Compact Securities Markets, Common Independence Knowledge, and the Complexity of Arbitrage (extended abstract)

David M. Pennock and Michael P. Wellman
University of Michigan AI Laboratory
1101 Beal Avenue
Ann Arbor, MI 48109-2110 USA
{dpennock,wellman}@umich.edu

March 3, 1998

Abstract

In a world where uncertainty is ubiquitous, a rational, risk-averse agent will seek to establish contingency plans in case of undesirable outcomes. The fundamental tool in economics and finance for agents to manage risk is the *securities market*, which allows for the purchase of goods whose receipt is tied to the realization of some state of nature. In general, all agents can hedge all of their risk only if the securities market is *complete*, or spans all possible states of nature, and most general equilibrium results break down for incomplete markets. But this number of theoretically required securities is, simply put, impossibly huge. We explore the prospect of more compact securities markets, in which agents can still construct any desired hedge. We show that, if there is *common knowledge* of risk-neutral independencies among some events in equilibrium, then a significant reduction in the number of markets is possible. We also argue that pricing redundant securities—that is, recognizing arbitrage opportunities—is, in general, computationally intractable.

Keywords: securities markets, probabilistic reasoning, common knowledge, arbitrage, risk-neutral probability

Joint Session: This paper may be appropriate for one of the TARK/Microeconomic Theory joint sessions.

1 Introduction

Securities markets play a central role in the economics of uncertainty. The existence of securities essentially allows agents to arrange for the transfer of wealth, contingent on the outcome of some uncertain events (e.g., to purchase insurance). Suppose that agents are uncertain as to which of S disjoint states of nature will occur (each agent with, in general, heterogeneous beliefs and utilities for outcomes). In order to guarantee that all agents can optimally reallocate all of their risk, the market must contain at least S securities, in which case the market is termed *complete*. But the number of states S can typically be considered exponential in the number of random events that agents are concerned with (e.g., if it will rain tomorrow, if my house will burn down, if IBM's stock will plummet, etc.). For this reason, the existence of a truly complete securities market is, for all practical purposes, impossible.

Research in uncertain reasoning is largely concerned with representing beliefs over large state spaces, and one of the major ideas is that structure in this set may lead to more compact specifications. In particular, the state space is typically factored into $N = \lceil \log_2 S \rceil$ random variables, such that their 2^N possible joint outcomes span all states of nature. If no other structure exists, then this factored representation is no more succinct than the full disjoint state space. On the other hand, if there are enough independencies among variables, a substantial savings in encoding (e.g., in Bayesian networks) may be achieved.

A natural question, then, is whether ideas about structuring probabilistic relations among random variables can be exploited in the design of configurations of securities markets. In prior work [11], we presented some evidence for an answer in the affirmative, in the form of a mapping from Bayesian networks to economies, where the competitive equilibrium of the latter represented the probabilities of the former, though we essentially assumed that producer agents enforce independencies. In this paper, we address the issue more directly, identifying a condition under which independence is selfenforcing among rational agents. We first argue that drawing the obvious analogy, where independencies are assumed to be common knowledge among agents, is of little or no help. However, if instead we assume a state of common knowledge of *risk-neutral* independencies in equilibrium, then compact securities markets may be possible. In addition, we consider some implications of this approach for the computational problems faced by agents. Specifically, pricing a redundant security consistently with existing securities, or equivalently identifying arbitrage, is #P-hard.

Section 2 covers some background material from economics and probabilistic reasoning, and introduces notation. Section 3 presents our main result, linking common knowledge with the possibility of compact, complete securities markets. Section 4 analyzes the complexity of arbitrage, and Section 5 concludes.

2 Background and Notation

The space of all possible world outcomes is assumed finite and discrete, denoted $\Omega = \{\omega_1, \omega_2, \ldots \omega_S\}$, where the ω are disjoint, exhaustive events. Each market participant ("agent") i has a subjective probability distribution $\Pr_i(\cdot)$ over Ω; its probability for an event E is then just $\Pr_i(E) = \sum_{\omega \in E} \Pr_i(\omega)$. Each agent has a set of available actions $a_i \in A_i$ (e.g., buy/sell/produce quantities of goods). Each agent has a utility $u_i(a_i, \omega)$ defined over the cross-product of actions and outcomes, and chooses the action $a_i \in A_i$ that maximizes its expected utility. When only one agent is under consideration, we simplify by dropping subscripts.

The next three subsections introduce relevant background concepts, discussing securities markets, risk-neutral probability, and compact probabilistic representations, respectively.

2.1 Securities markets

Economics has developed a substantial body of theory addressing uncertainty in markets.[1] The standard theory posits *securities*—each of which pays off in a common good (usually money) only if some uncertain event is realized. One of the key results of this theory is that an economy with a *complete* securities

[1]See a standard text on microeconomic theory (e.g., Mas-Colell, et al. [6]) for the precise statement and detailed development of these concepts.

market, plus *spot* markets in standard (uncontingent) goods, possesses the classical properties of competitive equilibrium (i.e., existence, uniqueness, Pareto optimality) under conditions identical to that for the case of certainty. To be complete, a market must consist of S securities, one for every possible state of nature (or any set of S exhaustive, non-redundant securities). If the set of markets is *incomplete*, then the market's allocation of risk may be inefficient (i.e., Pareto dominated), even if the economy is otherwise well-behaved.

We assume that one unit of a security, denoted $\langle E \rangle$, pays off one dollar if the event E occurs, and nothing otherwise. We also assume that all other goods can be exchanged with money, and that agents may borrow money to buy securities, or borrow securities to sell short. We denote the price of security $\langle E \rangle$ as $p^{\langle E \rangle}$.

The equilibrium prices of securities actually form a probability distribution over their corresponding events [8, 14]. To see why this is the case, recall de Finetti's argument for the existence of subjective probabilities [2]. Imagine that a bettor must decide at what price to buy and sell each of two securities: one pays \$1 if E_1 occurs, and nothing otherwise; the other pays \$1 if E_2. If E_1 and E_2 are disjoint events, then the bettor's total price for both tickets should never exceed \$1—otherwise, he is subject to a guaranteed loss: a shrewd opponent will sell both tickets to him for more than \$1, paying out at most \$1 in all outcomes. More generally, if the bettor's internal prices for a set of such securities are not legal probabilities, then he is exposed to sure losses, or *Dutch books*; conversely, the opponent can arrange for risk-free profit or *arbitrage*. An arbitrage opportunity exists if a portfolio of securities yields nonnegative returns in all outcomes, and positive return in at least one. A market in equilibrium must be arbitrage free [8, 14] (intuitively, if opportunities for risk-free profit still exist, then rational agents must not be maximizing utility). It follows, from de Finetti's argument, that the equilibrium prices in a securities market form a legal probability distribution. Note that the converse is also true—a market that is arbitrage free must be in equilibrium [8].

A *conditional security* $\langle E_1 | E_2 \rangle$ pays off exactly as $\langle E_1 \rangle$, conditional on the event that E_2 occurs. If $\overline{E}_2$ occurs, then the bet is called off, and any initial price paid for the security is returned. If E_2 occurs, the security pays off \$1 if E_1 also occurs, and nothing otherwise. In an arbitrage-free market, conditional security prices also obey the laws of probability. For example, the price of $\langle E_1 | E_2 \rangle$ times the price of $\langle E_2 \rangle$ equals the price of $\langle E_1 E_2 \rangle$ [2]. A set of S conditional securities can also form a complete market.

A securities market serves two main purposes. First, it enables agents to transfer wealth across states of nature, thereby hedging the inevitable risks posed by an uncertain world. Second, it functions as an aggregator and disseminator of information about uncertain events. Agents may have different beliefs about the relative likelihood of the possible states, and so securities markets may enable trading opportunities due to differences in subjective expected returns. We have explored the use of securities markets for this purpose, analyzing the properties of equilibrium prices as a representation of aggregate belief [12].

Securities markets are not just theoretical constructs. Real financial markets—in stocks, futures, options, and other derivatives—serve to continuously reallocate risk around the globe, and their going prices disseminate a tremendous amount of aggregate information concerning uncertain events.

2.2 Risk-neutral probabilities, equilibrium, and common knowledge

The analysis in this subsection applies to a pure securities market, without spot markets in standard goods, and where agents' utilities are defined over dollars.

An agent's behavior is the product of both its probabilities and utility. Specifically, an agent chooses the action $a \in A$ that maximizes its expected utility:

$$U(a) = E[u(Y_{a,\omega})] = \sum_{\omega \in \Omega} \Pr(\omega) u(Y_{a,\omega}), \tag{1}$$

where $Y_{a,\omega}$ is its payoff if action a is taken and state ω is realized. Notice that anyone observing this agent's choices cannot fully determine either Pr or u. Any action can be explained equally well by a family of transformed probabilities $\kappa \Pr(\omega) f(\omega)$ and utilities $u(Y_{a,\omega})/f(\omega)$, for any function $f(\omega)$, where κ is a normalization constant. An agent's *risk-neutral* probabilities (RNPs) are the transformed probabilities corresponding to risk-neutral transformed utility. Mathematically, RNPs are:

$$\Pr^{RN}(\omega) \equiv \kappa \Pr(\omega) u'(Y_{a,\omega}),$$

where u' is the agent's marginal utility for money in state ω [7]. A hypothetical agent with transformed probabilities $\Pr^{RN}$ and transformed risk-neutral utility $u^{RN} = u/u'$ has the same expected utility, and thus the same optimizing behavior, as an agent with Pr and u.[2]

In a securities market, a risk-neutral agent would choose to buy (sell) $\langle E \rangle$ if $p^{\langle E \rangle} < (>)\ \Pr(E)$, since it simply maximizes expected payoff. Then *any* agent would choose to buy (sell) $\langle E \rangle$ if $p^{\langle E \rangle} < (>)\ \Pr^{RN}(E)$. A market is in *equilibrium* when all opportunities for trade have been exhausted. If two agents h and i have differing RNPs for E—that is, $\Pr_h^{RN}(E) \neq \Pr_i^{RN}(E)$—then they will desire to trade. It follows that, at equilibrium, all agents' RNPs agree. Furthermore, since offers to buy and sell must match, the equilibrium prices equal these consensus RNPs [8].

A fact F is *common knowledge* among a group if every member knows F, knows that everyone else knows F, knows that everyone else knows that everyone else knows F, etc. [3] Equilibrium can be characterized, in a sense, as a state of common knowledge of risk-neutral probabilities (CKORNP) over available securities.[3] At equilibrium in an *incomplete* market, all desirable trades in available securities have occurred, but there may exist acceptable trades in unavailable ones. Equilibrium in a *complete* securities market is a state of CKORNP over all states of nature, ensuring that all potentially desirable transfers have occurred, and thus yielding a Pareto optimal allocation of risk. In other words, agents' behaviors in a complete market equilibrium are equivalent to those of risk-neutral agents with identical beliefs over Ω: there are simply no differences of opinion or utility on which to trade.

2.3 Compact probability representations

Researchers in probabilistic reasoning use *Bayesian networks* to encode joint probability distributions, over all S states, more compactly, by exploiting independencies among random events [5, 9, 10, 13].

Any space of S disjoint states can be factored into a set of $N = \lceil \log_2 S \rceil$ exhaustive, orthogonal events, $\{E_1, E_2, \ldots, E_N\}$. A *joint* probability distribution, $\Pr(E_1, E_2, \ldots, E_N)$, assigns a probability to all 2^N possible combinations of event outcomes, and thus to each of the original S states. With no further assumptions, the two representations are equivalent in both expressivity and size, although the factored set may be more natural, since

[2] Expected utilities will differ only by the multiplicative factor κ.

[3] We suspect that equal RNPs are not enough to ensure equilibrium. In any case, CKORNP is certainly required for *common knowledge* of equilibrium.

it is usually the outcomes of random events that agents are most directly interested in.

If, in addition, some events are conditionally independent of others, we can represent the joint distribution more compactly. Consider the event E_j. Suppose that, given a set of preceding events $\mathbf{pa}(E_j) \subseteq \{E_1, \ldots, E_{j-1}\}$, E_j is conditionally independent of all other preceding events. We call $\mathbf{pa}(E_j)$ the *parents* of E_j. Then we can rewrite the joint probability distribution in a (usually) more compact form:

$$\Pr(E_1, E_2, \ldots, E_N) = \prod_{j=1}^{N} \Pr(E_j | \mathbf{pa}(E_j)). \tag{2}$$

For each variable E_j, we need only record the conditional probabilities $\Pr(E_j|\mathbf{pa}(E_j))$, for all possible combinations of outcomes of E_j and $\mathbf{pa}(E_j)$. If k is the maximum number of parents per variable, then a specification of $O(N2^k)$ probabilities implicitly represents the full joint distribution.

All of the conditional independencies implicit in (2) can be represented graphically as a *directed acyclic graph* (DAG). Each event is a node or vertex in the graph, and directed edges encode parent relationships: there is a directed edge from E_i to E_j if and only if E_i is parent of E_j. Notice that recording all of the conditional independencies in this way requires $O(N^2)$ space.

Bayesian inference algorithms can compute any of the probabilities implicit in a Bayesian network, though the task is worst-case computationally intractable [1, 9].

3 Compact securities markets and common knowledge

Can conditional independence be exploited to reduce the size of securities markets? A natural approach might be to assume that all of the conditional independencies in a Bayesian network representation are common knowledge among agents. However, in general, the independencies represented in the equilibrium prices will *not* reflect commonly held independencies among agents, nor vice versa. In fact, researchers in *belief aggregation* have not found any reasonable function that can both combine a group's probability distributions and preserve unanimously held independencies, proving that this *independence preservation property* is incompatible with other desirable properties [4]. If, instead, we assume that the agents' *risk-neutral* conditional independencies are commonly known in equilibrium, then a more compact set of securities may be sufficient to support all desirable hedges.

In Section 2.2, we characterized an equilibrium in a complete securities market as a state of common knowledge of risk-neutral probabilities (CKORNP) over all states of nature $\omega \in \Omega$, in which case all acceptable trades have been exhausted, and the allocation is Pareto optimal. One possible method for eliminating securities is to somehow identify a subset of K securities such that, at equilibrium, all agents' RNPs also agree on the remaining $S - K$ states, and this agreement is commonly known. Then equilibrium ensures that there is no further possibility for trade in the K securities, and CKORNP ensures that there is no possibility for mutually beneficial trade in the remaining $S - K$ states. However, even if this were possible, the statement of common knowledge itself would presumedly consist of $S - K$ numbers, and the underlying size of the representation would be no smaller.

We believe that an assumption of common knowledge of risk-neutral independence (CKORNI, pronounced "corny") may be more reasonable. For example, consider a space of four states, $\{E_1E_2, E_1\overline{E}_2, \overline{E}_1E_2, \overline{E}_1\overline{E}_2\}$, factored into two events, and a market with only two securities, $\langle E_1 \rangle$ and $\langle E_2 \rangle$. If, at equilibrium, all agents RNPs agree that the events are independent, then their RNPs must agree across all four states. Even if another security were introduced, say $\langle E_1E_2 \rangle$, there would be no possiblity of mutually beneficial trade in it. Furthermore, its equilibrium price is already determined: it would be $p^{\langle E_1 \rangle}p^{\langle E_2 \rangle}$. In this sense, the smaller market is still complete, and succeeds in allocating risk in a Pareto optimal fashion, albeit under a fairly strong assumption of risk-neutral independence, and common knowledge of same. The main advantage of assuming CKORNI, instead of directly assuming CKORNP, is that an exponential reduction in securities may be possible with only a polynomial size statement of common knowledge.

We extend this reasoning to the more general case of S states factored into $N = \lceil \log_2 S \rceil$ events, and a market in only $K = O(N2^k)$ *conditional* securities, structured as a Bayesian network. That is, every conditional probability $\Pr(E_j|\mathbf{pa}(E_j))$ in the network is associated with a conditional security $\langle E_j|\mathbf{pa}(E_j) \rangle$ in the market. If, at equilibrium, all agents are in a state of CKORNI for those conditional independencies represented in the network's DAG, then they must be in a state of CKORNP over all S states. No further trades are mutually beneficial, and the current allocation is Pareto optimal. The statement of common knowledge is of size $O(N^2)$, while the reduction in securities, from 2^N to $O(N2^k)$, can be tremendous. Of course, if there are no commonly held risk-neutral conditional independencies in equilibrium, then S securities are still needed.

The above arguments hold only in equilibrium, and we have made no mention of how a state of CKORNI might be reached. We suggest an admittedly sketchy mechanism. Begin with securities in only the N random events $\{E_1, \ldots, E_N\}$. If any agent's demand for $\langle E_j | E_i \rangle$ (for any $i < j$) at price $p^{\langle E_j \rangle}$ is nonzero, then it creates a new market in $\langle E_j | E_i \rangle$. If, at some future time, the agent has zero demand for its new security, then it may retract the security. An additional condition for equilibrium is that no agent desires to create any new markets. Then, in equilibrium, it should be the case that all agents' RNPs agree across all states. We might want to add a transaction cost for opening new markets, so that equilibrium only ensures that risks are hedged up to a threshold cost. In some loose sense, this may be how real financial markets operate. New derivative assets, and customized exotic options, are created to meet the demands of investors and institutions to hedge particular risks. The absence of a market may imply that any advantage of a more precise hedge is outweighed by the cost of opening the new market.

The reader may object, feeling that CKORNI is a rather unrealistic assumption. While we have no direct evidence to the contrary, we believe that it is at least as plausible as common knowledge of "true" independence. While true probabilities are internal to an agent and, as discussed in Section 2.2, indeterminate from the perspective of an outside observer, an agent's risk-neutral probabilities might be inferred from its observable behaviors.

4 Recognizing arbitrage is #P-hard

In this section, we examine the computational complexity of finding arbitrage opportunities in a compact securities market under CKORNI. First, we note that any derivation of time complexity is in terms of the problem's input size, and the "input" to an arbitrage-finding algorithm is not uniquely defined. If the agents' probability distributions over Ω are considered a part of the input, then identifying arbitrage is relatively trivial. Instead, we conceptually analyze the problem's complexity from the perspective of an outsider observer, searching for any source of risk-free profit within the market. Then the input consists of the price of every available security, plus any statement of common knowledge.

Suppose that an equilibrium is reached in a compact market of $O(N2^k)$ securities, and that the participating agents announce a state of CKORNI, with a description length of $O(N^2)$ (as in Section 3), the implication of which is that all agents are in a state of CKORNP over Ω. Now imagine that a redundant security is introduced, say $\langle E_N \rangle$. In the given situation, the equilibrium price of $\langle E_N \rangle$ is already determined: it must equal $\Pr^{RN}(E_N)$; furthermore, if the current price does *not* equal $\Pr^{RN}(E_N)$, then the market is not in equilibrium, and an arbitrage is possible (recall that a market is in equilibrium if and only if it is arbitrage free [8]). For example, if $p^{\langle E_N \rangle} < \Pr^{RN}(E_N)$, then the observer could purchase it at the going price and sell it to any of the agents at price $p*$ such that $p^{\langle E_N \rangle} < p* < \Pr^{RN}(E_N)$. Although the observer does not have direct access to $\Pr^{RN}(E_N)$, it is uniquely computable from the original prices and the statement of common knowledge.

If the observer can identify an arbitrage by correctly pricing the redundant security, then it can perform Bayesian inference, which is known to be #P-complete [1, 9]. We therefore argue that, given our definition of the input, and allowing for statements of common knowledge, finding arbitrage is #P-hard. A problem is #P-hard (pronounced "number-P hard") if the problem of counting the number of solutions to a satisfiability problem can be reduced to it. Simply determining whether a solution to a satisfiability problem exists is NP-hard, so any problem that is #P-hard is also NP-hard.

5 Conclusion

The realization that large joint probability distributions can be represented compactly, by exploiting conditional independence, was a breakthrough in probabilistic reasoning, essentially shifting the task from the realm of theoretical desirability but practical impossibility, to one of significant real-world value and wide applicability [5, 9, 10, 13]. Complete securities markets, a requirement for Pareto optimal allocation of risk in economic settings, suffer from the same unmanageable size of representation. We have identified one condition—common knowledge of risk-neutral independence (CKORNI) in equilibrium—for which more compact securities markets may be possible. The question of whether CKORNI is a realistic assumption in any sense remains open. Though the concept certainly seems less plausible or natural than individual subjective independence, we believe that it may be at least as attainable as common knowledge of "true" independence among a group.

References

[1] G. Cooper. The computational complexity of probabilistic inference using Bayes belief networks. *Artificial Intelligence*, 42:393–405, 1990.

[2] Bruno de Finetti. *Theory of Probability: A Critical Introductory Treatment*, volume 1. J. Wiley, New York, 1974.

[3] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge.* MIT Press, Cambridge, MA, 1996.

[4] Christian Genest and James V. Zidek. Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 1(1):114–148, 1986.

[5] Finn V. Jensen. *An Introduction to Bayesian Networks.* Springer, New York, 1996.

[6] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory.* Oxford University Press, New York, 1995.

[7] Robert F. Nau. Coherent decision analysis with inseparable probabilities and utilities. *Journal of Risk and Uncertainty*, 10(1):71–91, January 1995.

[8] Robert F. Nau and Kevin F. McCardle. Arbitrage, rationality, and equilibrium. *Theory and Decision*, 31:199–240, 1991.

[9] Richard E. Neapolitan. *Probabilistic Reasoning in Expert Systems: Theory and Algorithms.* John Wiley and Sons, New York, 1990.

[10] J. Pearl. *Probabilistic Reasoning in Intelligent Systems.* Morgan Kaufmann, 1988.

[11] David M. Pennock and Michael P. Wellman. Toward a market model for Bayesian inference. In *Proceedings of the 12th Annual Conference on Uncertainty in Artificial Intelligence (UAI-96)*, pages 405–413, Portland, OR, USA, August 1996.

[12] David M. Pennock and Michael P. Wellman. Representing aggregate belief through the competitive equilibrium of a securities market. In *Proceedings of the 13th Annual Conference on Uncertainty in Artificial Intelligence (UAI–97)*, pages 392–400, Providence, RI, USA, July 1997.

[13] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice-Hall, New Jersey, 1995.

[14] Hal R. Varian. The arbitrage principle in financial economics. *Journal of Economic Perspectives*, 1(2):55–72, Fall 1987.