



**COLUMBIA
BUSINESS
SCHOOL**

PaineWebber

WORKING PAPER SERIES

in

*MONEY, ECONOMICS
AND FINANCE*

PRICE UNCERTAINTY AND DERIVATIVE
SECURITIES

Geoffrey Heal
Columbia Business School
G.Chichilnisky
Columbia University
J.Dutta
Cambridge University

January 1995
PW-94-05

**COLUMBIA BUSINESS SCHOOL
COLUMBIA UNIVERSITY**

[http:// www.columbia.edu/cu/business/wp/](http://www.columbia.edu/cu/business/wp/)

PRICE UNCERTAINTY AND DERIVATIVE SECURITIES

G. CHICHILNISKY* J. DUTTA† G. M. HEAL‡

July 1991 revised December 1993, August 1994, January 1995.

ABSTRACT. We consider trade in contracts which provide insurance against price uncertainty. This uncertainty results from the presence of multiple equilibria. Rational traders thus have an intrinsic inability to predict the functioning of the economic system. We assume that they know, and agree on, the objective probabilities with which each equilibrium price vector realizes, and can trade in commodities contingent on the equilibrium. With an extension of the market structure in Arrow [1], these markets allow traders to insure fully against the risk stemming from uncertainty about prices. However, they introduce further uncertainty because there may be several equilibrium prices for price-contingent commodities. The introduction of higher-order derivative products removes this uncertainty, but in turn introduces uncertainty about the prices of these products. This process converges in a finite number of steps to a unique fully-insured Pareto efficient allocation. The introduction of derivative price-contingent securities removes all uncertainty associated with inability to predict equilibrium prices. We thus provide a mechanism for resolving indeterminacy in economies with multiple equilibria and give an important resource-allocation role to derivative securities. Limitation of the feedback between derivative markets and underlying markets emerges as important in establishing a positive role for derivatives.

JEL Classification: D 80, G 10, G 13, G 22.

Keywords: endogenous uncertainty, price uncertainty, derivative securities, index securities, multiple equilibria.

* Program on Information and Resources, Columbia University. Chichilnisky acknowledges support from the NSF, SITE and Monte Dei Paschi di Siena.

† Cambridge University. Dutta acknowledges support from the ESRC under grant no. R00023285 and from the Commission of the European Communities under SPES grant no. CT 91-0057.

‡ Program on Information and Resources and Graduate School of Business, Columbia University. Heal acknowledges support from SITE, the NSF and Monte Dei Paschi di Siena. We are grateful to M. Adler, K. Arrow, O. Compte, J. Drèze, F. Hahn, M. Kurz, K. Shell, P. Streufert, J. Swinkels, H-M. Wu, Y. Zhou, the referees and A. Mas-Colell for valuable comments. This paper was first presented at the July 1991 meetings of the Stanford Institute for Theoretical Economics.

1. UNCERTAINTY ABOUT ENDOGENOUS VARIABLES

The unpredictability of states of nature, such as the weather or the occurrence of earthquakes,¹ drives classical theories of allocation under uncertainty. As the term indicates, these are factors determined outside the economic system. However, many of the uncertainties facing economic agents today are not about states of nature: they are often about endogenous variables such as interest rates, exchange rates and securities prices, which are determined as part of a market clearing equilibrium, and so are affected by the actions of individuals. Corporate activity in risk markets is often to hedge against uncertainty about these endogenous economic variables. In this paper we analyze the outcomes of trading securities which allow individuals to insure against such price uncertainty. Our concern is with mechanisms for efficient resource allocation in these circumstances.²

We consider a very pure form of price uncertainty. There are competitive markets in goods. Individuals take prices as given, and markets clear. The outcome is an equilibrium vector of prices: it is important that this need not be unique. As a result, even those individuals who know, or have observed, the structure of the economy, cannot predict the prices which will arise. This, we believe, is the simplest representation of the claim that individual agents, no matter how rational and well-informed, are unable fully to understand the functioning of the economy. Risk-averse traders will wish to hedge the uncertainty which arises from this source.

A natural response to uncertainty about prices, is to trade goods contingent on prices, just as goods contingent on states of nature are traded in the Arrow-Debreu framework of uncertainty about exogenous states. Contracts for goods contingent on prices, or for price-contingent securities, are “derivatives” in the sense of financial assets whose values depends on the price of another underlying good. The relationship between these assets and the derivatives usually traded on financial markets is set out below in section 4.

In its naïve form, the approach of trading price-contingent goods does not work. Suppose that some goods are traded on the same market as commodities contingent on their prices. This would mean, of course, that individuals can observe the price of the good itself at the same time as they trade the price-contingent commodity. Because there are no arbitrage possibilities at a competitive equilibrium, only one such contingent commodity can trade at a non-zero price. In fact there is no trade in this commodity at equilibrium as there are no other contingencies to trade against. As a result, the market for price-contingent commodities collapses, and in particular cannot provide insurance against price uncertainty. Nevertheless, if the possibility of

¹Note that anthropogenic influences on climate, such as carbon dioxide emission and the release of CFCs, suggest that even uncertainty about the weather is not truly exogenous.

²Associated papers (Chichilnisky and Wu [10], Chichilnisky Heal and Tsomocos [11]) review uncertainty about whether agents will honor their contracts, i.e. counterparty or default risk.

multiple equilibria and hence the unpredictability of prices is intrinsic, there remains a need for price insurance.

To trade goods and price-contingent securities together is clearly not the right market structure. Nontrivial trading in such securities and in goods needs to be sequential rather than simultaneous. Specifically, the approach that markets for all contingent trading be open at the same time as in Arrow-Debreu theory cannot accommodate insurance against price uncertainty. The sequential structure of markets proposed by Arrow [1] and extended by Radner [23] is better suited to this problem. We consider a refinement of that market structure which allows traders to insure fully against certain types of price uncertainty and leads to a unique Pareto efficient allocation of resources.

1.1. An island analogy. To demonstrate the need for insurance, and the nature of insurance which is achievable, consider the following parable. There are a number of markets in different physical locations, call them islands. At date 0, individuals have to choose to locate themselves on one, and only one, such island. They carry their endowments with them. At date 1, trade takes place in competitive markets on each such island, prices are announced, and markets clear. The uncertainty arises because individuals must choose the island before observing which price is called; and risk-averse traders would like to insure themselves against this uncertainty. Normally, of course, the market clearing price vector on each island will depend on the composition of the population of traders. We want to abstract from this consideration for the moment, assuming that the distribution of the population is identical on each island, and known to be so by all.³ The point we wish to make is that even this symmetry is no guarantee that the same prices will arise at competitive equilibria in separate but identical locations. If the typical economy has multiple competitive equilibria, different vectors of prices could be called at different locations. Let us suppose, then, that this is indeed the case, and that there is an objective distribution of equilibrium prices, known to all. Individuals choose a location before observing the choice of equilibrium at that or any other location. They are risk-averse, and would like to insure themselves against the choice of equilibrium on their island. What means of insurance are possible?

At the initial date, individuals would be willing to write insurance contracts which ensure that their final consumption is invariant to the equilibrium chosen on their island. Imagine, then, an insurance firm which operates at date 0. It knows the identities of individuals, and so the final consumption that they would get at each equilibrium. The simplest scheme that it could offer is to promise actuarially fair insurance, that is, offer the statistical average of the consumption allocations across equilibrium realizations to each individual. Such a scheme is feasible, because it is an average of equilibrium allocations which are themselves feasible. It is desirable, because individuals prefer the average with certainty as long as they are risk-averse.

³This could be made precise by assuming that there are in fact a continuum of individuals of each type, and concentrating on symmetric location decisions.

Given any objective distribution on equilibrium prices, the resulting insured allocation is unique. However, this allocation is typically not Pareto-efficient. After collecting their insurance claims, individuals would be willing to re-trade; and it is perfectly possible that this trade itself has more than one equilibrium outcome. The same argument suggests that the insurance firm can make positive profits with certainty, by offering bundles which every individual will accept and which add up to less than the aggregate endowment. This is not to say that efficient insurance is impossible. There are allocations which are feasible, efficient, and dominate actuarially fair price insurance. Any such allocation can be achieved by a set of insurance contracts which are tailored to individuals, such that every trader prefers the insured allocation to that associated with actuarially fair price-insurance, and therefore, necessarily better than the uninsured random allocation associated with price uncertainty. Interestingly, most such allocations cannot themselves be competitive equilibria, and cannot be reached in the absence of insurance. The problem is that there are many such, and it is necessary to understand how a determinate final allocation could be achieved.

We have argued before that this problem cannot be solved within the Arrow-Debreu framework. We make this claim precise later on. We also propose one way of achieving insurance against price-uncertainty. We consider efficient and insured allocations, which can be reached as the limit of a process of competitive trade in derivative securities. In the tale of islands, different locations correspond to different choices of market clearing prices. This gives a concrete physical form to the thought experiments of traders, which is how we would like to analyze the issue. Exactly as in the classical theory of resource allocation under uncertainty, the possibility that one of several realizations may occur creates a need for insurance. Assets allow transfers of income across states, and hence lead to the possibility that traders can ensure that their final consumption is invariant to which outcome actually occurs. A particular implication of this analysis is that the variability of prices is a veil, as individuals are willing and able to insure themselves against pure price risks, and that some kinds of assets traded in financial markets perform precisely that function.

This approach integrates the institutional structure of price-contingent securities with the more abstract economic problem illustrated in our parable of islands. We do not claim that this is the only possible mechanism which achieves this. Our framework extends an aspect of Arrow's 1953 model of securities, and incorporates this into a sequential trading process. The presence of multiple equilibria provides a natural justification for the hypothesis that rational individuals cannot predict the outcomes of the economic system with perfect certainty. Such uncertainty must be a more general phenomenon: it will occur whenever outcomes are sensitively dependent on initial conditions (as is very general with dynamical systems), or when the system is too complex for its outcomes to be computable. We have imposed no bounds, however natural, on individual rationality with respect to expectations about prices or other equilibrium outcomes. Nevertheless, our analysis suggests that such "endogenous uncertainty" needs to be analyzed somewhat differently, and specifically that the institutional structure of markets and trading possibilities needs to be integrated

explicitly in exploring its relevance to resource allocation.

2. INSURING PRICE UNCERTAINTY - AN EXAMPLE

2.1. The framework. We introduce price uncertainty in the simplest possible way, by supposing a competitive exchange economy to have a finite number of equilibria. There are common and accurate expectations both of the set of possible equilibrium prices and of their probabilities, but traders do not to know which of these equilibria will actually be selected. So traders first trade goods contingent on the equilibrium price, and then an equilibrium price is selected.^{4 5}

Consider then an exchange economy with two goods, x and y and two traders, A and B . Traders are characterized by strictly concave utilities $U_i(x_i, y_i) : \mathbb{R}^2 \rightarrow \mathbb{R}$ and by endowments $w_i = (\bar{x}_i, \bar{y}_i) \in \mathbb{R}^2$; $i = A, B$. The aggregate endowment is $X = \bar{x}_A + \bar{x}_B$ and $Y = \bar{y}_A + \bar{y}_B$. There are several possible competitive equilibria in this economy. For the purpose of illustration, we suppose individuals expect one of two such equilibria to realize.⁶ Call these E_1 and E_2 . Let p_1 and p_2 the relative prices of good y to good x at these equilibria, which are assumed to be distinct. Traders accurately expect the prices to be those associated with these equilibria, and believe that these have probabilities of π and $(1 - \pi)$. Importantly, the probability is non-degenerate, i.e. $\pi \in (0, 1)$, so that the uncertainty is genuine. In this situation, traders face uncertainty which is not imposed on them by external states such as the weather: *the uncertainty is intrinsic to the economy and derives purely from an inability to predict the functioning of the economic system and in particular the selection of equilibrium prices.* Our main results can be illustrated by a series of claims, which we hope are clear in the context of this example.

Claim 1 Price-uncertainty is costly in welfare terms.

The argument here is simple, and illustrated by the Pareto frontier in Figure 1. Before a price is selected, traders have an expected utility vector that is inside the Pareto frontier. This is irrespective of whether indirect utilities are convex or otherwise in prices: uncertainty about equilibrium prices necessarily translates into

⁴This introduces price uncertainty in a very pure form in a framework close to that of Arrow-Debreu. Any other approach to introducing price uncertainty would require a sequence economy with incomplete markets, introducing a set of additional complications not germane to price uncertainty. Such a framework would also make it impossible to derive general welfare theorems, which we are able to do in the framework we have selected. The assumption of multiple equilibria is of course easily justified by reference to the strength of the assumptions needed to ensure global uniqueness, and the genericity of economies with a finite number of locally unique equilibria.

⁵We are not taking a position on how the set of possible equilibrium prices, or their probabilities, become common knowledge: the whole issue of how equilibrium prices are determined in competitive markets is a very open one. One might, just as an example, think of traders who trade repeatedly under identical circumstances (i.e., with the same preferences and endowments). Suppose these to give rise to multiple equilibria, one of which is chosen by a random mechanism. The traders learn over time the set of equilibria of their economy and their probabilities.

⁶Normally the number of equilibria will be odd. We take it to be two only to simplify the diagram.

uncertainty about final consumption allocations, which is undesirable for risk-averse individuals. There is therefore scope for mutually beneficial trades.




Figure 1: the Pareto frontier of the exchange economy. With uncertainty about which equilibrium will be realized, the expected utilities of the agents are inside the frontier. The final allocation is on the frontier and is Pareto superior to the expected utilities of the equilibria.

Claim 2 Insurance against price-uncertainty is feasible and desirable.

This too is directly illustrated by Figure 1; note that there are feasible allocations which can achieve utility allocations to the northeast of the expected utility value. More directly, let $(x_A^*(1), y_A^*(1))$ be the demand of individual A at price p_1 , and $(x_A^*(2), y_A^*(2))$ her demand at price p_2 . The allocations at the two equilibria are $C(j) = (x_A^*(j), x_B^*(j) \equiv X - x_A^*(j), y_A^*(j), y_B^*(j) \equiv Y - y_A^*(j))$ for $j = 1, 2$. The equilibrium allocations, $C(1)$ and $C(2)$, are feasible, and any convex combination of these must also be feasible, including $\bar{C} = \pi C(1) + (1 - \pi)C(2)$, the statistical average of the equilibrium allocations. Risk-averse traders prefer to get the average with certainty, by definition.

Claim 3 Insurance against price-uncertainty cannot be achieved in Arrow-Debreu economies.

To appreciate this statement, consider the structure of markets in Arrow-Debreu economies. All trade occurs at an initial date, where traders observe all prices,

including contingent-goods prices, and choose their preferred trades at these prices. Insurance against price-uncertainty can be achieved by trading commodities contingent on prices p_1 or p_2 . If traders observe all prices, they also observe whether goods prices are in fact p_1 or p_2 . As a result, one of the contingent commodities will have a zero price. The other will have a positive price, but there will be no trade in it, as there is no other contingency to trade it against. This last claim suggests that the structure of markets needs to be sequential, if markets for price-contingencies are to operate at all. We turn to this next.

2.2. Equilibria with price insurance. Suppose now that there are three trading dates, call them $t = 0, 1, 2$. Assume now that goods x and y can be stored from period 1 to period 2; that individuals receive their endowments in period 1, and that they consume only at date 2. At date 0, individuals A and B trade contracts, or securities, which will pay off at time 2. The payoff of these securities depends on the realization of prices in spot markets, which take place at date 1. They are promises to pay or take delivery of goods at date 2 contingent on which of the two prices p_1 and p_2 are called at date 1. At that date, goods x and y are traded in a competitive market. Importantly, in this market, there is trade of goods only, and every trader must break even. They cannot, for example, use their securities holdings as a means of payment. Given their endowments, they trade, taking prices as given. An equilibrium realizes. At date 2, they redeem their securities. This involves a transfer of goods across individuals, and then consumption occurs. As in the theory of competitive trading under uncertainty, we assume that traded contracts can be enforced, so that there is no default at date 2.

We have specified the sequencing of markets, but not the rules of operation. In what follows, we consider the consequences of competitive trade in securities at date 0, when markets are assumed to be complete. Any portfolio of securities is equivalent to making price-contingent consumption plans for date 2, or of trading price-contingent goods at date 0. The fact that the alternative allocations C_1 and C_2 arise from competitive equilibria is irrelevant to the pure insurance problem. Imagine another economy where the aggregate endowment is X, Y , and the two individuals have exactly the same preferences $U_i(x_i, y_i)$ as before. However, in this economy, the distribution of endowments is random. Individual A receives endowment $e_A(1) = (x_A^*(1), y_A^*(1))$ with probability π and $e_A(2) = (x_A^*(2), y_A(2))$ with probability $(1-\pi)$. In each case, B receives the remainder of the state-independent aggregate endowment. In this economy, the distribution of endowments is ex-post Pareto efficient. If no assets trade in period 0, there will be no trade in goods in period 1. However, if assets do trade, individuals will choose to insure. With complete, competitive asset markets in period 0, individuals' final consumption levels will be different from their endowments. The aggregate endowment is non-stochastic, and individuals strictly risk-averse. It can be shown (Lemma 2 of the Appendix) that the competitive equilibrium of this extended economy must be state-independent. Asset prices and equilibrium allocations cannot depend on whether $C(1)$ or $C(2)$ is the actual distribution of endowments. For the

purpose of this illustration only, let us assume that this equilibrium is unique, and let C^* be the final allocation associated with it. We claim that equilibrium with competitive securities markets for price uncertainty must coincide with the equilibrium C^* of this artificial economy. For the moment, it is important to note that the fully insured allocation, C^* , is feasible, Pareto-optimal, and preferred to the allocations associated with equilibrium uncertainty. It can be achieved by trade in price-contingent goods, or in assets denominated in goods whose payoffs are contingent on prices in period 1. In such a market, there are four economic goods – two physical commodities, each contingent on one of two possible equilibria. The nature of the price-contingent trades is as follows: an equilibrium, say E_1 , is realized in period 1, and each trader receives her consumption vector associated with that equilibrium. Subsequent to the receipt of those consumption vectors trader 1 will deliver to trader 2 the vector of “goods contingent on price p_1 ” which trader 2 purchased from trader 1 in the market for price-contingent goods. Trader 1’s final consumption vector if equilibrium 1 is realized will therefore be her equilibrium 1 consumption vector minus the goods delivered to trader 2. Trader 2’s final consumption is computed similarly. This is similar to the way in which second-period endowments are augmented by securities purchased in the first period in Arrow’s 1953 model.

Figure 2 shows the price-contingent trades made: trader 1 sells to trader 2 a vector $T_1 \in \mathfrak{R}^2$ of “goods if p_1 is the price vector”, and reciprocally trader 2 sells to 1 the vector $T_2 \in \mathfrak{R}^2$ of “goods if p_2 is the price vector”. Whatever the equilibrium realized in the exchange economy, the final allocation of consumption is the point F on the contract curve: it may be reached by the realization of equilibrium 1 followed by trade T_1 or by the realization of equilibrium 2 followed by trade T_2 .⁷

It is now time to take stock of this example, and see what it indicates. It illustrates a situation where uncertainty about equilibrium prices can be insured fully by trading goods contingent on the equilibrium realized (represented by its prices). Trades in these contingent goods are of course based on expectations of the set of equilibrium prices, but are made before one is realized. These trades lead to a unique Pareto efficient allocation, a very satisfactory outcome.

⁷Note that for trader 1’s budget to balance, the value of T_1 at prices for goods contingent on price 1, must equal the value of T_2 at prices contingent on p_2 . This implies $\langle T_1, (p_{1/1}^2, p_{2/1}^2) \rangle = \langle T_2, (p_{1/1}^2 \frac{\pi_2}{\pi_1}, p_{2/1}^2 \frac{\pi_2}{\pi_1}) \rangle$, where $\langle x, y \rangle$ denotes the inner product of x and y . This can be rewritten as $\langle (p_{1/1}^2, p_{2/1}^2), (T_1 - T_2 \frac{\pi_2}{\pi_1}) \rangle = 0$. If as we assume just for convenience that the two probabilities are equal, then this means geometrically that the vector of relative prices for price-contingent goods is orthogonal to the difference between the trades T_1 and T_2 . Figure 2 reflects this.

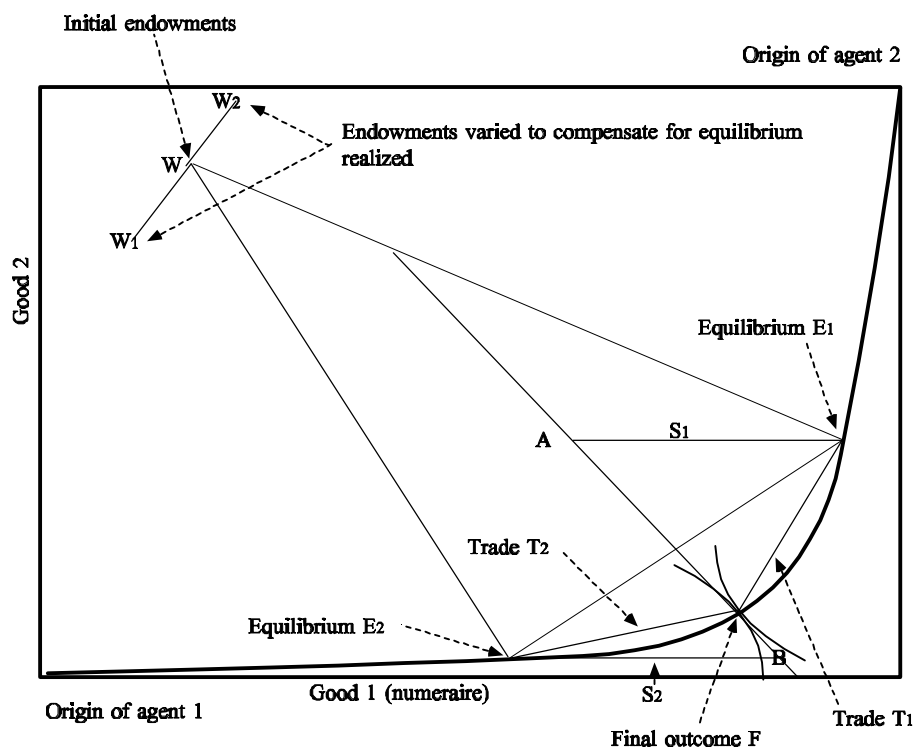


Figure 2: an exchange economy with 2 equilibria. Trading price-contingent goods leads via trades T_1 and T_2 to the fully insured Pareto efficient outcome F . This can also be reached by trading price-contingent securities via the trades S_1 and S_2 . Changing endowments to W_1 or W_2 to compensate for the equilibrium chosen will not lead to the same outcome.

Key to this result was the institutional structure assumed, which has the following components:

1. trading in goods and in price-contingent goods occurs in markets that cannot be accessed simultaneously.
2. there are common and accurate expectations of the equilibrium prices and associated probabilities. Positions are taken in price-contingent goods markets on the basis of expectations of equilibrium prices but before one is selected.
3. endowments in the market for price-contingent goods are the consumption vectors associated with the equilibrium prices in the market for goods.

Note several points of similarity to the structure in Arrow's 1953 paper [1], and its extension by Radner [23], in which:

1. trading occurs in two distinct rounds.
2. when trading securities in the first period, traders know the state-contingent prices that will rule in the second period. Knowing the prices linked to each state is immediate, in our model, as the prices *are* states.
3. endowments in the second trading round are determined by the purchases made in the first round.

This institutional structure provides the basis for the framework and the concept of equilibrium that we introduce in the next section.

The simplicity of our example owes much to the assumption that the market for price-contingent goods has a unique equilibrium, and in the rest of the paper we relax this assumption. Without it, there would be several possible final allocations, each of which would be on the Pareto frontier, with uncertainty about which would ultimately result. This residual uncertainty would have a welfare cost illustrated in figure 3, which could in principle be reduced by allowing traders to insure against uncertainty about the outcome in the market for price-contingent goods. This would require trading goods contingent on the prices of price-contingent goods. Below we formalize and generalize the process illustrated in this example, and show *inter alia* that even if all markets have multiple equilibria, a finite number of “levels” will always suffice to provide full insurance against price uncertainty.⁸ ⁹ Note however that in Proposition 1 below we show that if the probabilities of different equilibria occurring

⁸One further insight (due to Drèze [13]) can be extracted from the example. One way of motivating the trading process we use, is to rule out some obvious alternatives, as we have already done. Another such alternative is to change traders’ initial endowments in a way that depends on the equilibrium realized. After an equilibrium is realized, initial endowments might be removed from the trader best off at the realized equilibrium and transferred to the trader worst off. Traders would agree to insure each other by making changes in their endowments after an equilibrium has been selected but before it has been implemented, with the new endowments leading to a more equitable equilibrium. In our example, endowment would be transferred from trader 1 to trader 2 in the event of equilibrium 1 being realized, and vice versa. Geometrically, this would replace the initial endowment point W in the Edgeworth box by a pair of price-contingent endowments such as W_1 and W_2 in figure 2. Clearly a single efficient allocation such as F could not in general be supported from both members of such a pair: this approach would make it impossible to realize a single equilibrium-independent final allocation between traders, and therefore impossible to provide full insurance.

⁹A final comment (also due to Drèze [13]) on the example and figure 1: we have assumed that to insure, traders trade price-contingent goods. They could in fact trade price-contingent securities. The equivalence between contingent goods and securities, established by Arrow for the case of exogenous states, still applies here. Let good 1, the horizontal good in figure 2, be the numeraire, and let traders trade securities which pay units of the numeraire depending on the equilibrium realized in the exchange economy. In this case insurance trades would be represented by moves such as S_1 and S_2 in figure 1. These show that if equilibrium 1 is realized, trader 1 transfers numeraire to trader 2, leading to a new endowment point A. Similarly, if equilibrium 2 is realized, trader 2 transfers numeraire to trader 1 according to the trade S_2 . In either event, the opening of goods markets after these transfers will lead to a final equilibrium at F .

are concentrated at one equilibrium, the market for price-contingent goods will have a unique equilibrium.

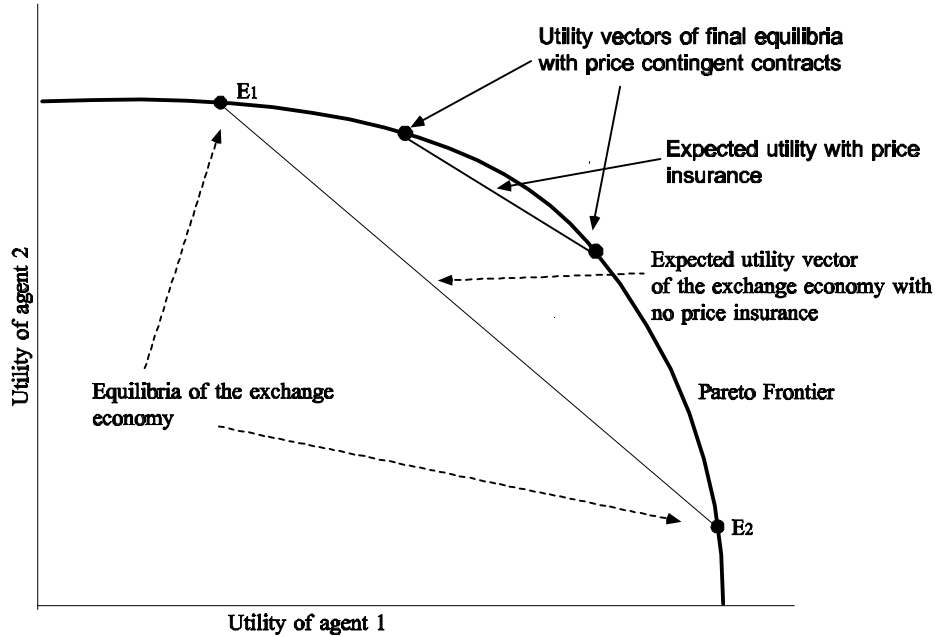


Figure 3: the Pareto frontier of the exchange economy. With several equilibria in the economy with price-contingent contracts, the outcome is a pair of utility vectors on the frontier and there is a cost to risk bearing from this.

2.3. The impossibility of completing the market against price uncertainty.

Intuitively it is clear that arbitrage possibilities will force a separation of markets for goods and for securities contingent on the prices of these goods. Chichilnisky Hahn and Heal (CHH) [7] and Chichilnisky [6] in fact prove formally that in an Arrow-Debreu framework one cannot use markets for price-contingent securities to hedge uncertainty about which equilibrium prices will rule. Roughly the argument is as follows.

Let an economy have two competitive equilibria, E_1 and E_2 , whose probabilities are known to agents. Introduce two securities S_1 and S_2 which pay a unit of numeraire if E_1 or E_2 occurs respectively. Agents maximize expected utility and are able to augment their incomes in E_1 or E_2 by trading securities. There are two states (equilibria) and two independent securities, so that the market is complete and the outcome should be efficient. What happens?

By Lemma 2 of the appendix, as the economy faces no aggregate uncertainty but only uncertainty about the distribution of resources between agents, an equilibrium will be fully insured, in the sense that each agent will have the same consumption

whatever the state (i.e., whatever the equilibrium of the exchange economy). This implies that the prices and consumptions will be independent of the state (equilibrium), so there will be only one equilibrium price vector. This contradicts the assumption of two equilibria. Hence the market cannot be completed in this way against price uncertainty.

3. INSURANCE AND PRICE-CONTINGENT CONTRACTS

We begin by formalizing the example of the previous section, and showing that the introduction of price-contingent contracts within the trading framework of that example will always lead to a Pareto improvement in traders' welfares, even if it does not fully remove endogenous uncertainty. We then show that in fact endogenous uncertainty can always be fully removed by a finite number of types of price-contingent contract or derivative.

We consider a competitive exchange economy with a finite number of competitive equilibria, a generalization of the example of section 2. This is an Arrow-Debreu economy except for the structure of information. Traders have accurate expectations of the set of market clearing prices,¹⁰ but do not know which will be selected. This uncertainty has a welfare cost to the traders. To reduce this, traders trade on markets for commodities contingent on the equilibrium price vector. The introduction of such price-contingent contracts (also called level 1 securities) allows traders to insure fully against the uncertainty about the equilibrium to be chosen.

However, further uncertainty arises if there are several possible equilibrium prices in the market for price-contingent commodities. This second round of uncertainty can in turn be removed by introducing higher-order or derivative securities, i.e., securities that pay off according to the prices of the level 1 securities: these we call level 2 securities. Again, this introduction of new securities allows traders to remove one source of uncertainty only to introduce another. In general, though not always, as we shall see below, *the introduction of securities may both remove one type of uncertainty and introduce another type.*

We show that *generically on utility functions this process of introducing successive levels of derivative securities will remove all endogenous uncertainty in a finite number of steps, and lead to a unique¹¹ fully-insured and risk-free Pareto efficient allocation.* The introduction of the "last" level of securities removes the uncertainty from the penultimate level, and introduces no further uncertainty. It is an implication of Theorem 1 below that there is always such a "last" level. Hence the introduction of price-contingent commodities (or equivalently securities), and further derivative securities, will remove all endogenous uncertainty associated with inability to predict equilibrium prices. We thus provide a mechanism for resolving indeterminacy in

¹⁰From these they can calculate the associated consumption vectors, which are just their demand vectors.

¹¹The final allocation is unique given the probabilities over equilibria. However, different probability distributions over equilibria will in general lead to different outcomes.

economies with multiple equilibria, and at the same time give an important resource-allocation role to derivative securities.

In addition, we show in section 4 that the payoff functions of the derivative securities that we introduce, can be replicated as the limit of payoff patterns emerging from trading combinations of options. In fact they are the payoff patterns of what Rubinstein [24] terms exotic options.

3.1. A framework for endogenous uncertainty. We now formalize a framework which generalizes the example of section 2. We consider an economy with i traders indexed by $i \in I = \{1, \dots, I\}$ and J goods indexed by $j \in J = \{1, \dots, J\}$. w_{ij} is trader i 's endowment of good j and w_i is trader i 's endowment vector in \mathfrak{R}^J . Preferences are represented by utility functions $U_i : \mathfrak{R}^J \rightarrow \mathfrak{R}$, and consumption vectors are $c_i \in \mathfrak{R}^J$. We make the following assumptions:

A1. $\forall i$, U_i is strictly concave, C^2 (twice continuously differentiable), monotonically increasing and has non-zero gradients.

A2. If $\{\pi^k, c_i^k\}$ is a lottery over consumption vectors c_i^k for trader i with probabilities π^k then i 's utility from this lottery is $\sum_k \pi^k U_i(c_i^k)$.

The exchange economy defined thus far will be denoted E^1 and referred to as the *underlying economy*: in E^1 endowment are w_{ij} , preferences are U_i , and the commodity space is \mathfrak{R}^J . $CE(E^1)$ will denote the set of competitive equilibria of E^1 , with p_k^1 being the price vector at the k -th. equilibrium and $c_{i,k}^1$ being trader i 's consumption at the k -th. equilibrium.

A3. The number of equilibria in E^1 is finite.

This assumption is satisfied by typical exchange economies. More precisely, the family of utility functions of which a residual set give finitely many equilibria for any endowments is the family \mathcal{U} of C^∞ functions whose bordered Hessians are non-zero everywhere.¹²

A4. Traders have accurate expectations about $CE(E^1)$, the set of possible equilibria of the underlying economy. They also expect that one of the equilibria of E^1 will be realized randomly according to a commonly-known exogenous probability distribution $\pi^1 = \{\pi_k^1\}$, $k = 1, \dots, \mathcal{N}^1$.

¹²This is known to be true generically on endowments for smooth preferences (Debreu [14]): it is also true for any endowments provided that the utility functions are selected from within a generic (formally, residual: a residual set is a countable intersection of open dense sets) set of utility functions meeting certain conditions. This follows from the work of Mas-Colell and Nachbar [22] and Herman [18] in extending Debreu's results on regular economies.

Traders hedge against the risks associated with inability to predict the equilibrium by trading goods contingent on the equilibrium selected in E^1 . The market for these contingent goods will typically have in its turn multiple equilibria, and commodities contingent on the prices of these contingent goods will be needed to remove uncertainty thus introduced. In order to define this construction concisely, we introduce the concept of a *multi-level economy* E , in which the underlying economy E^1 forms the first level. Y denotes the set of levels in E , with $y \in Y$ denoting a typical level. Levels 1 and 2 are defined as follows:

L1 Level 1 is the underlying exchange economy E^1 .

L2 Level 2, denoted E^2 , is a set of markets on which traders trade goods contingent on which element of $CE(E^1)$ is realized:

- the number of states in E^2 is \mathcal{N}^1 , the number of equilibria in E^1 .
- endowments in E^2 are consumption vectors at the equilibria of E^1 , so that $c_{i|k}^1$ is trader i 's endowment at the k -th. state of level 2, i.e., $w_{i|k}^2 = c_{i|k}^1 \in \mathfrak{R}^J$. The overall endowment vector of trader i is $w_i^2 = (c_{i|k}^1)_{k=1, \dots, \mathcal{N}^1} \in \mathfrak{R}^{Jx\mathcal{N}^1}$.
- trader i chooses a consumption vector $c_i^2 = \{c_{i|k}^2\}_{k=1, \dots, \mathcal{N}^1} \in \mathfrak{R}^{Jx\mathcal{N}^1}$ to maximize $\sum_k \pi_k^1 U_i(c_{i|k}^2)$ subject to $c_i^2 \cdot p = c_i^1 \cdot p$ where $c_i^1 = \{c_{i|k}^1\}_{k=1, \dots, \mathcal{N}^1}$.
- $CE(E^2)$ is the set of competitive equilibria of E^2 , which has cardinality \mathcal{N}^2 , with a typical price vector being $p_i^2 \in \mathfrak{R}^{Jx\mathcal{N}^1}$.

This formalizes the example of section 2. In that example $Y = \{1, 2\}$ as there were two layers, and $\mathcal{N}^1 = 2$, $\mathcal{N}^2 = 1$. Hence the number of states in level 2 was 2, and the commodity space in that level was $\mathfrak{R}^{Jx\mathcal{N}^1} = \mathfrak{R}^{2x2}$.

3.2. One type of derivative. We can now state formally the results illustrated by the example of the previous section. Given the realization of an equilibrium in L1 and an equilibrium in L2, agents' *consumptions* are given by the sum of their consumption vectors at the realized equilibria in L1 and in L2. In a two-level economy defined by preferences, endowments and expectations over market-clearing goods prices, an *equilibrium* is a set of trades for each agent in the markets for goods and for price-contingent goods such that all markets clear and utilities are maximized subject to budget constraints as in the definitions of L1 and L2. An equilibrium is *fully insured* if each agent's consumption is independent of which equilibrium is realized in the market for goods. Then:

Theorem 1. *Consider an exchange economy in which agents trade goods and goods contingent on the prices of goods, i.e., consisting of levels L1 and L2 as defined above. Then (a) any equilibrium of this economy is fully insured and Pareto efficient and (b) the introduction of markets for price contingent goods makes some agents better off and none worse off.*

Proof. Full insurance of the equilibria is established in Lemma 2 of the Appendix. Pareto improvement follows from Theorem 2 below. ■

Note that although each equilibrium is fully insured and Pareto efficient, there may be many equilibria because L2 will typically have many equilibria. Hence agents will again face the kind of uncertainty that started this analysis, namely uncertainty about equilibrium prices, this time in the market for derivatives. Figure 3 represents the situation for two agents: there are two equilibria in L1 and L2, so that finally agents face a distribution over two utility vectors on the utility frontier, with a consequent loss of welfare. Clearly the process of introducing derivatives can be iterated to remove this uncertainty, and we investigate this fully below. If L2 has a unique equilibrium then all uncertainty is removed: otherwise, the introduction of price-contingent contracts removes fully the risks associated with not knowing goods prices, but introduces a new type uncertainty, about the prices of price-contingent contracts. With the framework we work with here, the new uncertainty is “less important than” the initial uncertainty, in the sense that there is a Pareto improvement. Note that if one of the equilibria of the exchange economy is much more likely than any other, then the market for price-contingent goods *will* have a unique equilibrium:

Proposition 1. *If the probabilities of the equilibria in the underlying exchange economy are sufficiently concentrated around one equilibrium, then the market for price-contingent contracts will have a unique equilibrium. Formally, let $CE(E^1)$ be the set of competitive equilibria of the underlying exchange economy, and $\pi_i^1, i \in CE(E^1)$, be the probabilities of these equilibria being realized. Let $\#^1$ be the number of equilibria in $CE(E^1)$. Then $\exists \epsilon > 0$: if for some i , $\pi_i^1 > 1 - \epsilon$ and $\pi_j^1 < \epsilon / \#^1 \forall j \neq i$, the economy E^2 has a unique competitive equilibrium.*

Proof. A proof is given in the Appendix. ■

3.3. The general case. We now extend the analysis to the most general case in which there are several equilibria in the markets for derivatives and it is therefore necessary to introduce further derivatives to insure against this. We introduce markets for goods contingent on the prices of goods contingent on the market-clearing prices, and so on. For this, we need to extend the definition of a multi-level economy given above: we define additional levels inductively, extending the definition of level 2, L2, given above.

L3 Level y , denoted E^y , is a set of markets on which traders trade goods contingent on which element of $CE(E^{y-1})$ is realized:

- the number of states in E^y is \mathcal{N}^{y-1} , the number of equilibria in E^{y-1} .
- endowments in E^y are consumption vectors at the equilibria of E^{y-1} , so that $c_{i|k}^{y-1}$ is trader i 's endowment at the k -th. state of level y , i.e., $w_{i|k}^y = c_{i|k}^{y-1}$. The overall endowment vector of trader i is $w_i^y = \left(c_{i|k}^{y-1} \right)_{k=1, \dots, \mathcal{N}^{y-1}} \in \mathfrak{R}^{Jx\mathcal{N}^{y-1}}$.
- trader i chooses a consumption vector $c_i^y = \{c_{i|k}^y\}_{k=1, \dots, \mathcal{N}^{y-1}} \in \mathfrak{R}^{Jx\mathcal{N}^{y-1}}$ to maximize $\sum_k \pi_k^{y-1} U_i(c_{i|k}^y)$ subject to $c_i^y \cdot p = c_i^{y-1} \cdot p$ where $c_i^{y-1} = \{c_{i|k}^{y-1}\}_{k=1, \dots, \mathcal{N}^{y-1}}$.
- $CE(E^y)$ is the set of competitive equilibria of E^y , which has cardinality \mathcal{N}^y , with a typical price vector being $p_k^y \in \mathfrak{R}^{Jx\mathcal{N}^{y-1}}$.
- At every level y , traders hold accurate expectations about the set $CE(E^y)$ of competitive equilibria of that level. They also know that one of these will be realized according to a commonly-known exogenous probability distribution $\pi^y = \pi_k^y, k = 1, \dots, \mathcal{N}^y$.

Definition 1. A realization s is the selection of an equilibrium at every level $y \in Y$. A realization at level y is the selection of an equilibrium at every level up to and including y , and is a list of y integers.

A realization is a path through a tree whose nodes at each level are the equilibria at that level, as shown in figure 4. In our model, level 1 realizations will be the equilibria of the underlying exchange economy E^1 , level 2 realizations are pairs of equilibria, one from E^1 and one from E^2 , the markets for goods contingent on equilibria in E^1 . Each level of the economy corresponds to a different class of derivative security. These derivative securities are introduced in an order of logical priority, so that the payoff of each depends on the prices of the prior ones.

Within the multi-level economy E , a realization s_y at level y specifies price vectors for commodities and derivative securities up to level y : it also specifies for each trader a consumption vector for commodities and vectors for all types of contingent commodities up to level y . Geometrically, it specifies a vector at each node of the tree in figure 4 lying on the path from the base to the equilibrium realized at level y . The overall consumption of trader i in realization s is the sum of the consumption vectors of trader i at every equilibrium selected in this realization. It is thus a consumption vector corresponding to the equilibrium selected in E^1 , plus a vector of goods chosen in E^2 contingent on the equilibrium of E^1 , plus a vector chosen in E^3 contingent on the equilibrium in E^2 , etc. It is the sum of the initial consumption vector plus a series of price-contingent transfers of goods along the path through the tree in figure 4 implied by the realization selected. Posterior levels of price-contingent contracts entitle traders to delivery of goods vectors modifying their prior positions, and the overall consumption vector is the sum of all of these.

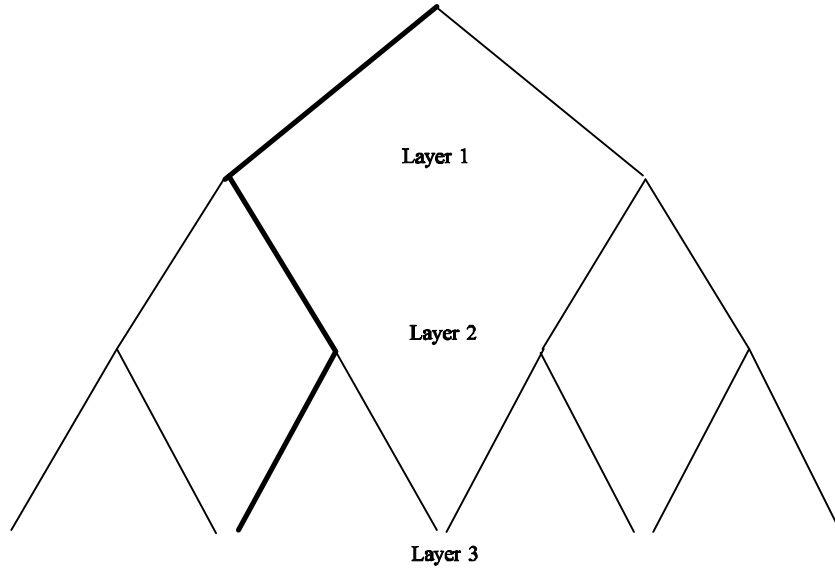


Figure 4: there are two equilibria at each of three levels. A realization is a path through the tree, as illustrated.

Definition 2. Let $c_i(s)$ be trader i 's overall consumption in realization s . Then

$$c_i(s) = \sum_{y=1}^{\mathcal{N}(Y)} c_{i|k(y,s)}^y$$

where $c_{i|k(y,s)}^y$ is trader i 's consumption in equilibrium $k(y, s)$ of level y and $k(y, s)$ is the state chosen at level y in realization s .

In the construction formalized here, each level is treated as a separate stochastic exchange economy, with state-dependent endowments and preferences over uncertain consumption vectors. The levels are related in that the states and endowments of level y depend on the equilibria of level $(y - 1)$. The dimensions of the commodity spaces of the E^y , $y \in Y$, depend on y . At each level traders maximize expected utility, with the expectation over a set of states defined by the equilibria of the previous level, and with consumption defined as the consumption vector realized in the exchange economy modified by a series of price-contingent transactions.

The next step is to give a definition of an equilibrium for the extended economy with multiple levels of price-contingent contracts. Consider the tree in figure 4: this

is analogous to the usual tree depicting the resolution of uncertainty in a multi-period Arrow-Debreu model, except that the states at each level are determined by the prior level, and that the evolution is not through time but along a logical sequence. An equilibrium is a set of price and consumption vectors for each node of the tree such that at a particular node, if traders take as given prices and all decisions made at prior levels, then the associated consumption vectors at that node maximize utilities subject to traders' budget constraints, and all markets clear.¹³ This is formalized in the following

Definition 3. *Consider the multi-level economy E , characterized by preferences, endowments and expectations over equilibrium prices. An equilibrium is*

1. *a finite set of levels Y^* and*
2. *a set $CE(E_y)$ of competitive equilibria (utility-maximizing and market-clearing trades for each agent) for the exchange economies E_y for each level $y \in Y^*$, where the levels E_y satisfy the relationships specified in points **L1** to **L3** above such that*
3. *the overall consumption of a trader in any realization, which is the sum of contingent trades defined by that realization as in definition 2, is independent of the realization (“full insurance”).*

We now establish that the introduction of a finite number of derivative securities at a finite number of different levels in the economy will suffice to remove all endogenous uncertainty and to provide a unique Pareto efficient and fully-insured allocation of resources in the underlying economy E^1 . Note that this allocation is *not* one of the competitive equilibria of the underlying economy E^1 .

Theorem 2. *Let the underlying exchange economy E^1 satisfy (A1) to (A4). Then for a generic set of utility functions there is a finite number N such that an N -level economy will have a unique equilibrium in which consumption vectors are independent of the realization selected and which is Pareto efficient in E^1 . The addition of extra levels of derivative security markets up to level N leads to Pareto improvements.*

Proof. The proof of this theorem is given in the Appendix.

We have established that introducing a finite number of levels of securities, and their derivatives, will provide full insurance against the endogenous uncertainty arising from lack of knowledge of the equilibrium price vector. Many levels are needed

¹³Note that this definition is not equivalent to traders maximizing the expectation of utility over all possible realizations, i.e., over the tree. It can be shown that such behavior will not in general lead to an outcome that is Pareto efficient.

because the introduction of each level of securities will in general remove the uncertainty associated with not knowing the equilibrium price of the previous level, but will introduce further uncertainty arising from lack of knowledge of the price at this new level. The addition of each level produces a Pareto improvement in the welfare levels of the economy.

This result has relevance to the debate about whether markets for derivative securities increase the risk and uncertainty to which the economy is subject. It indicates that, while it is true that a new derivative market introduces additional uncertainty, one has to set against this the fact that it provides insurance against preexisting uncertainty. The fact that the addition of extra derivative markets is always Pareto improving, implies that the gains from the new market outweigh the loss, at least in the framework analyzed here. In the next section, we turn to an institutional interpretation of Theorems 1 and 2. We relate it to securities based on price indices and to trading strategies based on options.

4. INSTITUTIONAL INTERPRETATIONS

4.1. Interactions between markets. In the discussions so far, there has been no allowance for feedback from derivative markets to the underlying goods markets, i.e., we have not allowed for the possibility that the hedges made by agents against uncertainty about equilibrium goods prices will affect the market clearing prices in the goods market. One might think of the following chain of argument: a set of possible market-clearing goods prices in an exchange economy induces agents to make certain price-contingent trades. The existence of these trades and the insurance which they provide modifies agents' demands in the exchange economy so that the market clearing prices in this are now different. Hence a new set of price-contingent trades is chosen, and so on. Does this process have an equilibrium? If so, what is it?

In fact the process just described *cannot* converge: convergence would contradict the theorem of Chichilnisky Hahn and Heal cited above. For convergence would imply the existence of a set of market-clearing prices and a set of trades in goods contingent on these prices such that all markets clear and utilities are maximized given market prices and price-contingent contracts, and agents are fully insured. This is precisely what the CHH theorem states cannot exist. It is therefore of the essence that feedback from the derivatives market to the goods market should be limited.

It is instructive to look in detail at what happens if agents take positions in the underlying markets in anticipation of being able to trade later in securities markets, and enquire exactly how this is destabilizing. It emerges from such enquiry that if agents are to be allowed to optimize with respect to all markets simultaneously, then certain institutional restrictions on their trades seem needed to preserve the desirable outcomes described in Theorem 2. In fact, such constraints arise very naturally in derivative markets. One can interpret the derivative markets as insurance markets: in insurance markets it is standard to limit the extent of the trades an agent can make. To be precise, one cannot buy \$30,000 of car insurance unless one can demonstrate ownership of a car whose value is at least \$30,000. Again, in markets for

derivative securities, collateral requirements (either directly or acting through credit restrictions) place limits on the trades open to agents.




Figure 5: prices and allocations here correspond to those in the Edgeworth box in figure 2. Agents who anticipate when trading goods that they will be able to trade securities then retrade goods at different prices may take extreme positions such as x_1 and x_2 . Their behavior will be stabilized if wealth is always evaluated at the vectors c_1 and c_2 .

Figure 5 shows the alternatives facing an agent who is trading in the underlying market when this market has two possible equilibrium prices and there is a unique equilibrium price for price-contingent securities. In this case, although the equilibrium price vector for price-contingent goods is four dimensional (two goods in two states), the relative prices of goods in the two states are the same so that this can be represented by a conventional budget constraint in a two-dimensional diagram. In figure 5, the two possible equilibrium prices are p_1 and p_2 , and p^* is the vector of relative prices for price-contingent goods. c_1 and c_2 are the agent's competitive equilibrium consumption vectors at the two possible equilibria of the exchange economy. In the second level of trading, agents trade securities in order to shift income between the two states given by the two possible equilibrium price vectors. In the case shown, the agent will shift income from state 1 to state 2. Agents then trade goods again after the state is realized and income transfers have been implemented: this trading is at

prices p^* . Clearly, if the agent is aware of the possibility of shifting income between states and trading subsequently at prices p^* , he or she will choose an extreme position such as x_1 or x_2 rather than c_1 or c_2 when facing prices p_1 or p_2 : such positions lead to higher pre-transfer income levels than the equilibrium consumptions c_1 or c_2 when evaluated at prices p^* . In this case, the introduction of level two markets, coupled with the ability to take positions in level one markets in anticipation of level two, will imply that c_1 and c_2 are no longer the equilibria at level one.

How do we prevent anticipation of derivative trading from destabilizing the system? The key is to use an appropriate rule for the evaluation of pre-transfer wealth. Define this as the inner product of a price vector with the equilibrium consumptions, either c_1 or c_2 . So when computing the wealth base from which transfers across states are subtracted, or to which they are added, it must be understood that these are the values of equilibrium consumptions in the exchange economy, irrespective of what positions agents may claim. c_1 and c_2 are the agent's legal entitlements in states 1 and 2, the most that she can lay claim to in these states. They are therefore all that could be offered as collateral in derivative transactions. We can also think of trading price-contingent securities as insurance against unfavorable prices: the agent is ensuring against favorable consumption c_1 being replaced by less favorable c_2 . Then it is natural to insist that the agent can only insure against the loss of a favorable position (such as c_1), and its replacement by one that is less favorable (such as c_2), if she can show that these are the positions to which she is legally entitled. For c_1 and c_2 this can be shown: for other positions, this entitlement cannot be demonstrated. So the interpretation of c_1 and c_2 as the only legal property rights in states 1 and 2, and the rule that we value wealth in these states by applying prices to these consumption vectors, will stabilize the multi-level economy and ensure that even if the possibility of trading posterior to the exchange economy is anticipated, the equilibria of this economy will be unchanged.

4.2. Options and price-contingent contracts. There are two results that help give an interpretation of Theorem 1 in terms of securities actually traded. One states that contracts which are contingent on the value of a price vector, can also be designed to be contingent on the value of a price index and can be interpreted as index contracts. The significance of this lies in the fact that index-based securities and contracts are now widespread in the financial world. The second result states that the pattern of payoffs as a function of goods prices exhibited by our price-contingent securities, can be replicated by a limit of trading strategies based on index options. Overall, these two results imply that the securities (price-contingent commodities) traded in our multi-level economy, can be understood as options on price indices and their derivatives.

Lemma 1. *Consider a set of distinct price vectors $p^1, p^2, \dots, p^k \in \mathfrak{R}^J$. Then for an open dense set of $- \in \mathfrak{R}^J$, each price vector gives a different value to the index p^i .- .*

Proof. Assume $x \neq 0$. Let p^1, \dots, p^g be such that

$$p^i \cdot x \neq p^j \cdot x, \forall i \neq j \text{ and } i, j \in \{1, \dots, g\} \tag{1}$$

For p^{g+1}, \dots, p^k we have

$$p^{g+1} \cdot x = p^{g+2} \cdot x = \dots = p^k \cdot x \tag{2}$$

Property (1) is an open property, so that there is a neighborhood N_x of x within which it holds. Property (2) is a closed property and implies that $(p^i - p^j) \cdot x = 0 \forall i, j$ in $\{g+1, \dots, k\}$. Now for any $x \neq 0$, $\{x : x \cdot p^i = 0\}$ has measure zero and has no interior in R^J . Hence in the neighborhood N_x of x there exists a x' such that (2) fails and (1) still holds. The set of such x' is open. This proves the theorem. ■

This result assures us that distinct equilibrium price vectors map into distinct index numbers using almost any set of weights. Hence *we can refer interchangeably to contracts that pay contingent on the value of a price vector or to contracts that pay contingent on the value of a price index.*

Finally, we relate the payoff patterns of the contracts traded in the multi-level economy to those that can be realized by trading index options. The contracts traded in level 2 are goods contingent on the equilibrium price vector selected in E^1 . By the Lemma above, they can be interpreted as goods contingent on the value of a price index. Alternatively in level two we could trade securities that pay if and only if a price index attains a particular value, i.e. their payoff as a function of the index number is zero everywhere except at a single point. This payoff structure can be replicated by the limit of a sequence of option trading strategies. Assume that the critical value of the price index is V . The basic option trading strategy is as follows:

Buy n index call options with exercise price $(V - \frac{1}{n})$.

Sell $2n$ index call options with exercise price V .

Buy n index call options with exercise price $(V + \frac{1}{n})$.

It is routine to verify that as $n \rightarrow \infty$, the payoff function from this strategy converges pointwise to a function that is zero everywhere except at V , where it assumes value 1. Hence the desired payoff structure can be approximated arbitrarily closely by the above strategy for n sufficiently large¹⁴. It is also worth noting that the payoff functions that characterize our index-contingent securities, are those associated with what Rubinstein [24] calls “binary options”: these are also discussed in Cox and Rubinstein [12]. It follows from this that the contracts traded in level 2, goods contingent on the equilibrium price vector selected in level 1, can be replaced by a combination of spot markets and markets for options based on the price index in level 1. The same reinterpretation can of course be carried out at other levels.

¹⁴This analysis of course ignores transaction costs. If the cost of each option trade are positive, then they will become infinite in the limit. The relative payoffs at different values of V are not affected by transaction costs, but their absolute values are affected and may be made negative

5. CONCLUDING COMMENTS

5.1. Related literature. There is some precedent for asking questions about uncertainty arising from economic variables. Hahn in 1973 identified our inability to address this as a key weakness of equilibrium theory [15], and Kurz [19] coined the phrase *endogenous uncertainty* for this uncertainty. Svensson [25] introduced aspects of price uncertainty into a temporary general equilibrium model.¹⁵ Radner [23] has also noted uncertainty about market prices as an important and unstudied problem. Working in a rather different framework, Chichilnisky and Wu [10] formalize endogenous uncertainty within a general equilibrium model with individual and collective risks. They show that this type of uncertainty is crucially dependent on the information structure and can be generated by financial innovation. Hahn [16] has addressed issues related to the present paper in an incomplete market model: his concern is to exhibit the logical inconsistency of rational expectations in a sequence economy where there are multiple equilibria in the second period. Chichilnisky, Hahn and Heal [7] and Chichilnisky [6] develop similar issues in a more general framework. Chichilnisky, Heal and Tsomocos [11] introduce the possibility of default by traders as part of an optimal response to the realization of exogenous variables: the possibility of this default is a source of endogenous uncertainty to other traders.

There is some connection between our results and those on sunspots (Cass and Shell [4]). Both literatures study uncertainty which does not affect the economy's total endowments:¹⁶ indeed sunspots by assumption do not affect any of its attributes. There is therefore a common element in the motivation of the studies. However, there are also big differences in the way the analysis is conducted. Sunspots by assumption do *not* directly affect any real variables: any effect that they have is via traders' beliefs and their impact on traders' behavior. Here, however, the state selected *by definition* has an effect on real variables because it determines the equilibrium chosen. A clear contrast is with Proposition 3 of Cass and Shell [4], which states that with complete markets ("unrestricted market participation") and agreement on probabilities, sunspots do not matter. All of our results occur in the context of complete markets and agreement on probabilities, precisely the case in which sunspot phenomena are not important. In addition, sunspots can matter in economies with unique equilibria ([4], appendix), a case about which we have nothing to say. So the two strands of literature are complementary.

5.2. Further research. A natural development of our present analysis, is in the direction of a model of asset pricing. Derivative securities play a natural and integral role in our model: it would be of interest to investigate the relationship between their prices at equilibrium, the equilibrium prices of goods in the underlying economy E^1 , and the probability distributions over equilibria. It would be of particular interest to have formulae for equilibrium derivative securities prices emerge from a model in

¹⁵Henrotte [17] has also considered options as a method of reducing endogenous price uncertainty.

¹⁶Cass and Shell call this extrinsic uncertainty.

which these securities are not redundant.¹⁷

6. APPENDIX

We have assumed the underlying exchange economy E^1 to possess a finite set of competitive equilibria for any initial endowments. In fact there seems to be no reason why our analysis should not extend certainly to countable sets of equilibria and possibly to continua of equilibria. In these cases, there would be infinite numbers of equilibrium-contingent commodities, in which case the framework of Chichilnisky and Heal [8] could be adopted.

An important preliminary step in our argument is establishing that all equilibria at any level of the multi-level economy E are *fully insured*, i.e., they give consumption vectors which are independent of the equilibria selected in all previous levels.

Definition 4. *We say that E achieves full insurance and its equilibria at any level are fully insured if for any level y , and for any equilibrium k selected at level y , trader i 's consumption vector at level y in equilibrium k in state s_y , $c_{i,k}^{s_y}$, is independent of the state s_y , i.e., independent of which equilibria are selected in all levels prior to y . Equivalently, consumption in equilibrium k at level y is independent of the realization by which the chosen equilibrium at level y is reached.*

Lemma 2. *Under assumptions (A1) to (A3) all equilibria at all levels of E are fully insured.*

Proof. The strategy of the proof is as follows. We show that by strict concavity of utility functions and the fact that the total endowment of the economy is realization-independent, any realization-dependent allocation is dominated by one consisting of its expected values. Hence any Pareto efficient allocation must give realization-independent consumption vectors. Of course, the equilibria at each level, being competitive equilibria of an exchange economy, are Pareto efficient in that economy.

Let p_k^y be the equilibrium price vector of the k -th. equilibrium of the y -th. level, $c_{i,k}^y$ be the associated consumption vector of trader i , and $c_{i,k}^{y_s}$ be the consumption vector of trader i at the k -th. equilibrium of level y in the s -th. realization at that level. Note that for any realization s , $\sum_i c_{i,k}^{y_s} = \sum_i w_i$ as the total endowment of the economy is the same at all realizations and all levels. Define $Ec_{i,k}^y = \sum_{y_s} \pi_{y_s} c_{i,k}^{y_s}$ as the expected consumption of trader i at level y in equilibrium k where the expectation is taken over realizations y_s . By strict concavity of utility functions (A1), we have $U_i(Ec_{i,k}^y) > \sum_{y_s} \pi_{y_s} U_i(c_{i,k}^{y_s})$ provided that the equilibrium consumption vector is not fully insured. It is feasible for trader i to consume $Ec_{i,k}^y$ in each realization, since

¹⁷In the usual arbitrage pricing models, the securities being priced are spanned by others and so from a risk-bearing perspective are redundant.

$\sum_i E c_{i,k}^y = \sum_i \sum_{y_s} \pi_{y_s} c_{i,k}^{y_s} = \sum_i w_i$. Hence $E c_{i,k}^y$ forms a feasible allocation that is Pareto superior to $c_{i,k}^{y_s}$, proving that the equilibrium must be fully insured. ■¹⁸

We shall use the concepts of *utility possibility set* and *Pareto frontier* for the economy E^1 . The *utility possibility set (UPS)* is a subset of \mathfrak{R}^I consisting of utility vectors $\{U_1(c_1), U_2(c_2), \dots, U_I(c_I)\}$ corresponding to feasible allocations in E^1 , i.e.,

$$UPS = \left\{ (U_1(c_1), U_2(c_2), \dots, U_I(c_I)) \in \mathfrak{R}^I : \sum_i c_i \leq \sum_i w_i \right\}$$

The *Pareto frontier (PF)* is the efficient frontier of the utility possibility set, i.e.,

$$PF = \left\{ y \in \mathfrak{R}^I : \sim \exists [U_1(c_1), U_2(c_2), \dots, U_I(c_I)] \geq y^{19} \text{ and } \sum_i c_i \leq \sum_i w_i \right\}$$

We assume that *the UPS is a compact set in R^I* . Closedness is automatic: boundedness requires extra conditions. For boundedness it would suffice if consumption sets were bounded below, or in the case that they are unbounded, that preferences satisfy for example the *limited arbitrage* condition of Chichilnisky [5], which is necessary and sufficient for the compactness of the Pareto frontier.

Lemma 3 establishes that if we add an infinite sequence of levels to the economy E^1 then in the limit the resulting equilibrium consumption allocations are Pareto efficient, and the associated utility vector is in the PF. In Theorem 1 we tighten this result to show that in regular economies it in fact holds after the addition of only a finite number of levels. Recall that $c_{i,k}^y \in \mathfrak{R}^{J \times N^{y-1}}$ is trader i 's consumption at the k -th. equilibrium of the y -th. level: $\sum_{y_s} \pi_{y_s} U_i(c_{i,k}^{y_s})$ is the expected utility of this vector, where the expectation is taken over all equilibria at level y . We abbreviate it to $EU_i[y]$: the vector of expected utility levels for all traders is then $EU[y] = [EU_i[y]]_{i=1, \dots, I}$. Finally, let $d(x, X)$ be the distance between the point x and the nearest element of the set X .

Lemma 3. $\lim_{y \rightarrow \infty} d(EU[y], PF) = 0$. *In words, the vector of traders' expected utilities of consumption across equilibria of the y -th. level converges to the Pareto frontier as the number of levels becomes infinite. Furthermore, the sequence of expected utility vectors $EU[y]_{y=1, 2, \dots}$ is Pareto improving, i.e., $EU[y+1] \geq EU[y]$ (where \geq denotes greater than or equal to in all coordinates and greater in some).*

Proof. $U_i(c_{i,k}^{y_s})$ denotes trader i 's utility from the k -th. equilibrium in the realization y_s : $\sum_{y_s} \pi_{y_s} U_i(c_{i,k}^{y_s})$ is the expected utility from this equilibrium over all realizations. By Lemma 2, $c_{i,k}^{y_s}$ is realization-independent. So we let $c_{i,k}^y$ stand for the

¹⁸Analogous results about fully insured equilibria were established in Malinvaud [20] [21] and in Cass and Shell [4].

consumption of trader i at equilibrium k of level y , without specifying the realization. Define the following subsets of the UPS:

$$I^y = \left\{ x \in UPS \subset R^I : \forall i, x_i \geq \min_k U_i \left(c_{i|k}^y \right) \right\}$$

I^y is the set of utility vectors that give each trader a utility level at least as great as that which the trader obtains at the equilibrium which is worst for that trader at level y . We define B^y similarly, except that the minimum utility level across equilibria is replaced by the expected utility level across equilibria. π_k^y is as usual the probability of equilibrium k being selected at level y .

$$B^y = \left\{ x \in UPS \subset R^I : \forall i, x_i \geq \sum_k \pi_k^y U_i \left(c_{i|k}^y \right) \right\}$$

This is the set of utility vectors in the UPS that give each trader at least his or her expected utility associated with level y . By construction of the levels, $\sum_k \pi_k^y U_i \left(c_{i|k}^y \right)$ equals the expected utility of the endowment vectors of trader i at level $y+1$. Clearly we have $B^y \subset I^y$ unless minimum and expected utility levels are equal. In this case all equilibria give the same utility values and we have a unique equilibrium. By Lemma 1 this is fully insured. As each equilibrium is Pareto efficient by normal arguments this gives a utility vector in the PF and we are done. From now on we assume that there are multiple equilibria at each level. As the equilibria of level $y+1$ are weakly Pareto superior to the endowments of this level, $I^{y+1} \subseteq B^y$. Hence we have the sequence:

$$I^1 \supset B^1 \supseteq I^2 \supset B^2 \supseteq I^3 \supset B^3 \dots I^y \supset B^y ..$$

The subsequence $\{I^y\}, y = 1, 2, \dots$ defines a strictly nested sequence of subsets of the UPS, which is itself a compact set bounded above by the PF. In each dimension i of R^I the greatest lower bound of I^{y+1} exceeds that of I^y (as we have assumed that the equilibrium is not unique), as

$$\forall i, \min_k U_i \left(c_{i|k}^y \right) < \sum_k \pi_k^y U_i \left(c_{i|k}^y \right) \leq \min_k U_i \left(c_{i|k}^{y+1} \right) \quad (3)$$

But $\lim_{y \rightarrow \infty} \left\{ \min_k U_i \left(c_{i|k}^y \right) - \min_k U_i \left(c_{i|k}^{y+1} \right) \right\} = 0$, because the sets I^y are bounded above. Now $\min_k U_i \left(c_{i|k}^y \right) - \min_k U_i \left(c_{i|k}^{y+1} \right) = 0$ implies $\min_k U_i \left(c_{i|k}^y \right) = \sum_k \pi_k^y U_i \left(c_{i|k}^y \right)$. This is true only if the equilibrium at level y is unique. Hence in the limit the equilibria at level y converge to a unique equilibrium, which is Pareto efficient. By (3) above expected utility vectors are Pareto improving as y increases. This completes the proof. ■

We now give the proof of Theorem 2:

Proof of Theorem 2. If an initial endowment vector $W^1 \in \mathfrak{R}^{LJ}$ is Pareto efficient in E^1 , then there is a unique no-trade equilibrium associated with this, which is just the initial endowment vector. By Debreu's theorem [14], there is a neighborhood of this initial endowment within which the equilibrium is still unique. Hence there is a neighborhood, denoted Ψ , of the set of Pareto efficient allocations in E^1 , such that if the initial endowment is in Ψ then the equilibrium is unique. Similarly there is such a neighborhood Ψ^y in E^y for any level y .

By Lemma 3, and the fact that the Pareto frontier is bounded, we know that for any agent i and any $\epsilon > 0$, $\exists y(\epsilon)$ such that $\min_k U_i(c_{i,k}^y) - \max_k U(c_{i,k}^y) < \epsilon$, whenever $y \geq y(\epsilon)$, i.e., the utility differences across equilibria go to zero for all agents as the number of levels goes to infinity. Hence the welfare loss due to uncertainty about which equilibrium (which state) will be selected also goes to zero. From this it follows that for some $y + 1$ the initial endowments (the equilibria of level y) will be in the neighborhood Ψ^{y+1} of the set of efficient allocations within which there is a unique competitive equilibrium. At this stage the competitive equilibrium of level $y + 1$ will be unique, and no more layers will be needed. ■

We now give a proof of proposition 1 and show that there will be a unique equilibrium at the first level of derivatives, E^2 , if there is one equilibrium which is clearly the most likely, in the sense that its probability is close to unity. Recall Proposition 1:

Proposition 1. *If the probabilities of the equilibria in the underlying exchange economy are sufficiently concentrated around one equilibrium, then the market for price-contingent contracts will have a unique equilibrium. Formally, let $CE(E^1)$ be the set of competitive equilibria of the underlying exchange economy, and $\pi_i^1, i \in CE(E^1)$, be the probabilities of these equilibria being realized. Let $\#^1$ be the number of equilibria in $CE(E^1)$. Then $\exists \epsilon > 0$: if for some i , $\pi_i^1 > 1 - \epsilon$ and $\pi_j^1 < \epsilon/\#^1 \forall j \neq i$, the economy E^2 has a unique competitive equilibrium.*

Proof. We give an outline of the proof, which is an obvious modification of that given above for Theorem 2. It also uses the fact that there is a neighborhood of the contract curve within which initial endowments lead to a unique equilibrium. If the statistical expectation of the equilibrium consumption vectors is close to one equilibrium, then the cost of risk-bearing arising from uncertainty about the equilibrium to be chosen is small. Hence the equilibrium consumptions in E^1 , which are the endowments of E^2 , are close to the set of efficient allocations in E^2 . Hence E^2 has a unique equilibrium. ■

Note that this result could be proved via a different route. One can show that the number of equilibria in the market for price-contingent commodities is a locally constant function of the probabilities of the various equilibria in the underlying exchange economy. (This is a modification of the standard argument due to Debreu showing that this number is locally constant with respect to endowments.) The number of equilibria in the contingent market is clearly one if for some i , $\pi_i = 1$ and $\pi_j = 0 \forall j \neq i$. This will then complete a proof.

REFERENCES

- [1] K.J. ARROW (1953). "The role of securities in an optimal allocation of risk-bearing." *Review of Economic Studies*, 1964, 31, pp. 91-96. Reprinted and translated from *Econometrie*, 1953.
- [2] Y. BALASKO (1975). "The graph of the Walras correspondence." *Econometrica*, Vol 43, pp. 907-912.
- [3] Y. BALASKO (1988). *Foundations of the theory of general equilibrium*. Academic Press, Boston.
- [4] D. CASS and K. SHELL (1983). "Do sunspots matter?" *Journal of Political Economy*, 91, pp. 193-227.
- [5] G. CHICHILNISKY (1993). "Arbitrage, gains from trade and social diversity: a unified perspective on resource allocation." Working Paper, Columbia University Department of Economics. Presented at the 1994 AEA meetings, *American Economic Association Papers and Proceedings*, May 1994.
- [6] G. CHICHILNISKY (1993). "Existence and optimality of equilibrium with endogenous uncertainty." Working Paper, Columbia University Department of Economics.
- [7] G. CHICHILNISKY, F.H. HAHN and G.M. HEAL (1992). "Price uncertainty and incomplete markets." Columbia Business School First Boston Working Paper, New York, forthcoming in G. Chichilnisky (Editor) *The Challenge of Economic Theory: Essays in Honor of Kenneth Arrow*, Cambridge University Press, forthcoming.
- [8] G. CHICHILNISKY and G.M. HEAL (1992). "Competitive equilibrium in Sobolev spaces without bounds on short sales." *Journal of Economic Theory*, 1993.
- [9] G. CHICHILNISKY, G.M. HEAL, P. STREUFERT and J. SWINKELS (1991). "Believing in multiple equilibria". Columbia Business School First Boston Working Paper, New York.
- [10] G. CHICHILNISKY and H.-M. WU (1991). "Financial innovation and default in incomplete markets". Columbia Business School First Boston Working Paper, New York.
- [11] G. CHICHILNISKY, G.M. HEAL and D.P. TSOMOCOS. "Option values and endogenous uncertainty in ESOPs, MBOs and asset-backed loans". Working paper, Columbia Business School, 1993.

- [12] J. COX and M. RUBINSTEIN (1985). *Options Markets*, Prentice Hall.
- [13] J.H. DREZE (1993). "Uncertainty in economics" in G. Chichilnisky (Editor) *The Challenge of Economic Theory: Essays in Honor of Kenneth Arrow*, Cambridge University Press, forthcoming.
- [14] G. DEBREU (1970). "Economies with a finite set of equilibria." *Econometrica*, 38, pp. 387-92.
- [15] F.H. HAHN. "On the notion of equilibrium". Inaugural Lecture, Cambridge University, 1973. Published by Cambridge University Press.
- [16] F.H. HAHN (1991). "A remark on incomplete market equilibrium." Working paper circulated at Stanford Institute for Theoretical Economics, July, 1991, forthcoming in G. Chichilnisky (Editor) *The Challenge of Economic Theory: Essays in Honor of Kenneth Arrow*, Cambridge University Press, forthcoming.
- [17] P. HENROTTE (1992). "Existence and optimality of equilibria in markets with tradeable derivative securities." Technical Report No. 48, August 1992, Stanford Institute for Theoretical Economics.
- [18] J.D. HERRMAN (1992). "Genericity analysis from the fibre bundle viewpoint." Mimeograph, Department of Mathematics, Rutgers University, 1992.
- [19] M. KURZ (1974). "The Kesten-Stigum model and the treatment of uncertainty in equilibrium theory." In *Essays on Economic Behavior Under Uncertainty*, Eds. M.S. Balch *et al*, North Holland, Amsterdam, pp. 389-399.
- [20] E. MALINVAUD (1972). "The allocation of small risks in large markets." *Journal of Economic Theory*, 4, pp. 312-28.
- [21] E. MALINVAUD (1973). "Markets for an exchange economy with individual risks." *Econometrica*, 41, 383-412.
- [22] A. MAS-COLELL and J.H. NACHBAR (1991). "On the finiteness of the number of critical equilibria, with an application to random selections." *Journal of Mathematical Economics*, Vol. 20, 397-409.
- [23] R. RADNER. "Equilibrium under uncertainty." Chapter 20 of *Handbook of Mathematical Economics*, Vol. II, North Holland, 1982 (eds. K.J. Arrow and M.D. Intriligator).
- [24] M. RUBINSTEIN (1991). "Exotic options." Mimeo, Department of Finance, Haas School of Business, University of California at Berkeley.
- [25] L.E.O. SVENSSON (1981). "Efficiency and speculation in a model with price-contingent contracts." *Econometrica*, Vol. 49, No. 1, pp. 131-51.

