The role of beliefs in economic theory
Morris, Stephen Edward
ProQuest Dissertations and Theses; 1991; ProQuest Dissertations & Theses Global

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

U·M·I

University Microfilms International
A Bell & Howell Information Company
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
313/761-4700 800/521-0600

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Order Number 9221426

The role of beliefs in economic theory

Morris, Stephen Edward, Ph.D.

Yale University, 1991

U·M·I
300 N. Zeeb Rd.
Ann Arbor, MI 48106

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

The Role of Beliefs in Economic Theory

A Dissertation

Presented to the Faculty of the Graduate School

of

Yale University

in Candidacy for the Degree of

Doctor of Philosophy

by

Stephen Edward Morris

December 1991

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

ABSTRACT

The Role of Beliefs in Economic Theory

Stephen Edward Morris

Yale University

1991

The dissertation extends standard analyses of trading with asymmetric information, "agreeing to disagree" and the value of information. I assume that rational agents maximize expected utility subject to some beliefs, but do not make two other standard, but conceptually and empirically independent claims: that agents update beliefs correctly in the light of new information and that agents with the same information have the same beliefs. In chapter two, I extend Harsanyi's notion of consistency of priors to derive necessary and sufficient conditions on the heterogeneity of beliefs to preclude trade, when agents have both asymmetric information and different prior beliefs. In chapter three, I present a general framework for analyzing the value of information, when agents may misinterpret the information they receive. I derive a generalization of Blackwell's Theorem and relate it to Geanakoplos' (1989) results for possibility correspondences. In chapter four, I show the relationship between "no trade", "no speculation" and "agreeing to disagree" results for agents with common priors who misinterpret information and agents with heterogeneous priors who process information correctly.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Preface

I would like to thank my advisor John Geanakoplos both for pivotal support and guidance in writing this dissertation and for nurturing my interest in economic theory. Truman Bewley and David Pearce played a similar dual role, and I thank them too.

Chapters II and V benefitted from comments by Jim Robinson, Torben Andersen, Curtis Taylor and seminar participants at Northwestern University, Boston University, Harvard Business School, the Kennedy School, the California Institute of Technology, and the Universities of Pennsylvania, Michigan, California at San Diego, California at Los Angeles and Iowa. Hal Varian provided me with some extremely valuable references.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Contents

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Chapter I

INTRODUCTION

There are three components to economists' standard models of rational choice under uncertainty. First, agents make consistent choices: this implies subjective expected utility maximization under the now standard assumptions of Savage (1972)[1]. Second, agents update their beliefs rationally, given new information (Bayesian updating). Thirdly, agents with the same information have the same beliefs (the common prior assumption). Whatever the merits of these three kinds of assumptions, as empirical facts or theoretical constructs, the three are logically and conceptually separate. Only the first - the idea of consistent choices - is part of the analytic core of neoclassical theory[2]. This dissertation is concerned with agents who make mistakes processing information, and agents who have different prior beliefs.

If different prior beliefs about the world were responsible for all differences in opinion (i.e. agents have the same objective information), then agents would be prepared to bet against each other. Yet people do not tend to bet on the basis of every difference of opinion[3]. But if all differences of opinion are explained by differences of information (i.e. agents have common prior beliefs), and there is common knowledge of rationality, then no trade can be explained by the differences of opinion alone, because other agents' willingness to trade reveals

---

1. Although the assumptions generating subjective expected utility maximization can be relaxed, e.g. the sure thing principle [Machina and Schmeidler (1990)] and completeness [Bewley (1986)].

2. Geanakoplos (1989) notes that errors in information processing do not make behavior "any less goal oriented and purposeful, and therefore less predictable."

3. This is a traditional argument for the common prior assumption; see Aumann (1987), Bewley (1986), Bernheim (1987).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

sufficient information to preclude trade[4]. Yet we do observe trades generated by asymmetric information. In chapter two, I look at the interaction between heterogeneous priors and asymmetric information in trading decisions. In such environments, it is possible that sometimes agents rationally do not trade with each other despite differences in prior beliefs, while other times agents may be prepared to trade on the basis of asymmetric information, but not if that information was public.

I identify necessary and sufficient conditions on the similarity of agents' beliefs to preclude trade in different environments. The key variable here is what kind of trades can be carried out. If trades are allowed to depend on agents' private information, then I say that *unrestricted* trade takes place. But typically it may not be possible to verify agents' private information ex post. If trades are allowed to depend on agents' private information, subject to incentives to truthfully reveal their information, then I say that *incentive compatible* trade takes place. Finally, if trades cannot be made contingent on agents' private information at all, I say that *public* trade takes place.

In section 2 of chapter II, I look at the case where agents are risk-neutral and there is a single good: in other words, a *betting* environment. In this case, unrestricted trade is precluded if agents have *consistent* beliefs, i.e. their posterior beliefs could have been derived from some "common prior" (even if, in fact, they weren't). Incentive compatible trade is precluded if agents have *reconcilable* beliefs. "Reconcilability" has the interpretation that agents may consider their own signals to be of less significance than others do. Public trade is precluded if agents have *publicly consistent* beliefs. Beliefs are publicly consistent if each agent's prior could be replaced by another prior with the same posterior beliefs conditional on

4. There is a literature of such "no speculation" results. See Geanakoplos (1988) for a survey.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

his private information, such that the replacement priors agree on the probability of publicly known and verifiable events (but may disagree about the probability of private signals).

But risk-neutral agents may be of little empirical relevance. So I consider a standard question from the literature. Suppose that, in the absence of or before observing their private information, risk-averse agents are not prepared to trade with each other. Suppose then that new (possibly asymmetric) information arrives, which is not payoff-relevant, but which could not be contracted on before. Does the arrival of new information lead them to trade with each other? Under the common prior assumption the answer is again no: in fact, Milgrom and Stokey (1982) showed that it is sufficient that agents' have the same beliefs about the distribution of the new information, conditional on the payoff-relevant events. In section 4 of chapter II, I show that properties analogous to those of section 2, but placing restrictions only on the conditional beliefs, are necessary and sufficient to preclude trade with risk-averse agents. The result is demonstrated by first showing (in section 3) that the same conditions hold for risk-neutral agents, and then showing (in section 4) that utility functions do not matter in ex ante efficient environments.

In chapter III, I derive a generalization of Blackwell's theorem ranking the informativeness of experiments[5] to agents who make mistakes processing information. An experiment is said to be more valuable than another if ex ante utility from observing one experiment is higher than from the other, for every decision problem. This definition can be extended to agents who make mistakes, but ambiguities may arise when different optimal decision rules imply different ex ante utilities. Section 1 of chapter III proposes alternative ways around the ambiguity: it can be resolved by requiring only that ex ante utility is higher for "almost all" decision problems.

5. Blackwell (1952).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

I show that an experiment is more valuable than another if and only if the difference between the two distributions of signals could have been derived from two other experiments where beliefs about one experiment are correct and sufficient for the other. This result specializes to Blackwell's theorem for agents who do not make mistakes.

A natural way to model agents who misinterpret information has been introduced in the literature by generalizing the notion of partition information[6]. Geanakoplos (1989) has applied this framework to a range of decision theoretic problems; critical to such applications are conditions showing when one possibility correspondence [PC] is ex ante more valuable than another, for all decision problems and all decision rules that are optimal given the agent's (possibly incomplete) interpretation of his information.

Chapter III thus provides a reinterpretion Geanakoplos' results in a framework which allows all kinds of conditional beliefs after observing information. When errors in information processing are only those generated in the PC framework, I show the theorem gives a somewhat strengthened version of Geanakoplos' results. This re-interpretation brings out the relation to standard decision theory, and allows us to solve problems with a dual approach. The multipliers have an interpretation as probabilities.

Brandenburger, Dekel and Geanakoplos [BDG] (1989) showed that for a given decision problem (and thus given distribution of information) facing a set of agents with heterogeneous priors who process information correctly, there is a decision theoretically equivalent decision problem where the agents have a common prior but make information processing errors; and vica-versa. The differences between the two cases are thus seen as one of interpretation. Geanakoplos (1989) used results about the value of information to prove "no speculation" results for possibility correspondences and notes that, in the light of the BDG result, these must

---

6. Bacharach (1985), Samet (1990), Shin (1987,1989), Geanakoplos (1989).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

also have an interpretation as results about heterogeneous prior beliefs. In chapter II, I show the relation between chapters II and III and the PC "no speculation results", as well as extending results of chapter II to agents who do not agree about which states are possible and to "agreeing to disagree" problems.

In the concluding chapter V, there is an extended discussion of the common prior assumption in economic theory, and applications that explicitly use heterogeneous prior beliefs; there is also a discussion of the relative merits of interpreting results in terms of heterogeneous prior beliefs and information processing errors. The relation of the dissertation to the literature is discussed within the chapters.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Chapter II

TRADING WITH HETEROGENEOUS PRIOR BELIEFS AND ASYMMETRIC INFORMATION

"It is frequently said of two bettors that one is a thief and the other is an imbecile; this is true in certain cases, where one of the two bettors is better informed than the other, and knows it; but it can also happen that two men in good faith in complex situations where they possess exactly the same elements of information will arrive at different conclusions on the probabilities of an event and that betting together, each one figures.. that he is the thief and the other the imbecile."

"If the adversary proposes a large bet, this tends to make people believe that he has a good hand.... the fact alone that the bet is proposed modifies the judgement which the bet is about." [Borel (1924)]

What is the origin of peoples' different opinions, or "posterior beliefs", about the world? We make a distinction in ordinary discourse between differences of opinion caused by different information, and differences of opinion arrived at on the basis of the same information. Decision theory makes the analogous distinction between asymmetric information and heterogeneous prior beliefs as explanations for different posterior beliefs[7]. In many contexts - theoretical and practical -it doesn't matter where the differences of opinion come from. For example, in a standard Walrasian equilibrium, where agents do not infer information from prices, only posterior beliefs matter.

The distinction becomes critical when agents are able to learn each others' information from each others' actions. Suppose you want to sell me I.B.M. shares at what seems to me to

---

7. If "rational" beliefs must be based on some prior information, there is a meaningless infinite regress (which the common prior assumption does not help to resolve). So how far back must we "unravel" beliefs to reach a truly "subjective" prior? An operational distinction for interactive decision theoretic purposes (and this paper) might go as follows: any opinion of yours that affects my beliefs has information content. We have unravelled beliefs far enough only when my "prior" beliefs are not changed by being told your "prior" beliefs. See discussion in Varian (1989) of the "credibility" of other agents' views.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

be a bargain price. I wonder whether you have reached a different assessment of I.B.M.'s prospects from the same information that I have (in which case, I should still consider buying), or whether you know something I don't (in which case, I should exercise extreme caution). This question arises whenever I trade in an uncertain environment where "objective" probabilities do not exist[8].

The particular case where all differences of opinion derive from differences in information (and agents have common prior beliefs) has been extensively studied in the economics literature[9]. "No trade" results have shown that, when agents' rationality is common knowledge, differences in information alone cannot explain trade. Agents' willingness to trade would reveal enough of their private information to each other to preclude trade.

But consider the other extreme, where all differences of opinion derive from heterogeneous prior beliefs (and there is no asymmetric information). We can certainly explain trade. But people should then be prepared to bet about every difference of opinion, and since differences of opinion seem to be more common than betting, this has been used as argument for supporting the common prior assumption as an empirical claim and a modelling tool[10].

In this chapter, I consider the intermediate case where differences of opinion are explained by both heterogeneous prior beliefs and asymmetric information. I can thus dispense

8. Of course, there are other reasons why valuations will differ in the absence of asymmetric information. You may have to sell your I.B.M. shares fast, because you have a liquidity crunch. In theoretical terms, different endowments and utility functions explain trade despite common priors. Or you may just be "irrational" (your actions are not coherently based on any set of beliefs and objectives). But, as long as "objective" probabilities do not exist, different subjective beliefs will be a possible explanation of your different valuation.

9. Milgrom and Stokey (1982) and Geanakoplos and Sebenius (1983) for the abstract trading environments considered here. Geanakoplos (1988) has shown how "no speculation" in competitive and rational expectations equilibrium results [e.g. Grossman (1981), Tirole (1982)] can be understood as special cases of such abstract problem. Other versions of the result include Rubinstein (1975) and Hakansson et al. (1982).

10. Aumann (1987), Bewley (1986).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

with both extreme results outlined above: i.e. I can explain how agents may sometimes rationally not trade with each other, despite differences in priors[11]; and I can explain how private information alone can lead to trade despite agents making rational inferences from other agents' willingness to trade. We must have both elements in order to make sense of my thought experiment about trading I.B.M. shares[12]. This interaction is the focus of this chapter[13].

But when will trade occur between asymmetrically informed agents? I consider three cases reflecting different possible contractual arrangements. An unrestricted trade allows complete contingent contracting. In particular, trades can depend on agents' private information. An incentive compatible trade cannot depend on private information directly, but can depend on some set of "verifiable" events and messages that agents are allowed to report; agents know how their payoffs depend on their messages, and choose messages optimally. A public trade cannot depend on agents' private information in any way, but only on the "verifiable" events and publicly known events[14]. Each agent is prepared to accept a trade only if his expected gain is non-negative, conditional on his information, and trade occurs only if all

11. Bewley (1986) noted that asymmetric information was a possible explanation of the absence of extensive betting among people with different prior beliefs (he develops another explanation by relaxing the completeness restriction on agents' preferences under uncertainty).

12. Without relaxing assumptions generating subjective expected utility maximization e.g. independence [Machina (1982)], completeness [Bewley (1986)], or correct interpretation of private information [Geanakoplos (1989)] and chpaters III and IV.

13. Varian (1989) looks at this interaction in a mean-variance model with a single risky asset. This is the only formal study of the interaction I know of. But asymmetric information has long been recognized as a problem in the elicitation of subjective probabilities; for example, it is explicitly recognized by Borel (1924) quoted above, Ramsey (1926), de Finetti (1937), Hartigan (1983). Since subjective probabilities are elicited by willingness to bet, the ideas in this paper directly address that problem. Where heterogeneous beliefs are assumed in the economics literature, they are typically open to either the interpretation that agents have observed different information, and fail to learn from other agents' actions, or that they have different priors [e.g. Lintner (1969), Harrison and Kreps (1978), Varian (1985)]. Heterogeneous beliefs in game theory are discussed in chapter V.

14. Incentive compatible and public trade are equivalent to equilibrium and common knowledge speculation, respectively, in Geanakoplos (1989). This relationship is discussed in chapter IV.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

agents accept, and at least one agent has strictly positive expected gain. For the moment I will assume, for each type of trade, that it must be common knowledge that the trade is accepted; each agent knows that other agents have accepted the trade, that other agents know that he has accepted the trade, and so on (this assumption will be weakened later). It will be shown that, under the common prior assumption, there cannot be any kind of acceptable trade, even with asymmetric information. On the other hand, if all agents are uninformed, the three definitions of trade are equivalent, and there will exist an acceptable trade whenever agents have different priors. But with both heterogeneous prior beliefs and asymmetric information, the different assumptions about contracting lead to different outcomes[15]. I extend (and weaken) Harsanyi's (1967) concept of consistency of beliefs to characterize the different necessary and sufficient conditions on priors to preclude trades of each type. Before giving an outline of the chapter, I discuss an example which introduces the extended concept of consistency and the differences between the three types of trade. In the formal treatment, I will define "betting" as the special case of "trading" when agents are risk neutral and there is a single good[16]. The example is a two person betting problem.

Suppose the Open Market Committee is to meet tomorrow to decide whether the discount rate will go up [U] or down [D]. You and I are having dinner and considering making a bet. But we both know that you had lunch with a member of the committee, and he made a prediction as to whether the meeting would decide to go up or down (we write "u" and "d" for his prediction; he didn't tell you anything else). Thus uncertainty can be represented by a four-state space, Ω = {uU, uD, dU, dD}, where you have observed a partition $\underline{P}$ = ({uU, uD},

15. As in Geanakoplos' (1989) results for Possibility Correspondences, discussed in chpater IV.

16. Geanakoplos and Sebenius (1983) proved a "no trade" theorem under the common prior assumption in this framework.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

{dU, dD}). Suppose your prior is π_y = (1/3, 1/6, 1/6, 1/3) and my prior is π_m = (2/9, 1/9, 2/9, 4/9). An unrestricted bet specifies how much you must pay me, or I must pay you, contingent on the actual realization of Ω. Thus the bet may be quite complicated (for example, you pay me 6 is uU occurs, 3 if dD occurs; I pay you 1 if uD occurs, 7 if dU occurs). Acceptance of the bet must be common knowledge. In this example, this means that you must be prepared to accept the bet whatever the prediction was (u or d).

It might appear easy to specify an unrestricted bet that both agents would accept. Your probabilities of observing U are 2/3 and 1/3, conditional on u and d respectively. Whereas I, who haven't observed u or d, have (unconditional) probability 4/9 of observing U. But no such bet exists[17], because agents have the same probabilities of observing U, conditional on the private information (I, also, have probabilities 2/3 and 1/3 of U, conditional on u and d respectively). More generally (with many agents, more than one of whom is informed), the necessary and sufficient condition to preclude trade is consistent beliefs. Following Harsanyi (1967), agents have consistent beliefs if their posterior beliefs, conditional on their own information, might have been derived from the same prior (even if, in fact, they weren't).

But now suppose trades cannot depend on agents' private information, because private information is not verifiable for contractual purposes. Thus it might not be possible for me to confirm with your lunch companion what prediction he made. You might then have an incentive to lie about which signal you had observed, even if you are required to commit at the time we make the bet - say, by writing down your prediction and putting it in a sealed envelope - before the actual discount rate decision is revealed.

17. Suppose there existed an unrestricted bet, with you receiving and I paying, (x_a, x_b, x_c, x_d) respectively in states (uU, uD, dU, dD). The x_i may be positive or negative. For you to accept, we must have $2x_a + x_b \geq 0$ and $x_c + 2x_d \geq 0$. For me to accept, we must have that $2x_a + x_b + 2x_c + 4x_d \leq 0$. But it is not possible to satisfy these inequalities with at least one inequality holding as a strict inequality.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

To see why this might happen, suppose in the discount rate example, my prior is replaced by $\pi_m' = \{3/8, 1/8, 1/8, 3/8\}$ and all else is unchanged. In this case, notice that we both believe that the unconditional probability of U is 1/2, and we both believe that the prediction is correlated with the eventual outcome. But I value the prediction more: I think there is a 3/4 chance that it is right, while you think there is only a 2/3 chance. There exists an unrestricted bet which we accept: suppose you pay me \$2 if the prediction is right and I pay you \$5 if the prediction is wrong. If you observed u, you are prepared to accept the bet (it has expected value \$1/3); if you observed d, you are also prepared to accept the bet (it again has expected value \$1/3). The bet has expected value \$ 1/4 for me. Suppose now that you cannot prove to me what the prediction was, but that you must write down your prediction (without telling me) at the time we make the bet (and before the actual discount rate change is announced). What will you write down? Suppose the prediction was actually u. If you write down u, your expected value under the bet is \$1/3. But if you write down d (i.e. you lie), you receive \$5 if the actual prediction (u) is right (i.e. uU occurs) and pay \$2 if the actual prediction is wrong (i.e. uD occurs). The expected value of the bet if you lie is thus \$ 8/3. But knowing this, I wouldn't believe you.

Now suppose agents can send messages, but (unlike in the previous example) the messages must be credible. Consider the natural mechanism design problem[18]: the bet specifies payoffs as a function of verifiable events (i.e. U or D) and some set of messages that agents can send. Agents decide, after observing their private information, both whether to accept the bet and what message to send. They do not observe other agents' messages, but acceptance of the bet is common knowledge. When U or D occurs, payment is made. By the

18. Hurwicz (1973).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

"revelation principle"[19], I can, without loss of generality, restrict attention to "truth-telling" mechanisms where agents' message spaces are the set of signals they observe, and they report their true signal in equilibrium. In the example, does there exist a bet, satisfying the conditions for unrestricted betting, which also satisfies the property that you don't have an incentive to lie? The answer is no[20]. Because I value your signal more than you do, any bet which satisfies the conditions for unrestricted betting will fail to satisfy incentive compatibility constraints. Only if we disagreed about which prediction (u or d) made U most likely, or if we agreed which made U most likely and you (the informed agent) valued the signal more highly than I (the uninformed agent), would incentive compatible betting occur. More generally, I will show that incentive compatible betting is precluded if and only if agents have reconcilable beliefs, where reconcilability is a weakening of consistency, allowing agents to have "averaged" beliefs about their own signals. It is not possible to exploit the differences in beliefs implied by the averaging because of the incentive compatibility problem.

Now consider one final version of the discount rate example. Suppose my prior is π_m'' = (3/10, 1/5, 1/5, 3/10), and again everything else is unchanged. Now you value your signal more than I do: you still believe there is a 2/3 chance that the prediction is right, I now believe there is only a 3/5 chance. So there exists an incentive compatible bet: you pay me $5 if the prediction is wrong (i.e. {uD, dU}) and I pay you $3 if the prediction is right (i.e. {uU, dD}). Now you have an incentive to report your true type.

19. Dasgupta, Hammond and Maskin (1979), Holmstrom and Myerson (1983).

20. Suppose I pay you (x_a, x_b, x_c, x_d) in states (uU, uD, dU, dD), respectively. For me to accept the bet, we must have $3x_a + x_b + x_c + 3x_d \leq 0$. For you to accept the bet on each of your information sets, we must have $2x_a + x_b \geq 0$ and $x_c + 2x_d \geq 0$. For you to truthfully reveal the prediction, we must have $2x_a + x_b \geq 2x_c + x_d$ and $x_c + 2x_d \geq x_a + 2x_b$. Multiplying your acceptance constraints by 4, and adding them and the two incentive compatibility constraints, we get $9x_a + 3x_b + 3x_c + 9x_d \geq 0$. But for me to accept the bet, we must have $9x_a + 3x_b + 3x_c + 9x_d \leq 0$. Therefore, it is impossible that either of us could have a strictly positive expectation of the bet on any information set.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

A public bet must depend only on verifiable events and publicly known events. In this example, there are no publicly known events and {uU, dU} and {uD, dD} are the only verifiable events. Common knowledge acceptance of the bet requires that the bet is accepted whatever you observed. But agents' unconditional probabilities of U are both 1/2, so there can't exist an acceptable common knowledge public bet. I will show that, in general, common knowledge public betting is precluded if and only if beliefs are publicly consistent. If each agent's prior could be replaced with some other prior with the same posterior beliefs conditional on his private information, such that the replacement priors agree on the probability of verifiable events and public events (events which we both know), then their beliefs are publicly consistent. In the example, beliefs would be publicly consistent only if my prior (unconditional) probability of U was either greater than or less than both your possible posterior probabilities of U.

The assumption that acceptance of bets must be common knowledge is actually stronger than required. Consider the "acceptance game" where each agent's strategy specifies when to accept and when to reject a bet. The bet is implemented only if all agents accept. The bet is "Nash accepted" if there is a Nash equilibrium of the acceptance game where some agent has strictly positive expected gain. This is a weaker requirement than common knowledge acceptance, but it will be shown that for both unrestricted and incentive compatible betting (but not public betting), conditions for the existence of an acceptable bet are the same under both acceptance rules.

In section 1, I introduce the uncertain environment which I will use in the chapter. A distinctive contribution of the chapter is that I work with both a "type space" representation and a "partition" representation of private information. Incentive compatibility is a difficult concept to motivate in "partition" notation, while common knowledge is an equally difficult concept to

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

motivate in "type space" notation[21]. By using both equivalent representations side by side, I am able to motivate my results better and relate apparently distinct approaches in the literature.

In section 2, I formally define the solution concepts for the betting case, discuss consistency, reconcilability and public consistency and prove they are the necessary and sufficient conditions to preclude unrestricted, incentive compatible and common knowledge public betting, respectively. Because I am not primarily concerned with the relation between solution concepts as such, the equivalence of different acceptance rules is proved in an appendix.

The advantage of the betting case is that there is no interaction with "insurance" motives for trade. But I wish to develop these ideas in a model more general than the simple betting case I use for motivation. If there are risk-averse agents trading many goods, the natural question to ask [as in Milgrom and Stokey (1982)] is whether agents with private information will trade with each other in situations where they would not be prepared to trade with each other in the absence of private information. Formally, I say that information leads to trade if agents with private information would trade when endowments are constrained efficient with respect to trades depending on verifiable events (but not depending on private information). If I assume in addition that the private information is not payoff-relevant (i.e. utility is independent of private information), then I am able to give conditions for trade as a function of agents' beliefs and information only (i.e. not their endowments and utility functions).

This argument proceeds in two stages: in section 3, I remain with the betting case, and ask when does information lead to betting? There is no uninformed betting if and only if agents have common beliefs about verifiable events. I show that under this assumption, conditions for no (informed) betting are independent of the common beliefs about verifiable events and depend

21. Holmstrom and Myerson (1983) give common knowledge results in type space notation.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

only on agents' beliefs about private information signals, conditional on verifiable events. Consistent interpretation, reconcilable interpretation, and publicly consistent interpretation are the necessary and sufficient conditions to preclude information leading to, respectively, unrestricted, incentive compatible and common knowledge public betting[22].

In section 4, I define the general trading problem with risk averse agents and many goods. I show that under the assumptions of "no uninformed trade" given above, conditions for trade in the general case reduce to those for the simple betting case. By allowing agents to trade to a constrained ex ante efficient outcome, and restricting signals to have no direct effect on utility, trade depends only, as in the betting case, on the interaction between beliefs and information.

The concluding chapter V considers applications.

Section 1: The Environment

I introduce notation for uncertainty which makes it easy to see the relation between a "partition" representation and a "type space" representation of private information; this is useful both in presenting the results in this chapter, and in explaining their relation to the literature.

There are H agents. I also write H for the set of agents. Each agent $h \in H$ observes a private signal $t_h \in T_h$[23]. The product of agents' signal spaces, $T = T_1 \times .. \times T_H$, thus incorporates all private information. I write $t = (t_1, .. , t_H)$ for a typical element of T. There

22. Milgrom and Stokey's (1982) "concordant" prior assumption similarly depends only on the "interpretation" of signals, but is stronger than all the conditions here.

23. An interpretation of the signal t_h is that it is agent h's "type" [see Harsanyi (1967)].

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

is also a set of "verifiable" events $\underline{Q}$, with typical element q. All relevant uncertainty is thus described by $\Omega = T \times \underline{Q}$. Each agent has a prior $\pi_h: \Omega \rightarrow \mathbb{R}_+$, $\Sigma_{\omega\in\Omega}\, \pi_h(\omega) = 1$.

I wish to study the case where agents' priors $\{\pi_h\}_{h\in H}$ differ. But I assume that agents at least agree about which combinations of private signals and verifiable events are possible[24]. Thus I assume that there exists $\Omega^* \subset \Omega$ such that $\{\omega\in\Omega \mid \pi_h(\omega) > 0\} = \Omega^*$, for all $h\in H$. This assumption enables me to give an alternative partition representation of uncertainty: define agents' private information partitions, $\{P_h\}_{h\in H}$, and verifiability partition, Q, on Ω^*, by:-

$$P_h(\{t,q\}) = \{\{t',q'\}\in\Omega^* | t_h' = t_h\} \quad \forall \{t,q\}\in\Omega^*$$

$$Q(\{t,q\}) = \{\{t',q'\}\in\Omega^* | q' = q\} \quad \forall \{t,q\}\in\Omega^*$$

Now the environment is completely described by uncertainty space Ω^*, private information partitions $\{P_h\}_{h\in H}$, verifiability partition Q and strictly positive priors on Ω^*, $\{\pi_h\}_{h\in H}$. Thus a betting environment $E = [\Omega^*, Q, \{\pi_h, P_h\}_{h\in H}]$.

What is the interpretation of the verifiability partition Q? If complete contingent claims markets do not exist, because certain events are not verifiable for contractual purposes, $Q(\omega)$ can be interpreted as the set of states which are not distinguishable from ω for contractual purposes; the range of Q, $\underline{Q} = \{C\in 2^{\Omega^*} \mid C=Q(\omega) \text{ for some } \omega\in\Omega^*\}$, is the set of verifiable events.

The primary motivation for verifiability restrictions is that agents' priors and partitions must be assumed to be common knowledge. Aumann has pointed out that this need not be a restriction as long as Ω^* is understood to be a complete description of all relevant uncertainty:

24. Holmstrom and Myerson (1983) make this common support assumption in a type space representation of uncertainty. This assumption is relaxed in chapter IV.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

"it is not an assumption, but a theorem, a tautology; it is implicit in the model itself"[25]. But then Ω^* must be understood as a description not only of some set of "natural" events, but also of agents' beliefs about other agents' priors and partitions, together with beliefs about beliefs, and so on in an infinite regress. While it may be possible to contract on whether it will rain tomorrow, it is presumably not possible to contract on what your prior beliefs about whether it will rain tomorrow are. Thus whenever common knowledge of the structure of the environment is assumed on the basis of Aumann's classic justification, verifiability restrictions must also be assumed.

The formal role of the partition Q in the theory is as a restriction on the type of bets that can be made; "verifiability" is one particular motivation for such restrictions that I will use through most of the chapter. But in section 3, I will give an example where events not in $\underline{Q}$ may be verifiable, but only bets contingent on events in $\underline{Q}$ are possible in the absence of private information, perhaps because of transaction costs. In section 2.3, I show that it would also be possible to incorporate a more general representation of incomplete asset markets (as in the General Equilibrium with incomplete markets literature).

What are the relative merits of the partition and type space representations of uncertainty? Many proofs (including some in this chapter) can be more easily stated in the type space notation. But the partition representation has a number of advantages for my purposes. Firstly, having defined uncertainty in terms of the product space, it is less easy to state comparative static results about changes in information structure, verifiability of contingent claims and, as I will show in section 4.2 below, the payoff-relevance of agents' signals. Secondly, the partition representation makes transparent the sometimes implicit assumptions about incomplete markets and payoff-relevance which drive results in the literature. Thirdly,

25. Aumann (1987) p9.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

the partition representation allows for easy comparison of different solution concepts. Fourthly, there are some concepts which are much easier to understand in the partition representation. For example, in the partition notation, we say that agent h knows event $E \subset \Omega^*$ at ω if $P_h(\omega) \subset E$. In the type space notation, h knows event $E \subset \Omega^*$ at $(t,q) \in E$ if $(\{t_{-h},t_h'\},q)$[26] $\in \Omega^*$ implies $(\{t_{-h},t_h'\},q) \in E$, for all $t_h' \in T_h$. In the partition notation, we say Q is finer than each P_h if $Q(\omega) \subset P_h(\omega)$ for all $h \in H$, $\omega \in \Omega^*$. The equivalent claim in the type space notation is that $(t,q) \in \Omega^*$ implies $(t',q) \notin \Omega^*$ for all $t' \neq t$. On the other hand, the idea of incentive compatibility is much more easily described in the type space notation, as I will show in section 2.2[27]. I will use both notations: in particular, proofs concerning incentive compatibility are given in type space notation and for common knowledge public trade in partition notation. Unrestricted trade proofs could easily be given in either.

I will need some additional notation for the interaction between agents' private information. Let P be the join (coarsest common refinement) of agents' private information partitions; let P_{-h} be the join of all but agent h's private information partitions. Notice that $P(\omega) \cap Q(\omega) = \{\omega\}$, for all $\omega \in \Omega^*$, by construction of Ω^*. Thus Ω^* can be thought of as a minimal description of uncertainty with no unnecessary states. I write M for the meet (finest common coarsening) of agents' private information; I will sometimes refer to M as the partition of public events.

---

26. Throughout the dissertation, the following notation is used:-

$$t = (t_1, \ldots, t_H) \in T_1 \times \ldots \times T_H = T$$
$$t_{-h} = (t_1, \ldots, t_{h-1}, t_{h+1}, \ldots, t_H) \in T_1 \times \ldots \times T_{h-1} \times T_{h+1} \times \ldots \times T_H = T_{-h}$$
$$\{t_{-h}, t_h'\} = (t_1, \ldots, t_{h-1}, t_h', t_{h+1}, \ldots, t_H) \in T.$$

27. In Blackwell's Theorem [Blackwell (1951), Cremer (1982)], uncertainty is usually represented by a cross-product of verifiable states (the states that the decision problem depends on) and private signals (from the "experiment"). In chapter III, I show how the equivalent partition representation version of Blackwell's Theorem shows when one information partition is always ex ante more valuable than another, for some restricted class of decision problems (thus the answer is not simply when one partition is finer than another). See also Green and Stokey (1978).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Section 2 : Betting

When I formally define the general trading problem in section 4, "betting" will be the special case of "trading" where there is a single good, and agents are risk neutral. Why treat betting first and separately? The betting case is the simplest way of abstracting from motivations for trade (e.g. insurance, state dependent utilities) other than the prior beliefs and information of agents. The results are particularly elegant in this simple case: all the theorems are straightforward applications of Farkas' lemma, as we ask if there exists a bet satisfying some set of linear inequalities. The linear algebraic conditions on beliefs then all have natural interpretations. Most of the intuition of the general trading problem becomes clear in the betting problem. In fact, the results for the general trading problem are proved by reduction to the betting case.

First, what is a bet? A bet is a set of random variables, one for each agent i.e. $\{x_h\}_{h\in H}$ with each x_h: $\Omega \to \mathbf{R}$. A bet is feasible if $\Sigma_{h\in H}\, x_h(\omega) \leq 0,\ \forall\ \omega\in\Omega$ [2.1]

It is useful to think of the $\{x_h\}_{h\in H}$ being defined on all of Ω (i.e. on all combinations of agents' signals and verifiable events), although for many purposes only their value on Ω^* (i.e. on those combinations which are believed to be possible) will be relevant.

To motivate the solution concepts here, think of there being an "acceptance game" for any given bet. I give a formal treatment of such an acceptance game in the appendix, here I am just seeking to motivate the formal description of the solution concepts, which do not have a game theoretic structure. Suppose an agent's strategy specifies on which information sets he will accept the bet and on which information sets he will reject the bet. "Nash acceptance" occurs if there is a Nash equilibrium of the acceptance game, with some agent strictly better off (there is also a trivial Nash equilibrium where all agents reject the bet, and thus no agent is

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

strictly better off)[28]. "Common Knowledge acceptance" occurs if there is Nash acceptance with common knowledge that each agent has accepted[29]. "Unconditional acceptance" occurs if there is Nash acceptance with each agent always accepting.

Nash and common knowledge acceptance rules both seem natural in different settings. If acceptance of the bet can somehow be signalled simultaneously, in such a way that agents cannot revise their acceptance strategies in the light of observing other agents' decisions, then the Nash rule makes sense. Under the Nash rule, agents know the strategies of other agents, but not their actions. Thus I may know that you accept if and only if event E occurs, without knowing if, in fact, E has occurred.

But if acceptance of the bet takes place publicly, without some special mechanism to ensure simultaneous, irrevocable, commitment to acceptance, then it is natural to require the stronger common knowledge acceptance. To invoke an image of Geanakoplos and Sebenius (1983), as the agents extend their hands to shake on the deal, each agent will learn that the other agent is accepting, and perhaps revise his strategy. It is important to note that in this chapter, agents' information is fixed (so that learning does not take place in the trading process); this assumption is discussed in section 2.4.

Since I have three acceptance rules (Nash, common knowledge and unconditional) and three contractual arrangements (unrestricted, incentive compatible and public), it might appear that I need nine solution concepts. But I will show that, when considering existence, all three acceptance rules (Nash, common knowledge and unconditional) are equivalent for unrestricted betting and incentive compatible betting. Common knowledge and unconditional acceptance (but not Nash acceptance) are equivalent for public betting.

---

28. This is closely related to "durability" in Holmstrom and Myerson (1983).

29. Geanakoplos (1988) discusses the common knowledge of strategies assumption in game theory.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

In this section, I give both partition (P) and type space (T) definitions. This redundant procedure is dropped in later sections, but should be extremely helpful for understanding the results.

2.1 Unrestricted Betting

In this section, I argue that if there is Nash acceptance of some unrestricted bet, then there is another unrestricted bet with unconditional (and thus common knowledge) acceptance. Then I define "interim rationality" which is the condition for unconditional acceptance. Next, I define and discuss, with examples, the concept of "consistency" of beliefs. Finally, I state and prove the theorem.

Suppose there is Nash acceptance of an unrestricted bet. Consider a new unrestricted bet which is the same as the old unrestricted bet at all states of the world where the old bet was accepted (by everybody), but equal to zero where the old bet was rejected (by anybody). Now there is not only "Nash acceptance" of the new bet, but also unconditional acceptance, since each agent is indifferent between accepting or rejecting if his expected payoff is zero. Thus it is not restrictive to require unconditional acceptance of an unrestricted bet[30].

P - Definition $\{x_h\}_{h \in H}$ is interim rational if

$$\sum_{\omega \in P_h(\omega')} \pi_h(\omega)\, x_h(\omega) \geq 0 \qquad \forall\, h \in H,\ \omega' \in \Omega^*$$

with strict inequality for at least one h, ω′

30. This is proved in the appendix.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

<u>T - Definition</u> $\{x_h\}_{h\in H}$ is <u>interim rational</u> if

$$\sum_{t_{-h}\in T_{-h}} \sum_{q\in Q} \pi_h(\{t,q\})\, x_h(\{t,q\}) \geq 0 \qquad \forall\, h\in H, t_h'\in T_h \qquad [2.2]$$

with strict inequality for at least one h, t_h

Interim rationality captures the idea of unconditional acceptance. Note that if strict inequality was not required somewhere, then a bet equal to zero everywhere would be classified as interim rational.

<u>Definition</u> $\{x_h\}_{h\in H}$ is an <u>unrestricted bet</u> if it is feasible and interim rational.

Thus we assume there is complete contingent contracting; in particular, the bet may depend on agents' private information. Thus either agents are compulsive truth-tellers, or agents' information must also be verifiable (as well as events in <u>Q</u>).

Harsanyi (1967) introduced the idea of "consistent" priors. The key observation is that once agents have observed their signals, all that matters is their posterior beliefs. If their posterior beliefs could have been derived from the same prior, then in any problem where only posterior beliefs matter, results which are true under the common prior assumption must also be true with "consistent" priors.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

P - Definition Beliefs $\{\pi_h\}_{h\in H}$ are consistent if there exists some "common prior" π on Ω^* such that $\pi_h[\,\omega \mid P_h(\omega)] = \pi[\,\omega \mid P_h(\omega)]$[31] $\forall\, h\in H,\ \omega\in\Omega^*$

T - Definition Beliefs $\{\pi_h\}_{h\in H}$ are consistent if there exists $\{\lambda_h\}_{h\in H}$, λ_h: $T_h \rightarrow \mathbb{R}_{++}$, and some "common prior" π on Ω, such that $\lambda_h(t_h)\ \pi_h(\{t,q\}) = \pi(\{t,q\})$, $\forall\, h\in H$, $\{t,q\}\in\Omega$

Proof of equivalence Substitute the following expression for $\lambda_h(t_h)$ in the T - definition

$$\lambda_h(t_h) = \frac{\sum_{t_{-h}\in T_{-h}} \sum_{q\in Q} \pi(\{t,q\})}{\sum_{t_{-h}\in T_{-h}} \sum_{q\in Q} \pi_h(\{t,q\})}$$

The T - definition shows how consistency arises as a critical condition from a set of linear inequalities.

Consider the example, discussed in the introduction, where there are two verifiable events (U and D) and two signals (u and d), perhaps correlated with U and D, which agent 1 has observed, but agent 2 has not. Suppose the agents' unconditional probabilities of agent 1 observing u are q and r, respectively; and the agents' probabilities of observing U and D are α_u and β_u, respectively, conditional on 1 observing u, and α_d and β_d respectively, conditional on 1 observing d. This is summarized below:-

31. For any probability distribution, ϕ,

$$\phi[\omega \mid P_h(\omega)] = \frac{\phi(\omega)}{\sum_{\omega'\in P_h(\omega)} \phi(\omega')}$$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Discount Rate Example

$\Omega^* = \{uU, uD, dU, dD\}$
$\pi_1 = \{q\alpha_u, q(1-\alpha_u), (1-q)\alpha_d, (1-q)(1-\alpha_d)\}$ $\quad$ $\pi_2 = \{r\beta_u, r(1-\beta_u), (1-r)\beta_d, (1-r)(1-\beta_d)\}$
$\underline{P}_1 = (\{uU,uD\}, \{dU,dD\})$ $\quad$ $\underline{P}_2 = (\Omega^*)$

Beliefs are consistent if and only if $\alpha_u = \beta_u$ and $\alpha_d = \beta_d$. This illustrates one general implication of consistency, that agents must agree about the probabilities of all events conditional on the join of private information i.e. $\pi_h[\ \omega \mid P(\omega)] = \pi[\ \omega \mid P(\omega)\ ]$, for some π, for all $h \in H$, $\omega \in \Omega^*$. But this is not always (as in this example) a sufficient condition: consistency also imposes restrictions on beliefs about other agents' signals. Consider another example:-

$\Omega^* = \{a, b, c, d\}$
$\pi_1 = (1/4, 1/4, 1/4, 1/4)$ $\quad$ $\pi_2 = (8/25, 12/25, 2/25, 3/25)$
$\underline{P}_1 = (\{a,b\}, \{c,d\})$ $\quad$ $\underline{P}_2 = (\{a,c\}, \{b,d\})$

Here the join of agents' private information is fully informative [$P(\omega) = \{\omega\}$, for all $\omega \in \Omega^*$], so it is trivially the case that agents agree about the probability of every event, conditional on the join of their beliefs (either 0 or 1). But agents disagree about the unconditional probability of every event in Ω^* that it is possible to disagree about[32]. To see why agents' beliefs are nonetheless consistent, consider prior $\pi = (2/5, 2/5, 1/10, 1/10)$. Each agent h's beliefs, conditional on his partition P_h, would be the same if his prior was replaced by π.

Theorem 1a There is no unrestricted betting if and only if agents' beliefs are consistent.

32. They agree about the probability of the empty set and universal set.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Proof By Farkas' Lemma, there do not exist $\{x_h\}_{h\in H}$ satisfying [2.1] and [2.2] if and only if there exist $\{\lambda_h\}_{h\in H}$, each $\lambda_h:T_h \to \mathbb{R}_{++}$, and $\phi: \Omega \to \mathbb{R}_+$, such that:-

$$\lambda_h(t_h)\pi_h(\{t,q\}) = \phi(\{t,q\}) \quad \forall\ h\in H,\ \{t,q\}\in\Omega$$

But now if we let $\pi(\omega) = \phi(\omega) / \Sigma_{\omega'\in\Omega}\, \phi(\omega')$, and $\lambda_h'(t_h) = \lambda_h(t_h) / \Sigma_{\omega'\in\Omega}\, \phi(\omega')$, the $\{\lambda_h'\}_{h\in H}$ and prior π satisfy the type space definition of consistency[33].

2.2 Incentive Compatible Betting

This section is structured as the previous section: first, I show why it is not restrictive to require unconditional acceptance of an incentive compatible bet. The revelation principle then implies that it is sufficient to consider mechanisms where agents truthfully reveal their information. I define the formal incentive compatibility restriction on bets: this is most easily stated in type space notation, but I also give a partition statement. Then I define and discuss reconcilability, the weakening of consistency which is the necessary and sufficient condition to preclude incentive compatible betting. Finally, I state and prove the theorem.

A mechanism specifies a set of possible messages for each agent, and a payoff to each agent as a function of verifiable events and all agents' messages. For a given verifiable event, and a given set of messages, the sum of agents' payoffs is less than or equal to zero. There is Nash acceptance of this mechanism bet if all agents accept the bet at some information set, with each agent choosing his message optimally as a function of his private information. But another message, "reject", could be added to each agent's set of possible messages, with all agents

33. In the rest of this paper, I will use this normalization of the multipliers from Farkas' Lemma to give a probability distribution without further comment.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

receiving zero if any agent sends the message "reject". Thus there is no loss of generality in assuming unconditional acceptance. But now by the revelation principle, there is also no loss of generality in restricting attention to "truth-telling" mechanisms.

In the appendix, I prove this "revelation principle with a twist". In what follows, I define incentive compatible trade in terms of the existence of a truth-telling mechanism bet with unconditional acceptance.

Incentive compatibility is much easier to define in type space notation:-

T - Definition Bet $\{x_h\}_{h \in H}$ is incentive compatible if

$$\sum_{t_{-h} \in T_{-h}} \sum_{q \in Q} \pi_h(\{t,q\})\, x_h(\{t,q\}) \geq \sum_{t_{-h} \in T_{-h}} \sum_{q \in Q} \pi_h(\{t,q\})\, x_h(\{(t_{-h},t_h'),q\}) \qquad \forall\, h \in H,\ \{t_h,t_h'\} \subseteq T_h$$

Thus if agent h's signal is t_h, then he is better off revealing his type truthfully than lying. Notice that I use the fact that the $\{x_h\}_{h \in H}$ are defined on all of $\Omega = T \times Q$, not just on Ω^*. What would be the equivalent definition in partition notation using only values of the $\{x_h\}_{h \in H}$ defined on Ω^*?

P - Definition A bet $\{x_h\}_{h \in H}$ is support incentive compatible if

$$\sum_{\omega \in P_h(\omega')} \pi_h(\omega)\, x_h(\omega) \geq \sum_{\omega \in P_h(\omega')} \pi_h(\omega)\, z_h(\omega,\omega^*) \qquad \forall\, h \in H,\ \omega' \in \Omega^*$$

$$\forall\, \omega^* \in \Omega^* \text{ such that } P_h(\omega^*) \cap P_{-h}(\omega'') \cap Q(\omega'') \neq \emptyset,\ \forall\, \omega'' \in P_h(\omega'),$$

$$\text{and } z_h(\omega,\omega^*) = x_h(\omega'') \text{ where } \omega'' \in P_h(\omega^*) \cap P_{-h}(\omega) \cap Q(\omega)$$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

"Support incentive compatibility" depends only on the values the $\{x_h\}_{h\in H}$ take on Ω^*, the support of the $\{\pi_h\}_{h\in H}$. It says that agents can only consider lying if there is no possibility of them being found out to be lying. The definitions are almost equivalent in the following sense: if there exists $\{x_h\}_{h\in H}$ incentive compatible, then $\{x_h\}_{h\in H}$ is support incentive compatible; and if $\{x_h\}_{h\in H}$ is support incentive compatible, then there exists a real number k such that for all $\{y_h\}_{h\in H}$ such that $y_h(\omega) = x_h(\omega)$, for all $\omega\in\Omega^*$, and $y_h(\omega) \leq k$, for all $\omega\notin\Omega^*$, $\{y_h\}_{h\in H}$ is incentive compatible. This means that if $\{x_h\}_{h\in H}$ is support incentive compatible, then for sufficiently unpleasant payoffs ($<k$) on impossible combinations of signals and verifiable events, $\{x_h\}_{h\in H}$ will be incentive compatible.

Definition $\{x_h\}_{h\in H}$ is an incentive compatible bet if $\{x_h\}_{h\in H}$ is feasible, interim rational and incentive compatible.

Reconcilability is a complicated property to state, but it has a natural interpretation.

P - Definition Beliefs $\{\pi_h\}_{h\in H}$ are reconcilable if there exist strictly positive priors $\{\zeta_h\}_{h\in H}$, priors $\{\xi_h\}_{h\in H}$, Markov matrices $\{\Theta_h\}_{h\in H}$ (with the dimension of each Θ_h equal to the number of elements of T_h), non-negative constants $\{k_h\}_{h\in H}$ and "common prior" π, such that:-

$$[1]\ \zeta_h[\omega|P_h(\omega)] = \pi_h[\omega|P_h(\omega)] \quad \forall\ h\in H,\ \omega\in\Omega^*$$

$$[2]\ \xi_h[\omega|P_h(\omega)] = \pi_h[\omega|P_h(\omega)] \quad \forall\ \omega\in\Omega^*,\ \textit{such that}\ \xi_h[P_h(\omega)] > 0,\ h\in H.$$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

$$[3]\quad \zeta_h(\omega) = \pi(\omega) - k_h[\xi_h(\omega) - \sum_{\omega' \in P_{-h}(\omega)\cap Q(\omega)} \theta_h(P_h(\omega'), P_h(\omega))\xi_h(\omega')\,] \quad \forall\, h \in H,\ \omega \in \Omega^*$$

where the $\theta_h(.,.)$ are elements of the matrix Θ_h.

<u>T - Definition</u> Beliefs $\{\pi_h\}_{h\in H}$ are <u>reconcilable</u> if there exist $\{\lambda_h\}_{h\in H}$, with $\lambda_h{:}T_h \rightarrow \mathbf{R}_{++}$, and $\{\mu_h\}_{h\in H}$, with μ_h: $T_h^2 \rightarrow \mathbf{R}_+$, and "common prior" π, such that[34]:-

$$\lambda_h(t_h)\pi_h(\{t,q\}) + \sum_{t'_h \in T_h} [\mu_h(t_h,t'_h)\pi_h(\{t,q\}) - \mu_h(t'_h,t_h)\pi_h(\{(t_{-h},t'_h),q\})] = \pi(\{t,q\}) \quad \forall\, h \in H,\ t \in$$

<u>Proof of equivalence</u> Note that summing over $t_h \in T_h$ in the T - definition gives:-

$$\sum_{t_h \in T_h} \lambda_h(t_h)\pi_h(\{t,q\}) = \sum_{t_h \in T_h} \pi(\{t,q\}) \quad \forall\, h \in H,\ t_{-h} \in T_{-h},\ q \in Q.$$

and that summing over the $\omega \in P_{-h}(\omega^*) \cap Q(\omega^*)$ in the P - definition gives:-

$$\zeta_h[P_{-h}(\omega^*)\cap Q(\omega^*)] = \pi[P_{-h}(\omega^*)\cap Q(\omega^*)] \quad \forall\, h \in H,\ \omega^* \in \Omega^*.$$

Now the following substitutions can be used to prove the equivalence (i.e. in both directions):

$$\alpha_h(t_h) = \sum_{t'_h \in T_h} \mu_h(t_h,t'_h) \qquad\qquad k_h = \sum_{t\in T}\sum_{q \in Q} \alpha_h(t_h)\pi_h(\{t,q\})$$

$$\zeta_h(\{t,q\}) = \lambda_h(t_h)\pi_h(\{t,q\})$$

$$\xi_h(\{t,q\}) = \frac{\alpha_h(t_h)\pi_h(\{t,q\})}{k_h} \qquad\qquad \theta_h(t_h,t'_h) = \frac{\mu_h(t_h,t'_h)}{\alpha_h(t_h)}$$

---

34. The structure of the linear algebraic restrictions is of the same form as those derived in a more general class of mechanism design problems, with a disagreement outcome, considered by Myerson (1989). He introduces the terminology that $\mu_h(t_h,t_h') > 0$ implies type t_h "jeopardizes" type t_h'. Thus type t_h can jeopardize type t_h' only if the constraint that type t_h should not be tempted to pretend to be type t_h' is binding. "Reconcilability" is also closely related to Shin's (1989) "dynamic representation" of beliefs; this is discussed in chapter III.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Note that if all the k_h are equal to zero, then reconcilability reduces to consistency. In general, reconcilability continues to require that agents' posterior-equivalent priors $\{\zeta_h\}_{h \in H}$ agree with the common prior π about other agents' signals and verifiable events. But agents' beliefs about their own signals, conditional on others' signals and verifiable events, under the $\{\zeta_h\}_{h \in H}$, may differ from the common prior π. With k_h non-zero, it is as if each agent believes he has observed a noisy version of his own signal; he puts less weight on his signal than other agents do.

Consider again the discount rate example. The "verifiability" partition is $\underline{Q}$ = ({uU, dU}, {uD, dD}). Agents' beliefs are reconcilable if and only if $\beta_u \geq \alpha_u > b > \alpha_d \geq \beta_d$, $\beta_u = \alpha_u = b = \alpha_d = \beta_d$, or $\beta_u \leq \alpha_u < b < \alpha_d \leq \beta_d$, where $b = r\beta_u + (1-r)\beta_d$ is agent 2's unconditional probability of U. To see why, note that according to the type space definition, beliefs are reconcilable if there exist λ_{11}, λ_{12} and λ_2 all strictly positive, and μ_1 and μ_2 non-negative such that:-

$$(\lambda_{11}+\mu_1)q\alpha_u-\mu_2(1-q)\alpha_d = \lambda_2 r\beta_u \qquad (\lambda_{11}+\mu_1)q(1-\alpha_u)-\mu_2(1-q)(1-\alpha_d) = \lambda_2 r(1-\beta_u$$

$$(\lambda_{12}+\mu_2)(1-q)\alpha_d-\mu_1 q\alpha_u = \lambda_2(1-r)\beta_d \qquad (\lambda_{12}+\mu_2)(1-q)(1-\alpha_d)-\mu_1 q(1-\alpha_u) = \lambda_2(1-$$

A straightforward manipulation shows that this is true if and only if $\lambda_2 r[\alpha_u - \beta_u] + \mu_2(1-q)[\alpha_u - \alpha_d] = 0$ and $\lambda_2(1-r)[\alpha_d - \beta_d] + \mu_1 q[\alpha_d - \alpha_u] = 0$, and thus $\beta_u \geq \alpha_u > \alpha_d \geq \beta_d$, $\beta_u = \alpha_u = \alpha_d = \beta_d$ or $\beta_u \leq \alpha_u < \alpha_d \leq \beta_d$. The requirement that λ_{11} and λ_{12} are strictly more than 0 give the extra conditions on b.

Theorem 1b There is no incentive compatible betting if and only if beliefs are reconcilable.

Proof When do there not exist $\{x_h\}_{h \in H}$, each $x_h: \Omega \rightarrow \mathbf{R}$, such that:-

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

$$\sum_{t_{-h}\in T_{-h}}\sum_{q\in Q}\pi_h(\{t,q\})x_h(\{t,q\}) \geq 0 \quad \forall\, h\in H,\ t_h\in T_h$$
$$\textit{with strict inequality for at least one } h,t_h$$

$$\sum_{t_{-h}\in T_{-h}}\sum_{q\in Q}\pi_h(\{t,q\})x_h(\{t,q\}) \geq \sum_{t_{-h}\in T_{-h}}\sum_{q\in Q}\pi_h(\{t,q\})x_h(\{(t_{-h},t_h'),q\}) \quad \forall\, h\in H,\ \{t_h,t_h'\}\subset T_h.$$

$$\sum_{h\in H}x_h(\{t,q\}) \leq 0 \quad \forall\, t\in T,\ q\in Q$$

By Farkas' Lemma, there do not exist such $\{x_h\}_{h\in H}$ if and only if there exist $\{\lambda_h\}_{h\in H}$, with λ_h:T_h $\rightarrow \mathbf{R}_{++}$, and $\{\mu_h\}_{h\in H}$, with μ_h: $T_h^2 \rightarrow \mathbf{R}_+$, and prior π, such that:-:-

$$\lambda_h(t_h)\pi_h(\{t,q\})+\sum_{t_h'\in T_h}[\mu_h(t_h,t_h')\pi_h(\{t,q\})-\mu_h(t_h',t_h)\pi_h(\{(t_{-h},t_h'),q\})] = \pi(\{t,q\}) \quad \forall\, h\in H,\ t\in T$$

2.3 Common Knowledge Public Betting

In this section, I first discuss the idea of public events in the partition framework. Public bets can depend on verifiable events and public events but not on events which are known to some agents and not to others. For public betting (unlike for unrestricted betting and incentive compatible betting), it matters whether we consider Nash acceptance or common knowledge acceptance. In this chapter, I'm concerned with common knowledge acceptance of public bets, but if there is common knowledge acceptance of some public bet, then there must also be another public bet which is unconditionally accepted. I use this in my formal definition of common knowledge public betting. I define and discuss the condition of publicly consistent beliefs, and then state and prove the theorem about common knowledge public betting.

Aumann (1976) introduced the notion of common knowledge into economics. An event E is said to be common knowledge at ω if $M(\omega) \subset E$, where M is the meet (finest common

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

coarsening) of agents' private information partitions. Aumann showed how this definition corresponds to the intuitive notion that an event is common knowledge if agents know that it's true, they know that other agents know that it is true, and so on[35]. An event E is a public event if E is common knowledge at ω, for all $\omega \in E$. By the definition of common knowledge, an event is a public event if and only if it is measurable with respect to M. For this reason, M will also be called the partition of public events.

It is natural to assume that even if bets cannot depend on private information, they should be able to depend on public events - events that all agents know and all agents know that all other agents know and so on. Public bets can depend on both public events and verifiable events. It may seem natural to assume that public events are verifiable [i.e. partition Q is as fine as the partition M]. But economists have considered situations where events publicly known to both parties to a contract are not verifiable[36]. Also, recall that the restrictions on trade implied by Q can be motivated by reasons other than verifiability (e.g. transaction costs): an example of this is given in section 3.

Definition Bet $\{x_h\}_{h\in H}$ is a public bet if each x_h is measurable with respect to the join of Q, the partition of verifiable events, and M, the partition of public events.

For public bets, there is a difference between Nash acceptance and common knowledge acceptance. Thus in the discount rate example, there exists Nash acceptance of a public bet if and only if there exists incentive compatible betting, while common knowledge acceptance of

35. Milgrom (1981) gives a formal proof of the equivalence.

36. For example, an employee's effort may be common knowledge among employer and employee, but it may not be possible to contract on effort, because a court of law cannot verify effort [e.g. Shapiro and Stiglitz (1984)].

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

a public bet is much more restrictive, as I will show. But suppose there is common knowledge acceptance of some public bet. Then if that bet is replaced by another equal to zero on those events where it was not accepted before, then the replacement bet will be unconditionally accepted. The appendix contains more detailed discussion of these issues.

Definition $\{x_h\}_{h\in H}$ is a common knowledge public bet if $\{x_h\}_{h\in H}$ is feasible, interim rational and public.

Lemma If there exists a common knowledge public bet, $\{x_h\}_{h\in H}$, then there exists an incentive compatible bet, $\{x_h^*\}_{h\in H}$.

Proof Let $x_h^*(\omega) = x_h(\omega)$, for all $\omega \in \Omega^*$, and $x_h^*(\omega) = k$, for all $\omega \notin \Omega^*$. For k sufficiently negative, all incentive compatibility constraints will hold. But feasibility and interim rationality must hold because $\{x_h\}_{h\in H}$ is a common knowledge public bet.

The lemma shows that conditions for common knowledge public betting are stronger than the conditions for incentive compatible betting. So the restriction on beliefs required to preclude common knowledge public betting, public consistency, is weaker than reconcilability.

P - Definition Beliefs $\{\pi_h\}_{h\in H}$ are publicly consistent if there exist positive priors $\{\zeta_h\}_{h\in H}$, and "common prior" π, such that:-

$$[1]\ \zeta_h[\omega | P_h(\omega)] = \pi_h[\omega | P_h(\omega)] \quad \forall\ h \in H,\ \omega \in \Omega^*$$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

$$[2]\ \zeta_h[Q(\omega)\cap M(\omega)] = \pi[Q(\omega)\cap M(\omega)] \quad \forall\, h\in H,\ \omega\in\Omega^*$$

T - Definition Beliefs $\{\pi_h\}_{h\in H}$ are publicly consistent if there exist $\{\lambda_h\}_{h\in H}$, each $\lambda_h: T_h \rightarrow \mathbf{R}_{++}$, and "common prior" π such that:-

$$\sum_{t\in E}\lambda_h(t_h)\pi_h(\{t,q\}) = \sum_{t\in E}\pi(\{t,q\}), \quad \forall\, h\in H, \textit{ for each public event } E\subset T^*.$$

where $T^* = \{t\in T \mid \{t,q\}\in\Omega^*$, for some $q\in \underline{Q}\}$ and $E \subset T^*$ is a public event if and only if $t\in E$ and $\{t_{-h},t_h'\}\in T^*$ implies $\{t_{-h},t_h'\}\in E$.

In words, beliefs are publicly consistent if each agent h's prior (π_h) can be replaced with another prior with the same posterior beliefs (ζ_h), such that the replacement priors ($\{\zeta_h\}_{h\in H}$) agree on the probability of public and verifiable events (the join of M and Q).

In the discount rate example, beliefs are publicly consistent if and only if $\alpha_u > b > \alpha_d$, $\alpha_u = b = \alpha_d$, or $\alpha_u < b < \alpha_d$ [recall that b is the uninformed agent's (unconditional) probability of U]. Thus common knowledge public betting occurs only if min $(\alpha_u,\alpha_d) \geq b$ and max $(\alpha_u,\alpha_d) > b$ or if max $(\alpha_u,\alpha_d) \leq b$ and min $(\alpha_u,\alpha_d) < b$. The uninformed agent's unconditional probability of U must be either greater than or less than both α_u and α_d.

In the example, public consistency is equivalent to the requirement that there exists π such that $\pi_h [\, Q(\omega) \mid P_h(\omega)\,] = \pi [\, Q(\omega) \mid P_h(\omega)\,] \quad \forall\, h\in H,\ \omega\in\Omega^*$.

This seems a simpler definition of public consistency, but in fact it turns out, in general, to be too strong a requirement, as the following example illustrates:-

$\Omega^* = \{a,b,c,d,e\}$
$\underline{Q} = (\{a,d\}, \{b,c,e\})$
$\pi_1 = (1/7, 1/7, 2/7, 2/7, 1/7)$ $\qquad \pi_2 = (1/6, 1/3, 1/6, 1/6, 1/6)$
$\underline{P}_1 = (\{a,b\}, \{c,d,e\})$ $\qquad \underline{P}_2 = (\{a,c\}, \{b,d\}, \{e\})$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Notice that Ω^* is the only public event [$M(\omega) = \Omega^*, \forall\omega\in\Omega^*$]. Beliefs are publicly consistent, since each agent's prior can be replaced with priors with the same posteriors: $\zeta_1 = \pi_1$ and $\zeta_2 = (5/14, 1/7, 5/14, 1/14, 1/14)$, respectively[37], such that the unconditional probabilities of event {a,d} in Q are 3/7 under both ζ_1 (i.e. π_1) and ζ_2. However, there does not exist a prior π such that agents' conditional probabilities of events in Q agree with π[38].

Theorem 1c There is no common knowledge public betting if and only if beliefs are publicly consistent.

Proof This is easiest to show in partition notation. Do there exist $\{x_h\}_{h\in H}$, each $x_h: \Omega^* \rightarrow \mathbf{R}$, measurable with respect to the join of Q and M, such that:-

$$\sum_{\omega\in P_h(\omega^*)} \pi_h(\omega)x_h(\omega) \geq 0 \quad \forall\, h\in H,\ \omega^*\in\Omega^*, \textit{ with strict inequality for at least one } h,\omega^*.$$

$$\sum_{h\in H} x_h(\omega) \leq 0 \quad \forall\, \omega\in\Omega^*.$$

By Farkas' Lemma, this is true if and only if there do not exist $\{\lambda_h\}_{h\in H}$, each $\lambda_h: \Omega^* \rightarrow \mathbf{R}_{++}$, each λ_h measurable with respect to P_h, and a common prior π such that:-

37. Agent 2's posterior beliefs, given partition P_2, are (1/2, 2/3, 1/2, 1/3, 1) under both π_2 and ζ_2.

38. Suppose there exists such a $\pi = (\pi_a, \pi_b, \pi_c, \pi_d, \pi_e)$. To agree with 1's conditional probabilities, we require $\pi_a/\pi_b = 1$ and $\pi_d/(\pi_c+\pi_e) = 2/3$; to agree with 2's conditional probabilities, we require $\pi_a/\pi_c = 1$ and $\pi_d/\pi_b = 1/2$; but now if $\pi_a = 2k$, we get $\pi_b = 2k$, $\pi_c = 2k$, $\pi_d = k$ and thus the contradiction $\pi_e = (3/2)\pi_d - \pi_c = -(1/2)k$.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

$$\sum_{\omega \in M(\omega^*) \cap Q(\omega^*)} \lambda_h(\omega)\pi_h(\omega) = \sum_{\omega \in M(\omega^*) \cap Q(\omega^*)} \pi(\omega) \quad \forall\, h \in H,\ \omega^* \in \Omega^*.$$

Now set $\zeta_h(\omega) = \lambda_h(\omega)\, \pi_h(\omega) \quad \forall\, h \in H,\ \omega \in \Omega^*$.

2.4 Discussion

The assumption that restrictions on trade take the form of a partition and derive from restrictions on verifiability is not necessary. I could have assumed, as in the general equilibrium with incomplete assets literature[39], that there was some arbitrary set of basic "assets" (random variables on Ω) and that all trades had to be linear combinations of those assets. Instead, I assumed, in effect, that the set of basic assets could be written in a particularly simple form: i.e. that for each $q \in \underline{Q}$, there exists an asset paying 1 if $\omega \in q$, 0 otherwise. In the more general case, instead of deriving restrictions on agents' beliefs, conditions about agents' valuations of different assets would be derived. The qualitative structure of the results is unchanged[40,41].

A more subtle problem is the assumption that agents' information $\{P_h\}_{h \in H}$ is fixed. But trading processes do reveal information. I have asked, for given information, whether bets satisfying certain properties exist. How could the idea that learning is part of the trading

39. Geanakoplos (1990).

40. But we do use the special form of incomplete asset markets in section 4 to argue that competitive equilibria (with incomplete markets) will be constrained efficient, in the general trading setting.

41. An interesting special case is when there is one good and two assets, a constant asset (paying, say, 1 in every state) and a variable asset. Would agents trade? The "private values" model in the auctioning and bargaining literature can be interpreted as an example of such an environment when agents have heterogeneous priors.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

process be incorporated in a general way[42]? Suppose initially agents have private information $\{P_h\}_{h\in H}$, but during the trading process, these become refined to $\{P_h^+\}_{h\in H}$. For example, if the trading process revealed all agents' private information[43], we would have $P_h^+ = P$, for all $h\in H$[44]. Say that $\{P_h^+\}_{h\in H}$ is a feasible refinement from $\{P_h\}_{h\in H}$ if each P_h^+ is as fine as P_h (h does not forget what he knew before) and as course as P (h cannot know something which no one knew before). Then, given $\{P_h\}_{h\in H}$, does there exist a feasible refinement $\{P_h^+\}_{h\in H}$ such that common knowledge public betting can occur. The conditions for no such "post revelation common knowledge betting" (depending on initial information $\{P_h\}_{h\in H}$) do not have a simple representation, but they are strictly weaker than consistency, strictly stronger than public consistency and cannot in general be ranked with respect to reconcilability.

Even when allowing for information revelation, necessary and sufficient conditions for betting are given as a function of beliefs and the initial private information. But for a given uncertainty space and verifiability partition, what condition on priors is required to preclude betting for all possible private information partitions? The question is non-trivial: for example, if all agents are uninformed, agents need agree only on the probability of verifiable events; if all agents are fully informed, they will not bet at all, whatever their priors. But as long as some events are distinguishable for contractual purposes [i.e. $Q(\omega) \neq \Omega^*$ for some (and thus all) $\omega\in\Omega^*$], there is no unrestricted / incentive compatible / common knowledge betting for all private information $\{P_h\}_{h\in H}$ if and only if agents have common prior beliefs about Ω^*[45].

42. Dubey, Geanakoplos and Shubik (1977) give an explicit model of information revelation in a trading process.

43. As in Radner (1979).

44. "Full Revelation Common Knowledge Betting" occurs if and only agents have common beliefs conditional on the join of private information i.e. if $\pi_h[\ \omega\ |\ P(\omega)\ ] = \pi[\ \omega\ |\ P(\omega)\ ]$ for all $\omega\in\Omega^*$, for some "common prior" π. This condition is strictly weaker than consistency, and but cannot in general be ranked with respect to reconcilability and public consistency.

45. These results about information revelation processes were proved in an earlier draft.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Section 3: When does information lead to betting?

When would agents bet in the absence of any private information? If and only if they have different beliefs about verifiable events. So information leads to betting, in the sense that the arrival of new information enables betting that would not otherwise have occurred, if betting occurs despite agents having common beliefs about verifiable events. In this section, we give necessary and sufficient conditions to preclude betting when agents have common beliefs about verifiable events; these conditions - consistent interpretation, reconcilable interpretation and publicly consistent interpretation - are independent of what the common beliefs about verifiable events are. Now it is agents' interpretation of the signals which matters.

Also in this section, a strengthening of the solution concepts of the previous section is introduced. Strict interim rationality requires that every agent is always better off under the bet conditional on some public event. Strict bets require strict interim rationality. This strengthening is required in the next section when we introduce risk-aversion. It is the weakest strengthening that will give the equivalence between betting and the general trading problem that I require.

Definition Agents have common beliefs about verifiable events if there exists a "common prior" π such that $\pi_h[\, Q(\omega)\,] = \pi[\, Q(\omega)\,]$, for all $h \in H$, $\omega \in \Omega^*$.

Definition Beliefs $\{\pi_h\}_{h \in H}$ satisfy consistent interpretation if there exists "common prior" π, such that $\pi_h\,[\,\omega\;|Q(\omega)] \,/\, \pi_h\,[\,\omega'\;|Q(\omega')] = \pi(\omega) \,/\, \pi(\omega') \quad \forall\, h \in H,\ \omega' \in P_h(\omega)$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

<u>Definition</u> Beliefs $\{\pi_h\}_{h\in H}$ satisfy (weakly) <u>reconcilable interpretation</u> if there exist $\{\lambda_h\}_{h\in H}$, each $\lambda_h: T_h \rightarrow \mathbb{R}_{(+)++}$, $\{\mu_h\}_{h\in H}$, each $\mu_h: T_h^2 \rightarrow \mathbb{R}_+$, and "common prior" π satisfying, $\Sigma_{t\in E}\,\Sigma_{q\in Q}\,\pi(\{t,q\}) > 0$, for all public events $E \subset T^*$, such that[46]:-

$$\lambda_h(t_h)\pi_h(t|q) + \sum_{t_h'\in T_h}[\mu_h(t_h,t_h')\pi_h(t|q) - \mu_h(t_h',t_h)\pi_h(t_{-h},t_h'|q)] = \pi(t,q) \quad \forall\, h\in H,\ t\in T,\ q\in Q.$$

<u>Definition</u> Beliefs $\{\pi_h\}_{h\in H}$ satisfy (weakly) <u>publicly consistent interpretation</u> if there exist positive (non-negative) priors $\{\zeta_h\}_{h\in H}$, and "common prior" π, with $\pi[\,M(\omega)\,] > 0$, for all $\omega\in\Omega^*$, such that:-

$$\zeta_h[\omega|P_h(\omega)] = \pi_h[\omega|P_h(\omega)] \quad \forall\, h\in H,\ \omega\in\Omega^*,\ \textit{such that}\ \zeta_h[P_h(\omega)] > 0.$$

$$\zeta_h[M(\omega)|Q(\omega)] = \pi[M(\omega)\cap Q(\omega)] \quad \forall\, h\in H,\ \omega\in\Omega^*$$

Consistent interpretation, reconcilable interpretation and publicly consistent interpretation are equivalent to consistency, reconcilability and public consistency when agents have common beliefs about verifiable events.

<u>Definition</u> Bet $\{x_h\}_{h\in H}$ is strictly interim rational if

$$\sum_{\omega\in P_h(\omega^*)} \pi_h(\omega)x_h(\omega) \geq 0 \quad \forall\, h\in H,\ \omega^*\in\Omega^*,$$
with strict inequality for all $h\in H$, $\omega^*\in E$, *for some public event* $E\subset\Omega^*$.

<u>Definitions</u> A bet is a <u>strict unrestricted bet</u> if it satisfies feasibility and strict interim rationality. A bet is a <u>strict incentive compatible bet</u> if it satisfies feasibility, incentive

---

46. I use notation $\pi_h(t|q) = \pi_h(\{t,q\}) \,/\, \Sigma_{t\in T}\,\pi_h(\{t,q\})$.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

compatibility and strict interim rationality. A bet is a strict common knowledge public bet if it satisfies feasibility, strict interim rationality and is public.

Theorem 2 (a) If agents have common beliefs about verifiable events, there is no (strict) unrestricted betting if and only if beliefs satisfy consistent interpretation.

(b) If agents have common beliefs about verifiable events, there is no (strict) incentive compatible betting if and only if beliefs satisfy (weakly) reconcilable interpretation.

(c) If agents have common beliefs about verifiable events, there is no (strict) common knowledge public betting if and only if beliefs satisfy (weakly) publicly consistent interpretation.

The proof is a straightforward extension of the proofs in the previous section, and is omitted. Note that the conditions for existence of strict and standard unrestricted betting are the same, but there is a difference between conditions for existence of strict and standard incentive compatible and public betting. In the following example, any common knowledge public bet must have agent 1 paying agent 2 \$x in event {a,b} and agent 2 paying agent 1 \$x in event {c,d}. But now 1 must have zero expected gain if he observes {a,c}, so there cannot be strict common knowledge public betting.

$\Omega^* = \{a, b, c, d\}$	
$\underline{Q} = (\{a,b\}, \{c,d\})$	
$\pi_1 = (1/4, 1/6, 1/4, 1/3)$	$\pi_2 = (1/3, 1/4, 1/4, 1/6)$
$\underline{P}_1 = (\{a,c\}, \{b,d\})$	$\underline{P}_2 = (\{a,d\}, \{b,c\})$
$x_1 = (-1, -1, 1, 1)$	$x_2 = (1, 1, -1, -1)$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Notice that if all events are verifiable [$Q(\omega) = \{\omega\}$, for all $\omega \in \Omega^*$], then information never leads to betting, under any solution concept, since beliefs must then satisfy consistent interpretation, reconcilable interpretation and publicly consistent interpretation. If all public events are verifiable [$Q(\omega) \subset M(\omega)$, for all $\omega \in \Omega^*$], then beliefs satisfy publicly consistent interpretation, and information cannot lead to common knowledge public betting. But if Q is interpreted as representing restrictions on betting not necessarily deriving from verifiability, then there are natural cases where Q is not finer than M. Before the Iraqi invasion of Kuwait, it was not possible to buy assets contingent on possible future Iraqi invasions of Kuwait (perhaps because transaction costs were high and demand was low). The day following the invasion (a very public event) was one of the busiest trading days of the year on the world's stock markets. Traders were able to make trades reflecting their different beliefs about the world, conditional on Iraq having invaded Kuwait, which they couldn't make before.

4. Trading

In this section, I generalize the results of the previous section to risk-averse agents trading in many goods. There is minimal discussion of the solution concepts as extension from the betting case is straightforward, and the reader can refer to section 2. The key conceptual issue in this section is the restriction to "no uninformed trade" environments. With risk averse agents, and many goods, there are motives for trade which have nothing to do with agents' beliefs and information (e.g. insurance, different preferences among goods, state dependence of utility functions). I am interested in environments where other motives for trade are not present, so I restrict attention to the case where endowments are constrained efficient, with respect to trades depending only on verifiable events, and where agents' private information is

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

not payoff-relevant (i.e. utility functions are independent of private information). In this case, conditions for no trade under the different solution concepts depend only on agents' information and beliefs (and not on their endowments and utility functions). This framework is very close to that used in Milgrom and Stokey's (1982) "no trade" result. Milgrom and Stokey observed that in this framework, it is agents' interpretation of signals - their beliefs about the distribution of agents' private signals conditional on verifiable events - which matter, not their beliefs about verifiable events, since they can make trades contingent on the verifiable events prior to observing private information.

4.1 The general trading problem

Consider the uncertain environment described in section 1, with $E = [\Omega^*, Q, \{P_h, \pi_h\}_{h\in H}]$, but now suppose there are a finite set L of different goods. I also write L for the number of goods. For each agent h, we must specify his utility function, $u_h: \mathbf{R}^L \times \Omega \rightarrow \mathbf{R}$ and endowment, $e_h: \Omega \rightarrow \mathbf{R}^L_+$, with each e_h measurable with respect to P_h. A feasible trade now involves bundles of commodities: $\{x_h\}_{h\in H}$ is a feasible trade if each $x_h: \Omega \rightarrow \mathbf{R}^L$ and $\Sigma_{h\in H}\, x_h(\omega) \leq 0$.

Solution concepts are defined by natural extension from the betting case:-

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Definition $\{x_h\}_{h\in H}$ is [strictly] interim rational if

$$\sum_{\omega\in P_h(\omega')} \pi_h(\omega)\, u_h[e_h(\omega)+x_h(\omega),\omega] \geq \sum_{\omega\in P_h(\omega')} \pi_h(\omega)\, u_h[e_h(\omega),\omega] \quad \forall\ h\in H,\ \omega'\in\Omega^*.$$

and with strict inequality for at least one h,ω'.

[with strict inequality for all h, and all $\omega'\in E$ for some public event $E \subset \Omega^*$]

Definition $\{x_h\}_{h\in H}$ is incentive compatible if

$$\sum_{t_{-h}\in T_{-h}} \sum_{q\in Q} \pi_h(\{t,q\})u_h[e_h(\{t,q\})+x_h(\{t,q\})] \geq \sum_{t_{-h}\in T_{-h}} \sum_{q\in Q} \pi_h(\{t,q\})u_h[e_h(\{t,q\})+x_h(\{(t_{-h},t_h'),q\})]$$

Definition $\{x_h\}_{h\in H}$ is public if each x_h is measurable with respect to the join of Q, the partition of verifiable events, and M, the partition of public events.

Definition $\{x_h\}_{h\in H}$ is a (strict) unrestricted trade if it is feasible and (strictly) interim rational.

Definition $\{x_h\}_{h\in H}$ is a (strict) incentive compatible trade if it is feasible, (strictly) interim rational and incentive compatible.

Definition $\{x_h\}_{h\in H}$ is a (strict) common knowledge public trade if it is feasible, (strictly) interim rational and public.

Now "betting" is the special case of trading where $L=1$ and $u_h(x,\omega) = x$, for all $x\in\mathbf{R}^L$, $\omega\in\Omega$. Note that if there is no unrestricted trade, then endowments satisfy "interim classical efficiency" in the sense of Holmstrom and Myerson (1983). But if there is no incentive

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

compatible trade, endowments are not necessarily "interim incentive compatible efficient", since there may not be an incentive compatible mechanism capable of implementing the endowments.

I could derive some general conditions for no trade but typically they would involve endowments and utility functions (in fact, marginal utilities), and the conditions have no obvious interpretation[47]. By looking at the "no uninformed trade" case described below, I can derive conditions for information leading to trade that are independent of endowments and utility functions.

4.2 No Uninformed Trade

Imagine that risk-averse agents were able to undertake trades, conditional only on verifiable events, before observing their own information. Suppose they traded to an ex ante constrained efficient allocation - constrained by the fact that trades can depend on verifiable events, not on (as yet unobserved and perhaps never verifiable) private information. Now suppose, in addition, that verifiable events incorporate all "payoff-relevant" uncertainty. Thus private signals are valuable to agents because they are correlated with payoff-relevant events, but they do not directly affect utility (they are not arguments of the utility function), nor do endowments depend on them. In this case, a constrained efficient allocation would arise in a competitive General Equilibrium with incomplete markets[48]. Or it might be reached by some

47. If I assume agents are risk-averse, then the interim rationality condition is convex, and thus I can derive linear algebraic conditions, analogous to the betting case, but involving marginal utilities, to preclude unrestricted and common knowledge public trade. Incentive compatibility constraints are not necessarily convex. But if randomization were allowed over possible trades (and thus incentive compatibility constraints became linear), results for incentive compatible trade could be derived.

48. It is not generally the case that general equilibria with incomplete markets [GEI] are constrained efficient [Geanakoplos and Polemarchakis (1986)]. But consider the Arrow-Debreu (complete markets) equilibria of the economy where uncertainty is specified only by verifiable events. Such equilibria are

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

other prior trading process[49]. It will turn out that when these three properties are satisfied, conditions for trade are independent of endowments and utility functions. Intuitively, this is because gains from trade because of different marginal utilities across states can be efficiently exploited before private information is available. Private information permits further (subjective) gains from trade by exploiting differences in prior beliefs.

Assumption 1: Each u_h is concave, differentiable and strictly increasing. This gives us risk-aversion.

Assumption 2: Constrained ex ante efficiency

There does not exist feasible $\{x_h\}_{h\in H}$, with each x_h measurable with respect to Q, such that

$$\sum_{\omega\in\Omega^*} \pi_h(\omega)\{u_h[e_h(\omega)+x_h(\omega),\omega]-u_h[e_h(\omega),\omega]\} \geq 0 \quad \forall\, h\in H$$

with strict inequality for at least one $h\in H$.

Thus we require endowments to be constrained ex ante efficient with respect to the verifiability partition Q.

efficient with respect to verifiable events and but are also the same as GEI equilibria of the larger economy. Note that this argument depends on the assumption that payoff-relevant events are verifiable. Note also that (for expository reasons) I have not made assumptions to ensure existence of GEI equilibria.

49. If agents anticipated that they would later be observing their private information signals, the rationale for trading to ex ante constrained efficient allocations is not clear; this is discussed in section 4.3. However, there is always remains a "counterfactual" interpretation of the result: when will agents trade with each other in situations where they would not have been prepared to trade with each other in the absence of the private information?

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Assumption 3: Payoff-relevant events are verifiable

First, define h's "payoff relevance partition", Z_h, by $Z_h(\omega) = \{\omega' \in \Omega^* \mid e_h(\omega') = e_h(\omega)$ and $u_h(x,\omega') = u_h(x,\omega)\ \forall x \in \mathbf{R}^L\}$. For given x, h cares only about which event in $\underline{Z}_h$ occurs, and not about any more refined events. But before x is known, information about other events may increase expected utility (because those events are correlated with events in $\underline{Z}_h$, and you can act accordingly). Notice that, in general, private information, $\{P_h\}_{h \in H}$, verifiability restrictions, Q, and payoff-relevance partitions, $\{Z_h\}_{h \in H}$, are independent properties, except for the requirement that, for each h, P_h is finer than Z_h[50]. You may know that a member of the Open Market Committee has predicted an increase in the discount rate tomorrow, but we cannot verify the prediction for contractual purposes, nor does the prediction directly affect our ex post utilities (though it may be correlated with events which do). Tomorrow's discount rate is verifiable, but we don't know what it will be, and it does not directly affect our ex post utilities. My bank manager's propensity to foreclose on my debts is payoff-relevant for me (and not for you), but it may be unverifiable for contractual purposes, and you don't know what it is[51].

In what follows, I assume that the verifiability partition, Q, is finer than each Z_h, thus $\omega' \in Q(\omega)$ implies $e_h(\omega') = e_h(\omega)$ and $u_h(x,\omega') = u_h(x,\omega)$, for all $h \in H$, $x \in \mathbf{R}^L$. Thus all payoff-relevant events are verifiable. I touched on one environment where this might be true in section 1. Suppose $\underline{Q}$ represents the set of "natural events", where natural events represent "substantive happenings that are not themselves described in terms of people knowing something"[52]. All

---

50. It was already assumed (in section 4.1) that $\omega' \in P_h(\omega) \Rightarrow e_h(\omega') = e_h(\omega)$. It is perhaps natural to assume also that $\omega' \in P_h(\omega) \Rightarrow u_h(x,\omega') = u_h(x,\omega)\ \forall\ x \in \mathbf{R}^L$. Neither assumption affects the formal argument.

51. Following Savage (1972), the treatment of uncertainty requires that we can describe Ω independently of the acts available to agents.

52. Aumann (1989). He notes that "the distinction between 'natural events' and 'events' is analogous to that between 'states of nature' and 'states of the world' that one sees in the literature".

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

"natural events" are verifiable, and so there are no uncontractable events in the conventional sense. But trading decisions with asymmetric information depend also on agents' beliefs about other agents' beliefs about natural events, and beliefs about what information other agents have observed, together with beliefs about other agents' beliefs about your beliefs etc... Our description of the uncertain environment must be expanded from $\underline{\Omega}$ to Ω^*[53]. Contracts cannot be written contingent on events concerning beliefs. Private information then includes your knowledge of your "type" (beliefs about $\underline{\Omega}$, partition of $\underline{\Omega}$, second order beliefs, etc...)[54]. Presumably, only "natural events" are payoff-relevant.

The "payoff-relevant events are verifiable" assumption is true in this interpretation, but not restricted to it.

Definition $[\Omega^*, Q, \{\pi_h, P_h, e_h, u_h\}_{h\in H}]$ is a no uninformed trade (NUT) environment if payoff-relevant events are verifiable, agents are risk-averse and endowments are constrained ex ante efficient.

The NUT environment is roughly that in Milgrom and Stokey (1982)[55]. In the NUT environment, I can use a reduced form notation without confusion. I write $e_h(q) = e_h(\{t,q\})$ and $u_h(x,q) = u_h(x,\{t,q\})$ for all $x \in \mathbb{R}^L$, $q \in \underline{Q}$, $t \in T$. Also I write $\pi_h(q) = \Sigma_{t\in T}\, \pi_h(\{t,q\})$.

53. Mertens and Zamir (1985), Brandenburger and Dekel (1985) and Aumann (1989) derive formal models of such hierarchies of beliefs.

54. Note that under this interpretation of agents private signals as "types", it does not make sense to imagine that agents had a prior opportunity to trade. But the counterfactual interpretation of the results (see note 49) remains.

55. Some authors have interpreted Milgrom and Stokey [MS] as assuming that the initial allocation is (unconstrained) efficient, leading to the conclusion that MS do not require a restriction on priors at all for their theorem. In fact, the MS theorem assumes only constrained efficiency (with respect to payoff-relevant and implicitly verifiable events), but their restriction on prior beliefs and the payoff-relevance assumption guarantee that constrained efficiency is equivalent to unconstrained efficiency.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

This notation to can be used to characterize NUT environments:-

Lemma If E is a NUT environment, there exist positive priors $\{\phi_l\}_{l\in L}$ on $\underline{Q}$, and positive constants $\{k_h\}_{h\in H}$ such that $k_h\ \partial u_h/\partial x_l[e_h(q),q]\ \pi_h(q) = \phi_l(q)\quad \forall\ h\in H,\ q\in \underline{Q},\ l\in L$.

Intuitively, when trading the lth good, agents will behave as if they have common prior, ϕ_l, on verifiable events. This will drive the results in the next section.

4.3 When does information lead to trade?

I say that "information leads to trade" if agents trade in a "no uninformed trade" environment. The conditions are identical to the betting case. The theorems are proved by reduction to the betting case.

Theorem 3 (a) Information leads to strict unrestricted trade if and only if beliefs do not satisfy consistent interpretation; (b) information leads to strict incentive compatible trade if and only if beliefs do not satisfy weakly reconcilable interpretation; (c) information leads to strict common knowledge public trade if and only if beliefs do not satisfy weakly publicly consistent interpretation.[56]

All three results are proved by reducing the trading case to the betting case of the previous section. I give the proof for (a) only; proofs for (b) and (c) take the same form. I

56. In the special case when all information is public, the three kinds of trade, and thus the three kinds of necessary and sufficeint conditions, are equivalent. Both Hakansson et al. (1982) and Milgrom and Stokey (1982) noted that consistent interpretation is the necessary and sufficient condition for no trade in this case.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

prove there exists strict unrestricted trading, in a "no uninformed trade" environment, if and only if there exists strict unrestricted betting in an environment where agents have the same beliefs conditional on verifiable events, and some arbitrary common beliefs about verifiable events. But we know from the previous section that consistent/reconcilable/publicly consistent interpretation are the necessary and sufficient conditions to preclude unrestricted/incentive compatible/common knowledge public betting with common beliefs about verifiable events.

Proof of (a)

1. Recall that by lemma of the previous section, in a NUT environment, there exist positive constants $\{k_h\}_{h\in H}$ and positive priors on $\underline{Q}$, $\{\phi_l\}_{l\in L}$, such that

$$k_h \frac{\partial u_h}{\partial x_l}[e_h(q),q]\,\pi_h(q) = \phi_l(q) \quad \forall\; h\in H,\; l\in L,\; q\in Q. \qquad [4.1]$$

2. Now suppose that there exists an unrestricted trade, $\{x_h\}_{h\in H}$. Consider the bet $\{y_h\}_{h\in H}$ defined by:-

$$y_h(\{t,q\}) = \frac{k_h \pi_h(q)\{u_h[e_h(q)+x_h(\{t,q\}),q]-u_h[e_h(q),q]\}}{\phi^*(q)}, \qquad [4.2]$$

where ϕ^* is any positive prior on $\underline{Q}$.

3. Now $\{x_h\}_{h\in H}$ feasible $\Rightarrow$ $\Sigma_{h\in H}\, x_{hl}(\{t,q\}) \leq 0 \quad \forall\, t\in T,\, q\in \underline{Q},\, l\in L.$ [4.3]

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

$$\sum_{h\in H} y_h(\{t,q\}) \leq \frac{\sum_{h\in H}\sum_{l\in L} k_h \pi_h(q) \frac{\partial u_h}{\partial x_l}[e_h(q),q]\, x_{hl}(\{t,q\})}{\phi^*(q)} \quad \forall\, t\in T,\ q\in Q,$$

by [4.2] *and risk-aversion.*

$$\leq \frac{1}{\phi^*(q)} \sum_{l\in L} \phi_l(q) \sum_{h\in H} x_{hl}(\{t,q\}) \quad \forall\, t\in T,\ q\in Q$$

by [4.1]

$$\leq 0 \qquad \forall\, t\in T,\ q\in Q$$

by [4.3]

$\Rightarrow$ $\{y_h\}_{h\in H}$ feasible

4. Write $\pi_h(t|q) = \pi_h(\{t,q\}) / \pi_h(q)$; now for any positive prior ϕ^* on Q, and $\forall\, h\in H$, $t_h\in T_h$

$$\sum_{t_{-h}\in T_{-h}} \sum_{q\in Q} \pi_h(t|q)\phi^*(q) y_h(t,q) - k_h \sum_{t_{-h}\in T_{-h}} \sum_{q\in Q} \pi_h(\{t,q\})\{u_h[e_h(q)+x_h(\{t,q\}), q] - u_h[e_h(q), q]\}$$

thus $\{x_h\}_{h\in H}$ strictly interim rational implies $\{y_h\}_{h\in H}$ strictly interim rational.

5. Conversely, suppose that $\{y_h\}_{h\in H}$ is a strictly unrestricted bet and define trade $\{x_h\}_{h\in H}$ by:-

$x_{h1}(\{t,q\}) = \epsilon y_h(\{t,q\})$, for some $\epsilon > 0$, $\forall\, t\in T$, $q\in Q$, $\{t,q\}\in E$, where E is the public event where interim rationality constraints hold with strict inequality.

$x_{hl}(t,q) = 0$ if $l \neq 1$ or if $\{t,q\} \notin E$. [4.4]

6. Now $\{y_h\}_{h\in H}$ feasible implies $\{x_h\}_{h\in H}$ feasible by construction.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

7. Suppose $\{y_h\}_{h\in H}$ is strictly interim rational. Then

Let $E_h = \{t_h \in T_h \mid (t,q)\in E \text{ for some } t_{-h}\in T_{-h},\ q\in \underline{Q}\}$

$$\text{if } t_h \in E_h, \text{ then } \frac{1}{\epsilon}\sum_{t_{-h}\in T_{-h}}\sum_{q\in Q}\pi_h(\{t,q\})\{u_h[e_h(q)+x_h(\{t,q\}),q]-u_h[e_h(q),q]\}$$

$$\underset{as\ \epsilon\to 0}{\to}\sum_{t_{-h}\in T_{-h}}\sum_{q\in Q}\pi_h(\{t,q\})\frac{\partial u_h}{\partial x_1}[e_h(q),q]y_h(\{t,q\}) \qquad by\ [4.4]$$

$$\underset{as\ \epsilon\to 0}{\to}\frac{\sum_{t_{-h}\in T_{-h}}\sum_{q\in Q}\pi_h(\{t,q\})\phi_i(q)y_h(\{t,q\})}{k_h\pi_h(q)} \qquad by\ [4.1]$$

$$\underset{as\ \epsilon\to 0}{\to}\frac{1}{k_h}\sum_{t_{-h}\in T_{-h}}\sum_{q\in Q}\pi_h(t|q)\phi_i(q)y_h(\{t,q\})$$

$$\text{if } t_h \notin E_h, \text{ then } \frac{1}{\epsilon}\sum_{t_{-h}\in T_{-h}}\sum_{q\in Q}\pi_h(\{t,q\})\{u_h[e_h(q)+x_h(\{t,q\}),q]-u_h[e_h(q),q]\} \to 0$$

therefore, for some $\epsilon>0$, $\{y_h\}_{h\in H}$ strictly interim rational implies $\{x_h\}_{h\in H}$ is strictly interim rational. Notice there are two senses in which the results in this section answer the question "when does information lead to trade?" One is counterfactual: is trade occurring in a situation in which agents would not have traded in the absence of their private information? Another is temporal: when would the arrival of information lead to trade? The latter is more problematic. If agents did not anticipate receiving information signals in the future, agents would indeed trade to an ex ante constrained efficient allocation in a co-operative environment (also, we noted, in a General Equilibrium with incomplete markets). But what if the agents anticipated receiving those signals in the future? Could they commit to a particular trade after

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

observing their information? There is no simple answer without describing how trades might be chosen, an issue that has been deliberately avoided[57].

Section 5: Conclusion

I have referred during the course of the chapter to related work in both the "partition" framework and the "type space" framework; here, I consider three additional issues.

Geanakoplos (1989) considers when trade occurs between agents who have common prior beliefs, but who "misinterpret" information. Each agent's information is represented by a "possibility correspondence" [PC] mapping $P_h: \Omega \rightarrow 2^{\Omega}/\varnothing$, which does not necessarily satisfy the characteristic property of partitions, that $P_h(\omega) = \{\omega' \in \Omega \mid P_h(\omega') = P_h(\omega)\}$[58]. Geanakoplos shows that common knowledge trade is precluded under weaker assumptions than (Nash) equilibrium trade. The result is of interest both because of the natural interpretations of the conditions for trade and because it clarifies the relation between solution concepts. Brandenburger, Dekel and Geanakoplos (1989) showed that, for any decision problem facing a set of agents with common priors and possibility correspondences, there is a decision theoretically equivalent problem with heterogeneous priors and standard partitions. This chapter gives the heterogeneous prior analogues of Geanakoplos' results. The relation is discussed in detail in chapter IV.

---

57. In an earlier draft, I considered one scenario; starting at an arbitrary allocation, does there exist a "prior trade", contingent only on verifiable events, and an unrestricted / incentive compatible / common knowledge trade "interim trade", such that the combined trade satisfies ex ante rationality, and the interim trade satisfies the conditions for unrestricted/ incentive compatible/ common knowledge trade? Theorems 3a, 3b and 3c are as before, with "consistent" beliefs replaced by "common" beliefs. [This case is thus related to ex ante incentive compatible efficiency, Holmstrom and Myerson (1983)].

58. This framework has been studied by Bacharach (1985), Samet (1990) and Shin (1987).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Milgrom and Stokey (1982) is perhaps the most cited "no trade" result and also the closest to the framework presented in section 4. Milgrom and Stokey identified a condition on priors, weaker than the common prior assumption, sufficient to preclude trade. Their "concordant prior" assumption is equivalent (in my terminology) to assuming common beliefs about private information, conditional on verifiable/payoff-relevant events. This chapter showed that this assumption can be weakened in two more ways to derive a necessary condition (as a function of agents' final information) for no common knowledge public trade (which is equivalent to the solution concept they use): agents need only have consistent, not common, beliefs; and they need have consistent beliefs only about public events, conditional on verifiable events (i.e. publicly consistent interpretation). In fact, their "concordant prior" assumption is sufficient to preclude not only common knowledge public trade (as their theorem states), and common knowledge unrestricted trade (as their discussion makes clear), but also unrestricted trade with Nash acceptance.

In this chapter, I have been concerned with agents who have exogenously given heterogeneous beliefs about events which are exogenous to the model. This should be contrasted with the very different issues that arise in game theory and other economic environments where agents' beliefs about endogenous events are themselves endogenous, and the question arises whether it is natural to impose the common prior assumption[59].

The common prior assumption has been criticized as empirically and philosophically unreasonable. Yet it underlies many important results in economics, not just the "no trade" results discussed here. The ultimate argument in favor is always: if you don't assume it, you

59. The literature on Bayesian foundations of game thoery shows that Nash equilibrium can be justified by assuming a common prior on agents' strategy spaces [Aumann (1987), Brandenburger and Dekel (1987b)]. On the other hand, one interpretation of the same formal results is that agents are choosing strategies conditional on some exogenous (not payoff-relevant) events about which they may or may not disagree [Aumann (1974)].

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

can't say anything. Chapter V includes a general discussion of the common prior assumption in economic theory, and applications with heterogeneous prior beliefs. This chapter shows that the common prior assumption is not necessary to discuss the existence of rational trading with asymmetric information, and that it is possible to use models with heterogeneous prior beliefs to gain new insights into trading behavior. Symmetric information and complete contingent contracting are not natural assumptions in a trading environment where the uncertainty space must incorporate not only "natural events", but also agents' beliefs about other agents' beliefs about "natural events", and so on. Therefore I examined the interaction between heterogeneous prior beliefs, asymmetric information, and restrictions on verifiability. For three natural assumptions about contracting, I extended Harsanyi's (1967) concept of consistency to characterize how much agreement is required to preclude trade.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Chapter II Appendix: Acceptance Rules and Contracting Rules

I use notation from the main text. There are two components of a solution concept: first, when is a bet accepted? Second, what kind of contracts can be implemented? I discuss each in turn.

A1. The Acceptance Game

Suppose $\{x_h\}_{h\in H}$ represent agents' expected payoffs under some bet (I haven't specified what kind of bets yet). Define an "acceptance game" as follows: each agent h chooses a strategy a_h, where $a_h: \Omega^* \rightarrow [0,1]$, a_h measurable with respect to P_h. Interpret $a_h(\omega) = 1$ to mean "agent h accepts the bet at ω" and $a_h(\omega) = 0$ means "agent h rejects the bet at ω". Bets are accepted only if all agents accept. Define utilities in the natural way:-

$$u_{h'}[\ \{a_h\}_{h\in H}\ ] = \Sigma_{\omega\in\Omega^*}\ \pi_{h'}(\omega)\ a(\omega)\ x_{h'}(\omega), \quad \text{where } a(\omega) = \Pi_{h\in H}\ a_h(\omega)$$

$\{x_h\}_{h\in H}$ is Nash accepted if there exists a Nash equilibrium $\{a_h^*\}_{h\in H}$ with $u_{h'}[\ \{a_h^*\}_{h\in H}\ ] > 0$ for some $h'\in H$. $\{x_h\}_{h\in H}$ is pure Nash accepted if there is Nash acceptance, with $a_h(\omega) = 0$ or 1, for all $h\in H$, $\omega\in\Omega^*$. $\{x_h\}_{h\in H}$ is common knowledge accepted if there is pure Nash acceptance with $\{\omega \mid a^*(\omega) = 1\}$ a public event. $\{x_h\}_{h\in H}$ is unconditionally accepted if there is common knowledge acceptance with $a^*(\omega) = 1$, for all $\omega\in\Omega^*$.

The different acceptance rules are discussed in the main text.

A2 Contracting rules

$\{x_h\}_{h\in H}$ is feasible if $\Sigma_{h\in H}\ x_h(\omega) \leq 0$, for all $\omega\in\Omega$.

$\{x_h\}_{h\in H}$ is implementable if there exist message spaces $\{M_h\}_{h\in H}$, message reporting distributions, $\{\sigma_h\}_{h\in H}$, where $\sigma_h(.\,|t_h)$ is a probability distribution over agent h's message space M_h, and a feasible bet $\{y_h\}_{h\in H}$, depending only on verifiable events and messages, such that:-

$$x_h(\{t,q\}) = \sum_{m\in M}\left\{\prod_{h'\in H}\sigma_{h'}(m_{h'}|t_{h'})\right\}y_h(q,m)$$

with the $\{\sigma_h\}_{h\in H}$ chosen optimally [the exact meaning of this only becomes clear once an acceptance rule is chosen].

$\{x_h\}_{h\in H}$ is public if each x_h is measurable with respect to the join of public event [M] and verifiable events [Q].

$\{x_h\}_{h\in H}$ is verifiable if each x_h is measurable with respect to the partition of verifiable events [Q].

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

A3 Combining acceptance rules and contractual rules

Four acceptance rules and four contracting rules give sixteen potential solution concepts. But if we are asking about existence (e.g. does there exist an implementable bet with pure Nash acceptance?), and if contractual rules are rich enough, the acceptance rules don't matter. By definition, if bet $\{x_h\}_{h\in H}$ is unconditionally accepted, it is Nash accepted. But it is also the case that if there exists a feasible bet with Nash acceptance, then there exists a feasible bet with unconditional acceptance (proved as result 1 below). Similarly, by definition, if implementable bet $\{x_h\}_{h\in H}$ is unconditionally accepted, it also Nash accepted. But if there exists an implementable bet with Nash acceptance, then there exists implementable bet with unconditional acceptance (proved as result 2 below). If there exists a public bet with common knowledge acceptance, then there exists a public bet with unconditional acceptance. The table below shows how the sixteen candidate solution concepts reduce, for existence questions, to six. The three cases with stars (*) are those considered in the main text. The examples in section A4 show that there is a difference between the conditions for existence of public and verifiable betting under the different acceptance rules.

<table>
<tr><th></th><th>Nash</th><th>Pure Nash</th><th>Common Knowledge</th><th>Unconditional</th></tr>
<tr><td>Unrestricted Feasible</td><td colspan="4">Unrestricted Betting*</td></tr>
<tr><td>Implementable Feasible</td><td colspan="4">Incentive Compatible Betting*</td></tr>
<tr><td>Public Feasible</td><td rowspan="2">Nash Public Betting</td><td rowspan="2">Pure Nash Public Betting</td><td colspan="2">Common Knowledge Public Betting*</td></tr>
<tr><td>Verifiable Feasible</td><td></td><td>Unconditional Verifiable Betting</td></tr>
</table>

Result 1 Suppose there is Nash acceptance (with strategies $\{a_h\}_{h\in H}$) of bet $\{x_h\}_{h\in H}$. Then there exists a feasible bet $\{x_h^*\}_{h\in H}$ with unconditional acceptance.

Proof We know, by optimality of agent h's strategy in Nash equilibrium, that

$$\sum_{\omega\in P_h(\omega^*)} \pi_h(\omega)a_{-h}(\omega)x_h(\omega) > 0 \Rightarrow a_h(\omega^*) = 1 \qquad \forall\, h\in H,\ \omega^*\in\Omega^*$$

$$\sum_{\omega\in P_h(\omega^*)} \pi_h(\omega)a_{-h}(\omega)x_h(\omega) < 0 \Rightarrow a_h(\omega^*) = 0 \qquad \textit{where } a_{-h}(\omega) = \prod_{h'\neq h} a_{h'}(\omega)$$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

We know by Nash acceptance that $\Sigma_{h\in H}\ \Sigma_{\omega\in\Omega^*}\ \pi_h(\omega)\ a(\omega)\ x_h(\omega) > 0$

Now define $\{x_h^*\}_{h\in H}$ by $x_h^*(\omega) = a(\omega)\ x_h(\omega)$, for all $\omega\in\Omega^*$. Each agent h is prepared to unconditionally accept $\{x_h^*\}_{h\in H}$, since

$$\sum_{\omega\in P_h(\omega^*)} \pi_h(\omega)x_h(\omega) \geq 0 \quad \textit{if} \quad a_h(\omega^*) = 1$$
$$\sum_{\omega\in P_h(\omega^*)} \pi_h(\omega)x_h(\omega) = 0 \quad \textit{if} \quad a_h(\omega^*) < 0$$

$$\textit{and} \sum_{h\in H}\sum_{\omega\in\Omega^*} \pi_h(\omega)x_h^*(\omega) = \sum_{h\in H}\sum_{\omega\in\Omega^*} \pi_h(\omega)a(\omega)x_h(\omega) > 0$$

Result 2 If there is Nash acceptance of an implementable bet, then there is unconditional acceptance of a "truth-telling" implementable bet.

Thus this result is a statement of the "revelation principle" with the twist that we show Nash acceptance is equivalent to unconditional acceptance.

Proof There exists a game described by a set of messages $\{M_h\}_{h\in H}$ and a feasible bet depending only on verifiable events and messages, $\{y_h\}_{h\in H}$. Write $M = M_1$ x.. x M_H; now each y_h: $\underline{Q}$ x $M \to \mathbf{R}$, with the sum of the y_h less than or equal to zero. Each agent's strategy consists of his message reporting rule, σ_h, and his acceptance rule, a_h. Agent h's utility is:-

$$u_h(\{a_h,\sigma_h\}_{h\in H}) = \sum_{q\in Q}\sum_{t\in T}\sum_{m\in M} \pi_h(\{t,q\})\left\{\prod_{h'\in H}\sigma_{h'}(m_{h'}|t_{h'})\right\}a(\{t,q\})y_h(q,m)$$

If there is Nash acceptance of an implementable bet, then there exist $\{a_h^*,\ \sigma_h^*\}_{h\in H}$ such that:-

$$[1]\ u_h[\{a_h^*,\sigma_h^*\}_{h\in H}] \geq u_h[(a_h,\sigma_h),\{a_{h'}^*,\sigma_{h'}^*\}_{h'\neq h}] \quad \forall\ (a_h,\sigma_h) \neq (a_h^*,\sigma_h^*)$$

$$[2]\ \sum_{h\in H} u_h[\{a_h^*,\sigma_h^*\}_{h\in H}] > 0$$

$$\textit{Now let } x_h^*(\{t,q\}) = \sum_{m\in M}\left\{\prod_{h'\in H}\sigma_{h'}(m_{h'}|t_{h'})\right\}a(\{t,q\})y_h(q,m)$$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

For this new bet, we can set $M_h = T_h$, $\sigma_h(t_h | t_h) = 1$ and $a_h(t_h) = 1$ for all t_h, and then there is unconditional acceptance.

A4. Examples: different acceptance rules for public and verifiable trades

In example A1, there is an incentive compatible bet as shown, but no public bet with Nash acceptance.

Example A1
$\Omega^* = \{a,b,c,d\}$
$\underline{\Omega} = (\{a,b,c,d\})$
$\pi_1 = \{1/6, 1/3, 1/3, 1/6\}$ $\pi_2 = \{1/3, 1/6, 1/6, 1/3\}$
$\underline{P}_1 = (\{a,b\}, \{c,d\})$ $\underline{P}_2 = (\{a,c\}, \{b,d\})$
$x_1 = (-1, 1, 1, -1)$ $x_2 = (1, -1, -1, 1)$

In example A2, there exists a public bet with Nash acceptance as shown, but there no public bet with pure Nash acceptance.

Example A2
$\Omega^* = \{a,b,c,d,e\}$
$\underline{\Omega} = (\{a,c,e\}, \{b,d\})$
$\pi_1 = \{3/8, 1/8, 1/8, 1/4, 1/8\}$ $\pi_2 = \{1/8, 3/8, 1/8, 1/4, 1/8\}$
$\underline{P}_1 = (\{a,b,e\}, \{c,d\})$ $\underline{P}_2 = (\{a,b,c\}, \{d,e\})$
$a_1 = \{1, 1, 1/2, 1/2, 1\}$ $a_2 = \{1, 1, 1, 1/2, 1/2\}$
$x_1 = \{1, -1, 1, -1, 1\}$ $x_2 = \{-1, 1, -1, 1, -1\}$

In example A3, there exists a public bet with pure Nash acceptance as shown, but no public bet with common knowledge acceptance.

Example A3
$\Omega^* = \{a,b,c,d\}$
$\underline{\Omega} = (\{a,c\}, \{b,d\})$
$\pi_1 = \{1/6, 1/3, 1/3, 1/6\}$ $\pi_2 = \{1/3, 1/6, 1/6, 1/3\}$
$\underline{P}_1 = (\{a,b\}, \{c,d\})$ $\underline{P}_2 = (\{a,b,c,d\})$
$a_1 = (1, 1, 0, 0)$ $a_2 = (1, 1, 1, 1)$
$x_1 = (-1, 1, -1, 1)$ $x_2 = (1, -1, 1, -1)$

In example A4, there is a verifiable bet with common knowledge acceptance as shown, but there is no verifiable bet with unconditional acceptance.

Example A4
$\Omega^* = \{a,b,c,d\}$
$\underline{\Omega} = (\{a,c\}, \{b,d\})$
$\pi_1 = \{1/6, 1/3, 1/3, 1/6\}$ $\pi_2 = \{1/3, 1/6, 1/6, 1/3\}$
$\underline{P}_1 = (\{a,b\}, \{c,d\})$ $\underline{P}_2 = (\{a,b\},\{c,d\})$
$a_1 = (1, 1, 0, 0)$ $a_2 = (1, 1, 0, 0)$
$x_1 = (-1, 1, -1, 1)$ $x_2 = (1, -1, 1, -1)$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Chapter III

THE VALUE OF MISINTERPRETED INFORMATION

Blackwell proved a remarkable theorem showing an equivalence between decision theoretic and statistical notions of informativeness[60]. Suppose you face a decision problem and you have a choice between two "experiments" to observe. An "experiment" is a distribution of signals correlated with payoff-relevant uncertainty. Experiment 1 is said to be more valuable than experiment 2 if, for every decision problem, you would choose experiment 1 over experiment 2. Blackwell showed that experiment 1 is more valuable than experiment 2 if and only if experiment 1 is a sufficient statistic for experiment 2, i.e. if experiment 2 could have been generated by adding noise to experiment 1.

Blackwell supposed that beliefs are correctly (i.e. Bayesian) updated after observing the experiment. I will be concerned with the case where agents' beliefs after observing the experiment are not necessarily correct. The concluding chapter V has a discussion of the axiomatic basis for such assumptions. It is still possible to characterize an agent's ex ante utility, for a given decision problem, where he chooses optimally given his (incorrect) beliefs. I show (in theorem 1 of section 2) that a linear algebraic condition called "sufficiency in differences" is the necessary and sufficient condition for one experiment to be preferred to another. I give a couple of interpretations of this condition: one is that the difference between the true distributions under the two experiments could have been derived from the difference between two other experiments, where the preferred experiment has correct beliefs and is sufficient for the second. I show how this specializes in the case of rationally interpreted experiments to Blackwell's Theorem. Another question of interest is when some (possibly

60. Blackwell (1951); see also Cremer (1982) and Laffont (1989).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

misinterpreted) information is better than no information. An experiment is better than no experiment if and only if beliefs are reconcilable with the true distribution. Beliefs are reconcilable if they look like a noisy version of the true distribution.

One natural way to model agents who make mistakes has been introduced into the literature by generalizing the notion of partition information[61]. Agents' beliefs may be represented by a mapping from the set of states of the world to subsets of the set of states of the world, but without requiring that the mapping satisfies the usual partition property. Geanakoplos (1989) has proved a remarkable Blackwell's theorem type property for such "possibility correspondences" (PCs): a possibility correspondence is preferred to a courser partition if and only if the possibility correspondence satisfies three natural properties: non-delusion (the true state of the world is always considered possible); "knowing that you know" (if you know that an event has occurred, you know that you know); and nestedness (at any two states of the world, either you know strictly more in one state than the other, or else possible states at the two states are mutually exclusive). In section 3, I show how Geanakoplos' result can be obtained and generalized to comparisons of any pair of non-deluded PCs by considering possibility correspondences as a special case of the framework of section 2.

The presentation thus reverses the order of research, which started from a desire to understand Geanakoplos' result. More important than the extensions of the PC results is the introduction of a dual framework for analyzing agents who make mistakes. Multipliers of the non-existence conditions have interpretations as probabilities. This dual approach is closely related to Shin's (1989) dynamic representation of Geanakoplos' conditions for ranking informativeness. The parameters of Shin's dynamic representation are in fact multipliers of the

---

61. Bacharach (1985), Brandenburger, Dekel and Geanakoplos (1989), Rubinstein and Wolinsky (1990), Samet (1990), Shin (1987, 1989).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

dual representation. Thus while Shin showed that his dynamic representation was equivalent to the logical properties of PCs which Geanakoplos showed are necessary and sufficient conditions for ranking informativeness, I give a direct argument showing that dynamic representation is the necessary and sufficient condition.

But the discussion thus far skipped a critical problem. There may exist many optimal decision rules, for a given decision problem and given beliefs. If beliefs are rational, this presents no problem: any of the optimal decision rules would give the same ex ante utility. This is not true if beliefs are not rational. Alternative criteria for ranking experiments arise depending on whether an experiment must have higher ex ante utility for all or only some optimal decision rules; whether the same action must be taken at the same signal in different experiments; whether we require higher ex ante for all, or just for "almost all", decision problems. Section 1 defines the different orderings but identifies a single natural ordering to study in the rest of the chapter. Readers prepared to take on trust the ordering used in sections 2 and 3 can skim section 1.

1. Ordering Experiments

This section introduces a variety of ways of ordering experiments when agents make mistakes processing information. There is an ambiguity in defining the value of an experiment because different optimal decision rules may lead to different ex ante utilities. Note that this problem does not arise when agents interpret information correctly, because while there may be many optimal rules for a given decision problem, each will lead to the same ex ante utility. The different orderings introduced in this section constitute different ways of dealing with that

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

ambiguity. The results of this section are proved in the appendix, partly because it is natural to prove them in a very different order than they are presented.

There is a finite uncertainty space, Ω, with "known", strictly positive[62], prior probability distribution, π, over states. A "decision problem" consists a finite set of actions A, and a utility function, u: A x $\Omega \rightarrow \mathbb{R}$. Thus a decision problem δ = [u,A] and we write Δ for the set of all decision problems. There is also a signal space σ = [S, γ], where S is a finite set of signals and $\gamma(.|s)$ gives beliefs on Ω, conditional on observing signal $s \in S$. Finally, β is an experiment, where $\beta(.|\omega)$ is a probability distribution on S, conditional on ω. It is notationally convenient to think of σ as fixed i.e. that different experiments involve distributions over the same set of signals. This entails no loss of generality: for any pair of experiments, let S include the union of the sets of signals observed under the two experiments. In this way, all information about the experiment is captured in the distribution β. Crucially, it is <u>not</u> assumed that the experiment is "rational" in the following Bayesian sense:-

<u>Definition</u> Experiment β is rational if and only if:-

$$\pi(\omega)\ \beta(s|\omega) = \gamma(\omega|s) \sum_{\omega' \in \Omega} \pi(\omega')\ \beta(s|\omega') \qquad \forall\ s \in S,\ \omega \in \Omega$$

<u>1.1 Simple Orderings</u>

I proceed by analogy to the case with rational beliefs, where experiment β is preferred to experiment β' if ex ante utility under experiment β is higher than ex ante utility under

62. This assumption is without loss of generality; we could simply ignore zero probability events in what follows.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

experiment β', given that agents make optimal decisions conditional on the signals they observe.

Define the set of optimal decision rules for decision problem δ by:-

$$F(\delta) = \left\{ f{:}S \to A \mid \sum_{\omega\in\Omega} \gamma(\omega|s)\ u[f(s),\omega] \geq \sum_{\omega\in\Omega} \gamma(\omega|s)\ u[a,\omega] \quad \forall\ a\in A \right\}$$

Now a problem arises in defining a partial ordering for the value of experiments: different optimal decision rules will give different ex ante utilities. Thus define:-

$$v^*[\beta,\delta] = \max_{f\in F(\delta)} \left\{ \sum_{\omega\in\Omega} \sum_{s\in S} \pi(\omega)\ \beta(s|\omega)\ u[f(s),\omega] \right\}$$

$$v_*[\beta,\delta] = \min_{f\in F(\delta)} \left\{ \sum_{\omega\in\Omega} \sum_{s\in S} \pi(\omega)\ \beta(s|\omega)\ u[f(s),\omega] \right\}$$

For most β, there exist δ where v^* strictly exceeds v_*. To be precise, I have:-

<u>Result 1</u> $v^*[\beta,\delta] = v_*[\beta,\delta]$, $\forall\ \delta\in\Delta$, if and only if β is rational.

If beliefs are not rational, there are four natural ways of ranking experiments by value:-

<u>Definitions</u>

[1] β is "always preferred" to β' ($\beta \succeq_a \beta'$) $\Leftrightarrow v_*[\beta,\delta] \geq v^*[\beta',\delta]$, $\forall\ \delta\in\Delta$.

[2] β is "upper preferred" to β' ($\beta \succeq_u \beta'$) $\Leftrightarrow v^*[\beta,\delta] \geq v^*[\beta',\delta]$, $\forall\ \delta\in\Delta$.

[3] β is "lower preferred" to β' ($\beta \succeq_l \beta'$) $\Leftrightarrow v_*[\beta,\delta] \geq v_*[\beta',\delta]$, $\forall\ \delta\in\Delta$.

[4] β is "sometimes preferred" to β' ($\beta \succeq_s \beta'$) $\Leftrightarrow v^*[\beta,\delta] \geq v_*[\beta',\delta]$, $\forall\ \delta\in\Delta$.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Result 2 notes a few properties of the orderings that we may wish to take into account in deciding which to study:-

<u>Result 2</u> (i) $\beta \succeq_a \beta' \Rightarrow \beta \succeq_u \beta'$ and $\beta \succeq_l \beta'$.

(ii) $\beta \succeq_u \beta'$ or $\beta \succeq_l \beta' \Rightarrow \beta \succeq_s \beta'$.

(iii) $\succeq_a$, $\succeq_u$, $\succeq_l$, and $\succeq_s$ are all transitive[63].

(iv) $\succeq_u$, $\succeq_l$ and $\succeq_s$ are reflexive[64].

(v) $\beta \succeq_a \beta$ if and only if β is rational.

(vi) If β is rational, $\beta \succeq_a \beta' \Leftrightarrow \beta \succeq_u \beta'$ and $\beta \succeq_l \beta' \Leftrightarrow \beta \succeq_s \beta'$.

(vii) If β' is rational, $\beta \succeq_a \beta' \Leftrightarrow \beta \succeq_l \beta'$ and $\beta \succeq_u \beta' \Leftrightarrow \beta \succeq_s \beta'$.

$\succeq_a$ is the easiest to study mathematically, because we ask a single existence question (does there exist any decision problem and any pair of optimal decision rules for which ex ante utility under one experiment exceeds another?); for example, Geanakoplos (1989) uses $\succeq_a$ for possibility correspondences. But it has the bizarre property that an experiment is at least as good as itself if and only if the experiment is rational. The following two examples illuminate the differences between the orderings. Suppose $\Omega = \{\omega_1, \omega_2\}$, $\pi = (1/2, 1/2)$, $S = \{s_0, s_1, s_2\}$, and $\gamma(\omega|s) = 1/2$, $\forall\, s \in S$, $\omega \in \Omega$. Now consider β_0, a completely uninformative experiment: $\beta_0(s_0|\omega) = 1$, for $\omega = \omega_1, \omega_2$; and β_1, an experiment which in fact reveals the true state of the world (ω_1 or ω_2), but which is treated as uninformative: $\beta_1(s_1|\omega_1) = \beta_1(s_2|\omega_2) = 1$, $\beta_1(s_1|\omega_2) = \beta_1(s_2|\omega_1) = 0$.

63. $\succeq$ is transitive if $\beta \succeq \beta'$ and $\beta' \succeq \beta'' \Rightarrow \beta \succeq \beta''$.

64. $\succeq$ is reflexive if $\beta \succeq \beta$, for all β.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

It can be shown that:-

$\beta_0 \succeq_a \beta_0$	$\beta_0 \succeq_u \beta_0$	$\beta_0 \succeq_l \beta_0$	$\beta_0 \succeq_* \beta_0$
$\beta_1 \not\succeq_a \beta_1$	$\beta_1 \succeq_u \beta_1$	$\beta_1 \succeq_l \beta_1$	$\beta_1 \succeq_* \beta_1$
$\beta_0 \not\succeq_a \beta_1$	$\beta_0 \not\succeq_u \beta_1$	$\beta_0 \succeq_l \beta_1$	$\beta_0 \succeq_* \beta_1$
$\beta_1 \not\succeq_a \beta_0$	$\beta_1 \succeq_u \beta_0$	$\beta_1 \not\succeq_l \beta_0$	$\beta_1 \succeq_* \beta_0$

To see why, consider $\delta = (A,u)$, $A = (a_1, a_2)$, $u(a_1,\omega_1) = 1$, $u(a_1,\omega_2) = -1$ and $u(a_2,\omega) = 0$, for $\omega = \omega_1,\omega_2$. Then all decision rules are optimal (the expected payoff is zero, for all three signals, for any action). Therefore, $v^*[\beta_0,\delta] = v_*[\beta_0,\delta] = 0$, $v^*[\beta_1,\delta] = 1$, $v_*[\beta_1,\delta] = -1$.

The examples make clear what is in fact a more general property: ambiguity between the different orderings depends on particular decision problems where agents take different actions after observing signals with the same posterior beliefs. In section 1.2, I show the effect of requiring decision rules to take the same action at the same signal in the comparison of experiments. In section 1.3, I show the effect of requiring one experiment to be better only for "almost all" decision problems.

1.2 Joint orderings and signal labelling

The reason why $\beta_1 \not\succeq_a \beta_1$, in the example above, is that the definition allows the agent to take a different action conditional on observing the same signal. Suppose instead that the optimal decision rules required the same decision, conditional on the same signal (although note that, for now, there may be many signals, s, with the same beliefs, $\gamma(.|s)$).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Definition

[1] β is "always jointly preferred" to β' ($\beta \succsim_{aj} \beta'$) if

$$\sum_{\omega\in\Omega}\sum_{s\in S} \pi(\omega)\ \beta(s|\omega)\ u[f(s),\omega] \geq \sum_{\omega\in\Omega}\sum_{s\in S} \pi(\omega)\ \beta'(s|\omega)\ u[f(s),\omega] \quad \forall\ \delta\in\Delta, f\in F(\delta)$$

[2] β is "sometimes jointly preferred" to β' ($\beta \succsim_{sj} \beta'$) if

$$\sum_{\omega\in\Omega}\sum_{s\in S} \pi(\omega)\ \beta(s|\omega)\ u[f(s),\omega] \geq \sum_{\omega\in\Omega}\sum_{s\in S} \pi(\omega)\ \beta'(s|\omega)\ u[f(s),\omega] \quad \forall\ \delta\in\Delta, \textit{for some}.$$

Result 3 (i) $\beta \succsim_a \beta' \Rightarrow \beta \succsim_{aj} \beta'$

(ii) $\beta \succsim_{aj} \beta' \Rightarrow \beta \succsim_u \beta'$ and $\beta \succsim_l \beta'$

(iii) $\beta \succsim_{sj} \beta' \Leftrightarrow \beta \succsim_s \beta'$.

(iv) $\succsim_{aj}$ and $\succsim_{sj}$ are transitive and reflexive

The joint requirement doesn't change $\succsim_s$, but weakens $\succsim_a$ so that $\succsim_{aj}$ is reflexive. But we can go a step further, by requiring that the same action must be taken any time (in either experiment) posterior beliefs are the same. This perhaps excessive assumption also reduces ambiguity. The easiest way to model this constraint is to assume that S cannot contain more than one signal with the same beliefs.

Definition σ satisfies no irrelevant labelling if $\gamma(\omega|s) = \gamma(\omega|s')$, $\forall\omega\in\Omega \Rightarrow s=s'$.

Result 4 If σ satisfies no irrelevant labelling, $\beta \succsim_s \beta' \Leftrightarrow \beta \succsim_{sj} \beta' \Leftrightarrow \beta \succsim_u \beta' \Leftrightarrow \beta \succsim_l \beta' \Leftrightarrow \beta \succsim_{aj} \beta'$.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Thus if we are prepared to restrict attention to experiments on signal spaces where any two signals with the same beliefs are the same, then all orderings other than $\succeq_a$ are equivalent.

1.3 Generic preference

An alternative approach to the ambiguity in ordering experiments would be to consider when one experiment is preferred to another for "almost all" decision problems, since the ambiguity rests on apparently non-generic decision problems where there are multiple optimal decision rules. Below I give two different ways of excluding "special cases" where ambiguity arises: first, I simply exclude those decision problems where there is ambiguity about the ex ante value of the decision problem under either experiment; second, I require that for every decision problem, there is some other decision problem arbitrarily close where the preferred experiment always has higher ex ante utility.

Definition

[1] β is "well-defined preferred" to β' ($\beta \succeq_w \beta'$) if and only if $v_*[\beta,\delta] \geq v^*[\beta',\delta]$ $\forall\, \delta$, such that $v^*[\beta,\delta] = v_*[\beta,\delta]$ and $v^*[\beta',\delta] = v_*[\beta',\delta]$, i.e. for all δ such that ex ante utilities are under β and β' are well-defined.

[2] β is "generically preferred" to β' ($\beta \succeq_g \beta'$) if and only if $\forall\, \epsilon > 0$, $\forall\, \delta = (u,A) \in \Delta$, there exists $\delta' = (u',A) \in \Delta$ such that $|\, u(a,\omega) - u'(a,\omega) \,| < \epsilon$, $\forall\, a \in A$, $\omega \in \Omega$, and $v_*[\beta,\delta'] \geq v^*[\beta',\delta']$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Result 5 $\beta \succeq_g \beta' \Leftrightarrow \beta \succeq_w \beta' \Leftrightarrow \beta \succeq_s \beta'$

Note that it doesn't matter in the definitions $\succeq_w$ and $\succeq_g$ above if the "always preference" requirement is replaced by upper, lower or sometimes preference.

1.4 Summary

Results 1 through 5 give the relationships between the orderings. In the appendix, I give examples demonstrating that these results were exhaustive. For the rest of the chapter, I will assume no irrelevant labelling and say β is preferred to β' [$\beta \succeq \beta'$], without qualification or subscript if $\beta \succeq_{aj} \beta'$ (or, equivalently, $\beta \succeq_u \beta'$, $\beta \succeq_l \beta'$, $\beta \succeq_s \beta'$, $\beta \succeq_w \beta'$ or $\beta \succeq_g \beta'$). But it should be understood that the results generalise to equivalent problems with irrelevant labelling when $\succeq$ is understood to mean (equivalently) $\succeq_s$, $\succeq_w$ or $\succeq_g$.

2. The value of misinterpreted information

I will proceed straight to giving the main result and then consider more special cases.

Definitions

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

$$Let\ \Phi = \left\{ \phi: S \times \Omega \to \mathbb{R} \mid \sum_{\omega\in\Omega}\sum_{s\in S} \phi(s,\omega) = 1 \right\} \qquad D = \left\{ d: S \to \mathbb{R}_+ \mid \sum_{s\in S} d(s) = 1 \right\}$$

$$\Theta = \left\{ \theta: S^2 \to \mathbb{R}_+ \mid \sum_{s'\in S} \theta(s,s') = 1, \ \forall\, s\in S \right\}$$

Experiment β is <u>sufficient</u> for β' if and only if there exists a Markov process $\theta \in \Theta$ such that

$$\beta'(s|\omega) = \sum_{s'\in S} \theta(s',s)\ \beta(s'|\omega) \qquad \forall\ s\in S,\ \omega\in\Omega$$

Thus β' is a noisy version of experiment β.

Experiment β is <u>sufficient in differences</u> for β' if and only if there exist a non-negative constant $k \in \mathbb{R}_+$, a probability distribution on S, $d \in D$, a Markov process $\theta \in \Theta$ and a "pseudo-probability" distribution $\phi \in \Phi$, such that:-

$$\pi(\omega)\ \beta(s|\omega) = k\ d(s)\ \gamma(\omega|s) + (1-k)\ \phi(s,\omega)$$
$$and\quad \pi(\omega)\ \beta'(s|\omega) = k \sum_{s'\in S} d(s')\ \theta(s',s)\ \gamma(\omega|s') + (1-k)\ \phi(s,\omega), \qquad \forall\ s\in S,\ \omega\in\Omega$$

I have two interpretations of sufficiency in differences. First, β is sufficient in differences for β' if and only if there exist experiment β_*' and rational experiment β_* which is sufficient for β_*', such that the difference between β and β' is proportional to the difference between β_* and β_*'. To see why this is true, consider rational experiment β_*, where $\pi(\omega)$ $\beta_*(s|\omega) = d(s)\ \gamma(\omega|s)\ \forall\ s\in S,\ \omega\in\Omega$, and experiment β_*', where $\pi(\omega)\ \beta_*'(s|\omega) = \Sigma_{s'\in S}\ d(s')$ $\theta(s',s)\ \gamma(\omega|s')$. Now the difference between β and β' is proportional (with factor k) to the difference between β_* and β_*', and β_* is sufficient for β_*'.

Secondly, consider the case where $k \leq 1$ and ϕ is non-negative. Now the condition says that the true distribution of signals β is a convex combination of that generated by a rational experiment (with d(s)) and an "irrational" distribution ϕ. β' is the convex combination

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

with the same weights of a noisy version of the rational experiment and the same "irrational" experiment. But we must allow also for k > 0 and sometimes negative ϕ.

Theorem 1 β is preferred to β' if and only if β is sufficient in differences for β'.

Claim 1 $\beta \succeq \beta' \Leftrightarrow \nexists$ x: S x $\Omega \rightarrow \mathbf{R}$ such that:-

$$\sum_{\omega\in\Omega} \gamma(\omega|s)\ x(s,\omega) \geq \sum_{\omega\in\Omega} \gamma(\omega|s)\ x(s',\omega) \qquad \forall\ s,s'\in S$$

$$\sum_{\omega\in\Omega}\sum_{s\in S} \pi(\omega)\ \beta(s|\omega)\ x(s,\omega) < \sum_{\omega\in\Omega}\sum_{s\in S} \pi(\omega)\ \beta'(s|\omega)\ x(s,\omega)$$

Proof of claim 1 If there exists such an x, consider δ = (S,x), i.e. the decision problem where the set of actions equals the set of signals and utilities are given by x. But then f(s) = s is an optimal decision rule under which the ex ante utility of β' is higher than β. On the other hand, suppose there exists δ and optimal decision rule f showing that $\beta \succeq \beta'$ is violated. Then let x(s,ω) = u[f(s),ω] and x must satisfy conditions above.

Claim 2 $\beta \succeq \beta' \Leftrightarrow$ there exists λ: $S^2 \rightarrow \mathbf{R}_+$ such that:-

$$\pi(\omega)\,\beta(s|\omega) - \pi(\omega)\,\beta'(s|\omega) = \sum_{s'\in S} \{\lambda(s,s')\gamma(\omega|s) - \lambda(s',s)\gamma(\omega|s')\} \qquad \forall\ s\in S,\ \omega\in\Omega$$

Proof of claim 2 By Farkas' Lemma from claim 1.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Proof of theorem 1 Now if there exist λ satisfying claim 2, then any (k,d,θ,ϕ) satisfying

$$k = \sum_{s\in S}\sum_{s'\in S}\lambda(s,s') \qquad kd(s) = \sum_{s'\in S}\lambda(s,s') \qquad kd(s)\theta(s,s') = \lambda(s,s')$$

$$(1-k)\phi(s,\omega) = \pi(\omega)\beta(s|\omega) - kd(s)\gamma(\omega|s) \qquad \forall\ \omega\in\Omega, s\in S$$

satisfy conditions for sufficiency in differences. Conversely, if there exist (k,d,θ,ϕ) satisfying conditions for sufficiency in differences, then $\lambda(s,s') = k\ d(s)\ \theta(s,s')$ will satisfy conditions of claim 2.

Theorem 1 is applied to certain special cases below - where one or both experiments are rational (corollaries 1.1 - 1.3) and where β' is uninformative (corollary 1.4). Proofs are in the appendix.

Corollary 1.1 (Blackwell's Theorem, 1951). If β and β' are rational, β is preferred to β' if and only if β is sufficient for β'.

Blackwell's theorem is proved from two corollaries which also strengthen the result:-

Corollary 1.2 If β is rational and β is sufficient for β', then $\beta \geq \beta'$.

Corollary 1.3 If β' is rational and $\beta \geq \beta'$, then β is sufficient for β'.

Definitions For any signal space $\sigma = [S, \gamma]$ with no irrelevant labelling, write s_0 for the (unique) signal with $\gamma(\omega|s_0) = \pi(\omega)$, $\forall\ \omega\in\Omega$. Now the zero (or uninformative) experiment β_0 satisfies $\beta_0(s_0|\omega) = 1$, $\forall\ \omega\in\Omega$.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Definition Experiment β is reconcilable if there exist $k \in (0,1]$, $c \in \mathbb{R}_+$, $\alpha, d, \mu \in D$ and $\theta \in \Theta$ such that

$$\pi(\omega)\{k\beta(s|\omega) + (1-k)\alpha(s)\} = \mu(s)\gamma(\omega|s) + c\left\{d(s)\gamma(\omega|s) - \sum_{s' \in S} d(s')\theta(s',s)\gamma(\omega|s')\right\}$$

Notice that beliefs are an average of the true distribution of signals β, the zero signal and a noisy version of beliefs. This property was discussed in chapter II and is closely related to Shin's (1989) "dynamic representation".

Corollary 1.4 Experiment β is preferred to the zero experiment (β_0) if and only if β is reconcilable.

3. Possibility correspondences

A possibility correspondence is a mapping $P: \Omega \rightarrow 2^{\Omega}/\{\varnothing\}$ with the interpretation that Ω is a (finite) set of states of the world and $P(\omega)$ is the set of states that are thought possible when the true state is ω. Recent work[65] has examined agents whose information can be represented by such possibility correspondences, but who do not necessarily satisfy the full rationality entailed by partitions, i.e. the requirement that $P(\omega) = \{\omega' \in \Omega \mid P(\omega') = P(\omega)\}$. Observing such a possibility correspondence is a kind of experiment:-

Definition Experiment β with signal space $\sigma = [S,\gamma]$ is decision theoretically equivalent to possibility correspondence P if and only if

65. See references in footnote 61.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

$$[1]\ S = 2^{\Omega}/\{\emptyset\} \qquad [2]\ \gamma(\omega|s) = \begin{cases} \dfrac{\pi(\omega)}{\sum_{\omega' \in s} \pi(\omega')}, & \text{if } \omega \in s \\ 0, & \text{otherwise} \end{cases} \qquad [3]\ \beta(s|\omega) = \begin{cases} 1, & s = P(\omega) \\ 0, & \text{otherwise} \end{cases}$$

Notice that there is no irrelevant labelling by construction. I will use results of the previous section to derive and extend Geanakoplos' informativeness result for such possibility correspondences. First, I note the appropriate analogues to the ordering of section 2:

A decision rule f: S → A is <u>optimal</u> if $\Sigma_{\omega \in s}\ \pi(\omega)\ u[f(s),\omega] \geq \Sigma_{\omega \in s}\ \pi(\omega)\ u[a,\omega]$, ∀ s∈S, a∈A.

<u>P ≥ P'</u> if for every decision problem $\delta \in \Delta$, and every optimal decision rule $f \in F[\delta]$,

$$\sum_{\omega \in \Omega} \pi(\omega) u[f(P(\omega)), \omega] \geq \sum_{\omega \in \Omega} \pi(\omega) u[f(P'(\omega)), \omega]$$

The main result of this section (theorem 2) will be given in terms of logical, not linear algebraic, properties. I will use the following:-

<u>Definition</u> The disagreement set of two PC's is the set of states where beliefs differ, i.e. $L(P,P') = \{ \omega \in \Omega \mid P(\omega) \neq P'(\omega) \}$

Clearly, only behaviour on the disagreement set will matter in comparing PCs.

<u>Definition</u> P is "non-deluded" [with respect to E] if $\omega \in P(\omega)\ \forall\ \omega \in \Omega$ [$\forall\ \omega \in E$]

Non-delusion requires that the true state is always thought possible.

<u>Definition</u> P satisfies "knowing that you know" (KTYK) [with respect to E] if $\omega' \in P(\omega) \Rightarrow P(\omega') \subseteq P(\omega)\ \forall\ \omega \in \Omega$ [$\forall\ \omega \in E$]

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

"Knowing that you know" requires that if you know that something is true, then you know that you know. Shin (1987) showed that non-delusion and KTYK are the necessary and sufficient conditions for knowledge to represent provability from a set of axioms.

Definition P satisfies "knowing that you don't know" (KTYDK) [with respect to E] if $\omega' \in P(\omega) \Rightarrow P(\omega') \supseteq P(\omega)\ \forall\ \omega \in \Omega\ [\forall\ \omega \in E]$

"Knowing that you don't know" requires that if you don't know that something is true, then you know that you don't know.

Definition P is "nested" [with respect to E] if $P(\omega) \cap P(\omega') = \emptyset$, $P(\omega)$ or $P(\omega')$, $\forall\ \omega,\omega' \in \Omega$ $[\forall\ \omega,\omega' \in E]$

Nestedness requires that at any two states of the world, either you know strictly more in one state than the other, or else information sets at the two states are mutually exclusive. Geanakoplos (1989) showed if a PC is non-deluded and nested, it is as if the agent has had a list of propositions revealed to him, but he only remembers up a certain point in the list.

These properties can be seen as absolute requirements on the rationality of a single possibility correspondence (see Geanakoplos (1988) for a discussion). Three properties for the comparison of possibility correspondences will be important. First, the partition concept of refinement generalises to possibility correspondences:-

Definition P refines P' if and only if $P(\omega) \subseteq P'(\omega)$, $\forall\ \omega \in \Omega$.

Sufficiency was a key property is the previous section. This translates in the PC case to:-

Definition P is sufficient for P' if $P(\omega') = P(\omega) \Rightarrow P'(\omega') = P'(\omega)$, $\forall\ \omega,\omega' \in \Omega$.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

I will require a third distinct comparitive property:-

Definition P reduces P' [with respect to E] if $\omega' \in P(\omega) \Rightarrow P'(\omega') = P(\omega)$, $\forall\, \omega \in \Omega$ [$\forall\, \omega \in E$].

To see the relation between these properties, consider, for any possibility correspondence P, its rational partition counterpart, P_*, defined by $P_*(\omega) = \{\, \omega' \in \Omega \mid P(\omega') = P(\omega) \,\}$. Then the following properties all follow from the various definitions above:-

Result (i) P refines P_* if and only if P satisfies KTYK and KTYDK

(ii) P_* refines P if and only if P satisfies non-delusion

(iii) P is sufficient for P' if and only if P_* refines P_*'

(iv) P reduces P' if and only if P refines P_*'

(v) P refines P' implies P reduces P' if and only if P' satisfies KTYK and KTYDK

(vi) P reduces P' implies P refines P' if and only if P' is non-deluded

(vii) P reduces P' implies P sufficient for P' if and only if P is non-deluded

(viii) P sufficient for P' implies P reduces P' if and only if P satisfies KTYK and KTYDK

Now I can state the main theorem of this section, which is somewhat stronger than theorem 1 of Geanakoplos (1989)[66].

66. Geanakoplos proved that if P' is a partition and P refines P' then $P \geq P'$ if and only P satisfies non-delusion, KTYK and nestedness. His proof is a direct argument by induction, rather than using the dual method of theorem 1. It can easily be extended to give theorem 2.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

<u>Theorem 2</u> If P satisfies non-delusion, KTYK, nestedness and reduction of P', all with respect to L(P,P'), then $P \geq P'$. These conditions are necessary for $P \geq P'$ if P refines P' or if P' is non-deluded.

The proof is given in the appendix. It consists of showing that the linear algebraic properties of theorem 1 are equivalent in the PC case to the logical properties of theorem 2. It is related to a proof in Shin (1989). I note also the analogues to the corollaries of the previous section:-

<u>Corollary 2.1</u> If P and P' are partitions, $P \geq P'$ if and only if P is sufficient for (equivalently, refines or reduces) P'.

<u>Corollary 2.2</u> If P is a partition and P is sufficient for (equivalently, reduces) P', then $P \geq P'$.

<u>Corollary 2.3</u> If P' is a partition and $P \geq P'$, then P is sufficient for P'.

The analogue to the empty or uninformative experiment is the uninformative possibility correspondence, $P_0(\omega) = \Omega, \forall\, \omega \in \Omega$.

<u>Corollary 2.4</u> $P \geq P_0$ if and only if P satisfies non-delusion, KTYK and nestedness.

Theorem 2 does not give complete necessary and sufficient conditions for one PC to be preferred to another. Of course, theorem 1 gives such necessary and sufficient conditions and covers the PC case, but can the linear algebraic conditions of theorem 1 always be translated into the kind of logical properties of theorem 2? The rest of this section explores that question, and gives (inconclusive) arguments why it might not be possible.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

The first question must be whether the sufficient conditions of the theorem are also necessary, even when P' is deluded and P does not refine P'. The following example shows that this is not the case.

	$P(a) = \{a,b\}$	$P'(a) = \{b\}$
$\Omega = \{a,b,c\}$	$P(b) = \{a,b\}$	$P'(b) = \{c\}$
	$P(c) = \{b,c\}$	$P'(c) = \{a\}$

$P \succeq P'$, even though P does not satisfy KTYK, nestedness or reduction of P' w.r.t. L(P,P'). Note that P' is deluded and P does not refine P', so the example does not contradict theorem 2. I give an argument showing why $P \succeq P'$, which is rather suggestive; consider the following PC's:

$P(a) = \{a,b\}$	$R_1(a) = \{b\}$	$R_2(a) = \{b\}$	$R_3(a) = \{b\}$	$P'(a) = \{b\}$
$P(b) = \{a,b\}$	$R_1(b) = \{b\}$	$R_2(b) = \{b,c\}$	$R_3(b) = \{c\}$	$P'(b) = \{c\}$
$P(c) = \{b,c\}$	$R_1(c) = \{b,c\}$	$R_2(c) = \{b,c\}$	$R_3(c) = \{c\}$	$P'(c) = \{a\}$

Corollary 2.3 below shows that $P \succeq R_1 \succeq R_2 \succeq R_3 \succeq P'$. Here is a sequence of simple alterations, each of which makes ex ante utility worse under optimal decision rules, which explain, by transitivity, why $P \succeq P'$.

More generally, say that P is a "simple alteration" of P' ($P \leftrightarrow P'$) if and only if $P(\omega) = P(\omega')$ and $P'(\omega) = P'(\omega')$, $\forall\, \omega,\omega' \in L(P,P')$. A simple alteration only makes one qualitative change in the PC. In the above example, $P \leftrightarrow R_1 \leftrightarrow R_2 \leftrightarrow R_3 \leftrightarrow P'$. When is it the case that a simple alteration is improving (as in that sequence)?

<u>Corollary 2.5</u> If $P \leftrightarrow P'$, then $P \succeq P'$ if and only if <u>either</u> $P(\omega) = L(P,P')\ \forall\, \omega \in L(P,P')$ <u>or</u> $P(\omega) = P'(\omega) \cup L(P,P')$ and $P'(\omega) \cap L(P,P') = \phi$, $\forall\, \omega \in L(P,P')$

In words, a simple alteration is improving only if the disagreement set L(P,P') is replaced in the simple alteration with either the true information set [L(P,P')] or, if beliefs on

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

L(P,P') were always deluded in P' [i.e. $P'(\omega) \cap L(P,P') = \varnothing$], with the union of the old information set and the true information set. Will there always be a sequence of simple alterations, as in the example above, which explain why one PC is preferred to another?

Corollary 2.6[67] If P' is non-deluded, or P refines P', and $P \succeq P'$, then there exists a sequence of PCs, $(R_1,.. ,R_n)$, such that $P \succeq R_1 \succeq .. \succeq R_n \succeq P'$ and $P \leftrightarrow R_1 \leftrightarrow .. \leftrightarrow R_n \leftrightarrow P'$.

Unfortunately, this argument also does not extend to general P', as the following example shows:-

	$P(a) = \{a,c\}$	$P'(a) = \{b\}$
$\Omega = \{a,b,c\}$	$P(b) = \{a,b\}$	$P'(b) = \{c\}$
	$P(c) = \{b,c\}$	$P'(c) = \{a\}$

$P \succeq P'$ but there does not exist $(R_1,.. ,R_n)$ such that $P \succeq R_1 \succeq .. \succeq R_n \succeq P'$ and $P \leftrightarrow R_1 \leftrightarrow .. \leftrightarrow R_n \leftrightarrow P'$[68].

67. This is closely related to Geanakoplos' (1989) "memory property" result.

68. For any $\delta \in \Delta$, and any optimal decision rule f,

$$u[f(\{a,c\}), a] + u[f(\{a,c\}),c] \geq u[f(\{a\}), a] + u[f(\{a\}), c]$$
$$u[f(\{a,b\}), a] + u[f(\{a,b\}),b] \geq u[f(\{b\}), a] + u[f(\{b\}), b]$$
$$u[f(\{b,c\}), b] + u[f(\{b,c\}),c] \geq u[f(\{c\}), b] + u[f(\{c\}), c]$$
$$u[f(\{a\}), a] \geq u[f(\{a,b\}), a]$$
$$u[f(\{b\}), b] \geq u[f(\{b,c\}), b]$$
$$u[f(\{c\}), c] \geq u[f(\{a,c\}), c]$$

Summing the six inequalities gives:-

$$u[f(\{a,c\}), a] + u[f(\{a,b\}),b] + u[f(\{b,c\}),c] \geq u[f(\{b\}),a] + u[f(\{c\}),b] + u[f(\{a\}),c]$$

All sequences of improving simple alterations from P' can be constructed to show there does not exist one that reaches P.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Thus an obvious way of translating the linear algebraic properties of theorem 1 to logical properties, for all possibility correspondences, fails[69]. Part 3 of the appendix gives some additional results which extend theorem 2, but are still not completely general.

4. Conclusion

In this chapter, I have identified necessary and sufficient conditions on experiments to rank even misinterpreted information, and related it to earlier results for possibility correspondences. The value of information is a natural starting point for examining agents who make mistakes processing information. But it is also a question with many important economic applications, e.g. Holmstrom (1979). Geanakoplos (1989) showed that "speculation" can occur only if misinterpreted information is not more valuable than no information. The relation to this work is discussed in the next chapter. In any environment where there is one-sided asymmetric information, the higher ex ante value of "more" information is usually being implicitly used.

---

69. However, I do not know of examples where $P \geq P'$ and the multipliers of the necessary and sufficient linear algebraic conditions (see property A in proof of theorem 2 in appendix) are not zeros and ones. This suggests that a translation may exist after all.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Appendix to Chapter III

1: proofs

section 2

Result 1. If beliefs are rational, $\Sigma_{\omega\in\Omega} \pi(\omega)\, \beta(s|\omega)\, u[f(s),\omega] = k(s)$, for some k, all $\delta\in\Delta$, $f\in F[\delta]$, therefore $v^*[\beta,\delta] = v_*[\beta,\delta]$. If beliefs are not rational, then there exist s^*,ω^*, such that

$$\gamma(\omega^*|s^*) \sum_{\omega'\in\Omega} \pi(\omega')\, \beta(s^*|\omega') > \pi(\omega^*)\, \beta(s^*|\omega^*)$$

Now let $\delta=(A,u)$, $A = (a_1,a_2)$, $u(a_1,\omega^*) = 1 - \gamma(\omega^*,s^*)$, $u(a_1,\omega) = -\gamma(\omega^*,s^*)\ \forall\omega\neq\omega^*$, $u(a_2,\omega) = 0\ \forall\omega\in\Omega$. Observe that either a_1 or a_2 are optimal given s^*. For any $f\in F[\delta]$, consider $f_1,f_2\in F(\delta)$, where $f_1(s^*) = a_1$ and $f_1(s) = f(s)\ \forall s\neq s^*$, and $f_2(s^*) = a_2$ and $f_2(s) = f(s)\ \forall s\neq s^*$. Now

$$v^*[\beta,\delta] \geq \sum_{\omega\in\Omega}\sum_{s\in S} \pi(\omega)\,\beta(s|\omega)\,u[f_2(s),\omega] > \sum_{\omega\in\Omega}\sum_{s\in S} \pi(\omega)\,\beta(s|\omega)\,u[f_1(s),\omega] \geq v_*[\beta,\delta]$$

Result 2. (i), (ii), and the transitivity and reflexivity of $\succeq_a$, $\succeq_u$ and $\succeq_l$ in (iii) in (iv) follow directly from the definitions of the orderings and $v^*[\beta,\delta] \geq v_*[\beta,\delta]$, $\forall\ \delta\in\Delta$. The transitivity of $\succeq_s$ in (iii) is implied by result 4 below, since $\succeq_{aj}$ is clearly transitive. (v), (vi) and (vii) are consequences of result 1.

Result 3. (i), (iv) and $\beta \succeq_{aj} \beta' \Rightarrow \beta \succeq_s \beta'$ in (iii) follow from definitions. $\beta \succeq_s \beta' \Rightarrow \beta \succeq_{aj} \beta'$ follows from result 5: $\beta \succeq_w \beta' \Rightarrow \beta \succeq_{aj} \beta'$ by the same argument as $\beta \succeq_w \beta' \Rightarrow \beta \succeq_s \beta'$. A proof that $\beta \succeq_{aj} \beta' \Rightarrow \beta \succeq_u \beta'$ follows. An exactly symmetric proof shows that $\beta \succeq_{aj} \beta' \Rightarrow \beta \succeq_l \beta'$ and thus (ii).

Proof that $\beta \succeq_{aj} \beta' \Rightarrow \beta \succeq_u \beta'$

$\beta \not\succeq_u \beta' \Rightarrow$ there exists $\delta\in\Delta$, $f\in F(\delta)$, such that

$$\sum_{\omega\in\Omega}\sum_{s\in S} \pi(\omega)\,\beta'(s|\omega)\,u[f(s),\omega] - v^*[\beta',\delta] > v^*[\beta,\delta] \geq \sum_{\omega\in\Omega}\sum_{s\in S} \pi(\omega)\,\beta(s|\omega)\,u[f(s),\omega]$$

So $\beta \not\succeq_{aj} \beta'$.

Result 4. Note that given results 2 and 3, it is sufficient to show that, with no irrelevent labelling, $\beta \succeq_s \beta' \Rightarrow \beta \succeq_{aj} \beta'$. The only proof I have uses the dual characterization of the necessary and sufficient conditions for each result and shows that, under "no irrelevant labelling", they are equivalent.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Suppose $\beta \succeq_s \beta'$. Then there does not exist x: $S \times \Omega \to \mathbb{R}$ such that:-

$$\sum_{\omega\in\Omega} \gamma(\omega|s)\, x(s,\omega) > \sum_{\omega\in\Omega} \gamma(\omega|s)\, x(s',\omega) \qquad \forall\, s,s'\in S$$

$$\sum_{\omega\in\Omega}\sum_{s\in S} \pi(\omega)\, \beta(s|\omega)\, x(s,\omega) < \sum_{\omega\in\Omega}\sum_{s\in S} \pi(\omega)\, \beta'(s|\omega)\, x(s,\omega)$$

(note the strict inequality in the optimality constraints).
Therefore, by Farkas' Lemma, there exist k ≥ 0 and $\lambda: S \to \mathbb{R}_+$, not all zero, such that:-

$$k\pi(\omega)\{\beta(s|\omega) - \beta'(s|\omega)\} = \sum_{s'\in S} \{\lambda(s,s')\gamma(\omega|s) - \lambda(s',s)\gamma(\omega|s')\} \qquad \forall\, s\in S,\ \omega\in\Omega$$

Now suppose $\beta \not\succeq_{aj} \beta'$. Then there do not exist (k,λ) satisfying the above conditions with k > 0 (by theorem 1). So there exist $\lambda: S^2 \to \mathbb{R}_+$, not all zero, such that

$$\sum_{s'\in S} \{\lambda(s,s')\gamma(\omega|s) - \lambda(s',s)\gamma(\omega|s') = 0\} \qquad \forall\, s\in S,\ \omega\in\Omega$$

This requires $\gamma(\omega|s) = \gamma(\omega|s')$ $\forall \omega\in\Omega$, for some $s\neq s'$, $s,s'\in S$. But this contradicts no irrelevant labelling.

Result 5. $\beta \succeq_s \beta' \Rightarrow \beta \succeq_w \beta'$ follows from definitions. Therefore it is sufficient to show $\beta \succeq_w \beta' \Rightarrow \beta \succeq_g \beta'$ and $\beta \succeq_g \beta' \Rightarrow \beta \succeq_s \beta'$.

first, let $m(\delta)$ be the minimum change in expected utility required to make a non-optimal decision rule optimal:-

$$m(\delta) = \min_{s\in S, f\in F(\delta), g\notin F(\delta)} \left\{\sum_{\omega\in\Omega} \gamma(\omega|s)\, u[f(s),\omega] - \sum_{\omega\in\Omega} \gamma(\omega|s)\, u[g(s),\omega]\right\}$$

proof that $\beta \succeq_w \beta' \Rightarrow \beta \succeq_g \beta'$. Suppose $\beta \succeq_w \beta'$. For any $\delta=(A,u)$, let N = #A, and label the actions 1 to N and let n(a) be the label of action a. For any $\epsilon>0$, let $\delta'=(A,u')$ with $u'(a,\omega) = u(a,\omega) - (1/2)\,[n(a)/N] \min(\epsilon,m)$; notice that $F(\delta)$ is a singleton, and therefore has a well-defined ex ante utility; therefore $v_*[\beta,\delta'] \geq v^*[\beta,\delta']$.

proof that $\beta \succeq_g \beta' \Rightarrow \beta \succeq_s \beta'$. Suppose $\beta \not\succeq_s \beta'$. Then there exist $\delta\in\Delta$ such that $v_*[\beta',\delta] - v^*[\beta,\delta] = c > 0$. Consider $0 < \epsilon < (1/2) \min(m(\delta),c)$. Now, for any $\delta' = [A,u']$ such that $|\, u(a,\omega) - u'(a,\omega)\, | < \epsilon \;\forall\, a\in A, \omega\in\Omega$, $v^*[\beta',\delta'] \geq v_*[\beta',\delta] - \epsilon$ and $v_*[\beta,\delta'] \leq v^*[\beta,\delta] + \epsilon$. So:-
$v^*[\beta',\delta'] \geq v_*[\beta',\delta] - \epsilon = v^*[\beta,\delta] + c - \epsilon \geq v_*[\beta,\delta'] + c - 2\epsilon > v_*[\beta,\delta']$, and thus $\beta \not\succeq_g \beta'$.

Section 3

Corollary 1.2 In definition of sufficiency in differences, set k=1, θ equal to the Markov process showing β sufficient for β', and $d(s) = \Sigma_{\omega\in\Omega}\, \pi(\omega)\, \beta(s|\omega)$.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

<u>Corollary 1.3</u> let $d(s) = \Sigma_{\omega\in\Omega}\, \pi(\omega)\, \beta(s|\omega)$, $d'(s) = \Sigma_{\omega\in\Omega}\, \pi(\omega)\, \beta'(s|\omega)$; let $S^* = \{\, s\in S \mid d'(s) > 0 \,\}$, $S_0 = \{\, s\in S \mid d(s) > 0,\ s\notin S^*$; by rationality of β', $\pi(\omega)\, \beta'(s|\omega) = d'(s)\, \gamma(\omega|s)$; now $\beta \geq \beta'$ and claim 2 $\Rightarrow$ there exist λ: $S^2 \rightarrow \mathbf{R}_+$ such that:-

$$\pi(\omega)\,\beta'(s|\omega)\left\{1 + \frac{\sum_{s'\in S}\lambda(s,s')}{d'(s)}\right\} = \pi(\omega)\,\beta(s|\omega) + \sum_{s'\in S_0}\lambda(s',s)\gamma(\omega|s) + \sum_{s'\in S^*}\left\{\frac{\lambda(s',s)\,\pi(\omega)\,\beta'(s|\omega)}{d'(s)}\right\} \quad \forall\ s\in S^*, \omega\in\Omega$$

$$\sum_{s'\in S}\lambda(s,s')\,\gamma(\omega|s) = \pi(\omega)\,\beta(s|\omega) + \sum_{s'\in S^*}\left\{\frac{\lambda(s',s)\,\pi(\omega)\,\beta'(s|\omega)}{d'(s)}\right\} + \sum_{s'\in S_0}\lambda(s',s)\,\gamma(\omega|s) \qquad \omega\in\Omega, s\in$$

Now an expression for $\gamma(.|s)$, $s\in S_0$, as a linear combination of β and β' can be derived by progressive substitution in [2]. This can be substituted in [1], and then progressive substitutions give β' as a linear combination of β.

<u>Corollary 1.4</u> In claim 2, set

$$k = \frac{1}{1 + \sum_{s\neq s_0}\lambda(s_0,s)} \qquad c = \sum_{s\neq s_0}\sum_{s'\neq s_0}\lambda(s,s') \qquad (1-k)\,\alpha(s) = k\lambda(s_0,s)$$

$$\mu(s) = \left\{\frac{\lambda(s,s_0)}{\sum_{s\neq s_0}\lambda(s,s_0)}\right\} \qquad cd(s) = \sum_{s'\neq s_0}\lambda(s,s') \qquad d(s)\,\theta(s,s') = \lambda(s,s')$$

<u>Section 4</u>

<u>Theorem 2</u> Let $S^* = \{\, s\in S \mid s = P(\omega)$ for some $\omega\in L(P,P') \,\}$; $\zeta_s(\omega) = 1$ if $P(\omega) = s$, 0 otherwise; $\zeta_s'(\omega) = 1$ if $P'(\omega) = s$, 0 otherwise; $\chi_s(\omega) = 1$ if $\omega\in s$, 0 otherwise.

<u>Step 1.</u> Condition A is necessary and sufficient for $P \geq P'$

<u>Condition A</u> there exists λ: $S^2 \rightarrow \mathbf{R}_+$, such that $\zeta_s(\omega) - \zeta_s'(\omega) = \Sigma_{s'\in S}\, \{\lambda(s,s')\, \chi_s(\omega) - \lambda(s',s)\, \chi_{s'}(\omega)\}$ $\forall\ s\in S, \omega\in\Omega$.

This can be seen by substituting the conditions defining a PC into theorem 1, and incorporating within the λ the denominator of the PC expressions for γ.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Step 2. If $P \geq P'$, then P is non-deluded with respect L(P,P').

$\zeta_s(\omega) \geq \zeta_s'(\omega) + \Sigma_{s' \in S} \{\lambda(s,s')\ \chi_{s'}(\omega)\}$ (from A)

therefore if $s = P(\omega)$ and $\omega \in L(P,P')$ then $\omega \in s$.

Step 3. So if P refines P', and $P \geq P'$, then P' is non-deluded w.r.t. L(P,P')

Step 4. Suppose P is non-deluded, KTYK, nested and a reduction of P', w.r.t. L(P,P'), but P' is not non-deluded w.r.t. L(P,P'). Then consider R with $R(\omega) = P'(\omega)$, $\forall\ \omega \notin L(P,P')$ and $R(\omega) = \{\ \omega' \in L(P,P')\ |\ P'(\omega') = P(\omega)\ \}$, $\forall\ \omega \in L(P,P')$. Now, $R \geq P'$, R is non-deluded w.r.t. P, and P is non-deluded, KTYK, nested and a reduction R w.r.t. R.

Steps 1 to 4 imply that it is sufficient for the theorem to show that condition C below is equivalent to condition A, when P' is non-deluded w.r.t. L(P,P'). I show this by showing that conditions A, B and C are equivalent.

Condition B there exist T, $I \subset S^*$, one-to-one "precedence" mapping α: $S^*/I \rightarrow S^*/T$, $\alpha(s) \neq s\ \forall\ s \in S^*/I$, and a partition Q of L(P,P') such that:-

$$I = \{\ s \in S\ |\ s = Q(\omega), \text{ for some } \omega \in L(P,P')\ \} \qquad [\text{B.1}]$$

$$Q(\omega) \subset \{\ \omega' \in L(P,P')\ |\ P'(\omega') = P'(\omega)\ \} \quad \forall\ \omega \in L(P,P') \qquad [\text{B.2}]$$

$$\zeta_s(\omega) = \chi_s(\omega) - \Sigma_{s' \in f(s)}\ \chi_{s'}(\omega) \quad \forall\ s \in S^*, \omega \in L(P,P') \qquad [\text{B.3}]$$

where $f(s) = \{\ s' \in S^*\ |\ \alpha(s') = s\}$

Condition C P is non-deluded, KTYK, nested and a reduction of P', w.r.t. L(P,P')

$A \Rightarrow B$

L1. $\lambda(s,s') > 0 \Rightarrow s \subset s'$[70]

P and P' non-deluded, and multiplying [A] by $[1 - \chi_s(\omega)] \Rightarrow$

$$\Sigma_{s' \in S}\ \lambda(s',s)\ \chi_{s'}(\omega) = \Sigma_{s' \in S}\ \lambda(s',s)\ \chi_{s'}(\omega)\ \chi_s(\omega) \Rightarrow [\ \omega \in s, \omega \notin s' \Rightarrow \lambda(s',s) = 0\ ]$$

Note that w.l.o.g., we can assume $\lambda(s,s) = 0$, $\forall\ s \in S$, in A.

L2. $\lambda(s,s') = 0\ \forall\ s'$ such that $\omega \in s'$, $s' \not\subset P'(\omega)$

Proof by induction on the number of elements of s: L1 implies L2 true for $s = \Omega$. Suppose true for $\#s = n$. Then if $\#s = n-1$:-

$$\zeta_s(\omega) - \zeta_s'(\omega) = -\ \Sigma_{s' \neq s}\ \lambda(s',s)\ \chi_{s'}(\omega)$$

If $s \neq P'(\omega)$, then $\lambda(s',s) = 0\ \ \forall\ s'$.t. $\omega \in s'$, so L2 true for $\#s = n-1$.

70. Where $s \subset s'$ means strict inclusion i.e. $s \subseteq s'$ and $s \neq s'$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

L3. $P(\omega) \subset P'(\omega)\ \forall\ \omega \in L(P,P')$

$\zeta_s(\omega) \leq \Sigma_{s' \in S}\ \lambda(s,s')\ \chi_s(\omega) = 0,\ \sim \omega \in L(P,P'),\ \not\subset\ \subset P'(\omega)$ by L2

Let $F(s) = \{ s' \in S^* \mid s' \subset s\}$; let $f(s) = \{ s' \in S^* \mid s' \in F(s)$, and there does not exist s" such that $s" \in F(s),\ s' \in F(s")\}$.

L4. $\lambda(s,s') = 1$ if and only if $s \in f(s'),\ \forall\ s,s' \in S^*$
L5. $\Sigma_{s' \neq s}\ \lambda(s,s') = 1,\ \forall\ s \in S^*$

Note first that $\Sigma_{s' \neq s}\ \lambda(s,s') > 0,\ \forall\ s \in S^*$, from A

$$By\ L2,\ \sum_{s' \neq s} \lambda(s,s')\chi_s(\omega) - \zeta_s(\omega) + \sum_{s' \in F(s)} \lambda(s',s)\chi_{s'}(\omega) \qquad \forall\ s \in S^*$$

Now by progressive substitutions:-

$$\sum_{s' \neq s} \lambda(s,s')\chi_s(\omega) - \zeta_s(\omega) + \sum_{s' \in f(s)} \left\{ \frac{\lambda(s',s)\,\zeta_{s'}(\omega)}{\sum_{s'' \neq s} \lambda(s',s'')} \right\} + \sum_{s' \in F(s) \setminus f(s)} \nu(s',s)\,\zeta_{s'}(\omega) \qquad \forall\ s \in S^*$$

where ν is some positive function of the λ.

Therefore $\Sigma_{s' \neq s}\ \lambda(s,s') = 1,\ \forall\ s \in S^*$, and $\lambda(s,s') = 1$ if $s \in f(s)$.

But $\lambda(s,s') > 0$ for some $s \in F(s')$, $s \not\ni f(s')$ implies $\nu(s',s) > 1$, which is impossible.

So $\alpha^*(s) = \{ s' \in S^* \mid s \in f(s') \}$ has at most one element. Let $I = \{ s \in S^* \mid \alpha^*(s) = \varnothing \}$. Let $\alpha : S^*/I \rightarrow S^*$ uniquely satisfy $s \in f[\alpha(s)]$.

B $\Rightarrow$ C P satisfies non-delusion and reduction of P' w.r.t. P' follow immediately.

$\chi_s(\omega) = \Sigma_{s' \in F(s)}\ \zeta_{s'}(\omega)$. So $\omega \in s$, $s' = P(\omega) \Rightarrow s' \in F(s) \Rightarrow s' \subset s$, so KTYK holds.

$\omega \in s \cap s' \Rightarrow P(\omega) \in F(s),\ P(\omega) \in F(s') \Rightarrow s \in F(s')$ or $s' \in F(s)$, so nestedness holds.

C $\Rightarrow$ B Let $I = \{ s \in S^* \mid$ there does not exist $s' \in S^*$ such that $s \subset s'\}$. By nestedness w.r.t. P', $s,s' \in I \Rightarrow s=s'$ or $s \cap s' = \phi$. By non-delusion w.r.t. P', there exist $s \in I$ s.t. $\omega \in s$, $\forall\ \omega \in L(P,P')$. Therefore, the sets in I partition L(P,P'); call the partition Q. Now B2 is true because P satisfies reduction of P' w.r.t. P'. Now define $F(s) = \{ s' \in S^* \mid s' \subset s \}$, $f(s) = \{ s' \in F(s) \mid$ there does not exist s" s.t. $s' \in F(s),\ s' \in F(s")\}$, $\alpha(s) = \{ s' \in S^* \mid s \in f(s') \}$. By construction, $s \not\ni s'$. Suppose $s',s" \in \alpha(s)$. Then $s \in f(s')$ and $s \in f(s")$. But this contradicts nestedness, therefore $\alpha(s)$ has at most one element. It remains only to show B3. By nestedness, sets in f(s) are mutually exclusive. Now suppose $\omega \in s$, $\omega \not\ni s'$, $\forall s' \in f(s)$. Then by KTYK w.r.t. P', $s = P(\omega)$.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

B ⇒ A let $\lambda(s,s') = 1$ if $s'=\alpha(s)$ or if $s\in I$, $s=P(\omega)$, $s'=P'(\omega)$, $\omega\in L(P,P')$, 0 otherwise.

Corollary 2.5

Suppose $P(\omega) = A$ and $P'(\omega) = B$, $\forall\ \omega\in L(P,P')$. Then condition A above gives

$$\sum_{s'\in S}\{\lambda(A,s')\chi_A(\omega) - \lambda(s',A)\chi_{s'}(\omega)\} = \begin{cases}1, & \forall\ \omega\in L(P,P')\\ 0, & \forall\ \omega\notin L(P,P')\end{cases}$$

$$\sum_{s'\in S}\{\lambda(B,s')\chi_B(\omega) - \lambda(s',B)\chi_{s'}(\omega)\} = \begin{cases}-1, & \forall\ \omega\in L(P,P')\\ 0, & \forall\ \omega\notin L(P,P')\end{cases}$$

$$\sum_{s'\in S}\{\lambda(s,s')\chi_s(\omega) - \lambda(s',s)\chi_{s'}(\omega)\} = 0,\quad \forall\ \omega\in\Omega, s\neq A\ or\ B$$

These three equations are satisfied only if:-

[1] $L(P,P') = \varnothing$
[2] $\lambda(A,B) = 1$, $\lambda(s,s') = 0$ otherwise, and thus $A = L(P,P')$
[3] $\lambda(A,B) = \lambda(B,A) = 1$, $\lambda(s,s') = 0$ otherwise, and thus $A = B \cup L(P,P')$ and $B \cap L(P,P') = \varnothing$.

Corollary 2.6: see part 3 below

2: Examples for section 2

I give examples to show that every combinations of rankings not excluded by the results of section 2 are possible.

Let $\Omega = \{a,b\}$, $S = \{(s_1,\ldots,s_n), a',b'\}$, with $\pi(\omega) = \gamma(\omega|s_i) = 1/2$, $\forall\ i=1,n$, $\omega\in\Omega$, and $\gamma(b|b') = \gamma(a|a') = 1$. Thus the agent is which state has occurred if he observes a' or b'. But observing any signal s_i does not change his prior beliefs.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

First consider the case where n = 1, and thus there is no irrelevent labelling of signals; experiment β_0 with $\beta_0(s_1|\omega) = 1$, for ω = a,b and experiment β_1 with $\beta_1(t|\omega) = 1$, for ω = a,b.

Now $\beta_0 \succeq_a \beta_0$ (and thus $\beta_0 \succeq_{aj} \beta_0$) but $\beta_1 \not\succeq_a \beta_1$, although $\beta_1 \succeq_{aj} \beta_1$.

Now consider n > 1, so there is irrelevant labelling of signals. Let $\beta_2(s_2,a) = \beta_2(s_3,b) = 1$. $\beta_2 \not\succeq_a \beta_2$, but $\beta_2 \succeq_{aj} \beta_2$; $\beta_0 \not\succeq_{aj} \beta_2$, $\beta_0 \not\succeq_u \beta_2$ but $\beta_0 \succeq_l \beta_l$; $\beta_2 \not\succeq_{aj} \beta_0$, $\beta_2 \not\succeq_l \beta_0$ but $\beta_2 \succeq_u \beta_0$.

Let β_3 satisfy $\beta_3(a'|a) = \beta_3(b'|b) = 1$. Thus β_3 is completely informative. Now $\beta_3 \succeq \beta_2$, even though there is irrelevent labelling.

Let β_4 satisfy $\beta_4(s_3|a) = \beta_4(b'|b) = 1$. Now $\beta_4 \not\succeq_{aj} \beta_2$ but $\beta_4 \succeq_u \beta_2$ and $\beta_4 \succeq_l \beta_2$.

Finally, I need a larger state space for my final example. Let Ω = {a,b,c,d}. Let S = {t_1, t_2, t_3, t_4}, with $\pi(\omega) = \gamma(\omega|t_i) = 1/4$, $\forall\ \omega \in \Omega$, i=1,4. Let $\beta_5(t_1|a) = \beta_5(t_1|b) = \beta_5(t_2|c) = \beta_5(t_2|d) = 1$ and let $\beta_6(t_3|a) = \beta_6(t_4|b) = \beta_6(t_3|c) = \beta_6(t_4|d) = 1$. Now $\beta_5 \not\succeq_u \beta_6$ and $\beta_5 \not\succeq_l \beta_6$, but $\beta_5 \succeq \beta_6$.

3: Further results about simple orderings

It is useful to introduced some additional language

Definition P is a type 1 improving simple alteration on P' (P' $\rightarrow_1$ P) if
[1] $P(\omega) = L(P,P')$, $\forall\ \omega \in L(P,P')$ [2] $P'(\omega') = P'(\omega)$, $\forall\ \omega,\omega' \in \Omega$.

Definition P is a type 2 improving simple alteration on P' (P' $\rightarrow_2$ P) if
[1] $P(\omega) = P'(\omega) \cup L(P,P')$, $\forall\ \omega \in L(P,P')$ [2] $P'(\omega) \cap L(P,P') = \varnothing$, $\forall\ \omega \in L(P,P')$
[3] $P'(\omega') = P'(\omega)$, $\forall\ \omega,\omega' \in \Omega$.

Definition P is [type 1 / type 2] reached by P' (P' $\Rightarrow_{[1/2]}$ P] if there exists a sequence of possibility correspondences, (R_1, R_2,.. ,R_n), such that

$$P' \rightarrow_{[1/2]} R_1 \rightarrow_{[1/2]} .. \rightarrow_{[1/2]} R_n \rightarrow_{[1/2]} P.$$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Result A3.1. If $P' \Rightarrow P$ and P' is non-deluded, then $P' \Rightarrow_1 P$.

If P' is non-deluded, there cannot exist any type 2 improving alterations.

Result A3.2. If P is non-deluded, KTYK, nested and a reduction of P', with respect to L(P,P'), then $P' \Rightarrow_1 P$.

This can be shown by construction. This implies corollary 2.6 in the text.

Definition P satisfies simple learning [w.r.t. E] if $P(\omega) \cap P(\omega') = \varnothing \nexists \omega'' \in \Omega$ s.t. $\{\omega,\omega'\} \subseteq P(\omega'')$, $\forall \omega,\omega' \in \Omega$ [$\forall \omega,\omega' \in E$]

Result A3.3. Suppose P and P' satisfy:-

[1] $P(\omega) = P(\omega') \Rightarrow P'(\omega) \cap \{ \omega'' \in L(P,P') \mid P(\omega'') = P(\omega) \} = P'(\omega') \cap \{ \omega'' \in L(P,P') \mid P(\omega'') = P(\omega')\}, \forall \omega,\omega' \in L(P,P') \}$.

[2] $Q(\omega) = \{ \omega' \in L(P,P') \mid P(\omega') = P(\omega), \omega' \in P(\omega), \omega \notin P'(\omega) \} \neq \varnothing, \forall \omega \in L(P,P')$

[3] Q is satisfies non-delusion, KTYK, nestednes, and simple learning w.r.t. L(P,P').

Then $P' \Rightarrow_2 P$.

This can be shown by construction.

Result A3.4. Let $P^{-}(\omega) = [P(\omega)]^{c}$. Let $E = \{ \omega' \in \Omega \mid P(\omega) = \Omega \}$. If P^{-} satisfies non-deludion, KTYK, nestedness and simple learning w.r.t. E, then $P_0 \geq P$.

This comes from applying result to the case when $P' = P_0$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Chapter IV

TRADING, SPECULATION AND CONSENSUS

When can "trade" occur between risk-neutral agents with asymmetric information? When can their posterior probabilities of an event be common knowledge, yet differ - i.e. when can they "agree to disagree"? When will agents take a "speculative" action, in an environment where it is known that it is impossible for everybody to gain and everybody is rational? The relationship between these questions is examined in this paper by identifying necessary and sufficient conditions on agents' prior beliefs for each of these results.

This relationship has been examined in the context of agents with common prior beliefs who process information using possibility correspondences[71]. A possibility correspondence specifies which states of the world the agents thinks possible in each state of the world, without requiring that the correspondence is a partition[72]. Brandenburger, Dekel and Geanakoplos [BDG] have shown that there is a decision theoretic equivalence between modelling agents as having common priors, but updating information with boundedly rational possibility correspondences, and modelling agents as having heterogeneous prior beliefs, and but updating information correctly. This paper uses this equivalence to show the relation between the possibility correspondence results and heterogeneous prior results.

In exploring these relationships, the presentation is necessarily somewhat taxonomic: I am trying to show how various old and new results are linked together. This introduction is intended to motivate this taxonomy, and identify the interesting new results.

---

71. Geanakoplos (1989), Rubinstein and Wolinsky (1990).

72. Such representations of information have been discussed in the economics literature by Bacharach (1985), Brandenburger, Dekel and Geanakoplos (1989), Shin (1987, 1989), Samet (1990), as well as the authors cited in note 71.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

In chapter II, I gave necessary and sufficient conditions on beliefs for no trade. Recall that three kinds of trade were considered: unrestricted, where trades can be made contingent on agents' private information; incentive compatible, where trades can depend on agents' private information only when it is incentive compatible for agents to truthfully report their signals; and public, where trades cannot depend on private information at all. Necessary and sufficient conditions to preclude the three types of trade are, respectively, consistency, reconcilability and public consistency of beliefs. Beliefs are consistent if agents posterior beliefs could have been derived from the same prior (even if, in fact, they weren't). Reconcilability has the interpretation that agents may consider their own signals to be of less importance than others do. Beliefs are publicly consistent if each agent's prior could be replaced by another prior with the same posterior beliefs conditional on his private information, such the replacement priors agree on the probability of publicly verifiable events. In section 2, I summarise these general results and consider the meaning of consistency, reconcilability and public consistency when agents have special kinds of beliefs. First, when beliefs are conditionally independent, so that agents believe that the signals they observe are independent, conditional on exogenous uncertainty. Secondly, when agents' beliefs are generated by possibility correspondences (PCs). In the latter case, the necessary and sufficient conditions for no trade are that all possibility correspondences are partitions (for unrestricted trade), all PCs satisfy non-delusion, "knowing that you know" and nestedness (for incentive compatible trade) and all PCs satisfy relative positive balance (for public trade). These results are closely related to Geanakoplos' results for speculation discussed below.

Aumann (1976) showed that under the common prior assumption, if agents' posterior probabilities of an event are common knowledge, then those posterior probabilities must be the same even if agents have different information. Thus agents cannot "agree to disagree". This

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

result has been extended to expected values of random variables[73]. But how different must their beliefs be to allow the possibility that agents can agree to disagree? It has long been recognized that this is closely related to the no trade result[74], the difference being that the inequalities representing acceptance of trade must be replaced by equalities. In section 3, I derive analogues to the three kinds of no trade results for agreeing to disagree: positive multipliers in the necessary and sufficient conditions are replaced by positive or negative multipliers. For unrestricted agreeing to disagree, this makes no difference, and the necessary and sufficient condition remains consistency. But suppose agents' private information in not observable, so that agents' ex post valuation of a random variable depends on what (other) agents choose to report; and suppose that one agent has an incentive to maximise the value of the random variable and the other agent has an incentive to minimise it. Does there then exist a random variable such that agents' different valuations can be common knowledge? When is there no incentive compatible agreeing to disagree? The necessary and sufficient condition is negative reconcilability. Finally, suppose we are concerned with random variables depending on some exogenous uncertainty (and not the signals agents have observed). There is no public agreeing to disagree if and only beliefs are negatively publicly consistent.

Again, these results can be related to possibility correspondence results. If beliefs are generated by possibility corespondences, then necessary and sufficient conditions for no agreeing to disagree are (1) all PCs are standard partitions (unrestricted), (2) all PCs satisfy non-delusion and "knowing that you know" (incentive compatible) and (3) all PCs satisfy relative balance (public). The last result is an extension of one in Geanakoplos (1989). What this framework makes clear is that Samet (1990) and Geanakoplos' (1989) results about agreeing

73. Nielson (1987)

74. For example, in Geanakoplos and Sebenius (1982).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

to disagree occur because the PC framework is implicitly restricting what the random variables can depend on (they can't depend directly on the signals agents have observed). The incentive compatible agreeing to disagree result is of particular interest, because non-delusion and "knowing that you know" together are natural restrictions to impose on knowledge, and this is a decision theoretic result for which they are jointly necessary and sufficient.

Geanakoplos (1989) introduced a powerful abstract definition of "speculation" for possibility correspondences. In section 4, I give a version of this for my general heterogeneous prior environment. Suppose there is a "true" probability distribution, π^*. Does there exist a Bayesian game, where there is some non-participation action available for each agent which gives him some reservation ex ante expected utility (evaluated with π^*), where it is not possible for any agent to strictly improve on his reservation utility if all other agents at least achieve their reservation utility, but where nonetheless some agent chooses an action other than the "safe" action and expects strictly positive gains from doing so? I show that there is no speculation if and only if there would be no incentive compatible trade even if an uninformed agent with prior beliefs π^* were added. This is another way of stating Geanakoplos' result that there is no speculation if and only if each agent's information increases ex ante utility for every one person decision problem.

I give results for no common knowledge speculation - with the additional requirement that actions are common knowledge in the speculation; and again, I can then specialise to the possibility correspondence case. Conditions for no incentive compatible trade and no speculation are equivalent (Geanakoplos gave a very different proof for the no speculation case). The necessary and sufficient condition to preclude common knowledge speculation is positive balance (this is a slight extension of Geanakoplos' result). This differs from the no public trade result where relative positive balance was required. The key difference between "no trade" and

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

"no speculation" results is that the existence of trade depends on differences in information processing errors, whereas the existence of speculation depends on an absolute standards of errors.

1. Framework and Common Knowledge with Probability One

I first summarise (section 1.1) the general framework I used in chapter II to represent the uncertain environment where there is both heterogeneous private information and exogenous uncertainty. It is important for the results of this paper to have a notion of common knowledge when agents do not agree on which states of the world are possible. So in section 1.2, I give such a notion of "common knowledge with probability 1", based on Brandenburger and Dekel (1987a). In section 1.3, I consider restrictions on agents' heterogeneous beliefs: first, that all agents believe that signals are independent, conditional on exogenous (non-informational) events; and secondly, restrictions that give beliefs as generated by possibility correspondences. I also derive for future use a simple characterisation of "common knowledge with probability 1" for possibility correspondences.

1.1 Framework

There are H agents. I also write H for the set of agents. Each agent $h \in H$ observes a private signal $t_h \in T_h$. The product of the agents' signals, $T = T_1 \times .. \times T_H$, thus incorporates all private information. I write $t = (t_1, .. , t_H)$ for a typical element of T. But not all uncertainty in the economy is reflected in agents' private information. There is also a set $\underline{Q}$ of exogenous events (i.e. not private information), with typical element q. All relevant uncertainty is thus

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

described by $\Omega = T \times Q$. Each agent has a prior $\pi_h : \Omega \rightarrow \mathbb{R}_+$, $\Sigma_{\omega \in \Omega}\pi_h(\omega) = 1$. But agents do not necessarily assign positive probability to every state: let agent h's "possibility set", $\Omega^+_h = \{ \omega \in \Omega \mid \pi_h(\omega) > 0 \}$. I assume that every agent thinks that each of his signals $t_h \in T_h$ is possible i.e. $\{t,q\} \in \Omega^+_h$, for all $t_h \in T_h$, for some $t_{-h} \in T_{-h}$, $q \in Q$[75]. It will be important for the analysis of possibility correspondences that I allow for the case where these possibility sets differ. But it is useful to at least get rid of those states that no one thinks possible: let $\Omega^+ = \cup_{h \in H}\{\Omega^+_h\}$. Now I can give a partition representation of uncertainty:-

$$P_h(\ \{t,q\}\ ) = \{\ \{t',q'\} \in \Omega^+ \mid t_h' = t_h\ \} \quad \forall \{t,q\} \in \Omega^+$$

$$P^+_h(\ \{t,q\}\ ) = \{\ \{t',q'\} \in \Omega^+_h \mid t_h' = t_h\ \} \quad \forall \{t,q\} \in \Omega^+$$

$$Q(\ \{t,q\}\ ) = \{\ \{t',q'\} \in \Omega^+ \mid q' = q\ \} \quad \forall \{t,q\} \in \Omega^+$$

$P_h(\ \{t,q\}\ )$ is the set of states indistinguishable from $\{t,q\}$, while $P^+_h(\ \{t,q\}\ )$ is the set of states that agent h believes possible (i.e. occur with probability 1) at state $\{t,q\}$. Note that I could have described uncertainty by starting with state space Ω^+, giving agents partitions $\{P_h\}_{h \in H}$, partition Q to represent other events, and specifying priors such that every state was thought possible by some agent, and each agent thinks every information set in his partition possible[76].

---

75. As in chapter II, the following notation is used:-

$$t = (t_1,.. ,t_H) \in T_1 \text{ x.. } T_H = T$$
$$t_{-h} = (t_1,.. ,t_{h-1},t_{h+1},.. ,t_H) \in T_1 \text{ x.. x } T_{h-1} \text{ x } T_{h+1} \text{ x.. x } T_H = T_{-h}$$
$$\{t_{-h},t_h'\} = (t_1,.. ,t_{h-1},t_h',t_{h+1},.. ,t_H) \in T$$

76. Or I could have started with Ω^+ and chosen "probability 1" correspondences, $\{P_h^+\}_{h \in H}$, each P^+_h: $\Omega^+ \rightarrow 2^{\Omega^+}/\{\emptyset\}$ satisfying [1] $\omega' \in P^+_h(\omega) \Rightarrow P^+_h(\omega') \subseteq P^+_h(\omega)$ ["knowing that you know"], [2] $\omega' \in P^+_h(\omega) \Rightarrow P^+_h(\omega') \supseteq P^+_h(\omega)$ ["knowing that you don't know"], but not necessarily the third condition on possibility correspondences necessary for a partition, $\omega \in P^+_h(\omega)\ \forall\ \omega \in \Omega^+$ ["non-delusion"]. See Geanakoplos (1988) for a discussion of these axioms.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

1.2 Common Knowledge with Probability One

I need both a definition of common knowledge and an easy way of verifying common knowledge in this environment. The following results are implicit in Brandenburger and Dekel's (1987a) general treatment of "common knowledge with probability one", but can be stated in a form exactly like the standard results (where agents agree on which states are possible, $\Omega^+_h = \Omega^+$, $\forall\, h \in H$)[77]. Throughout this discussion, I am concerned only with states and events in Ω^+.

Definition Agent h knows event A at ω if and only if $P_h^+(\omega) \subseteq A$.

$$K_h A = \{\omega \in \Omega^+ \mid P_h^+(\omega) \subseteq A\}$$

Definition Event A is common knowledge at ω if every agent knows event A, every agent knows that every agent knows event A, and so on.

$$CKA = (\ \omega \in \Omega^+ \mid \omega \in K_{h_1} .. K_{h_n} A \quad \forall\, (h_1,..,h_n) \in H^n,\ n \in \mathrm{N}_{++}\ )$$

Definition Event F is public if all agents know F whenever F occurs i.e. $F \subseteq K_h F$.

Note that it is not the case that $K_h A \subseteq A$, for all events A, so that if an agent knows event A, event A is true. This property is true if agents agree on which states are possible. Not being able to use this property adds some extra steps to the proof of the theorem.

77. Notice the following result would not hold if I had not assumed that $P_h^+(\omega)$ is always non-empty and restricted attention to Ω^+, where at least one agent thinks every state is possible. Aumann (1976) made the original argument that an event can be common knowledge at a given state if and only there exists a public event (an event measurable on the meet) containing the state and contained in the event.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Theorem Event A is common knowledge at ω [i.e. $\omega \in CKA$] if and only if there exists a public event F such that $\omega \in F$, $F \subseteq A$.

Proof (if)

$$Let \quad K^m A = (\omega \in \Omega^+ \mid \omega \in K_{h_1}..K_{h_m} A \quad \forall (h_1,..,h_m) \in H^m)$$

Now if F is a public event, $F \subseteq K_h F$, $\forall h \in H$, by definition. So if $F \subseteq K^m A$, then $K_h F \subseteq K_h K^m A$ $\forall h \in H$, so $F \subseteq K_h F \subseteq K^{m+1} A$. But if $F \subseteq A$, then $K_h F \subseteq K_h A$ and $F \subseteq K_h A$, $\forall h \in H$. So $F \subseteq K^1 A$. Now, by induction, $F \subseteq K^m A$, $\forall m \in N_+$, and thus $F \subseteq CKA$. So if $\omega \in F$, A is common knowledge at ω.

(only if) Let $F = CKA$; now $\omega \in F$; F is a public event by definition of common knowledge; $F \subseteq K_h A$, by definition of common knowledge; $K_h A \cap \Omega^+_h \subseteq A$, $\forall h \in H$, by definition of knowledge. Now if $\omega \in F$, $\omega \notin A$, it follows that $\omega \notin \Omega^+_h$, $\forall h \in H$, which is not possible for $\omega \in \Omega^+$. So $F \subseteq A$.

1.3 Restriction on beliefs

I will economize in notation in what follows by using π_h for agent h's beliefs about states, events and conditional events, i.e.

$$\pi_h(t_h) = \sum_{t_{-h} \in T_{-h}} \sum_{q \in Q} \pi_h(t,q) \qquad \pi_h(q|t_h)\ \pi_h(t_h) = \sum_{t_{-h} \in T_{-h}} \pi_h(t,q)$$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

$$\pi_h(q) = \sum_{t \in T} \pi_h(t,q) \qquad\qquad \pi_h(t_{h'}|q)\, \pi_h(q) = \sum_{t_{-h'} \in T_{-h'}} \pi_h(t,q) \quad \forall h' \in H$$

Suppose agents believe that they are observing independent signals $(t_1, \ldots, t_H)$ about which uncertain event $q \in \underline{Q}$ has occurred. Then I say their beliefs $\{\pi_h\}_{h \in H}$ are conditionally independent.

<u>Definition</u> Beliefs are conditionally independent if (equivalently)

$$\pi_h(t,q) = \pi_h(q)\left\{\prod_{h' \in H} \pi_h(t_{h'}|q)\right\} \qquad or \qquad \pi_h(t,q) = \pi_h(t_h)\, \pi_h(q|t_h)\left\{\prod_{h' \neq h} \pi_h(t_{h'}|q)\right\}$$

A number of authors have considered agents whose beliefs are generated by possibility correspondences: in each state of the world, $q \in \underline{Q}$, there is a set of states of the world, $R_h(q)$, which agent h thinks are possible, so that $R_h : \underline{Q} \rightarrow 2^{\underline{Q}} / \{\varnothing\}$[78]. There is a sense in which the agent is not being fully rational if R_h is not a partition i.e. if $R_h(q) \neq \{ q' \in \underline{Q} \mid R_h(q') = R_h(q)\}$, because the agent should be able to refine his possibility correspondence by noting what he does <u>not</u> know. Brandenburger, Dekel and Geanakoplos (1989) showed a decision theoretic equivalence between problems where agents have common priors, but make mistakes interpreting their information, and problems where agents have heterogeneous priors and interpret information correctly. This equivalence is illustrated and used below by modelling agents with possibility correspondences as a special case of conditionally independent beliefs.

Suppose each agent $h \in H$ has a possibility correspondence on $\underline{Q}$. Then the set of signals he observes, $T_h = \{ C \in 2^{\underline{Q}} \mid C = R_h(q) \text{, for some } q \in \underline{Q} \} \subseteq 2^{\underline{Q}}$. If all agents have the same prior on q, then the "true" distribution π^* on $\Omega = T \times \underline{Q}$ is:-

78. see references in footnotes 71 and 72.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

$$\pi^*(t,q) = \pi^*(q)\left\{\prod_{h\in H} \zeta(t_h|q)\right\} \qquad \text{where } \zeta_h(t_h|q) = \begin{cases} 1, & \text{if } t_h = R_h(q) \\ 0, & \text{otherwise} \end{cases}$$

But conditional on observing signal $t_h \in T_h$, agent h updates his beliefs conditional on being in set t_h, so that

$$\pi_h(q|t_h) = \frac{\pi^*(q)\ \chi_{t_h}(q)}{\sum_{q'\in t_h} \pi^*(q')} \qquad \text{where } \chi_{t_h}(q) = \begin{cases} 1, & \text{if } q\in t_h \\ 0, & \text{otherwise} \end{cases}$$

Now suppose that while agents make mistakes interpreting their own signal, they correctly infer which signal other agents will have observed.

<u>Definition</u> Agents beliefs, $\{\pi_h\}_{h\in H}$, are generated by possibility correspondences, $\{R_h\}_{h\in H}$, if they satisfy:-

$$\pi_h(t,q) = \frac{\pi_h(t_h)\ \pi^*(q)\ \chi_{t_h}(q)\prod_{h'\neq h} \zeta_{h'}(t_{h'}|q)}{\sum_{q'\in t_h} \pi^*(q')}$$

Agents' prior beliefs over their own signals ($\pi_h(t_h)$) will turn out to be irrelevant, so I make no special assumption about them. Also notice that agents disagree about which states are possible.

$$\Omega^+_h = (\ \{t,q\}\ |\ q\in t_h,\ t_{h'}=R_h(q)\ \forall\ h'\neq h)$$

$$P^+_h(\ \{t,q\}\ ) = (\ \{t',q'\} \in \Omega^+\ |\ t'_h = t_h,\ q'\in t_h,\ t_{h'}=R_h(q')\ \forall\ h' \neq h)$$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

It will be useful to have a simple characterisation of public events, and thus common knowledge, when beliefs are generated by PCs.

<u>Definition</u> Event $F \subseteq \underline{Q}$ is self-evident[79] if, for each $q \in F$, $h \in H$, $R_h(q) \subseteq F$.

<u>Definition</u> Event $F \subseteq \underline{Q}$ is strongly self-evident if, for each $q \in \underline{F}$, $R_h(q) \subseteq F$, $\forall\, h \in H$, and $q \in R_{h'}(q')$ for some $h' \in H$, $q' \in \underline{Q}$.

Thus an event is self-evident if each agent knows that F has occurred, whenever it occurs. F is strongly self-evident if, in addition, every state in F is considered possible by some agent at some (possibly different) state in F.

<u>Definition</u> R_h is non-deluded if $q \in R_h(q)$, $\forall\, q \in \underline{Q}$

Note that if each R_h is non-deluded, F is strongly self-evident whenever it is self-evident.

<u>Theorem</u> If beliefs are generated by PCs $\{R_h\}_{h \in H}$, an event $E \subseteq \Omega^+$ is public if and only if there exists a strongly self evident event $F \subseteq \underline{Q}$ such that

$$E = g(F) = \{ \{t,q\} \in \Omega^+ \mid q \in F,\ t_h = R_h(q') \text{ for some } q' \in F,\ \forall h \in H \}$$

<u>Proof</u> (if) $q \in F \Rightarrow R_h(q) \subseteq F$, so $\{t',q'\} \in P_h^+(\{t,q\})$ and $\{t,q\} \in E \Rightarrow \{t',q'\} \in E$

(only if) Let $F = \{ q \in \underline{Q} \mid \{t,q\} \in E$, for some $t \in T \}$, and $E_h = \{t_h \in T_h \mid \{t,q\} \in E$ for some $t_{-h} \in T_{-h}$, $q \in \underline{Q}\}$, for each $h \in H$. Now suppose $\{t',q'\} \in \Omega^+$ satisfies $t_h' \in$

---

79. This notion was introduced in Shin (1987) and is discussed in Geanakoplos (1989) and Brandenburger, Dekel and Geanakoplos (1989).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

E_h, for each $h \in H$, and $q' \in F$. $\{t',q'\} \in \Omega^+ \Rightarrow q' \in t_h'$, $t_{h'} = R_{h'}(q') \; \forall \; h' \neq h$, for some $h \in H$. Then there exists $\{t,q\} \in E$ such that $\{t',q'\} \in P_h^+(\{t,q\})$, so $\{t',q'\} \in E$, and $[E_1 \times .. \times E_H \times F] \cap \Omega^+_h$. But now F must be strongly self-evident because [1] $q \in F \Rightarrow R_h(q) \in E_h \Rightarrow q' \in F$ if $q' \in R_h(q)$ and [2] $q \in F \Rightarrow q \in t_h$, for some $t_h \in E_h$.

In making the BDG translation between thinking of agents making mistakes processing information (the possibility correspondences approach) and agents with heterogeneous prior beliefs, it is not obvious how to translate ideas like common knowledge. The theorem shows that "common knowledge with probability 1" in the heterogeneous prior environment translates well in the framework, since self-evident is a natural extension of common knowledge to non-deluded possibility correspondences.

2. Trading

This section summarises results from chapter II, with somewhat changed terminology, and applies them to the case of conditional independent beliefs and beliefs generated by possibility correspondences.

Definitions A trade is a set of random variables, $\{x_h\}_{h \in H}$, each $x_h : \Omega \rightarrow \mathbb{R}$. A trade is feasible if

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

$$\sum_{h \in H} x_h(\omega) \leq 0, \quad \forall \omega \in \Omega$$

A trade is publicly accepted if there exists a public event $E \subseteq \Omega^+$, such that:-

$$\sum_{\omega' \in P_h^+(\omega)} \pi_h(\omega')\, x_h(\omega') \geq 0, \quad \forall\, h \in H, \omega \in E$$

with strict inequality for some $h \in H$, $\omega \in E$.

A trade is incentive compatible if

$$\sum_{t_{-h} \in T_{-h}} \sum_{q \in Q} \pi_h(t,q)\, x_h(t,q) \geq \sum_{t_{-h} \in T_{-h}} \sum_{q \in Q} \pi_h(t,q)\, x_h(\{t_{-h}, t_h'\}, q), \quad \forall\, h \in H, t_h \in T_h, t_h' \in T_h$$

A trade is verifiable if each x_h is measurable with respect to partition Q. Unrestricted trade occurs if there exists x satisfying feasibility and public acceptance. Incentive compatible trade occurs is there exists x satisfying feasibility, public acceptance and incentive compatability. Public trade occurs if there exists x satisfying feasibility, public acceptance and verifiable.

The motivation of chapter II was that there are different kinds of contractual arrangments: if complete contingent contracting is possible, then unrestricted trade is the relevant definition. If agents' private information is not verifiable, but it is possible to make payoffs contingent on private information as long as revelation of the private information is incentive compatible, then incentive compatible trade is the relevant definition. If it is not possible to make trade contingent on private information at all, and common knowledge acceptance is required, then public trade is the relevant concept.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Definition Agent h's beliefs are consistent with π if

$$\pi_h [t_h, q \mid t_h] = \pi [t_h, q \mid t_h], \ \forall \{t,q\} \in \Omega^+$$

Theorem (i) There is no unrestricted trade if and only if there exists a "common prior" π such that each agents' beliefs are consistent with π.

(ii) If agents' beliefs are generated by possibility correspondences, then there is no unrestricted trade if and only if all agents' PCs are partitions.

Proof (i) is proved in chapter II; (ii) follows by substituting special restrictions on priors into consistency conditions.

Definitions Beliefs are [negatively] reconciliable with π if and only if there exist $\{\lambda_h\}_{h\in H}$, $\lambda_h : T_h \rightarrow \mathbf{R}_{++}$ $[\mathbf{R}]$, $\{\mu_h\}_{h\in H}$, each μ_h: $T_h^2 \rightarrow \mathbf{R}_+$, such that:-

$$\lambda_h(t_h)\ \pi_h(t,q) + \sum_{t'_h\in T_h} \{ \mu(t_h,t'_h)\ \pi_h(t,q) + \mu(t'_h,t_h)\ \pi_h(\{t_{-h},t'_h\},q)\} = \pi(t,q), \quad \forall\ t,q\in\Omega^+$$

R satisfies "knowing that you know" if $q' \in R(q) \Rightarrow R(q') \subseteq R(q), \ \forall\ q \in \underline{Q}$.

R satisfies nestedness if $R(q) \cap R(q') = \emptyset$, $R(q)$ and $R(q')$, $\forall\ q \in \underline{Q}$.

See Geanakoplos (1988) for a discussion of the PC properties.

Theorem (i) There is no incentive compatible trade if and only if there exists a "common prior" π such that each agents' beliefs are reconcilable with π.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

(ii) If agents' beliefs are generated by possibility correspondences, then there is no incentive compatible trade if and only if all agents' PCs are non-deluded, KTYK and nestedness.

Proof (i) is proved in chapter II; corollary 2.4 of chapter III showed that the logical properties are obtained by substitution of the PC beliefs. Geanakoplos (1989) proved that non-delusion, KTYK and nestedness were necessary and sufficient conditions for no speculation (see section 4). Shin (1989) showed that a "dynamic representation", which is essentially equivalent to agents' beliefs about Q, conditional on their own signals, being reconciliable with π, in necessary and sufficient for non-delusion, KTYK and nestedness.

Definition (i) Agent h's beliefs are [negatively] publicly consistent with π if, for every public event $E \subseteq \Omega^+$, there exist $\{\lambda_h\}_{h\in H}$, each λ_h: $T_h \rightarrow \mathbf{R}_+$ [$\mathbf{R}$], such that:-

$$\sum_{(t|(t,q)\in E)} \lambda_h(t_h)\, \pi_h(t,q) = \sum_{(t|(t,q)\in E)} \pi(t,q) \qquad \forall\, q\in Q$$

(ii) Possibility Correspendence R satisfies [positive] {relative} balance[80] if, for every strongly self-evident event $F \subseteq \Omega^+$, there exist $\{\lambda_h\}_{h\in H}$, each λ_h: $T_h \rightarrow \mathbf{R}$ [$\mathbf{R}_+$], {k: $F \rightarrow \mathbf{R}_{++}$, k non-zero},

$$\sum_{t_h\in T_h} \lambda_h(t_h)\, \chi_{t_h}(q) = 1 \;\{k(q)\}, \quad \forall\, q\in Q$$

Theorem (i) There is no public trade if and only if there exists a "common prior" π such that each agents' beliefs are publicly consistent with π.

80. This generalizes to deluded PCs a concept in Geanakoplos (1989).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

(ii) If agents' beliefs are generated by possibility correspondences, then there is no public trade if and only if all agents' PCs are satisfy relative positive balance.

3. Agreeing to Disagree

The relationship between no trade results and no "agreeing to disagree" results is most easy to see in the two person case[81]. Both kinds of results generalise easily to many agents, but they generalise in different ways. If there are two agents, we can think of trades being represented by a single random variable which one agent pays to the other, i.e. $x_a(\omega) = x(\omega) = -x_b(\omega)$, $\forall\ \omega \in \Omega$. This extra requirement never prevents trades from existing. In this case, I have:-

Definition Random variable x is publicly accepted if there exists a public event E such that

$$\sum_{\omega' \in P_a^*(\omega)} \pi_a(\omega')\, x(\omega') \geq 0, \quad \forall\ \omega \in E \qquad \sum_{\omega' \in P_b^*(\omega)} \pi_b(\omega')\, x(\omega') \leq 0, \quad \forall\ \omega \in E$$

with strict inequality for some $h \in H$, $\omega \in E$.

x is incentive compatible if

81. As in Geanakoplos and Sebenius (1982)

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

$$\sum_{t_b\in T_b}\sum_{q\in Q}\pi_a(t,q)\ x(t,q) \geq \sum_{t_b\in T_b}\sum_{q\in Q}\pi_a(t,q)\ x((t_b,t_a'),q), \quad \forall\ t_a\in T_a,\ t_a'\in T_a$$

$$\sum_{t_a\in T_a}\sum_{q\in Q}\pi_b(t,q)\ x(t,q) \leq \sum_{t_a\in T_a}\sum_{q\in Q}\pi_b(t,q)\ x((t_a,t_b'),q), \quad \forall\ t_b\in T_b,\ t_b'\in T_b$$

x is verifiable if it is measurable with respect to partition Q.

Definition Agents agree to disagree about x if there exists a public event E and $k_a \neq k_b$ such that

$$\sum_{\omega'\in P_a^*(\omega)} \pi_a(\omega')\ [\ x(\omega') - k_a\ ] = 0 \quad \forall\ \omega\in E \qquad \sum_{\omega'\in P_b^*(\omega)} \pi_b(\omega')\ [\ x(\omega') - k_b\ ] = 0,$$

Now if I replace public acceptance of trade by agreeing to disagree, I will get three sets of results (unrestricted, incentive compatible and public) which differ from the trade case only in requiring requiring equality rather than inequality in the expected value of x constraints.

Definition There is no unrestricted agreeing to disagree if there does not exist x satisfying agreeing to disagree.

If there is no unrestricted disagree to disagree, then it is impossible for the agents' expected value of the random variable to be common knowledge without it being the same.

Theorem (i) There is no unrestricted agreeing to disagree if and only if beliefs are consistent.

(ii) If beliefs are generated by PCs, there is no unrestricted agreeing to disagree if and only if all PC's are partitions.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Definition There is no incentive compatible agreeing to disagree if there does not exist x satisfying agreeing to disagree and incentive compatibility.

Suppose agents' private information is not verifiable, so the only way anyone else can know an agent's private information is if he reports it. Suppose one agent has an incentive to exaggerate the value of the random variable, while the other agent has an incentive to undervalue it. Does there exist a random variable, such that agents' have an incentive to truthfully report their signal, where posteriors are common knowledge, but where they differ?

Theorem (i) There is no incentive compatible agreeing to disagree if and only if beliefs are negatively reconcilable with some prior π.

(ii) If beliefs are generated by possibility correspondences, there is no incentive compatible agreeing to disagree if and only all PCs satisfy non-delusion and KTYK.

Proof The change of sign of the $\{\lambda_h\}_{h\in H}$ in (i) come replacing public acceptance with public agreeing to disagree. For (ii), I must show that there exist $\{\lambda_h\}_{h\in H}$, $\lambda_h : T_h \to \mathbb{R}$, $\{\mu_h\}_{h\in H}$, $\mu_h : T_h^2 \to \mathbb{R}_+$, such that:-

$$\zeta(t_h|q) = \lambda(t_h)\, \chi_{t_h}(q) + \sum_{t_h'\in T_h} \mu(t_h,t_h')\, \chi_{t_h}(q) - \mu(t_h',t_h)\, \chi_{t_h'}(q)$$

if and only R_h satisfies non-delusion and KTYK. To see this, let $\mu(t_h,t_h') = 1 \Leftrightarrow t_h \subset t_h'$ and $\nexists\, t_h''$ s.t. $t_h \subset t_h'' \subset t_h'$[82], and 0 otherwise. Now let

82. where $C \subset C'$ means strictly inclusion, i.e. $C \subseteq C'$ and $C \neq C'$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

$$\lambda(t_h) = \sum_{t'_h \in T_h} \mu(t_h, t'_h)$$

This result is of some interest because non-delusion and KYTK are the axioms on possibility correspondences that seem most natural from an epistemic point of view[83]; Shin (1987) showed that non-delusion and KTYK are the conditions for knowledge to represent "provability"; they have been shown to be necessary for a number of decision theoretic results: e.g. Samet (1990) for what I will define as no public agreeing to disagree, Geanakoplos (1989) for no speculation. Here we have a decision theoretic result for which they are necessary and sufficent conditions.

<u>Definition</u> There is public agreeing to disagree if there exists x satisfying agreeing to disagree and verifiability.

<u>Theorem</u> (i) There is no public agreeing to disagree if and only if beliefs are negatively publicly consistent with some prior π.

(ii) If beliefs are generated by possibility correspondences, there is no public agreeing to disagree if and only all PCs satisfy relative balance.

This public agreeing to disagree is most closely related to Samet's (1990) and Geanakoplos' (1989) results about agreeing to disagree.

Agreeing to disagree and no trade results generalise to many agents in different ways. There is agreeing to disagree among many agents if and only if there is agreeing to disagree

83. Kripke (1963), Hintikka (1962) and Halpern (1986).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

between any pair of agents. On the other hand, there may be trade between a group of three or more agents, when there is no trade between any pair of agents.

4. Speculation

Geanakoplos (1989) proposed a definition of speculation for agents with possibility correspondences which generalises as follows:-

Definitions Consider the game Bayesian G, where agent h has a simple action space A_h, and chooses a mapping, f_h, from the set of signals he observes, T_h, to his set of actions, A_h. His payoff depends on some exogenous uncertainty represented by Q, and the actions of other players i.e. $u_h : A \times Q \rightarrow \mathbf{R}$, where $A = A_1 \times .. \times A_H$, and his utility

$$v_h(f) = \sum_{t \in T} \sum_{q \in Q} \pi_h(t,q)\ u_h(f(t),q), \qquad \text{where } f: T \rightarrow A \text{ satisfies } f(t) = \{f_1(t_1),.. f_H(t_H)\}$$

G is a speculation game with respect to π^* if each player h has an action $z_h \in A_h$ such that $\Sigma_{t \in T} \Sigma_{q \in Q} \pi^*(t,q)\ u_h(\{z_h, f_{-h}(t_{-h})\}, q) = u_h^*$, for all f, and such that $\Sigma_{t \in T} \Sigma_{q \in Q} \pi^*(t,q)\ u_h(f(t),q) \geq u_h^*$, $\forall\, h \in H$, implies $\Sigma_{t \in T} \Sigma_{q \in Q} \pi^*(t,q)\ u_h(f(t),q) = u_h^*$, $\forall\, h \in H$.

Speculation occurs if there exists a Nash equilibrium of a speculation game where some action is strictly preferred to z_h by some agent in some state.

In other words, speculation occurs if, despite the fact that it is not possible for any agent to do strictly better (in ex ante expected utility terms, under π^*) than playing z_h everywhere, some agent believes he will do strictly better. Geanakoplos (1989) has shown how this notion speculation result underlies no speculation in competitive and rational expectations equilibrium results.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

How does this relate to the no trade results of section II above and chapter 2? No speculation is a stronger requirement than no trade. Consider the following strengthenings of speculation:-

Definition Speculation with no ex ante expected gains occurs if speculation occurs with $\Sigma_{h\in H} \Sigma_{t\in T} \Sigma_{q\in Q} \pi^*(t,q)\, u_h[f(t),q] \leq \Sigma_{h\in H} u_h^*$, for all f. Speculation with no ex post gains occurs if speculation occurs with $\Sigma_{h\in H} u_h[f(t),q] \leq \Sigma_{h\in H} u_h[\{z_h, f_{-h}(t_{-h})\},q]$, $\forall$ t, q, f.

Note that requiring no ex post gains is strictly stronger than requiring no ex ante gains which is strictly stronger than the basic definition of speculation. The easiest way to see the relation between these definitions is to think of "nature" as being another player with beliefs π^*, and no private information. Then I have:-

Theorem There is no speculation, and no speculation with no ex ante expected gains, if and only if there is no incentive compatible trade among agents in H and nature. There is no speculation with no ex post gains if and only if there is no incentive compatible trade among agents in H.

Proof If there is speculation, then, for some h, $\Sigma_{t\in T} \Sigma q \in Q\, \pi_h(t,q)\, u_h(t,q) > u_h^*$. Now if $u_{h'}$ is set equal to zero for all h' $\neq$ h, then there is speculation with no ex ante expected gains. Now if there is speculation with payoffs u and equilibrium strategies f, let

$$x_h(t,q) = u_h[f(t),q] - u_h[\{z_h, f_{-h}(t_{-h})\},q] \qquad x_n(t,q) = -\sum_{h\in H} x_h(t,q)$$

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

then the $\{x_h\}_{h \in H \cap \{n\}}$ satisfy feasibility, unconditional acceptance and incentive compatibility, where "n" is "nature", an uninformed agent with beliefs π^*. If there is speculation with no ex post gains, then the $\{x_h\}_{h \in H}$ satisfy feasibility. Conversely, if the $\{x_h\}_{h \in H \cap \{n\}}$ satisfy feasibility, unconditional acceptance and incentive compatibility, then

$$A_h = T_h \cup \{z_h\} \qquad u_h(f(t),q) = x_h(t,q) \qquad u_h(,\{z_h, f_{-h}(t_{-h})\},,q) = 0$$

satisfy the conditions for speculation. If the $\{x_h\}_{h \in H}$ satisfy feasibility, then u and f satisfy the conditions for speculation with no ex post gains.

The equivalence of speculation with no ex ante gains and incentive compatible trade is comes essentially from the revelation principle.

<u>Corollary</u> (i) There is no speculation if and only agents' beliefs are reconcilable with π^*.

(ii) If beliefs are generated by PCs, there is no speculation if and only if beliefs satisfy non-delusion, KTYK and nestedness.

Note that for possibility correspondences (but not in general) conditions for no speculation and no trade are the same.

<u>Definition</u> There is common knowledge speculation [with no ex ante expected gains / with no ex post gains] if there is speculation [with no exte expected gains / no ex post gains], and a public event E such that $f_h(t_h)$ is constant in E, for each h, and strictly preferred to z_h for some h.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Theorem (i) There is no common knowledge speculation if and only beliefs are publicly consistent with π^*.

(ii) If beliefs are generated by PCs, there is no common knowledge speculation if and only if beliefs satisfy positive balance.

Note that for no common knowledge speculation, positive balance is required - no public trade required only relative positive balance. For example, if all agents observe the same PC, then, however badly behaved, they satisfy relative positive balance and do not trade. But they may speculate.

Section 5: Conclusion

This chapter generalized the no trade results of chapter II in two directions: first, to "agreeing to disagree" type results and "speculation"; and, secondly, to see the relation to possibility correspondences. By doing so, the relation between existing results in different frameworks is clearer. A key idea was the translation of common knowledge between heterogeneous prior and possibility correspondence frameworks.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

CHAPTER V

CONCLUSION

In this chapter, I discuss the role of the common prior assumption in economic theory, the arguments given in favor of using it, and past and possible future work using the heterogeneity of prior beliefs (section 1). Section 2 is a discussion of misinterpreted information and its relation to heterogeneous priors. Section 3 re-examines the results of this dissertation in the light of those discussions.

1. The Common Prior Assumption in Economic Theory

> Why is (it that) common priors are implicit or explicit in the vast majority of the differential information literature in economics and game theory? Why has the economic community been unwilling, in practice, to accept and actually use the idea of truly personal probabilities in much the same way that it did accept the idea of personal utility functions? After all, in (Savage's expected utility theory), both the utilities and probabilities are derived separately for each decision maker. Why were the utilities accepted as personal, and the probabilities not?[84]

One way to "refute" the common prior assumption is to appeal to Aumann's (1976) "Agreeing to Disagree" paper. If agents had common prior beliefs then it would be impossible for rational agents to disagree with each other about anything when their beliefs are common knowledge, even if they started out with asymmetric information[85]. Since such "agreeing to

---

84. Aumann (1987).

85. Geanakoplos and Polemarchakis (1982) give a dynamic version of the "agreeing to disagree" result.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

disagree" seems to be a fundamental aspect of social life, not excluding economic life, surely any theory with such assumptions at its core is doomed to failure.

But this will not be enough. For a start, there is Aumann's response, discussed further below, that we ought to interpret observed "agreeing to disagree" as a failure of rationality in information processing, not as a result of differences in prior beliefs. More generally, and for good reason, it is not enough to argue the common prior assumption is not true. It is necessary to argue that by assuming the common prior assumption, important aspects of economic reality, which could be understood with the help of heterogeneous prior models, are missed. It is necessary (in section 1.2) to address the amalgam of philosophical, empirical and pragmatic reasons given in support of the common prior assumption (including the comprehensive set offered by Aumann in response to his own question quoted above). This is attempted on a point-by-point basis. Before doing so, it is necessary to cover (in section 1.1) some preliminary ground concerning the meaning of probabilities in philosophical, decision theoretic and economic contexts. Finally (in section 1.3), I discuss work that has been and can be done explicitly dropping the common prior assumption.

There exist some excellent discussions of the merits of the common prior assumption[86], and this paper repeats many familiar arguments. One point which this discussion emphasizes is the necessity of different justifications for the common prior assumption depending on whether we are concerned about agents' beliefs about events exogenous or endogenous to the economic model.

Aumann (1987), Bernheim (1987), Varian (1989).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

1.1 Probability

Before discussing the role of the common prior assumption (CPA) directly, it is useful to review what might be meant by probabilities in different contexts: in philosophy, in decision theory and in economic models, and to clarify the relation between different uses of the mathematical concept of probability in these different contexts.

"Probability" means different things to different people; Savage (1954) identified three underlying kinds of views of probability. "Objectivistic" or "frequentist" views[87]

> hold that some repetitive events, such as tosses of a penny, prove to be in reasonably close agreement with the mathematical concept of independently repeated random events, all with the same probability. According to such views, evidence for the quality of agreement between the behavior of a repetitive event and the mathematical concept, and for the magnitude of the probability that applies (in case any does), is to be obtained by observation of some repetitions of the event, and from no other source whatsoever.

Frequentist probabilities are the long run frequencies of repeated events.

"Personalistic" or "subjectivist Bayesian" views[88]

> hold that probability measures the confidence that a particular individual has in the truth of a particular proposition, for example, the proposition that it will rain tomorrow. These views postulate that the individual concerned is in some ways "reasonable", but they so not deny the possibility that two reasonable individuals faced with the same evidence may have different degrees of confidence in the truth of the same proposition.

Personalistic probabilities are measured by agents' willingness to bet.

Finally, "necessary", "logical" or "objective Bayesian" views[89]

---

87. von Mises (1957).

88. Borel (1924), Ramsey (1926), de Finetti (1974), Savage (1954).

89. Bayes (1763), Keynes (1921), Carnap (1950).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

> hold that probability measures the extent to which one set of propositions, out of logical necessity and apart from human opinion, confirms the truth of another. They are generally regarded by their holders as extensions of logic, which tells when one set of propositions necessitates the truth of another.

It is not clear how to measure logical probabilities.

Because probabilities in the frequentist and logical views are independent of any individual, they describe a possible origin for a common prior. Because probabilities in the personalistic view are a function of individual choices, there can be no presumption that they must be the same (but see below).

Economic theorists are presumably not interested in the philosophical meaning of probability, but in the role of probabilities in rational decision making. Natural (if stringent) axioms on choice (independence, continuity, asymmetry, negative transitivity) on "lotteries" (specified in terms of some exogenous probabilities) ensure that there exists a [von Neumann-Morgenstern] utility function such that choices maximize expected utility under the exogenous probabilities. Since personalist probabilities are defined in terms of decision theory, these exogenous probabilities must be understood to be logical or frequentist probabilities.

Savage showed that expected utility maximization (with endogenous "personal" probabilities) is implied by analogous axioms on choices over state-contingent lotteries, with no mention of any exogenous probabilities. Personal probabilities are defined in terms of this result. Thus they are a measure of peoples' willingness to bet.

Most peoples' intuition is that "rational" agents should agree on the probability of some events, but not necessarily on others. One way to combine approaches is to assume the Savage axioms but also that agents' choices contingent on events where objective (i.e. logical or frequentist) probabilities exist must conform to the von-Neumann Morgenstern axioms [e.g. Anscombe and Aumann (1963)]. But if objective probabilities are derived from a frequentist

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

view, a more compelling reconciliation uses de Finetti's (1937) exchangeability result. Suppose instead in the Savage framework, there are repeated trials of some random process; and suppose that agents are indifferent between receiving a dollar conditional on some sequence of outcomes and receiving a dollar conditional on any other sequence of outcomes with the same total number of outcomes of each type (i.e. regardless of the order); then if there exist limiting frequencies of different types of outcomes, then each agent conditional beliefs converge to those limiting relative frequencies.

I have one observation to make about the meaning of probabilities in decision theory. Either (logical or frequentist) probabilities exist logically prior to the agent's decision problem, in which case we can think of them as exogenous to the decision problem. Or else they do not, in which case they are in some sense endogenous to the decision problem, representing parameters of consistent choice. By applying mixed objective-subjective frameworks or exchangeability, we are imposing additional assumptions which ensure that the endogenous personal probabilities must be equal, in some cases, to logically prior, exogenous probabilities i.e. probabilities determined by something other than consistent choice. I have never heard of any reason why we would expect the endogenous probabilities to be equal unless there exist such exogenous probabilities that we expect them to be equal to[90]. Thus Aumann (1987) assumes "as in Savage (1954), that each player has a subjective probability distribution over the set of all states of the world" (p2) and then imposes the common prior assumption. My point is that it makes little sense to do so unless the common prior existed logically prior to the decision problem[91]. Thus Aumann argues that:

---

90. The only way to do so would to be impose some addition axiom on the choices of a group of agents e.g. that uninformed agents do not bet with each other. But this is dangerously close to assuming the conclusion.

91. I argue below why the existence of such a logically prior probability is a problem for his argument.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

> At one point, Savage wrote that "the personalistic view incorporates all.. criteria for reasonableness in judgement known to me, and.. when any criteria that may have been overlooked are brought forward, they will be welcomed...". It's just possible that he would have welcomed the CPA.

But if my argument above is correct, by welcoming the CPA, Savage would be recognising the need for some other (non-personalistic) view of probabilities to account for the logically prior exogenous probability that agents' personalistic priors turn out to equal. But he explicitly rejects alternative meanings of probability[92].

Having looked at the meaning of probability in general, and in the context of the decision theory, I move on to the meaning of probability in economic models. Models may involve uncertain events that are exogenous to the model (the actions of the agents whose behavior is modelled do not effect likelihood of these events), while other events are endogenous. For purposes of economic modelling, the weather may be thought of as exogenous, while agents' beliefs about the actions of other agents in game theory and about prices in rational expectations equilibria are endogenously determined. The distinction will be important in the argument that follows, but notice that it is sometimes a matter of interpretation. For example, in Aumann (1974), subjectively correlated equilibria are interpreted as arising when agents condition their actions on exogenous events about which agents have different subjective beliefs, while in Aumann (1987) the same formal solution is given a different interpretation where the subjectivity concerns endogenous events (the actions of other agents).

---

92. Indeed, in the sentence after that quoted by Aumann, Savage makes clear that he does not expect any new criteria for reasonableness to eliminate differences in priors:-

> The criteria incorporated in the personalistic view do not guarantee agreement on all questions among all honest and freely communicating people, even in principle. That incompleteness, if one will call it such, does not distress me, for I think that at least some the disagreement we see around us is due neither to dishonesty, to errors in reasoning, nor to friction in communication.... (p27)

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

A crucial question, for the argument that follows, will be whether the endogeneity and exogeneity of uncertainty in the model and the endogeneity and exogeneity of probabilities (to each agent's decision making) can be thought of as independent.

1.2 The arguments

I attempt a division of arguments for the common prior assumption in economic theory into four categories. The first two relate to logical and frequentist justifications for objective probabilities. Both are arguments that rationality - and, in the frequentist case, sufficient historical evidence - ensure that agents' behavior should be consistent with the common prior assumption. I give general arguments against these claims, but also argue that these rational arguments for the common prior assumption are particularly flawed when the beliefs concern endogenous events. But ultimately, perhaps, the issue should be determined by empirical evidence. I discuss some "stylized facts" that are held to support the common prior assumption. Finally, I discuss some pragmatic arguments in favor of the common prior assumption.

"Logical" justifications

Rationality by itself entails the common prior assumption. If agents have the same information and different beliefs, one of them has made a mistake and is therefore not rational.

I will not attempt to delve into the philosophical issues in any detail, since I argue below that they are not compelling for economic theory anyway. But note that a logical view

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

requires that there exists a probability relation, giving the true probability as a function of the sum of evidence. Ramsey (1926) replying to Keynes' "Treatise on Probability" wrote that a

> ... fundamental criticism of Mr. Keynes' views... is the obvious one that there really do not seem to be any such things as the probability relations he describes. He supposes that, at any rate in certain cases, they can be perceived; but speaking for myself I feel confident that this is not true. I do not perceive them, and if I am to be persuaded that they exist it must be by argument; moveover I shrewdly suspect that others do not perceive them either, because they are able to come to so very little agreement as to which of them relates any two propositions."

Logic tells us how to update a prior given new information, but not how to choose a prior, just as logic tells how to deduce new knowledge from new information, but not how to choose axioms.

Even if it is accepted that rationality entails the common prior assumption, there is still be a massive conceptual and empirical distinction between the traditional economists' notion of "rationality as consistency" - which, together with the axioms like the "sure-thing principle", implies expected utility maximization with respect to some probability distribution - and the common prior assumption. It would be certainly still be a sensible research agenda to examine the implications of dropping the common prior assumption from otherwise rational behavior.

Whatever the merits of a logical view of probability with respect to exogenous events, it cannot possibly make sense in the context of endogenous events. Consider an example of Savage:-

> I personally consider it more probable that a Republican president will be elected in 1996 [recall this was written in 1954!] than that it will snow in Chicago in the month of May 1994. But even this late spring snow seems to me more probable than that Adolf Hilter is still alive."

The point of the example is to illustrate that it is difficult to imagine the rules for deducing the probability of these events from a set of information. But it is at least possible to imagine what information might be used: "Hilter's body was never found", "it has never

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

snowed in May in Chicago the last fifty years" etc... But suppose in game theory, we attempted to use a logical view of probability to justify agents having a common prior over the strategy space and thus Nash equilibrium[93]. We must then be able to find the information that led agents to hold those common (logical) beliefs. But in a symmetric game with symmetric Nash equilibria - say a many agent matching pennies game[94] - no evidence favoring one outcome over another exists (by assumption of symmetry). The formal problem is that assuming a common, logical, prior about endogenous events makes the logical relation (if one existed) between information and priors self-referential. In fact, the Bayesian foundations of game theory literature gives a most compelling explanation of why truly personalistic probabilities must be required for at least some problems[95].

Differences in beliefs are explained by differences in information; therefore agents with the same information must have the same beliefs.[96]

This argument is a non-sequitur. Either we assume, by the logical view, the conclusion that agents with same information must have the same beliefs. Or else we do not, in which case we think of beliefs as not explained by information at some level. It is worth considering why the argument seems more convincing than simply asserting the conclusion. I suspect it is

---

93. Aumann (1987), Brandenburger and Dekel (1987b).

94. Players independently choose Heads or Tails and each receives $1 if all choose the same, 0 otherwise. Pearce (1984) has used this example to illustrate the inadequacy of Nash equilibrium.

95. The particular problems of assuming the common prior assumption in this literature are discussed in Shin (1988), Bernheim (1987), Gul (1990).

96. The argument is often referred to as the "Harsanyi doctrine" [e.g. Aumann (1987)], following Harsanyi (1967).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

because we believe many differences in (posterior) beliefs are explained by differences in information. The argument implicitly assumes that differences in information and differences in prior beliefs are mutually exclusive explanations for differences in (posterior) beliefs. This is further discussed below.

Frequentist justifications

We are justified is assuming common priors, because past experience will have removed differences of beliefs unexplained by differences in information.

Even if there exist some events where past experience has lead to common priors (life insurance), presumably there are many other events of economic importance for which we have not had an opportunity to learn (the collapse of communism in Eastern Europe, Iraqi invasions of Kuwait). We would like to examine the qualitative and empirical differences between the kinds of uncertainty.

The prevalence of finite parameter stationary economic models should not be confused with reality. Bayesian statisticians have found that "even in cases where well-behaved and non-degenerate priors have actually been formulated (over infinite dimensional spaces), there is an additional problem that a nonparametric Bayesian approach may be inconsistent: the posterior distribution need not converge on the "true" parameter vector as the sample size tends to infinite"[97]. Bewley (1988) shows that "not every stochastic law can be learned from the data it generates"; he gives a simple example of such an "undiscoverable" law and argues that "there

97. Rust (1988).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

seems to be no sound reason for believing that economic time series necessarily have discoverable laws, the popularity of time series methods notwithstanding".

Again, there is an additional argument against the common prior assumption relating exclusively to endogenous events. When you are learning about endogenous events, what you learn is also endogenously determined, and beliefs may not converge for this reason only[98]. Pearce (1984) makes the related point that in game theory, agents' learning from experience (of the same opponent) is naturally modelled by repeated games and should not be used as a justification for Nash equilibrium (where learning is assumed to have converged).

Finally, even if past experience does imply convergence to a common prior, there are good reasons for modelling that convergence explicitly and not assuming the CPA. Is it the case that as priors get arbitrarily close, economic outcomes get arbitrarily close to those under the CPA? Results of chapter II suggest otherwise. For both endogenous and exogenous events, the process of convergence of beliefs to the common prior assumption may be of importance.

Empirical justifications

If agents' different (posterior) beliefs were explained by differences in prior beliefs, then agents would be prepared to bet with each other on the basis of every difference in (posterior) beliefs. But betting seems to be restricted to special environments (horse racing, stock markets, etc..), suggesting that there cannot in general be differences in prior beliefs.

Bewley's (1986) Knightian decision theory aims to explain this, by allowing incompleteness and inertia in agents' choices. He notes that an alternative explanation would

98. Grandmont and Laroque (1990), Fudenberg and Kreps (1988).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

be the presence of asymmetric information, in which case it may well make sense not to bet despite the different prior beliefs and thus the ex ante subjective gains from trade. Chapter II showed conditions on information and beliefs to preclude betting. However different agents' prior beliefs, there is always some information sufficient to preclude betting. In economic environments, the relevant private information includes not only information about economic variables, but also about agents' utility functions.

If two agents had different prior beliefs, then it would be possible for a third agent to make certain positive gains brokering bets between them, so that there exists a money pump.

This argument is based on a confusion. One justification for subjective probabilities is that they are necessary and sufficient to ensure no Dutch book, i.e. no set of bets that guarantee a positive profit to the bookmaker (in all states of the world). The most compelling argument for the no Dutch book argument is that the victim of the Dutch book should realize that he has a guaranteed loss from the set of bets. This does not occur when many agents have different priors. It is true that if all agents are risk-neutral, a middleman can make infinite profits by brokering bets. On the other hand, if you prefer apples to bananas and I prefer bananas to apples and our utility functions are both linear in apples and bananas, then a middleman can also make infinite gains by brokering trades. Either lower bounds on consumption or strict convexity of preferences will preclude both kinds of "money pump". A more modest argument would be that differences of prior beliefs imply large, but finite, gains to be made from brokering trades in financial assets. This does not seem to contradict reality.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Differences in prior beliefs require that, in the absence of asymmetric information, agents do not want revise their beliefs when learning others' beliefs. Yet they do.

This relates back to the "Harsanyi doctrine" as discussed above, but now we can give it an "empirical" twist. People do actually want to revise their beliefs on learning others' beliefs, so differences in beliefs must after all be explained by differences in information. I would suggest that the problem here is with the notion of "information". If you are an "expert" on U.S. politics, I may well want to alter my beliefs, on learning your beliefs, about who will win the 1992 presidential election, even if you have already told me all relevant "information" in the usual sense. Your "expert prior" is an information signal for me, so we must interpret it as information. An operational distinction between priors and information might go as follows: any opinion of yours that affects my beliefs has information content. We have unravelled beliefs far enough only when my "prior" beliefs are not changed by being told your "prior" beliefs[99]. The key point is that while some peoples apparent prior beliefs implicitly convey information some of the time, it does not mean that all differences in beliefs are explained or can be modelled by differences in information.

Pragmatic justifications

Normative analysis becomes impossible without the common prior assumption

In particular, bets which are zero-sum ex post, lead to ex ante gains from trade and thus Pareto-improvements. Thus although there is no problem formally with welfare analysis, this

99. See a related discussion in Varian (1989) of the "credibility" of other agents' beliefs.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

seems paradoxical to some, antithetical to the whole notion of Pareto-improvement. One approach is to assess welfare gains with respect to some "true" probability distribution[100]. But it merely reflects a more general unwillingness to take the supposed logic of welfare analysis seriously. Consider an example: suppose workers are observed to accept dangerous jobs with apparently very small extra compensation for the risk, perhaps because of cognitive dissonance[101], despite being fully informed of all known evidence. A welfare economist offers to mediate between a paternalist who would like to interfere - say, by guaranteeing them higher wages - and a libertarian who would like to allow workers their free informed choices. The welfare economist discovers that half the workers accept the job because - despite the evidence - their consistent choices reveal a surprisingly low probability of accidents. The other half accept the job despite a more realistic (i.e. higher) assessment of the probability of accidents: their consistent choices reveal a surprisingly low disutility from having a maiming accident. The welfare economist recommends interfering in the choice of the first group - for they just made mistakes assessing the probability; but not interfering in the choice of the latter group, because welfare economists do not make value judgements about peoples' revealed preferences. Neither the paternalist nor the libertarian are happy.

100. Broome (1989), Hammond (1983); see also Machina (1989).

101. See Akerlof and Dickens (1982).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

If probabilities are just parameters of the utility function, we do not need explicit models of heterogeneous prior beliefs. Existing models without expected utility maximization encompass heterogeneous prior beliefs.

Thus general equilibrium theory (which does not use expected utility maximization) is entirely consistent with heterogeneous prior beliefs. But the result is that we do not know whether the large proportion of microeconomic theory which assumes, in one form or another, expected utility maximization with common priors, is being driven by expected utility maximization or the common prior assumption. In any case, there are many contexts when it is important for the interpretation to distinguish between beliefs and preferences over certain outcomes.

People make mistakes updating beliefs on the arrival of new information; apparent heterogeneous prior results are best modelled and understood as a consequence of information processing errors.[102]

Information processing errors are an important subject of study in their own right[103]. There is also a formal sense in which it can be assumed that at the beginning of time agents misinterpreted a signal and this is what gave them heterogeneous prior beliefs [Brandenburger, Dekel and Geanakoplos (1989)]. I will argue in section 2.2 below that there are reasons for studying both interpretations.

102. Aumann (1976), Geanakoplos (1988).

103. see Geanakoplos (1988) for discussion of work in this area.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Anything can happen with heterogeneous prior beliefs

Differences in prior beliefs are a component of differences in utilities. It is no doubt a valid criticism of neoclassical economics that anything can happen when people have different utility functions[104]; yet, for one reason or another, it does not seem to have brought neoclassical economics grinding to a halt. It is no doubt a valid criticism of information economics (and one more commonly made by economists) that anything can happen when people have different information[105]; information economics seems to be thriving. Without getting into a discussion of what makes explanations acceptable to economists, clearly there are de facto restrictions on the kinds of differences in utility and information that are found plausible or are empirically tested. There is some sense that ad hoc assumptions of differences in utilities / information / prior beliefs are not considered sufficient explanations, although how to draw the line is always a matter of controversy. But there is no difference in principle about prior beliefs, information and utility functions in this regard. To take an example that I find particularly puzzling, it is claimed that to say that some trading is explained by differences in beliefs not accounted for by differences in information is somehow illegitimate[106]. Yet a prevailing paradigm in financial economics explains trading by asymmetric information and the presence of noise traders[107]. Now the noise traders are controversial - they are in the grey

---

104. For example, Heckman and McCurdy (1988) argue that appropriate assumptions about heterogeneity of agents' preferences potentially explain evidence of disequilibrium in the labor market. See also Debreu (1974) and Sonnenschein (1973).

105. "With a little judicious selection here and there, it will turn out that the data are just barely consistent with your thesis adviser's hypothesis that money is neutral (or non-neutral) everywhere and always, modulo an information asymmetry, any old asymmetry, don't worry, you'll think of one." Solow (1976).

106. e.g. Aumann (1987)

107. Kyle (1985).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

area of "ad hoccery" outlined above. But what about the asymmetric information? The asymmetric information is no more modelled, no more observable and imposes no more restrictions on the outcomes than heterogeneity of prior beliefs[108].

1.3 Abandoning the common prior assumption

I am aware of the quixotic character of attempting to argue against what are essentially rationalizations of the common prior assumption. Economists will be persuaded that these rationalizations are just that around the time they find applications which abandon the common prior assumption which seem persuasive. Here I try and give an account of past and possible future work without the common prior assumption.

There are two key themes which I believe should guide applied work without the common prior assumption. First, the interaction between asymmetric information and heterogeneous prior beliefs is key. It is difficult to think of real world examples where differences of opinion do not reflect both asymmetry of information (so agents want to revise their beliefs in the light of others' beliefs), and also heterogeneous prior beliefs (so they can still agree to disagree)[109]. This was a key motivation for chapter II. Second, given reasonable skepticism about explanations relying on "ad hoc" differences in prior beliefs, a natural starting point for theoretical study is to contrast economic behavior in situations where agents' priors agree and situations where they differ[110]. This is analogous to Knight's (1921) between risk (where agents agree) and uncertainty (where they don't). We can make some safe

---

108. Of course, a particular form of information can be assumed with testable and interpretable results.

109. This point is made in Varian (1989).

110. See Varian (1985) and Bewley (1989)

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

generalizations about where agents are more likely to agree (life insurance, old, large and predictable stocks) and where they are less likely to (political prospects in Iraq, new, small and unpredictable stocks). With these two issues in mind, I consider three areas of research in economics and the implications of dropping the common prior assumption. I do not discuss dropping the common prior assumption in game theory which is extensively discussed elsewhere[111].

In the bargaining literature, a distinction is made between "common value" and "private value" models. One source of different valuations may be heterogeneous prior beliefs. When risk neutral agents are bargaining over financial assets or any asset with common ex post valuations, it is the sole source of different valuations.

Consider the following classroom illustration of the winner's curse: students are asked to submit sealed bids for a jar full of pennies. Rational students, the argument goes, should condition on winning as they choose their bids, and we are (rightly) confident that if they don't, the winning bid will exceed the number of pennies in the jar. Thus we are assuming that students have private information about the number of pennies in the jar (private information underlies the winner's curse argument), even though apparently all of them are equally informed - they publicly observe the jar. The private information that students have is, of course, their own beliefs about the number of pennies. Students would probably want to alter their beliefs if they were aware of other students' beliefs. Would students always want to alter their beliefs on observing other students' different beliefs? (i.e. is it impossible to agree to disagree about the number of pennies in the jar?) I suspect not. Therefore we must interpret the situation to be one where there is both asymmetric information and heterogeneous prior beliefs. Think of this as a paradigm for situations where agents may not have private information in a

111. Bernheim (1987), Gul (1990).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

conventional sense, but they form different beliefs on the basis of the same apparent public signals. It is natural to interpret those differences in beliefs as partly reflecting heterogeneous priors and partly reflecting private information. So the private values model of the bargaining can be applied more generally than its conventional interpretation.

There has been extensive work on asset pricing allowing for heterogeneous beliefs[112]. Such work has been criticized because at least some of the differences in beliefs we observe are surely explained by asymmetric information. Varian (1989) argued that differences in equilibrium beliefs in such models are those that persist after learning from other agents' behavior and equilibrium prices. In any case, it is important when looking at the role of heterogeneous beliefs in asset pricing to explicitly address the interaction with asymmetric information. Work on asset pricing shows that, with auxiliary assumptions as everywhere in economic theory, there are definite implications of heterogeneous prior beliefs, most clearly seen by comparison with the case of homogeneous prior beliefs. Harrison and Kreps (1978) showed that if risk neutral traders have heterogeneous prior beliefs, infinite endowment and can re-trade assets but cannot sell them short, then the asset price in each state is always greater than or equal to each agent's expected value of future payments from the asset, and typically (in particular, if the common prior assumption does not hold) strictly greater. The result can be seen as a critique of a "fundamental" theory of pricing and as a formalization of Keynes' notion of speculation in the *General Theory*: speculation is buying as asset for its short term capital gain, at a price higher than the value of the discounted stream of future returns. Varian (1985) showed that in an Arrow-Debreu single good economy with common preferences but heterogeneity of prior beliefs, increasing heterogeneity of beliefs decreases equilbrium asset prices if risk aversion does not decline with increasing income too rapidly.

---

112. Lintner (1969), Varian (1985). See Varian (1990) for a survey.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

The importance of the common prior assumption has been highlighted by no trade theorems: in the absence of other reasons for trade, asymmetric information does not lead to trade under the common prior assumption. Given the apparent absence of sufficient other reasons for trade to explain observed trading volumes, difference in heterogeneous prior beliefs must play a key role in explaining trading volume:-

> ... it is difficult to imagine that the volume of trade in security markets has very much to do with the modest amount of trading required to accomplish the continual and gradual portfolio balancing inherent in our current intertemporal models. It seems clear that the only way to explain the volume of trade is with a model that is at one and the same time appealingly rational and yet permits divergent and changing opinions is a fashion that is other than ad hoc....
>
> Surely there is nothing more embarrassing to an economist that the ability to explain the price in a market while being completely silent on the quantity.[113]

Chapter II gave a systematic treatment of when differences of beliefs lead to trade in the presence of asymmetric information and when the arrival of new information leads to trade. Varian (1989) extended the single risky asset framework in which Grossman studied rational expectations equilibria[114], to obtain a simple characterization of the determinants of trading volume and price changes. Trading volume is determined by parameters which reflect the heterogeneity of agents' different interpretation of new information, while price changes reflect an average of agents' different beliefs[115].

---

113. Ross (1989).

114. Grossman (1976).

115. See also Simpson (1990), Hindy (1989).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

2. Misinterpreted Information

2.1 A decision theoretic justification for agents making mistakes

It is often wrongly assumed that axiomatic subjective expected utility theory somehow implies Bayesian updating. Hacking (1967) notes that an implicit assumption - he calls it the "dynamic assumption" - is missing from most interpretations of subjective expected utility maximization[116]. This assumption would require that contingent choices equal actual choices so that beliefs after observing an event are equal to the mathematically defined conditional probability. This property is not formally entailed by anything in Savage (1972) or any static decision theory. It is quite possible for choices to satisfy Savage's axioms at each point in time, and violate the "dynamic assumption" i.e. the beliefs implicit in choices after observing some information are not equal to the appropriate conditional probabilities derived from beliefs implicit in initial choices.

But while static decision theory does not require Bayesian updating of beliefs, neither does it require a constant utility function. The results of this dissertation require the decision problem, including the utility function, to be the same before and after information is observed. What axioms on behavior would justify that combination of non-Bayesian updating of beliefs and constant utility functions? Consider a world with both events with "objectively known" probabilities (roulette wheels) and events with "unknown" probabilities (horse races)[117]. Suppose agents make the same choices through time between unconditional objective lotteries.

116. See also Kreps (1988), chapter 10, "Conditional Preference, Probability and Choice"; he calls the missing assumption the "planted axiom"; Weller (1978), Machina (1989).

117. Anscombe and Aumann (1963).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Then agents' utility functions are determined and constant, while they may "incorrectly" update their subjective beliefs about events with unknown probability[118].

2.2 Misinterpreted information or heterogeneous prior beliefs?

In the light of the formal decision theoretic equivalence between agents who make mistakes processing information and agents with heterogeneous information, discussed in chapter IV, is there any point in studying both variants on standard economic models? Brandenburger, Dekel and Geanakoplos [BDG] (1989) suggested that information processing errors are a less "ad hoc" explanation of deviations from standard models than heterogeneous prior beliefs. Aumann (1976) points to known biases such as those discussed in Kahnemann, Slovic and Tversky (1982).

One response is that when the formal equivalence is investigated in detail (as in chapter IV), results in one framework turn out to have very different (but equally intuitive) interpretation in the other. Whatever framework is "right", it makes sense to try and understand the issues arising in both interpretations.

But beyond that, there is a real operational distinction between heterogeneous prior beliefs and information processing errors. Consider a world where agents behave rationally[119] given their beliefs, have the same priors, but make mistakes processing information. Such agents would never accept a static Dutch book - a contemporaneous set of bets which together imply a sure loss - but they could be fooled by a dynamic Dutch book:- a set of bets spread

118. It need not be specified whether the "objective" beliefs about events with "known" probabilities can be assumed to be updated correctly or not.

119. in the sense of Savage (1954).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

over the period the agent receives (and misinterprets) information, which imply a sure loss. But uninformed, or identically informed agents, would never bet against each other. On the other hand, consider a world where agents behave rationally given their beliefs, interpret information correctly but have different priors. Such agents would never accept a static or dynamic Dutch book against them. But they would bet against each other when uninformed or identically informed.

3. Contributions of the Dissertation

We understand economic theory better when we understand the role that particular assumptions play in results. This dissertation can be seen as an investigation of relaxing assumptions about beliefs in well known results about no trade, no speculation, agreeing to disagree and the value of information.

But beyond that, I have argued that economic analysis can and should be extended to agents with heterogeneous prior beliefs and agents who make mistakes, and provided a framework for such analysis. The type-space and partition notation in chapter II to represent the uncertain environment clarifies the relation between results in the literature and provides a way of extending analysis beyond the standard, abstract issues discussed in the dissertation. Chapter III provided some foundational work (in section 1) on how to define the value of information when agents makes mistakes processing (and thus there is potential for ambiguity). The key properties concerning the similarity of beliefs (consistency, reconcilability and public consistency) and the relation between experiments (sufficiency in differences) can be expected to recur as key properties in applications[120].

120. As consistency already has since Harsanyi (1967).

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

References

Akerlof, G. and Dickens, W. (1982). "The Economic Consequences of Cognitive Dissonance," *American Economic Review*, **72:3**, 307-319.

Anscombe, F. and Aumann, R. (1963). "A Definition of Subjective Probability," *Annals of Mathematical Statistics*, **34**, 199-305.

Aumann, R. (1974). "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, **1**, 67-96.

Aumann, R. (1976). "Agreeing to Disagree," *The Annals of Statistics*, **4:6**, 1236-1239.

Aumann, R. (1987). "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, **55:1**, 1-18.

Aumann, R. (1989). "Notes on Interactive Epistemology".

Bacharach, M. (1985). "Some Extensions of a Claim of Aumann in an Axiomatic Model of Knowledge," *Journal of Economic Theory*, **37:1**, 167-190.

Bayes, T. (1763). "An Essay towards solving a Problem in the Doctrine of Chances," reprinted in *Biometrika* (1956), **45**, 293-315.

Bernheim, D. (1984). "Rationalizable Strategic Behavior," *Econometrica*, **52:4**, 1007-1028.

Bernheim, D. (1986). "Axiomatic Characterizations of Rational Choice in Strategic Environments," *Scandinavian Journal of Economics*, **88:3**, 473-488.

Bewley, T. (1986). "Knightian Decision Theory: Part I," Cowles Foundation Discussion Paper #807.

Bewley, T. (1987). "Knightian Decision Theory: Part II," Cowles Foundation Discussion Paper #835.

Bewley, T. (1988). "Knightian Decision Theory and Econometric Inference," Cowles Foundation Discussion Paper #868.

Bewley, T. (1989). "Market Innovation and Entrepreneurship: A Knightian View," Cowles Foundation Discussion Paper #905.

Blackwell, D. (1951). "The Comparison of Experiments" in *Proceedings, Second Berkeley Symposium on Mathematical Statistics and Probability*, 93-102. University of California Press.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Borel, E. (1924). "Apropos of a Treatise on Probability," in *Studies in Subjective Probability*, Kyburg and Smokler eds. (1st edition). John Wiley and Sons (1964).

Brandenburger, A. and E. Dekel (1985). "Hierarchies of Beliefs and Common Knowledge," Research paper, Graduate School of Business, Stanford.

Brandenburger, A. and E. Dekel (1987a). "Common Knowledge with Probability 1," *Journal of Mathematical Economics*, **17**, 237-245.

Brandenburger, A. and E. Dekel (1987b). "Rationalizability and Correlated Equilibria," *Econometrica*, **55:6**, 1391-1402.

Brandenburger, A., E. Dekel and J. Geanakoplos (1989). "Correlated Equilibrium with Generalized Information Structures," Cowles Foundation Discussion Paper #884R.

Broome, J. (1989). "Should Social Preferences be Consistent?," *Economics and Philosophy*, **5:1**, 7-18.

Carnap, R. (1950). *Logical Foundations of Probability*. Chicago: University of Chicago Press.

Cremer, J. (1982). "A simple proof of Blackwell's "Comparison of Experiments" Theorem," *Journal of Economic Theory*, **27**, 439-443.

Dasgupta, P., P. Hammond and E. Maskin (1979). "The Implementation of Social Choice Rules; Some General Results on Incentive Compatibility," *Review of Economic Studies*, **46:2**, 185-216.

de Finetti, B. (1937). "La Prevision: Ses Lois Logiques, Ses Sources Subjectives," *Annales de l'Institute Henri Poincaré*, **7**, 1-68, translated in English and reprinted in *Studies in Subjective Probability*, Kyburg and Smokler eds. (2nd edition). New York: Robert E. Krieger Publishing Co. (1980).

de Finetti, B. (1974). *Theory of Probability*. John Wiley and Sons.

De Long, J., A. Shleifer, L. Summers, and R. Waldmann (1990). "Noisy Trader Risk in Financial Markets," *Journal of Political Economy*, **98:4**, 703-738.

Debreu, G. (1974). "Excess Demand Functions," *Journal of Mathematical Economics*, **1**, 15-21.

Dubey, P., J. Geanakoplos and M. Shubik (1987). "Revelation of Information in Strategic Market Games: A Critique of Rational Expectations," *Journal of Mathematical Economics*, **16:2**, 105-138.

Ellsberg, D. (1961). "Risk, Ambiguity and the Savage Axioms," *Quarterly Journal of Economics*, **75:4**, 643-669.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Fudenberg, D. and D. Kreps (1988). "A Theory of Learning and Experimentation in Games."

Geanakoplos, J. (1988). "Common Knowledge, Bayesian Learning, and Market Speculation with Bounded Rationality," Yale University.

Geanakoplos, J. (1989). "Game Theory without Partitions, and Applications to Speculation and Consensus," Cowles Foundation Discussion Paper #914.

Geanakoplos, J. (1990). "An Introduction to General Equilibrium with Incomplete Assets Markets," *Journal of Mathematical Economics*, **19:1**, 1-38.

Geanakoplos, J. and H. Polemarchakis (1982). "We Can't Disagree Forever," *Journal of Economic Theory*, **28:1**, 192-200.

Geanakoplos, J. and H. Polemarchakis (1986). "Existence, Regularity, and Constrained Suboptimality of Competitive Allocations when Markets are Incomplete," in W. Heller, R. Starr, and D. Starrett (eds.), *Essays in Honor of Kenneth Arrow, Volume 3*. Cambridge: Cambridge University Press.

Geanakoplos, J. and J. Sebenius (1983). "Don't Bet On It: Contingent Agreements with Asymmetric Information," *Journal of the American Statistical Association*, **78**:382, 424-426.

Glosten, L. and P. Milgrom (1985). "Bid, Ask, and Transactions Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, **14:1**, 71-100.

Grandmont, J.-M. and G. Laroque (1990). "Economic Dynamics with Learning: some Instability Examples".

Green, J., and N. Stokey (1978). "Two Representations of Information Structure and their Comparisons," Stanford Institute for Mathematical Studies in the Social Sciences, Economics Series Technical Report #271.

Grossman, S. (1976). "On the Efficiency of Competitive Stock Markets where Agents have Diverse Expectations," *Journal of Finance*, **31**, 573-585.

Grossman, S. (1981). "An Introduction the Theory of Rational Expectations under Asymmetric Information" *Review of Economic Studies*, **48:4**, 541-559.

Gul, F. (1990). "Rational Strategic Behavior and the Notion of Equilibrium," Stanford University.

Hacking, I. (1967). "Slightly More Realistic Personal Probability," *Philosophy of Science*, **34**, 311-325.

Hakansson, N., J. Kunkel, and J. Ohlson (1982). "Sufficient and Necessary Conditions for Information to have Social Value in Pure Exchange," *Journal of Finance*, **37:5**, 1169-1181.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Halpern, J. (ed.) (1986). *Theoretical Aspects of Reasoning about Knowledge.* Morgan Kaufman, Los Altos, California.

Hammond, P. (1983). "Ex-post Optimality as a Dynamically Consistent Objective for Collective Choice under Uncertainty," in *Social Choice and Welfare*, edited by P. Pattanaik and M. Salles. Amsterdam: North-Holland.

Harrison, M. and D. Kreps (1978). "Speculative Investor Behaviour in a Stock Market with Heterogeneous Expectation," *Quarterly Journal of Economics*, **92:2**, 323-336.

Harsanyi, J. (1967/68). "Games with incomplete information played by "Bayesian" Players I-III," *Management Science*, **14:3**, 159-182; **14:5**, 320-334; **14:7**, 486-502.

Hartigan, J. (1983). *Bayes Theory*. New York: Springer-Verlag.

Heckman, J. and T. McCurdy (1988). "Empirical Tests of Labor-Market Equilibrium: An Evaluation," *Carnegie-Rochester Conference Series on Public Policy*, **28**, 231-258.

Hindy, A. (1989). "An Equilibrium Model of Futures Markets Dynamics," M.I.T.

Hintikka, J. (1962). *Knowledge and Belief.* Cornell University Press, Ithaca, New York.

Holmstrom, B. (1979). "Moral Hazard and Observability," *Bell Journal of Economics*, **10**, 74-91.

Holmstrom, B. and R. Myerson (1983). "Efficient and Durable Decision Rules with Incomplete Information," *Econometrica*, **51:6**, 1799-1819.

Hurwicz, L. (1973). "The Design of Mechanisms for Resource Allocation," *American Economic Review*, **63:2**, 1-30.

Kahneman, D., P. Slovic and A. Tversky (1982). *Judgement under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.

Keynes, J. (1921). *A Treatise on Probability*. London and New York: Macmillan & Co.

Kreps, D. (1988). *Notes on the Theory of Choice*. Boulder: Westview Press.

Knight, F. (1921). *Risk, Uncertainty and Profit*. Boston and New York: Hougton - Miller

Kripke, S. (1963). "Semantic analysis of modal logic," *Z. Math. Logik Grundlag. der Math.*, **9**, 67-96.

Kyle, A. (1985). "Continuous Auctions and Insider Trading," *Econometrica*, **53**, 1315-1335.

Laffont, J-J. (1989). *The Economics of Information and Uncertainty*. Boston: M.I.T. Press.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Lintner, J. (1969). "The Aggregation of Investors' Diverse Judgements and Preferences in Purely Competitive Markets," *Journal of Financial and Quantitative Analysis*, **4**, 347-400.

Machina, M. (1982). "Expected Utility Analysis without the Independence Axiom," *Econometrica*, **53:6**, 1315-1335.

Machina, M. (1989). "Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty," *Journal of Economic Literature*, **27**, 1622-1668.

Machina, M. and D. Schmeidler (1990). "A More Robust Definition of Subjective Probability," The Foerder Institute For Economic Research, Tel-Aviv University.

Mertens, J.-F. and S. Zamir (1985). "Formulation of Bayesian analysis for Games with Incomplete Information," *International Journal of Game Theory*, **14:1**, 1-29.

Milgrom, P. (1981). "An Axiomatic Characterization of Common Knowledge," *Econometrica*, **49:1**, 219-222.

Milgrom, P. and N. Stokey. (1982). "Information, Trade and Common Knowledge," *Journal of Economic Theory*, **26:1**, 17-27.

Myerson, R. (1989). *Analysis of Conflict*, chapter 10 draft.

Nielsen, L. (1984). "Aggregation of Expectations, Common Information and Market Efficiency," University of Texas at Austin.

Pearce, D. (1984). "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, **47:3**, 655-678.

Radner, R. (1979). "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices," *Econometrica*, **47:3**, 655-678.

Ramsey, F. (1926). "Truth and Probability," in *The Foundations of Mathematics and other Logical Essays*. London: Kegan Paul (1930).

Ross, S. (1989). "Intertemporal Asset Pricing," in *Theory of Valuation: Frontiers in Modern Finacnial Theory*, edited by S. Bhattacharya and G. Constantinides. Totowa, N.J.: Rowman and Littlefield.

Rubinstein, A. and A. Wolinsky (1990). "On the logic of "Agreeing to Disagree" type results," *Journal of Economic Theory*, **52:1**, 190-207.

Rubinstein, M. (1975). "Security Market Efficiency in an Arrow-Debreu Economy," *American Economic Review*, **65:5**, 812-824.

Rust (1988). "Comment on Poirier: The Subjectivist Perspective of a Spiritual Bayesian," *Journal of Economic Perspectives*, **2:1**, 145-151.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.

Samet, D. (1990). "Ignoring Ignorance and Agreeing to Disagree," *Journal of Economic Theory*, **52:1**, 190-207.

Savage, L. (1954). *The Foundations of Statistics*. John Wiley and Sons.

Shapiro, C. and J. Stiglitz (1984). "Equilibrium Unemployment as a worker discipline device," *American Economic Review*, **74:3**, 433-444.

Shin, H. (1987). "Logical Structure of Common Knowledge," Nuffield College, Oxford.

Shin, H. (1988). "Some Issues in Defining Bayes-Rationality for Games," Magdalen College, Oxford.

Shin, H. (1989). "Non-Partitional Information on Dynamic State Spaces and the Possibility of Speculation," University of Michigan.

Simpson, J. (1990). "Beliefs, Information and Trading Intensity," M.I.T.

Solow, R. (1976). "Economics. Is something missing?" in Parker, W., ed., *Economics History and the Modern Economist*. Oxford: Basil Blackwell.

Sonnenschein, H. (1973). "Do Walras' Identity and Continuity Characterize the Class of Community Excess Demand Functions?," *Journal of Economic Theory*, **6**, 345-354.

Tirole, J. (1982). "On the Possibility of Speculation under Rational Expectations," *Econometrica*, **50:5**, 1163-1181.

Varian, H. (1985). "Divergence of Opinion in Financial Markets: A Note," *Journal of Finance*, **40:1**, 309-317.

Varian, H. (1989). "Differences of Opinion in Financial Markets," in Stone, C., ed., *Financial Risk: Theory, Evidence and Implications*. Boston: Kluwer Academic Publications.

Varian, H. (1990). "Differences of Opinion," draft entry for *The New Palgrave Dictionary of Money and Finance*.

Verrecchia, R. (1981). "On the Relationship between Volume Reaction and Consensus of Investors: Implications for Tests of Information Content," *Journal of Accounting Research*, **19:1**, 271-283.

von Mises, R. (1957). *Probability, Statistics and Truth*. George Allen and Unwin Ltd.

Weller, P. (1978). "Consistent Intertemporal Decision-making Under Uncertainty," *Review of Economic Studies*, **45:2**, 263-266.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.