (s, k))$. We also say that N_k generates risk for process x . Let $Y_x(t)$ be the vector of all processes that generate risk for process x and refer to it as x 's risk process. Naturally, \mathbf{F}_x will be filtration generated by Y_x 's internal history¹³.

This definition determines the first part of risk, the probability of a loss: whenever a process ¹⁴ $N_k(s)$ alters the future probabilities of process x it generates risk for that process. It also allows one to identify the effects of the announcement as opposed to other changes that could be taking place due to other non-uncertain factors in a more general version of the model, such as income growth, retirement, etc. Let Y_i denote the risk process which arises from all the processes that are generators of risk for agent i 's endowment, $(e_i(t))_{t \in \mathbf{T}}$, and Y_e that arising from the generators of risk for the aggregate endowment, $(e(t))_{t \in \mathbf{T}}$. \mathbf{F}_i and \mathbf{F}_e are the corresponding filtrations, and $\mathbf{F}_{i \vee e}$ the filtration generated by both risks. Under Assumption 4.a, for example, $Y_i = N_i$ and $Y_e = N$.

Recall the definitions above for x , and $\omega(t)$. Define a more general history (a larger event) $\omega_{k \leftrightarrow k'}(t)$ as all the possible histories obtained by switching the indices between processes k and k' . That is, a history of \tilde{N} up to time $t \in \mathcal{T}$, $\omega(t)$ includes a (possibly empty) set of jump times

¹²For all (s, k) such that there exists $\omega' \in \Omega$ such that $(s, k) \in \tilde{N}(\omega')$, there exists $\omega'' \in \Omega$, such that $\omega'' = \omega' - (s, k)$. This implies that all counterfactuals are well-defined.

¹³This means that $Y_x(t)$ is what is left after deleting from \tilde{N} all the processes that do not generate risk for process x . Also, note this extends \mathbf{F}_x from \mathbf{T} to \mathcal{T} .

¹⁴Note we are imposing that a risk be described by an entire process and do not look for a more microscopic definition of risk that could involve jumps individually. One could potentially construct a single process that describes different risks, for example: $\tilde{N}(t) = N_1(t)$; risk is given by: the first jump implies agent i loses L and the second jump implies agent j loses L' . We rule these out not because we think such processes do not exist but because we believe that in such cases the modelling strategy should be different. In particular, one could maintain a direct relationship between the description of the underlying agent-specific uncertainty, \tilde{N} , and the risks it generates by constructing the model as follows: let $\tilde{N} = (N_1(t), N_2(t))$ be two-dimensional (one process for each agent) and make the hazard of the process for the second agent, $N_2(t)$, depend on the realizations of the first, $N_1(t)$. The new model would have $k(1) = 1$, $Y_1 = N_1$ and $k(2) = 2$, $Y_2 = (N_1, N_2)$. For agent 2, there are two types of risks, N_1 which puts her at risk, and N_2 which implies the realization of a loss. This makes the relationship between the risks and uncertainty transparent by construction and risks would be suitable for the type of analysis we propose.

for both processes $\{\tau_k\}_{\omega(t)}$ and $\{\tau_{k'}\}_{\omega(t)}$. The history $\omega_{k \leftrightarrow k'}$ is the set of all histories $\omega'(t)$ such that all $N_m \neq N_k, N_{k'}$, have the same number and timing of jumps as along $\omega(t)$ and $\{\tau_k\}_{\omega'(t)} \cup \{\tau_{k'}\}_{\omega'(t)} = \{\tau_k\}_{\omega(t)} \cup \{\tau_{k'}\}_{\omega(t)}$. With this define the equivalence between processes as:

DEFINITION 5.2. A process, $N_k(s)$, is *x-equivalent* to process $N_{k'}(s)$ iff for all $t \in \mathbf{T}$, $A \in \mathcal{F}_t^x$ $P(A|\omega(t)) = P(A|\omega_{k \leftrightarrow k'}(t))$.

This definition requires that substituting the index of two processes does not alter the future probabilities for process x . Note that our definition allow the inclusion of risks that do not directly imply losses but just changes in probabilities over future losses - hence including risk factors as well as loss announcements. For the major part of the analysis we will consider constant magnitudes and consider different risks in terms of their effect on the arrival of losses. We postpone introducing variable risk magnitudes for a single agent till Section 7, where we will see that they just increase the number of assets needed to optimally deal with risk.

Nevertheless, our notion of equivalence includes both the arrival and the magnitude, and hence can be used to classify risks in general. When we defined Y_x we used all the counting processes that generate risk for x . This number can be reduced, in particular, the following is true:

LEMMA 5.1. *If the two processes $N_k, N_j, k, j \in \{1, \dots, K\}$, are x-equivalent they can be summarized in \mathbf{F}_x by a the process $N_l = N_k + N_j$.*

Two things to note are: (1) we have allowed different classifications for aggregate and agent-specific risk. The reason is that the classification of risks we propose above could be quite different, depending on whether one was looking at aggregate or agent-specific risks, and this has substantial consequences for the number and types of assets considered (as we will see below). (2) Agent i 's risk is being referred to as agent-specific and not as idiosyncratic risk. We have seen above that this is an important aspect, specially in finite agent economies, because the subject of insurance is not uncertainty but agent-specific risk and that risk will have two components: a diversifiable one and an undiversifiable one. This agent-specific risk becomes idiosyncratic only in the context where the law of large numbers applies; formally, in continuum economies¹⁵ such as those analyzed by Al-Najjar(1995).

¹⁵There are some models using an uncountable infinity of agents but they are the exception.

5.2. Introducing Some Symmetry

The insurance literature, in particular Malinvaud(1972,73) and CCW(1996), suggest that an important aspect of well-functioning insurance markets is the existence of symmetry. They impose symmetry in terms of agents' endowments and preferences, but we have seen above that that kind of symmetry is not necessary. Nevertheless, it seems some symmetry is required in order to have a certain degree of parsimony in the number and types of contracts (insurance and otherwise) needed to decentralize a Walrasian equilibrium. Clearly from the above discussion, that symmetry has to come from the fundamental uncertainty and its effect on the economy, i.e. from the risk processes. We propose the symmetry assumptions considered by Malinvaud and CCW, plus the general class of Markovian risks. Recall we are assuming a fixed loss magnitudes, hence the assumptions apply to the properties of the distribution of loss announcements' arrivals. Later, we will allow for different magnitudes of losses and for losses to have consequences that persist over time.

5.2.1. Independence

The first, classic form of symmetry is assuming that agents' risks are risks of independent losses which are felt by one agent at a time. Formally, this condition is formalized as agent i 's endowment satisfying:

Assumption 4.b. (Independent Private Risk: InPR). The endowment is described by a fixed quantity, w_i , at each consumption date t , which is subject to potential losses:

- (1) Of a fixed magnitude L ,
- (2) With arrival distribution described by $k'(i) \subset \{1, \dots, K\}$, such that the processes $(N_{k'(i)})$ is the vector of processes that generate risk for e_i and (i) is equivalent to a single point process¹⁶ ($k(i) = 1$): $Y_i = (N_i)$, which admits an intensity, λ_i , measurable on the information generated by $N_i(t)$; and (ii) for any $j \in I$, $j \neq i$, $k'(i) \cap k'(j) = \emptyset$.

These processes are independent because (i) their hazard functions do not depend on information outside of that provided by the N_i process itself, and (ii) the process N_i does not affect any other agent. This first case already generalizes beyond time independence (the time homogenous Poisson case considered in the section above). It allows for non-time homogenous Poisson, i.e. $\lambda_i = \lambda_i(t)$, and for self-exciting processes, $\lambda_i = \lambda_i(\tau - t)$, where τ denotes the last jump time of the N_i

¹⁶The use of the index sets $k'(i)$ in this case is not necessary as by assumption $Y_i = N_i$ but it is important to keep it in order to make this assumption comparable to others that follow.

process. Basically, we allow for any hazard rate that depends on clock time t and the past of the process N_i .

We have distinguished between two index sets: $k'(i)$ and $k(i)$. This distinction is very important and arises from the dual definition of Y_i . In general the $k'(i)$ processes obtained from marks of the MPP \tilde{N} are equivalent to a smaller number of processes when defining Y_i denoted by the second index: the fundamental processes in \tilde{N} are indexed by $k'(i)$ and the more succinct description of the processes that affect agent i are indexed by $k(i)$. In a couple of paragraphs we will see an earthquake example that should help further clarify this distinction.

5.2.2. *Locally Independent Private Risks*

One could generalize the independence condition by assuming agent i 's endowment satisfies a weak version of independence:

Assumption 4.c. (Locally Independent Private Risks: LInPR). The endowment is described by a fixed quantity, w_i , at each consumption date t , which is subject to potential losses:

- (1) Of a fixed magnitude L ,
- (2) With arrival distribution described by $k'(i) \subset \{1, \dots, K\}$, such that the processes $(N_{k'(i)})$ is the vector of processes that generate risk for e_i and is equivalent to a single point process ($k(i) = 1$): $Y_i = (N_i)$, which admits an intensity, λ_i , measurable on the information generated by $N_i(t)$.

This condition implies independence locally - among the processes whose indices describe the risk for agent i - but by allowing an agent's risk process to share indices with other agents, independence across agents may fail. In particular, this allows there to be marks generating risk for a group of agents and hence allows for correlated losses in general and catastrophic losses in particular. Let us illustrate this (and the distinction between $k'(i)$ and $k(i)$) via an earthquake example:

An earthquake can be modelled as follows: suppose $n = 100$ and each agent has a fixed endowment of the consumption good each period $w_i = 5,000$. The risk of an earthquake is modelled as a massive endowment loss, say of 100,000 units of consumption. Suppose there can be at most one earthquake between two consumption dates (to avoid problems of negative endowments) and the arrival time of the earthquake is given by the hazard rate¹⁷ $\lambda = \lambda(t)$.

¹⁷This implies that if $\Delta N(s) = 1$, $s \in (t, t + 1]$, then for $s' \in (s, t + 1]$, $\lambda(s') = 0$.

All agents face the same risk: given an earthquake has occurred, they all have the same probability of loss: 20 names are picked at random and those 20 lose 5,000 units of consumption at the next consumption date.

The appropriate marked point process, \tilde{N} , contains $100!/(20!80!) > 10^{20}$ marks corresponding to independent Poisson processes: one for each possible combination of the 20 names. For any agent i , $k'(i) = 99!/(19!80!)$, those marks that contain i 's name and hence generate risk for agent i . But as the risk for i is the same (the loss of 5,000 units of consumption) they are all equivalent so that $k(i) = 1$, and the corresponding hazard, λ_i is 1/20th the one for an earthquake. Obviously, agents' losses are not independent accross agents, but the losses of any one list of 20 is independent of that of any other different list of 20 agents¹⁸.

Note that the martingale dimension of \tilde{N} is $100!/(20!80!)$ but the definition of equivalent risks above allows us to exploit symmetry and will have important consequences for the number of assets needed to effectively complete the market, as we will see below.

5.2.3. Exchangeability

CCW extend risks to include exchangeability (a notion we owe to Haag (1924) and De Finetti (1937)) within groups of ex-ante identical agents. Generally, events A_1, \dots, A_n from a probability space (Ω, \mathcal{F}, P) are exchangeable (interchangeable) if for all choices of $1 \leq i_1 < \dots < i_j \leq n$ and all $1 \leq j \leq n$, $P(A_{i_1} \cdot \dots \cdot A_{i_j}) = p_j$ - Chow and Teicher (1988, p.33).

CCW considered a static context and a joint distribution of the agent's private risk and the economy's aggregate endowment of the form:

$$P(Y_i = y | Y_e = x) = P(Y_j = y | Y_e = x) = \frac{\pi(y, x)}{\sum_{y'} \pi(y', x)},$$

for all agents i, j of the same type. They use exchangeability on the indices representing the endowment outcomes of agents of the same type. Dynamically, we can interpret their assumption as allowing the probability of agent's risks to depend on what happens to the aggregate endowment¹⁹.

Recall the definitions, Y_e , and the corresponding filtration, \mathbf{F}_e . Let $k(e)$ denote the index set of the process that make up Y_e and refer to the

¹⁸Except for the period between an earthquake and the consumption date just following it.

¹⁹Note that by restricting ourselves to the arrival process we are not considering different magnitudes so that this assumption alone is not a generalization of that made by CCW. Their results will only be subsumed in this framework after Section 7.

corresponding processes, $((N_j)_{j \in k(e)})$ as the factors generating aggregate risk, $N_{k(e)}$. Then, let agent i 's endowment satisfy:

Assumption 4.d. (Exchangeable Private Risks: ExPR). The endowment is described by a fixed quantity, w_i , at each consumption date t , which is subject to potential losses:

- (1) Of a fixed magnitude L ,
- (2) With arrival distribution described by $k'(i) \subset \{1, \dots, K\}$, such that the processes $(N_{k'(i)})$ is the vector of processes that generate risk for e_i and are equivalent to a single point process, $Y_i = N_i$, which admits an intensity, λ_i , measurable on the information generated by $N_i(t)$ and the factors generating aggregate risk, $N_{k(e)}$.

5.2.4. Markov and Semi-markov Processes

Another very general class of risks one could consider is that generated by the introduction of a state space and allowing the hazards of the process describing agents' losses to depend on the state at date t . By allowing a random process to describe transitions accross states we have a Markov process if the transition rates depend only on the current state and clock time, and a Semi-markov process if they also depend on the time spent at the current state. Let us just consider a finite state space²⁰ \mathcal{S} with $m(e)$ states and transition rates $\alpha_{hj}(t)$. For purposes of exposition we will speak only of Markov processes but the results apply to Semi-markov ones as well. The MPP describing the current state is denoted $M(t)$. Then, let agent i 's endowment satisfy:

Assumption 4.e. (Markovian Private Risks: MaPR). The endowment is described by a fixed quantity, w_i , at each consumption date t , which is subject to potential losses:

- (1) Of a fixed magnitude L ,
- (2) With arrival distribution described by $k'(i) \subset \{1, \dots, K\}$, such that the processes $(N_{k'(i)}, M(t))$ is the vector of processes that generate risk for e_i and are equivalent to the process $M(t)$ plus a single point process, N_i , which admits an intensity, λ_i , measurable on the information generated by $N_i(t)$ and $M(t)$.

Moreover, if $M(t)$ generates risk for agent i then it also generates risk for the aggregate endowment, e .

²⁰Naturally expand $(\Omega, \mathcal{F}, (\mathcal{F}_t))$ appropriately and P be the measure implied by λ and α on the new (Ω, \mathcal{F}) .

The latter part of the assumption to avoid the technical possibility that individual risks be Markovian but as you sum over all agents the Markovian effects wash out (which would only add further cases to consider).

In the same way we have allowed for an additional Markov process, we could add more risks by allowing for regression variables (multiplicative hazards, ...) and other factors. We have already considered the independent case (allowing for correlation across agents), the exchangeable case (to compare with CCW) and the very general Markov case. Further generalizations would take us too far from the aim of this paper: analyzing efficient insurance markets. Our results naturally extend to those additional cases.

6. INSURABILITY AND TRADING BEHAVIOUR

We now look at what is the number and type of assets required to optimally diversify risk. We have considered a very large number of possible types of risks. Fortunately, all those possibilities were taken into account in the definition of individual risk and risk equivalence. In order to prove our results we apply the powerful machinery of martingale theory.

We are primarily interested in establishing the number of insurance contracts needed to deal with agent-specific risks. When defining risks above we noted that the classification of aggregate and agent-specific risks could be quite different and will have important effects on the number and type of assets traded. In fact, the number and type of assets depends very strongly on whether the following condition on private risks holds: Let $N_{ij}(t)$ be any process that generates risk for e_i , $N_{ij} \in Y_i$. Let B_{ij} denote the indices of $\{1, 2, \dots, K\}$ that added together equal N_{ij} .

Local-Global Risk Condition (LoGRC). For all $N_{ij} \in Y_i$, there exists a process N_k that generates risk for the aggregate endowment, e , such that B_{ij} is a subset of the indices of $\{1, 2, \dots, K\}$ that added together equal N_k .

This definition seems a little complex but the intuition is relatively simple: if this condition holds, then any process that describes risk for agent i also describes (a single) risk for the aggregate process. In the earthquake example we saw a case where there was a large number of indices that were equivalent to a single risk for that agent, losing 5,000 units of consumption, and generated a single (yet different) risk for the aggregate economy: a loss of 100,000: all risks for agent i of the same type were also of the same type for the aggregate economy. On the

other hand we could have a case where the LoGRC condition does not hold. Consider the following example: an economy with three agents, A, B, C living on the Great Plains. All three are subject to a loss of 5,000 consumption units from a tornado. The arrival of a tornado is described by a Poisson process. As the tornado proceeds across the plains it can affect any subset of the three agents (each agent has a fixed independent probability of being hit conditional on a tornado). This is modelled by an 8-dimensional point process whose indices represent the possible combinations of agents the tornado affects. The indices are represented by the letters H, T such that HTH stands for agents A and C are hit and B is spared. From agent A's point of view, the indices (HTT, HHT, HTH, HHH) are equivalent but clearly they are not so from the aggregate endowment point of view. The agent only cares whether she got hit or not but the economy cares whether it was a small, medium or large tornado (in terms of losses). We will now see what the effect of this condition is. We start by assuming that the magnitude of an agent's loss (per announcement) is constant and equal to L_i - i.e. we allow for loss heterogeneity across agents.

THEOREM 6.1. *The economy \mathcal{E} together with individual risks such that for every agent $i \in I$, e_i satisfies one of assumptions 4b-e and the LoGRC condition, requires at most n insurance contracts, $k(e)$ mutual funds and a riskless bond to decentralize the Pareto optimal allocation from any Walrasian equilibrium. Every agent, $i \in I$, needs at most to trade only in her own private insurance contract, the $k(e)$ mutual funds and the bond.*

If on the other hand, condition LoGRC does not hold we have:

THEOREM 6.2. *The economy \mathcal{E} together with individual risks such that for every agent $i \in I$, e_i satisfies one of assumptions 4c-e, requires at most $\min\{K, 2nk(e)\}$ mutual insurance contracts and a riskless bond to decentralize the Pareto optimal allocation from any Walrasian equilibrium. Every agent, $i \in I$, needs at most to trade in $\min\{K, 2k(e)\}$ mutual insurance contracts and the bond.*

Note that assumption 4.b, independence across agents, implies the LoGRC condition. The intuition for the differences in the two theorems is relatively straight-forward. If condition LoGRC is not fulfilled then, paraphrasing Kreps(1982), the information flow is not nice enough. The LoGRC condition implies that a single piece of information revealed on an agent's loss is coupled with a single piece of information on the aggregate endowment. In the tornado example, an agent's loss could be

associated with multiple different aggregate losses and hence you need extra assets to ensure that you can deal with whatever tornado happens to occur.

We are defining risks as insurable if they can be optimally diversified, i.e. dependent on the asset structure of the economy. Whether a certain class of accidents is insurable or not depends on two things: one, the number of risk management tools available to deal with incoming loss information; and two, the suitability of those assets to deal with the characteristics of the risk (in our case, mutual funds or mutual insurance contracts).

If we look at the economy to try to deduce what kind of risks are being insured, the lack of the widespread use of mutual insurance contracts (as we have defined them) indicates that either risks are primarily of the LoGRC type, or the incompleteness from not having mutual insurance is not sufficiently important. On the other hand, if the LoGRC condition does not hold, agents cannot separate the problem of dealing with the aggregate risk and individual risk and they require a much larger set of assets. If these are not available, it could show up in practice as distorted optimal coverage at the individual level. Counting assets at the industry level may be misleading because the number of assets needed by insurance companies is much smaller than the is needed by the agent. This would lead one to conclude that markets are complete when there are enough assets to deal with aggregate risk but, at the individual agent level, agents face incomplete markets.

As for the role of financial innovations in insurance markets, it seems that the creation of new assets such as catastrophe bonds and actively traded insurance derivatives are fulfilling a latent need for a more diversified range of risk management tools than those provided by traditional reinsurance. This should be improving the diversification of insurance risks and expanding the range of insured risks (and hence effectively insurable risks). As these tools are used more efficiently one should observe a call for lower and/or more flexible capital requirements by insurance and reinsurance companies (and a greater flight towards more flexible legal environments like Bermuda while those calls are not met by national legislatures).

7. EXTENSIONS AND DISCUSSION

7.1. Spanning Aggregate Risk

Let us return to more technical issues: up until now we have paid little attention to the value of $k(e)$. One of the reasons is that by applying Lemma 5.1 to e one can directly obtain how many mutual funds one needs. Assumptions 4b-e allows us to say a little more.

In the case of independent risks, InPR, $k(e)$ depends on the number of loss magnitudes and the number of different hazard functions (two different loss magnitudes will imply different aggregate endowments). In particular, one should note that when agents' hazards are time varying and arise from independent self-exciting processes it can be impossible to add them, even if they are identical. This is because the hazard of the sum can be different depending on the identity of the last jump (which would automatically rule out e -equivalence).

If we allow agents to share risks, as in the LInPR, then it very much depends on how risks combine. Recall that the tornado example used above has LInPR risks. We can clearly see that such risks can quickly turn into a spanning nightmare. Nevertheless, working on the model used in the earthquake example, one could come up with a weaker version of the LoGRC condition under which one can still have a manageable number and type of assets and keep markets effectively complete.

As for exchangeability, the ExPR condition adds dependence of private risks on the aggregate risk. Apart from the additional assets required for heterogeneous losses which applies to all conditions, this assumption could be a way to reintroduce self-exciting processes as a more manageable risk in as far as the exciting comes from the aggregate process. Making hazard changes depend on aggregate information would reduce the complexity of the aggregated process and allow for fewer aggregate assets (a smaller $k(e)$).

The introduction of an additional source of uncertainty, the Markov process, necessarily implies the need for more assets. Nevertheless, it is important to have in mind that it is not the total number of states of the Markov process that determines how many new assets are needed, but the properties of the transition function. This is because the transition function determines how 'nicely' information flows - the fewer the possible transitions out of each state the fewer the number of assets needed. This is particularly obvious in the Brownian motion case, which has an uncountably infinite state space but a uni-dimensional martingale space.

7.2. Persistent Losses

We propose two ways to extend the description of loss magnitudes. Let us start by allowing losses to last for several (possibly an infinite number of) periods. Also, allow for the losses generated by an announcement to vary depending on the time interval they occurred.

Consider an announcement made at s that affects agent i , $\Delta N_i(s) = 1$, $s \in (t, t + 1]$. The first step was to consider that agent i would lose L units of consumption at date $t + 1$. Now, define the loss function $L_j(t, m) : \mathbf{T} \times \mathbf{N} \rightarrow \mathbf{R}$. This function tells us that after $\Delta N_i(s) = 1$ agent i will lose $L(t, 1)$ at $t + 1$, $L(t, 2)$ at $t + 2$, ...

Allowing these changes does not alter the above results in the following sense:

LEMMA 7.1. *For any given risk affecting an agent, i , whether the magnitude of the risk is given by a single loss at the next period, L , or a loss function $L(t, m)$ does not affect the number or type of assets needed to decentralize the Pareto optimal Walrasian equilibrium. The same is true for the pre-loss endowment e_i which can be any function of clock-time.*

The intuition behind Lemma 7.1 is that if future losses are known to occur with certainty at the time of the announcement, they are equivalent to the loss of their discounted value at the next period, and hence to a fixed loss. Then, any economic variable that changes deterministically with clock-time is known with certainty at any point in time so again its value can be discounted to the appropriate date.

7.2.1. Different Magnitudes

Another way to generalize the loss magnitude is to allow for the realization of the magnitude of the loss, L , to be random. This generalization is very important for insurance practice yet it poses problems for the existence of complete markets. This has already arisen in the literature on asset pricing where it is implicit in the problem of finding a unique martingale measure when one allows for asset prices to have discontinuous jumps (also modelled as MPPs).

We will limit ourselves to a finite number of different loss magnitudes. Let the set describing the possible loss outcomes be denoted by $\mathcal{L} \equiv \{\ell_1, \ell_2, \dots, \ell_m\}$ (wolg these could be m loss functions like the ones in the previous section).

When allowing for random magnitudes of losses generally (in insurance and actuarial circles) people assume loss magnitudes to be independent of the process describing the arrival of the accident. In our context the distribution of loss magnitudes could also depend on the same factors as the risk process, so that one could consider combinations of Assumptions 4.b-e for the magnitude as well as the arrival of the loss. Before analyzing this case we introduce the possibility of there being multiple risks (allowing, for example, auto and health insurance in the same model).

7.2.2. Multiple Risks

We have considered that each agent was exposed to a ‘single’ risk and that there were possible differences in risks between agents; both have primarily affected the aggregate variable $k(e)$. One could also introduce more risks and have different subsets of agents exposed to them. The

subset of agents exposed to different risks could be quite general, overlapping for some risks, non-overlapping for others, ... The catch is that the more heterogeneity one introduces the greater will be the number of assets needed to deal with such risks, but how many?

In general, introducing m new constant loss risks is equivalent (for our purposes) to allowing m different loss magnitudes, as each new risk requires a new process to describe it, both at the individual and the aggregate level. By the definition of \mathbf{F}_i , the number of risks the agent is exposed to is given by the number of processes (after applying equivalence) that make up \mathbf{F}_i . This implies that if each agent is exposed to m_i different risks (on top of the $k(e)$ or $m(e)$ global ones in Assumptions 4d,e) and each of the different risks does not depend on the specific risk of the others, and if the LoGRC condition holds for each of those risks, then Theorem 6.1 can be extended to:

THEOREM 7.1. *The economy \mathcal{E} together with individual risks such that for every agent $i \in I$, e_i is exposed to m_i risks satisfying assumptions 4b-e and the LoGRC condition, requires at most $\sum_i m_i$ insurance contracts, $k(e)$ mutual funds and a riskless bond to decentralize the Pareto optimal allocation from any Walrasian equilibrium. Every agent, $i \in I$, needs at most to trade only in m_i private insurance contracts on herself, the $k(e)$ mutual funds and the bond.*

Naturally, $k(e)$ will be greater in general for an economy with agents exposed to $m_i > 1$ risks. Note also that we allow for agent i 's risks to be correlated but only through aggregate risks, i.e. the $k(e)$ aggregate factors for Assumption 4.d and the $m(e)$ processes defining the Markov states for Assumption 4.e. Theorem 6.2 can be similarly extended:

THEOREM 7.2. *The economy \mathcal{E} together with individual risks such that for every agent $i \in I$, e_i satisfies one of assumptions 4c-e, requires at most $\min\{K, 2k(e)(\sum_i m_i)\}$ mutual insurance contracts and a riskless bond to decentralize the Pareto optimal allocation from any Walrasian equilibrium. Every agent, $i \in I$, needs at most to trade in $\min\{K, 2k(e)m_i\}$ mutual insurance contracts and the bond.*

Allowing for correlations between individual's risks could increase the number of assets needed even further.

8. CONCLUSION

We have presented a very general framework in which to analyze insurance problems in a context that allows for dynamic trading. This includes

general definitions of individual and aggregate risks, private insurance and mutual insurance contracts, and mutual funds. We have considered the independence framework analyzed by Malinvaud and exchangeability as analyzed by CCW. Both those models are clearly included as special cases within our framework: let $\mathbf{T} = \{0, 1\}$, let $\tilde{N}(s) = 0$ for all $s \in \mathcal{T} = [0, 1]$ and $\tilde{N}(1)$ reflect the random vector describing all the risk in the economy²¹. We have extended the work by Ellickson and Penalva(1997) with a deeper analysis of risk, and by looking at the number and types of assets needed to decentralize Pareto optimal Walrasian allocations. We have concluded that the most important determinant of the type and number of assets traded is whether the revelation of information about agent-specific risks is matched with a single aggregate risk, i.e. whether information flows are sufficiently nice.

In terms of the existing literature, we have extended the generalization of Malinvaud's classic analysis of insurance market by CCW in a number of ways: we have allowed for risks that have effects that persist over time, we have allowed for a dynamic description of the arrival of information about risks and have expanded the class of risks in the analysis, and we have allowed for much greater heterogeneity across agents.

In our more general framework we have analyzed the notion of risk and looked at the type and number of assets needed to optimally deal with those risks. We have concluded that the only type of heterogeneity that matters is heterogeneity in the types of risks agents are exposed to (both in terms of the description of the arrival of losses as well as magnitudes); we have shown that allowing dynamic trading reduces the number of assets needed to effectively complete markets (extending the results in DH) and have made a detailed analysis of what the number and types of assets are that will ensure markets are effectively complete. In our analysis we have explained why in the context analyzed by CCW it is not enough to consider standard insurance contracts but one needs to introduce the concept of mutual insurance contracts, as they did.

Nevertheless, I would like to add some caveats that apply to this, as well as CCW's analysis. These results rely strongly on the properties of Walrasian equilibrium. It has been shown elsewhere (Ostroy and Zame(1994), for example) that the use of Walrasian equilibrium may not be justified in a context where markets are not competitive. The specific issue of competitive markets under conditions of uncertainty have not been studied in detail and provide an interesting further direction of study. Possibly, the conditions for competition may revive the classical need for "a large number of exposure units". Also, we have assumed

²¹Strictly speaking this would not be within the above model because it imposes an infinite hazard at $t = 1-$ which is ruled out by assumption, but the intuition still holds.

common knowledge of the information available in the economy, leaving aside problems of moral hazard and adverse selection which are naturally very important in insurance markets. Nevertheless, we find our analysis enriches existing views of insurance markets.

In terms of more practical contributions, this analysis suggests how different types of insurance risks should fare in an economy with access to stock market-type institutions (which allow for dynamic trading) in a context of full information. Dynamic trading, which takes the form of active capital management, is a potential substitute for having large static capital reserves. This would imply that firms following dynamic risk management strategies would be looking to reduce the large capital reserves and asset management restrictions imposed by old-fashioned laws, either via lobbying or via the use of off-shore subsidiaries. Also, as financial markets are liberalized, we should see an increasing use of dynamic risk management. On an industry-wide scale, dynamic risk management has led to changes that range from more finance-like dynamic reserve management to the creation of very active trading markets in insurance derivatives such as the ones in Bermuda and Chicago. Financial innovations such as insurance derivatives, catastrophe bonds, or securitized insurance, add value in the economy in that they allow a more efficient use of insurance capital and reserves.

Our analysis also suggests that insurance practice is fundamentally sound in that it classifies risks according to actuarial parameters and ignores issues of risk preferences and endowments. Insurability is not a matter of agents' preferences or endowments; they determines prices but not whether risks can be efficiently distributed. We have shown that insurable risks are no longer those that are "socially removed by the operation of the law of large numbers" (Malinvaud, 1972).

REFERENCES

1. AASE, K. (1993): "Continuous Trading in an Exchange Economy under Discontinuous Dynamics: A Resolution of the Equity Premium Puzzle," *Scandinavian Journal of Management*, 9:Suppl., S3-S28.
2. AL-NAJJAR, NABIL I. (1995): "Decomposition and Characterization of Risk with a Continuum of Random Variables," *Econometrica*, 63(5), 1195-1224.
3. ARROW, KENNETH J. (1965): "Aspects of the Theory of Risk-Bearing, lecture 3," reprinted in *Collected Papers of Kenneth J. Arrow, Vol V*, (1984) 77-86.
4. ARROW, KENNETH J. AND ROBERT C. LIND (1970): "Uncertainty and the Evaluation of Public Investment Decisions," in *American Economic Review*, 60(3) 364-378.
5. BEWLEY, T. (1972): "Existence of Equilibria in Economies with Infinitely Many Commodities," *Journal of Economic Theory*, 43, 514-540.
6. BORCH, K. (1990): *Economics of Insurance*. Edited and completed by Knut K. Aase and Agnar Sandmo, North-Holland.
7. BRÉMAUD, P. (1981): *Point Processes and Queues: Martingale Dynamics.*, Springer-Verlag, New York.

8. CASS, DAVID, GRACIELA CHICHILNISKY AND HO-MOU WU (1996): "Individual Risk and Mutual Insurance," *Econometrica*, 64, 333–341.
9. CHOW, YUAN SHIH AND HENRY TEICHER (1988): *Probability Theory: Independence, Interchangeability, Martingales*, Second edition. Springer-Verlag.
10. CHRISTENSEN, PETER O., SVEND E. GRAVERSEN, AND KRISTIAN R. MILTERSEN (1996/1999): "Dynamic Spanning in the Consumption-based Capital Asset Pricing Model," *Publications from Department of Management, Odense University 2/1996, revised 1999*.
11. DE FINETTI, B. (1937): "La prévision, ses lois logiques, ses sources subjectives," *Annales de l'Institut Poincaré* 7, pp 1-68, reprinted and translated as "Foresight: Its Logical Laws, Its Subjective Sources," (1964) in *Studies in Subjective Probability*, ed. by H.E.Kyburg Jr. and H.E.Smokler. Wiley, New York.
12. DUFFIE, D. AND C. HUANG (1985): "Implementing Arrow-Debreu equilibria by continuous trading of few long-lived securities," *Econometrica*, 53, 1337–1356.
13. DUFFIE, D. AND WILLIAM R. ZAME (1989): "The Consumption-based Capital Asset Pricing Model," *Econometrica*, 57(6), 1279–1297.
14. Ellickson, B., and J. Penalva Zuasti (1997): "Intertemporal insurance," *Journal of Risk and Insurance* 64(4): 579–598.
15. HAAG, J. (1924): "Sur un probleme general de probabilites et ses diverses applications," *Proc. Inst. Cong. Math., Toronto, 1924, 1928*, pp 629–674.
16. HARRISON, J., AND D. KREPS (1979): "Martingales and arbitrage in multiperiod securities markets." *Journal of Economic Theory*, 20, 381–408.
17. HINDY, A., AND C. HUANG (1993): "Optimal consumption and portfolio rules with durability and local substitution," *Econometrica*, 61,85-121.
18. HUANG, CHI-FU (1987) "An Intertemporal General Equilibrium Asset Pricing Model: The Case of Diffusion Information," *Econometrica*, 55(1), 117–142.
19. KREPS, DAVID M. (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.
20. LAST, G. AND A. BRANDT (1995): *Marked Point Processes on the Real Line: The Dynamic Approach.*, Springer-Verlag. New York.
21. MALINVAUD, E. (1972): "The Allocation of Individual Risks in Large Markets," *Journal of Economic Theory*, 4, 312–328.
22. MALINVAUD, E. (1973): "Markets for an Exchange Economy with Individual Risks," *Econometrica*, 41(3), 383–410.
23. MERTON, ROBERT C. (1990) *Continuous-Time Finance*. Blackwell Publishers. Cambridge, Massachusetts.
24. PENALVA ZUASTI, JOSE S. (1997) *The Theory of Financial Insurance with an Application to Earthquakes and Catastrophe Bonds*, PhD Thesis, UCLA.
25. OSTROY, JOSEPH AND WILLIAM R. ZAME (1994) "Nonatomic Economies and the Boundaries of Perfect Competition," *Econometrica*, 62(3), 593–633.
26. REJDA, GEORGE E. (1995) *Principles of Risk Management and Insurance*, Fifth Edition. HarperCollins.
27. ZHOU, YUQING (1995): "Dynamic Equilibrium with Idiosyncratic Risk," *mimeo*.

APPENDIX

A.1. PROOF OF LEMMA 5.1

Proof.

Define the process \tilde{N}_{k+j} by altering \tilde{N} as follows: delete the elements N_j and N_k . Substitute N_j by the constant zero and N_k by $N_l = N_k + N_j$.

Let us proceed by contradiction. Suppose N_k and N_j are x -equivalent but there exists $s \in \mathcal{T}$, $t \in \mathbf{T}$, $s \leq t$, event $A \in \mathcal{F}_t^x$ and history $\omega(s)$ such that the lemma does not hold. That is, for the corresponding history for the process \tilde{N}_{k+j} , $\hat{\omega}(s) \supset \omega(s)$, $P(A|\hat{\omega}(s)) > P(A|\omega(s))$ (the direction of the inequality is without loss of generality). By definition, $\hat{\omega}(s)$ has the same non- k , j jumps as $\omega(s)$ but $\{\tau_k\}_{\omega(t)}$ and $\{\tau_j\}_{\omega(t)}$ have been substituted by $\{\tau_l\}_{\omega(t)}$, so that $\omega(s) \subset \hat{\omega}(s)$. The condition $P(A|\hat{\omega}(s)) > P(A|\omega(s))$ implies there is another element of $\hat{\omega}(s)$, $\omega'(s)$ such that $P(A|\omega'(s)) > P(A|\omega(s))$. But $\omega'(s)$ differs from $\omega(s)$ in the realizations of N_k and N_j . Yet this implies that $\omega'(s) \in \omega_{k \leftrightarrow j}$ - a contradiction. ■

A.2. PRELIMINARY LEMMAS

In order to simplify exposition of the proofs we first state a few Lemmas which summarize relevant results in the literature which are directly applicable in our context, which appropriate proofs or references.

LEMMA A.1. *There exists a price-allocation pair (π, x) which is a Walrasian equilibrium of economy \mathcal{E}*

Proof. See Bewley(1972). ■

LEMMA A.2. *Given any equilibrium pair (π, x) , there exists functions $g : \mathbf{T} \times \mathbf{R} \rightarrow \mathbf{R}$ and $f_i : \mathbf{T} \times \mathbf{R} \rightarrow \mathbf{R}$ for all $i \in I$ such that:*

(1) *there exists a measure Q whose Radon-Nikodym derivative (pricing kernel) at every $t \in \mathbf{T}$ is given by*

$$E \left[\frac{dQ}{dP} \middle| \mathcal{F}_t \right] = \prod_{s=1}^t \frac{g(s, e(s))}{E_P[g(s, e(s)) | \mathcal{F}_{s-1}]} \quad (\text{A.1})$$

(2) *Agent i 's optimal consumption process, x_i , is of the form*

$$x_i(t) = f_i(t, e(t)). \quad (\text{A.2})$$

(3) There is an interest rate process, $r(t)$, given by

$$1 + r(t) = \frac{g(t-1, e(t-1))}{E_P[g(t, e(t))|\mathcal{F}_{t-1}]}, \quad (\text{A.3})$$

such that for any $y \in L_+$

$$\pi(y) = \sum_{t \in \mathbf{T}} E_Q[y^*(t)], \quad (\text{A.4})$$

where $y^*(t)$ denotes the discounted value of $y(t)$

$$y^*(t) = \gamma(t)y(t); \quad \gamma(t) \equiv \frac{1}{\prod_{s=1}^t 1 + r(s)}.$$

Proof. Huang(1987) establishes that economies such as \mathcal{E} have a representative agent representation, which implies the existence of a function $g' : \mathbf{T} \times \mathbf{R} \rightarrow \mathbf{R}$ such that for any $y \in L_+$, $\pi(y) = \sum_{t \in \mathbf{T}} E_P[g'(t, e(t))y(t)]$. Normalize by setting the value of consumption at date zero to be equal to one and denote the renormalized function by g . Then define r and dQ/dP as is done in equations (A.1) and (A.3). Equation (A.4) follows directly from these redefinitions. That dQ/dP defines a measure comes from the positiveness of π and from discounting, which ensures $E_P[dQ/dP] = 1$ for every $t \in \mathbf{T}$. Part (2) is also demonstrated by Huang(1987) and is generally known as the property of optimal risk sharing. ■

LEMMA A.3. For economy \mathcal{E} , given an interest rate process r and a martingale measure Q defined by the Radon-Nikodym derivative $\nu(t)$ defined on $t \in \mathbf{T}$, there is a P -a.s. extension of Q onto \mathcal{T} defined by the Radon-Nikodym derivative $\xi(s)$, $s \in \mathcal{T}$ defined by $\xi(0) = \nu(0)$, and

$$\xi(s) = E_P[\nu(t+1)|\mathcal{F}_s], \quad s \in (t, t+1], t \in \mathbf{T}.$$

Proof. By definition $\xi(t) = \nu(t)$ for $t \in \mathbf{T}$. Hence, for $s \in (t, t+1]$, $t \in \mathbf{T}$, $\xi(s) = E_P[\xi(t+1)|\mathcal{F}_s]$. Using this and the fact that ν is a change of measure defined on \mathbf{T} , for $t \in \mathbf{T}$, so that $E_P[\nu(t+1)|\mathcal{F}_t] = \nu(t)$, implies that ξ is a P -martingale. As ν is a change of measure it will always be positive and, as conditional expectations preserve that property, so will be ξ . Also, $E_P[\nu(t)] = 1$ for $t \in \mathbf{T}$, hence for $s \in \mathcal{T}$, $s \in (t, t+1]$,

$$E_P[\xi(s)] = E_P[E_P[\xi(s)|\mathcal{F}_t]] = E_P[\xi(t)] = E_P[\nu(t)] = 1.$$

Hence, ξ defines a change of measure on \mathcal{T} that coincides with Q at every $t \in \mathbf{T}$. ■

For $s \in \mathcal{T}$, define the process

$$X_i(s) = E_Q \left[\sum_{t \in \mathbf{T}} \Delta_e x_i^*(t) \middle| \mathcal{F}_s \right] - E_Q \left[\sum_{t \in \mathbf{T}} \Delta_e x_i^*(t) \right]$$

where $\Delta_e x_i^*(t) = x_i^*(t) - e_i^*(t)$, $x_i^*(t) = \gamma(t)x_i(t)$ and $e_i^*(t) = \gamma(t)e_i(t)$. By construction, $(X_i(s))_{s \in \mathcal{T}}$ is a Q -martingale.

LEMMA A.4. *Given a martingale measure Q and an interest rate r derived from a Walrasian equilibrium price, π (as in Lemma A.2) and an equilibrium allocation for agent i , x_i , if there are sufficient assets to represent the Q -martingale X_i , then there exists a trading strategy using those assets and a riskless bond with will allow her to attain the allocation, x_i , and there does not exist an alternative allocation x^* that is strictly preferred to x_i within the agents' budget set.*

Proof. This is proven by Duffie and Huang(1985) in the context of an economy with $\mathbf{T} = \{0, T\}$ and the same method can be used to extend to the more general setting considered here (e.g. see Penalva(1997)). ■

A.3. PROOF OF THEOREMS 4.1, 6.1, AND 6.2

Theorem 4.1 is essentially a corollary of the more general Theorem 6.1. The order in the presentation required it to be stated as a Theorem. To prove Theorem 6.2 we use a procedure very similar to that for Theorem 6.1. So we start by proving Theorem 6.1, which is not vacuous by Lemma A.1.

Proof. (Theorem 6.1) By Lemma A.4 it suffices to show that for every agent $i \in I$, X_i is in the span of the discounted gains process of a private insurance contract and $k(e)$ mutual funds¹.

- *Step 1. Filtrations:* Choose an arbitrary $i \in I$. Recall the filtrations \mathbf{F}_i , \mathbf{F}_e and $\mathbf{F}_{i \vee e}$ from Section 3.2 and their generic extensions in Section 5.1.
- *Step 2. Martingale dimension:* If agent i 's endowment satisfies Assumptions 4.b-c, then $Y_i = N_i$, if it satisfies 4.d, Y_i includes either the

¹It is well-established in the finance literature that the discounted gains process of an asset is a Q -martingale.

$k(e)$ processes that generate Y_e or if it satisfies 4.e, the $m(e)$ process that generate $M(t)$.

If the LoGRC condition holds for e_i then there exists $N_l \in Y_e$ which can be decomposed into N_i and N_{-i} where N_i corresponds to all the indices in $\{1, \dots, K\}$ that generate risk for e_i and N_{-i} corresponds to the rest of indices that correspond to N_l and do not generate risk for e_i . The LoGRC condition and Lemma 5.1 applied to e imply that N_{-i} admits the intensity $\lambda_l - \lambda_i$.

The vector process $N' \equiv (Y_e \setminus N_l, N_i, N_{-i})$ generates the filtration $\mathbf{F}_{i \vee e}$. As the compensated version of N' is (by the martingale representation theorem) a basis for the space of $(P, \mathbf{F}_{i \vee e})$ -martingales then the martingale dimension, and hence the number of long-lived risky assets we need, is at most $k(e) + 1$ (where $k(e)$ includes $m(e)$ assets that account for the randomness generated by $M(t)$ in case of Assumption 4.e).

Note that by construction, $e_i(t) \in \mathbf{F}_i$ and $e(t) \in \mathbf{F}_e$.

• *Step 3. Martingale representation:* Let $C = E_Q[\sum_{t \in \mathbf{T}} \Delta_e x_i^*(t)]$. As C is a constant it can be safely ignored for the martingale analysis so we will omit it to reduce the burden on notation.

$$\begin{aligned} X_i(t) &= E_Q \left[\sum_{s=0}^T \Delta_e x_i^*(s) \middle| \mathcal{F}_t \right] \\ &= E_Q \left[\sum_{s=0}^T \gamma(s) (f_i(s, e(s)) - e_i(s)) \middle| \mathcal{F}_t \right] \\ &= E_Q \left[\sum_{s=0}^T \gamma(s) f_i(s, e(s)) \middle| \mathcal{F}_t \right] - E_Q \left[\sum_{s=0}^T \gamma(s) e_i(s) \middle| \mathcal{F}_t \right] \quad (\text{A.5}) \\ &= X_i^A(t) - X_i^I(t). \quad (\text{A.6}) \end{aligned}$$

Let X_i^A be the aggregate component and X_i^I the agent-specific component.

Both X_i^A and X_i^I are \mathbf{F} -martingales. For any $s > t$:

$$\begin{aligned} E_Q[X_i^A(s) | \mathcal{F}_t] &= E_Q \left[E_Q \left[\sum_{u=0}^T \gamma(u) f_i(u, e(u)) \middle| \mathcal{F}_s \right] \middle| \mathcal{F}_t \right] \\ &= E_Q \left[\sum_{u=0}^T \gamma(u) f_i(u, e(u)) \middle| \mathcal{F}_t \right] = X_i^A(t), \end{aligned}$$

by the law of iterated conditional expectations. As all martingales are predictable, and $f_i(t, e(t))$ is measurable with respect to \mathbf{F}_e for all t , $X_i^A(t)$ is also measurable on \mathbf{F}_e so that X_i^A is an \mathbf{F}_e martingale. As for $X_i^I(t)$, the change of measure which is a function of e , implies that one

need be concerned about both \mathbf{F}_i and \mathbf{F}_e , hence the best one can say in general is that X_i is $(Q, \mathbf{F}_{i \vee e})$ -martingale.

• *Step 4. Spanning with the assets:* Select the $k(e)$ mutual funds such that the discounted price processes, $S^*(t)$, are not linearly dependent. Applying the martingale representation theorem, there exists a $k(e) \times k(e)$ -dimensional predictable process $\mathbf{m}(t)$ such that

$$\begin{aligned} S^*(t) &= S^*(0) + \int_0^t \mathbf{m}(s) dY_e(s), \\ &= S^*(0) + \int_0^t \mathbf{m}_{-l}(s) d(Y_e \setminus N_l)(s) + \int_0^t m_l(s) dN_{-i}(s) + \int_0^t m_l(s) dN_i(s) \end{aligned}$$

where \mathbf{m}_{-l} is the $(k(e)) \times (k(e) - 1)$ -dimensional vector obtained by eliminating the l -th column, and m_l corresponds to that l -th column.

If $P(N_i(t) > 0) > 0$ and $P(N_i(t) = N_l(t)) < 1$ for some $t \in \mathcal{T}$ then $\mathbf{F}_i \neq \mathbf{F}_e$, otherwise \mathbf{F}_i is trivial or equal to \mathbf{F}_e . The latter implies that the agent only needs the $k(e)$ mutual funds and the bond to attain her optimal allocation.

Let $d_i(t)$ represent an insurance contract with different pay-offs for every realization of $N_i(t)$. The corresponding discounted gains process, S_i^* , is at worst a $(Q, \mathbf{F}_{i \vee e})$ -martingale (it will be a (Q, \mathbf{F}_i) -martingale if Assumptions 4.b-c hold). By the martingale representation theorem there exists a $(k(e) + 1)$ -dimensional vector process $(\mathbf{n}, n_{-i}, n_i)(t)$ such that

$$S_i^*(t) = S_i^*(0) + \int_0^t \mathbf{n}(s) d(Y_i \setminus N_l)(s) + \int_0^t n_{-i}(s) dN_{-i}(s) + \int_0^t n_i(s) dN_i(s)$$

where n_i and n_{-i} must be (generically) distinct: suppose $N_i(s-) = n$, then $dN_i(s) = 1$ makes the probability $\{N_i(t) = n\}$ equal to zero for all future $t \in \mathbf{T}$, while $\Delta N_{-i}(s) = 1$ does not, which generically implies that $S_i^*(s)$ is different after $\Delta N_i(s) = 1$ than after $\Delta N_{-i}(s)$ so that n_{-i} and n_i have to be different.

This implies, that S_i^* is not a linear combination of S^* so that (S^*, S_i^*) forms a martingale basis for $(Q, \mathbf{F}_{i \vee e})$ -martingales.

In Step 3 we showed that X_i is $(Q, \mathbf{F}_{i \vee e})$ -martingale, so that by Lemma A.4, there is a trading strategy using the bond, the $k(e)$ mutual fund and the insurance contract that lets agent i attain her optimal allocation. As we have chosen i arbitrarily, and the mutual funds and the bonds are common to all agents, the theorem holds. ■

Proof. (Theorem 4.1) Assumption 4.a is a special case of assumption 4.b. As $\tilde{N} = (N_i)_{i \in I}$ and for all i , $Y_i = N_i$ then the LoGRC condition is

trivially satisfied. As all agent's losses are the same and have the same hazard, they are all e -equivalent. Hence $k(e) = 1$. ■

Proof. (Theorem 6.2) The proof is exactly as that of Theorem 6.1 except that because the LoGRC condition does not hold then the best we can say about the martingale dimension of $\mathbf{F}_{i \vee e}$ is that it is at most $2k(e)$. This implies that each agent requires $2k(e)$ assets which in general would need to be mutual insurance contracts (i.e. written on both what happens to e_i and to e). Because these assets can differ from one agent to another (given possible heterogeneous losses) one potentially needs a different set of $2k(e)$ assets per agent. Of course, this number can never be greater than the martingale dimension generated by \tilde{N} , which is K . Also, in aggregate, the economy cannot need more than K assets. ■

A.4. PROOF OF LEMMA 7.1

Proof. (Lemma 7.1) Substituting a constant base endowment from the constant w_i at every t , to constant $w_i(t)$ does not add uncertainty.

For any $i \in I$, recall the definition of Y_i in Section 5.1 and consider changing the loss generated by $\Delta N_i(s) = 1$ of L_i at the next consumption date, t , to the loss $L_i(t, 1)$ at t , $L_i(t, 2)$ at $t + 1$, Apply the same logic as for making the endowment a function of consumption dates above but now on the endowment after a loss. ■

A.5. PROOF OF THEOREMS 7.1 AND 7.2

The proofs are exactly as those of Theorems 6.1 and 6.2. The difference is that N_i in those theorems is now substituted by $\tilde{N}_i = (N_{ij})_{j=1}^{m_i}$.

If the LoGRC condition holds (Theorem 6.1), for each N_{ij} there is an N_{lj} in Y_e such that the indices that make up N_{lj} include all the indices that make up N_{ij} for all $j = 1, \dots, m_i$. This is then naturally incorporated into the definition of N' in Step 3. The martingale dimension of $\mathbf{F}_{i \vee e}$ will at most be $m_i + k(e)$, and the same arguments go through with m_i insurance contracts for agent i instead of one.

If the LoGRC condition does not hold and because different risks do not depend on each other directly, then each risk requires in the worst possible case $2k(e)$ assets per risk, in total the agent needs $2m_i k(e)$ mutual insurance contracts.