Economics and Philosophy, 11 (1995), 227–253.

THE COMMON PRIOR ASSUMPTION IN ECONOMIC THEORY

Stephen Morris
University of Pennsylvania.

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?[1]

1. INTRODUCTION

It has become an article of faith among economists that differences in beliefs among rational individuals must be explained by different information. Thus it is assumed that it is possible to represent any uncertain environment by some state space and a common prior probability distribution. Individuals may have different information about the true state of the world, and thus different posterior beliefs, but those posteriors are derived by updating the common prior.

Perhaps the most compelling argument against the common prior assumption is the following *reductio* argument. If individuals had

This paper incorporates material from the author's dissertation at Yale University (Morris, 1991). An earlier version of this paper was circulated as University of Cambridge Economic Theory Discussion Paper No. 199. Financial Support from an ESRC research grant (R000232865) and an EEC contract (SPESCT910057) is gratefully acknowledged. I would like to thank James Robinson for valuable references and Michael McPherson, Philippe Mongin and two anonymous referees for detailed and useful comments.

[1] Aumann (1987, pp. 12 and 13).

common prior beliefs then it would be impossible for them to publicly disagree with each other about anything, even if they started out with asymmetric information.[2] Since such public 'agreeing to disagree' among apparently rational individuals seems to be common, in economic environments as elsewhere, an assumption which rules it out is surely going to fail to explain important features of the world.

But the rationale for using the common prior assumption in economic theory is more subtle. For the purposes of this introduction, this can be divided into three strands. First, there is the claim that rationality implies the common prior assumption. I review arguments in support of this claim but conclude, on the contrary, that the common prior assumption is inconsistent with economists' usual notion of rationality.

Second, there is the methodological claim that relaxing the common prior assumption makes economic theorizing too easy. Any outcome, the argument goes, is consistent with heterogeneous prior beliefs. I will argue that while this criticism is overstated (and could be equally directed at many accepted heterogeneities in economic theorizing), it has a grain of truth. Not *any* heterogeneous prior beliefs should be acceptable as explanations. We should resort to unmodelled heterogeneities of prior beliefs only when we can imagine an origin for the differences in beliefs and we can perform comparative static exercises, comparing the predictions of heterogeneous prior models with alternative explanations.

Finally, there is the argument that it is always more productive to interpret apparent failures of the common prior assumption as a failure of rationality in information processing. We should be seeking models of bounded rationality, not rational heterogeneous prior beliefs. I will argue that while in some cases, and perhaps most cases, explicit modelling of information processing is important, in other cases we can identify reasonable differences in prior beliefs which do not have their origin in information processing errors.

The paper is organized as follows. In Section 2, there is a brief discussion of the role of the common prior assumption in economic theory. In Section 3, some preliminary material on the interpretation of probabilities in decision theoretic and economic contexts is introduced. The core of the paper is Section 4, where the amalgam of philosophical, empirical and pragmatic reasons conventionally cited in support of the common prior assumption are reviewed and – largely – rejected. Finally,

[2] Aumann (1976). Formally, if a group of individuals have common prior beliefs and there is common knowledge of their posterior beliefs about an event, then those posterior beliefs must be the same. If individuals' publicly disagree about the posterior probabilities of an event (and fail to revise them), then there is common knowledge that their posteriors are different. Geanakoplos and Polemarchakis (1982) showed public agreement must eventually occur (if individuals exchange posterior beliefs), although it may take an arbitrarily long time.

in Section 5, I discuss the circumstances where heterogeneous prior beliefs might be used – selectively – in economic analysis. I conclude by discussing a number of areas of research where models with heterogeneous prior beliefs have been successfully used to develop economic insights, and where I believe there are no reasonable alternative explanations available.

There exist some excellent discussions of the merits of the common prior assumption.[3] I make no claim to originality for the arguments in this paper. The purpose is to give a comprehensive discussion of the somewhat confusing array of arguments often cited, to argue against each in turn and to give some positive recommendations on how heterogeneous prior beliefs can be incorporated into economic analysis.

2. A BRIEF HISTORY OF THE COMMON PRIOR ASSUMPTION

The standard economists' model of choice under uncertainty and asymmetric information can be represented as follows. Suppose there is some economically relevant 'state of the world' that an individual cares about. Suppose he observes a signal which is correlated with the state of the world. We assume that the individual has some *prior beliefs* over combinations of states of the world and signals (i.e., the cross product of the set of states of the world and the set of possible signals). Contingent on observing a particular signal, his posterior beliefs over states of the world can be generated by conditioning by Bayes rule. If there are many individuals, we must assume each individual has prior beliefs over states of the world, and all possible combinations of signals observed by all individuals. Again, his posterior beliefs are generated by conditioning on his own signal.[4]

If individuals end up having different posterior beliefs about states of the world, they may have observed identical signals (and know they have observed identical signals) but have different prior beliefs over the world; or they may have common prior beliefs, but have observed different signals. For many purposes, it does not matter if differences in posterior beliefs are explained by differences in information or differences in prior beliefs. The distinction becomes critical in modelling when individuals make deductions about others' information from their actions. So the common prior assumption becomes important when we are concerned with *interactive environments* and there is the possibility of *asymmetric*

[3] Including Bernheim (1986), Aumann (1987), Varian (1989), Gul (1993).

[4] This discussion is under the assumption that individuals understand and agree on the description of a standard information system, so that *by assumption* differences in prior beliefs are the only source of disagreement. In Section 4.4, there is a discussion of the relation between differences in prior beliefs and more general models of the information system.

information. This explains why Harsanyi's revolutionary work on games with incomplete information[5] became the first work in the economics literature to address the issue head on. Harsanyi showed that a game with incomplete information can be reduced to a standard game of imperfect information with an initial move by nature, if and only if individuals could share a common prior over payoffs in some state space (i.e., have 'consistent' beliefs).[6] Aumann (1976) showed that if individuals' posterior beliefs are derived from a common prior, and there is common knowledge of those posteriors, then the posteriors must be the same. This result suggested that asymmetric information, then a new but rapidly growing topic of microeconomic research, had less explanatory power than might be thought: in the absence of differences in prior beliefs, asymmetric information could not explain common knowledge differences in posterior beliefs.

Aumann's work stimulated work on *no trade* results which establish that, in the absence of *ex ante* gains from trade, asymmetric information cannot generate trade. In particular, Sebenius and Geanakoplos (1983) – extending Aumann's argument – showed that (under the common prior assumption) it cannot be common knowledge that risk neutral individuals are prepared to bet against each other, that is, that one individual's posterior beliefs exceed another's. Milgrom and Stokey (1982) showed an analogous result in a more general setting of risk averse traders. Since no-trade results can be shown to underlie many important results in microeconomic theory, it had by now become clear that the common prior assumption was critical. Arguments which started from differences in posterior beliefs, without explicitly modelling their origins, came to be viewed with increasing suspicion.

The common prior assumption re-emerged as a critical and controversial assumption in game theory, as a link between epistemic approaches, which seek to justify outcomes on the basis of common knowledge of rationality and the structure of the game alone, and equilibrium solution concepts. The common knowledge assumptions imply only that players use strategies which survive iterative deletion of dominated strategies.[7] The assumption of a common prior generating players' beliefs about each others' actions is necessary to restrict attention to correlated equilibria, and, in the absence of correlation in beliefs, Nash equilibria.[8]

[5] Harsanyi (1967, 1968a, 1968b).

[6] Harsanyi also explained how to carry out analysis of games of incomplete information without the common prior assumption. Indeed Harsanyi's approach can easily be extended to incorporate heterogeneous prior beliefs if the differences in beliefs are incorporated into payoffs (see Myerson 1991, pp. 72 and 73).

[7] Moulin (1979). Bernheim (1984) and Pearce's (1984) notion of rationalizability is equivalent in two person games, and is more restrictive in many player games as no two players may have correlated beliefs about a third player's behavior.

[8] Aumann (1987), Bernheim (1986), Brandenburger and Dekel (1987).

3. WHAT ARE PRIOR BELIEFS?

Before discussing the role of the common prior assumption (CPA) directly, it is useful to review what *might* be meant by prior probabilities in decision theoretic and economic contexts. Savage (1954) identified three underlying interpretations of probability. 'Frequentist' views[9]

> 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. (Savage, 1954, p. 3)

Frequentist probabilities are the long run frequencies of repeated events.

'Personalistic' or 'subjectivist Bayesian' views[10]

> 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 do 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. (Savage, 1954, p. 3)

Personalistic probabilities are measured by individuals' willingness to bet.

Finally, 'necessary', 'logical' or 'objective Bayesian' views[11]

> 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. (Savage, 1954, p. 3)

It is not clear how to measure logical probabilities.

Because frequentist and logical probabilities are independent of any individual, they provide a possible origin for a common prior. In the next section, logical and frequentist views of probability will be shown to underlie certain defenses of the common prior assumption. Because

[9] von Mises (1957).
[10] Borel (1924), Ramsey (1926), de Finetti (1974), Savage (1954).
[11] Bayes (1763), Keynes (1921), Carnap (1950).

probabilities in the personalistic view are a function of individual choices, there can be no presumption that they must be the same.

Probabilities enter economics via decision theory. There are two distinct ways in which the standard expected utility representation of preferences under uncertainty can be guaranteed. Given some exogenous probabilities, some natural axioms on choice[12] on 'lotteries' (specified in terms of the 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 themselves components of individuals' preferences, the exogenous probabilities in this von-Neumann–Morgenstern approach must be understood to be logical or frequentist probabilities.

Savage, on the other hand, did not assume the existence of objective probabilities. He considered individuals choosing among acts with uncertain outcomes. Specifically, there is a set of possible states and the outcome of each act depends on the state. Now natural axioms on preferences over acts – analogous to von-Neumann–Morgenstern axioms – ensure expected utility maximization, for some 'endogenous' personal probabilities.

The three views of probability are apparently mutually exclusive. The probability of a particular event must be defined as the long run frequency of some event *or* as the willingness to bet of some individual *or* as some measure of the logical relation between evidence and that event. However, there may be circumstances under which these different notions coincide. For example, there is a common intuition that rational individuals should agree on the probability of some events – where probabilities are objectively known, say because long run frequencies have been observed (see, for example, Allais, 1953). But they may disagree about the probability of unique events for which historical frequencies are not available. There are two ways in which one might combine the objective and subjective expected utility approaches. One way is to assume the Savage axioms but also that individuals' choices contingent on events where objective probabilities exist must conform to the von-Neumann–Morgenstern axioms (e.g., Anscombe and Aumann, 1963). De Finetti's (1937) exchangeability result provides a deeper reconciliation. Suppose in the Savage framework, there are repeated trials of some random process; and suppose that individuals 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); if there exist limiting frequencies of different types of outcomes, and individuals put strictly positive probability on the truth, then each

[12] Independence, continuity, asymmetry and negative transitivity.

individual's conditional beliefs converge to those limiting relative frequencies.

Before discussing the common prior assumption in economic analysis, I want to highlight a couple of issues which hinge crucially on the interpretive issues raised in this section. First, observe that either (logical or frequentist) probabilities exist logically prior to the individual'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 that is, probabilities determined by something other than consistent choice. But there is surely no reason why we would expect the endogenous probabilities to be equal unless there exist such exogenous probabilities to which we expect them to be equal or at least approximately equal.[13] 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' (p. 2) 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.[14] Aumann argues that:

> 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. (1987, p. 6)

But if my argument above is correct, by welcoming the CPA, Savage would be recognizing the need for some other (non-personalistic) view of probabilities to account for the logically prior exogenous probability that individuals' personalistic priors turn out to equal. But he explicitly rejects alternative interpretations of probability.[15]

[13] The only way to do so would be to impose some additional axiom on the choices of a group of agents, e.g., that uninformed agents do not bet with each other (see Nau and McCardle, 1990). But this is dangerously close to assuming the economically relevant conclusion.

[14] I argue below why the existence of such a logically prior probability is a problem for his argument.

[15] Indeed 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

This question – of whether probabilities are exogenously given to the decision maker, or are in some sense endogenous – is crucial in economic analysis. Models may involve uncertain events that are exogenous to the model (the actions of the individuals whose behavior is modelled do not affect the likelihood of these events), while other events are endogenous. For purposes of economic modelling, the weather may be thought of as exogenous, while individuals' beliefs about the actions of other individuals in game theory and about prices in rational expectations equilibria are endogenously determined. There are special problems defending the common prior assumption in a model where beliefs are endogenously determined.

4. JUSTIFYING THE COMMON PRIOR ASSUMPTION

Many arguments of different types are now part of the economics profession's folklore in support of the common prior assumption. It will be useful to try and categorize them into four broad categories, and present responses to each in turn. The first two categories relate to the two objective probability theories – logical and frequentist – discussed above. Both are arguments that rationality – and, in the frequentist case, sufficient historical evidence – ensure that individuals' behavior should be consistent with the common prior assumption. I give general arguments against these claims, but also argue that these objective arguments for the common prior assumption are particularly flawed when the beliefs concern endogenous events.

But the decisive arguments for economists may have less to do with the philosophy of probability, and more to do with the methodology of economic modelling. In Section 4.3 I discuss some stylized facts which are held to support the common prior assumption. Finally, in Section 4.4, I discuss some of the pragmatic arguments in favor of the common prior assumption.

4.1 Logical Justifications

Rationality by itself entails the common prior assumption. If individuals have the same information and different beliefs, one of them has made a mistake and is therefore not rational.

(a) The logical view of probability has been largely discredited in the philosophical literature (see Levi, 1980); it is beyond the scope of this paper to go into this debate in detail. It will suffice to mention Frank Ramsey's famous critique[16] of Keynes's version of the logical view of

> of the disagreement we see around us is due neither to dishonesty, to errors in reasoning, nor to friction in communication ... (1954, pp. 67—8)

[16] Ramsey (1926).

probability in his *Treatise on Probability*.[17] Recall that a logical view of probability requires that there exists a relation connecting the evidence concerning an event with its probability. Ramsey 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. (1926, pp. 65—6)

At best, logic tells us how to update a prior given new information, but not how to choose a prior, just as logic tells us how to deduce new knowledge from new information, but not how to choose axioms.

(b) Even if it were accepted that rationality entailed the common prior assumption, there is still a vast conceptual distinction between the traditional economists' notion of 'rationality as consistency' – which, together with axioms like the 'sure-thing principle', implies expected utility maximization with respect to some probability distribution – and the common prior assumption. It would certainly still be a sensible research agenda to examine the implications of dropping the common prior assumption from otherwise rational behavior.

(c) 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 published in the 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 Hitler is still alive. (1954, p. 27)

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: 'Hitler's body was never found', 'it has never snowed in May in Chicago in the last fifty years', etc. But suppose in game theory, we attempted to use a logical view of probability to justify individuals having a common prior over the strategy space and thus Nash equilibrium (as discussed in Section 2). We must then be able to find the

[17] Keynes (1921). Keynes's approach allowed qualitative probability judgments ('event A is more likely than event B') which could not be quantified; but it is subject to many of the same criticisms as a quantitative logical view.

information that led individuals to hold those common (logical) beliefs. But consider a symmetric game with multiple symmetric Nash equilibria. 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.[18]

Differences in beliefs are explained by differences in information; therefore individuals with the same information must have the same beliefs.

Thus Harsanyi (1968b, p. 497) writes that

> any inconsistency among the various players' (posterior) probability distributions is always a result of discrepancies among the (prior) probability distributions used by the different players. On the other hand, these discrepancies among the (prior) probability distributions will themselves often admit of explanation in terms of *differences in the information* available to different players.

It is tempting to go one step further and argue that individuals with the same information must have the same beliefs.[19] But this argument is a non sequitur. Either we assume, by the logical view, the conclusion that individuals with the 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.

The argument has its origins in a thought experiment of Harsanyi (1967, pp. 163–164). Suppose a group of individuals have different posterior beliefs about some event. Call these their 0th order beliefs. Learning another individual's posterior beliefs may well cause some individual to revise his beliefs. Taking a Bayesian view, we might think of each individual having a joint distribution over both the original event and also the beliefs of other individuals. Call these first order beliefs. We can imagine extending this construction *ad infinitum*, with (n+1)th order beliefs incorporating beliefs about the nth order beliefs of others. A number of authors have formalized this infinite hierarchy of beliefs.[20]

Yet despite the fact that the hierarchy can be shown to be mathematically well-behaved – in the sense that the infinite hierarchy of

[18] See Bernheim (1986).

[19] This argument, 'all differences in beliefs are explained by differences in information' is often referred to as the 'Harsanyi doctrine'.

[20] Böge and Eisele (1979), Mertens and Zamir (1985) and Brandenburger and Dekel (1993).

beliefs exhausts the uncertainty of each individual about others – the construction lends no support to the common prior assumption. Certain differences in 0th order beliefs may indeed turn out to be explained by differences in information in a precise sense: that is, it is possible to identify a subset of the universal beliefs space – the limit space of the infinite hierarchy of beliefs – where (1) everyone knows they are within that subset and (2) their beliefs could have been derived from a common prior. But there is no guarantee that this is the case. Nyarko (1991a) has shown that the set of states where the common prior assumption holds in the universal beliefs space are in some sense non-generic. Gul (1993) has emphasized that the hierarchical construction of beliefs is concerned with posterior beliefs alone and suggests that confusion arises because we introduce the construct of prior beliefs which are not really meaningful.

In conclusion, the argument implicitly assumes that differences in information and differences in prior beliefs are mutually exclusive explanations for differences in (posterior) beliefs. But they can equally well be seen as complementary explanations. This is further discussed in Section 4.3 below.

4.2 Frequentist Justifications

We are justified in assuming common priors, because past experience will have removed differences of beliefs unexplained by differences in information.

Thus, as discussed in Section 2, a common prior implies, together with other assumptions, Nash equilibrium behavior. Kalai and Lehrer (1993a) provided a formalization of the old alternative justification that learning must lead to convergence to Nash equilibrium. I'll give three types of reasons why this type of argument does not justify assuming the common prior assumption in general.

(a) *Consistency* results in Bayesian statistics do give sufficient conditions for posterior beliefs about the parameters of some data generating process to converge a point mass on the true parameters. However, those conditions are very strong, even in single person decision problems where the observations are independent of the decision maker's actions. In any case, consistency does not guarantee either correct conditional beliefs about the future or common conditional beliefs about the future. Such results use mutual absolute continuity conditions on individuals' priors: these require that all individuals agree on which events should be assigned positive probability. When applied to infinite sequences of observations, this requirement seems scarcely weaker than assuming the common prior assumption.[21]

[21] Thus Kalai and Lehrer (1993a) rely on a mutual absolute continuity assumption. See Blume and Easley (1993) for a survey of rational learning and well-behaved examples where convergence of posterior predictions does not occur.

Kurz (1994b) makes a related point in a frequentist setting. Suppose that individuals do not know structural relations in an economic system but compute beliefs from relative frequencies in the data. Suppose the system is stable in the sense that empirical frequencies always converge. A belief is said to be a rational belief if it generates the same long run empirical frequencies as the data. If the system is not stationary, there will typically be many different rational beliefs.

(b) 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 alone. A canonical example is the multi-armed bandit problem (see Rothschild, 1974). Consider a gambler choosing repeatedly between two slot machines in a casino, one with a known probability of payout, one with an unknown probability. The optimal strategy may allow experimentation with the unknown machine. But it will also require that the unknown machine be abandoned forever, after some finite number of experiments, if it does not perform well. Thus the gambler never learns the true probability of success on that machine. Learning may cease (and so differences in beliefs persist) because the decision to learn is endogenous and costly. Models of 'conjectural',[22] 'self confirming'[23] and 'subjective'[24] equilibria also have this property. Further examples where this issue is important are discussed in Section 5.4.

(c) Even if in the limit learning did remove differences in beliefs unexplained by differences in information, we presumably do not live in a world where learning has ceased. There are two exercises (without the common prior assumption) which might be interesting in the light of this. First, by analogy with Knight (1921), we can compare 'risky' situations where learning has ceased and therefore past experience has led to common priors and 'uncertain' situations where learning has not ceased, and thus there is still scope for different beliefs on the basis of the same information. We would like to examine the qualitative and quantitative differences between such 'risk' and 'uncertainty'. Second, 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? For both endogenous and exogenous events, the process of convergence of beliefs to the common prior assumption may be important. Applications of this type are discussed in Section 5.

[22] Hahn (1977).
[23] Fudenberg and Levine (1993).
[24] Kalai and Lehrer (1993b).

4.3 Empirical Justifications

If differences of posterior beliefs were explained by differences in prior beliefs, then individuals 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.

The argument implicitly assumes that information is symmetric. In the presence of asymmetric information, it may well make sense not to bet despite the different prior beliefs and thus the *ex ante* subjective gains from trade.[25] Morris (1993a) shows that individuals' willingness to bet will exhibit a bid ask spread property in the presence of heterogeneous prior beliefs and asymmetric information. Thus an individual will be prepared to accept a bet for an event at sufficiently favorable odds, will be prepared to bet against at sufficiently favorable odds, but there will exist a range of intermediate odds where he is not prepared to bet either way. However different individuals' 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 individuals' utility functions.

If two individuals had different prior beliefs, then it would be possible for a third individual to make certain positive profits brokering bets between them, so that there exists a money pump. But we don't observe money pumps.

This argument is based on a confusion. One justification for subjective probabilities is that they are necessary and sufficient to ensure no Dutch book, so that there is no set of bets which guarantees a positive profit to the bookmaker in all states of the world. The most compelling argument for no Dutch book is that the victim of the Dutch book should realize that he has a guaranteed loss from the set of bets.[26] This does not occur when many individuals have different priors. It is true that if all individuals 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 conclusion would be that differences of prior beliefs imply large, but finite, gains to be made

[25] Bewley (1986) suggests but rejects this explanation (in favor of a 'Knightian decision theory').

[26] Schick (1986) has questioned whether, even in the single person case, a no Dutch book argument provides any independent rationale for subjective probabilities beyond the normative appeal of transitivity.

from brokering trades in financial assets. This does not seem to contradict reality.[27]

Differences in prior beliefs require that, in the absence of asymmetric information, individuals do not want to revise their beliefs on learning others' beliefs. Yet they do.

This relates back to the 'Harsanyi doctrine' as discussed in Section 4.1, 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. 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 1996 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.[28] 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 by differences in information.

4.4 Pragmatic Justifications

Normative analysis becomes impossible without the common prior assumption.

When a group of risk averse individuals insure each other, a trade which is zero-sum *ex post* leads to *ex ante* gains from trade and thus an (*ex ante*) Pareto improvement. The same principle is at work when a group of risk-neutral individuals with different prior beliefs betting with each other achieve an *ex ante* Pareto improvement despite the absence of *ex post* gains from trade. Yet the latter seems more conceptually problematic to some.

A formal reason for being especially concerned about differences in prior beliefs has been identified by Hammond (1983), Broome (1989) and Mongin (1993). These authors examine when an *ex ante* application of the Pareto criterion can be reconciled with maximizing a Bayesian social welfare function. For sufficiently diverse utility functions, a common

[27] Nau and McCardle (1991) study whether 'no arbitrage' conditions can be used as a primitive assumption in economic models, replacing untestable assumptions about preferences and beliefs. But (consistent with the argument above) they note that if risk neutrality is not assumed, beliefs and utility functions are not separable, so it is not possible to deduce the common prior assumption from 'no arbitrage'.

[28] See a related discussion in Varian (1989) of the 'credibility' of individuals' beliefs.

prior is required for the reconciliation. Thus dropping the common prior assumption means that a Bayesian 'social planner' may never choose an *ex ante* efficient allocation.

There are two obvious responses to such results. One is that the idea of representing welfare via a Bayesian social planner will not work in such a context. An alternative view is that the 'right' welfare notion requires a Bayesian social planner and that *ex ante* efficiency should be abandoned.[29] But in either case, dropping the common prior assumption at most complicates welfare analysis: it certainly does not make it impossible.

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 as a result, we do not know whether the large proportion of micro-economic 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 beliefs are best modelled and understood as a consequence of information processing errors.[30]

Information processing errors are an important subject of study in their own right.[31] Brandenburger, Dekel and Geanakoplos (1992) showed that there is a formal sense in which heterogeneous prior beliefs can be interpreted as the outcome of information processing errors. Thus individuals may have misinterpreted a signal at the beginning of time,

[29] This requires that an individual's own prior is ignored in making welfare judgements. It could be argued that if an individual's prior is to be ignored, why not also ignore measures of his utility? Consider an example: suppose a group of workers are observed to accept dangerous jobs with apparently very small extra compensation for the risk, perhaps because of cognitive dissonance (Akerlof and Dickens, 1982), despite being fully informed of all known evidence. Detailed investigation reveals that half the workers accept the job because they assess a probability of accidents well below the historic empirical frequency. The other half accept the job despite a subjective probability of accidents equal to the historic empirical frequency: they reveal a surprisingly low disutility from having a maiming accident. It could be argued that using the first group of workers' subjective probabilities in welfare analysis is no more reasonable or unreasonable than using the second group of workers' 'unrealistic' utility functions.

[30] Aumann (1976), Geanakoplos (1989).

[31] See Rubinstein (1993) for a recent attempt to derive economic implications from information processing mistakes in a simple model.

and this is what gave them heterogeneous prior beliefs. Because of this formal connection,[32] we can always interpret one set of results in terms of the other. My own conjecture is that while many, perhaps most, improvements on our current models of rationality will come from models of bounded rationality, there are some situations where we have a very clear idea that there exist heterogeneous prior beliefs that have nothing to do with information processing errors in anything other than a tautological sense. In those cases, it will be more insightful to take the heterogeneous prior beliefs as primitive and not attempt to reduce them to information processing errors. Some of the examples I have in mind are discussed in Section 5.

Anything can happen with heterogeneous prior beliefs.

Differences in prior beliefs are a component of differences in utility functions. It is no doubt a valid criticism of neoclassical economics that anything can happen when people have different utility functions;[33] 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;[34] information economics seems to be thriving. Without getting into a detailed discussion of what makes explanations acceptable to economists, clearly there are *de facto* restrictions on the kinds of differences in utility functions and information that are found plausible or are empirically tested. There is some notion that *ad hoc* differences in utility functions, information and prior beliefs are not 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.[35] To take one puzzling example, it is claimed that it is somehow illegitimate to say that some trading is explained by differences in beliefs not accounted for by differences in information. Yet

[32] Explored further in Morris (1991), Chapter 4.

[33] For example, Heckman and McCurdy (1988) argue that many apparent disequilibrium features of labor markets are consistent with full market clearing, for appropriate assumptions about the heterogeneity of agents' preferences. Sonnenschein (1973) and Debreu (1974) showed that any aggregate demand behavior, consistent with basic accounting identities, could be generated by some collection of rational agents with heterogeneous utility functions.

[34] '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).

[35] Philippe Mongin has pointed out to me that a *practical* difference is that, for utility functions, there exist common properties of utility functions (monotonicity, concavity, etc) which economists have agreed to assume and which do not imply common utility functions. For prior beliefs, there does not exist a comparable middle ground.

a prevailing paradigm in financial economics explains trading by asymmetric information and the presence of noise traders (discussed further in Section 5.3). Now the noise traders are controversial – they are in the grey area of 'ad hoccery' outlined above. But what about the asymmetric information? The asymmetric information is no more modelled, no more observable and (in the presence of appropriate noise traders) may impose no more restrictions on the outcomes than heterogeneity of prior beliefs.[36]

5. ECONOMICS WITHOUT THE COMMON PRIOR ASSUMPTION

'Agreeing to disagree' – public differences in beliefs – seems to be an important aspect of the world around us and an important factor in understanding the economy. What is the origin of the differences in prior beliefs? Presumably they are not different in principle from the origin of differences in utility functions or ability of agents in economic models. The profession is typically comfortable in treating the latter as exogenous. Are there reasons why it is less reasonable to treat differences in prior beliefs as exogenous (i.e., not explained by the model) and heterogeneous? The discussion of the previous section suggests there is not. Moreover, to the extent that prior beliefs are restricted endogenously by rationality or other assumptions, there is every reason to think there will still be heterogeneity.

The conclusion of the previous section is that none of the arguments in support of the common prior assumption is compelling. But this claim will clearly be in vain if it is not possible to point to economic analysis which explicitly rejects the common prior assumption, but nonetheless generates original insights, refutable predictions, and interesting research agendas. That is the purpose of this section. Before doing so, however, it will be useful to summarize some useful lessons to be learnt from the criticisms of heterogeneous prior beliefs, and to lay down some guidelines for when an explanation relying on heterogeneous prior beliefs should be acceptable.

I draw two main conclusions from the discussion of the previous section. First, while rational learning does not justify the common prior assumption, it does help us identify situations where heterogeneity of prior beliefs will be important. Second, the most compelling arguments for the common prior assumption are methodological ones which focus on the degrees of freedom in modelling introduced by heterogeneous prior beliefs and the possibility that heterogeneity of prior beliefs is substituting for bounded rationality in information processing. In the

[36] Of course, a particular form of information may be assumed with testable and interpretable results.

light of these, I would propose a number of common sense preconditions for explanations involving heterogeneous prior beliefs.

Expect greater heterogeneity of prior beliefs when there has been little opportunity for learning.

In other words, we should always be able to carry out a comparative static exercise. Suppose we explain some phenomenon by heterogeneity of prior beliefs, because of a lack of learning opportunities. Then we should be able to verify that that phenomenon does not occur when learning opportunities exist.

Apparent heterogeneities which are associated with updating beliefs are best analyzed as bounded rationality.

For example, experimental psychologists and others have identified many systematic 'biases' in the way beliefs are updated in the light of new information.[37] It would be possible to interpret such systematic biases as a consequence of different prior beliefs on some larger state spaces. Such an interpretation would be misleading. But if the heterogeneity is not associated with some updating process, it may be equally forced and misleading to interpret the heterogeneity as a consequence of bounded rationality.

The assumption of heterogeneous prior beliefs should have some justification independent of the phenomenon being explained.

An analogy is useful. Economists do not feel any particular need to justify heterogeneity of consumers' taste in manufactured goods. Presumably there is an explanation for this at some level, but we feel happy studying the implications for market structure of heterogeneous consumer tastes, without explicitly modelling that heterogeneity.[38] On the other hand, the thesis that the northern hemisphere is richer than the southern hemisphere because there is a lower demand for leisure in the northern hemisphere is viewed with justified suspicion. The differences in tastes assumed are tailored to the phenomenon being explained, and do not seem to have any independent rationale. Similarly, it is not very useful to observe that an asset is traded because the buyer attaches higher probability to a price increase than the seller does. But below I will discuss heterogeneities of beliefs which have more independent justification.

In the next four sections, I review work that allows for heterogeneous prior beliefs that I believe satisfies these criteria.

[37] See, for example, Kahneman, Slovic and Tversky (1982).

[38] A more extreme view is that differences in tastes are *never* an acceptable explanation, e.g., Becker and Stigler (1977).

5.1 Game Theory

It was argued in Sections 4.1 and 4.2 that when individuals have beliefs about actions which are endogenous to the model, as in game theory, objective probabilities cannot make sense, so one set of arguments for the common prior assumption is automatically ruled out. One line of research has focused on those special situations when it is possible to justify equilibrium solution concepts (e.g., Nash) on the basis of iterated deletion of dominated strategies or rationalizability alone, so that without implicitly or explicitly assuming the common prior assumption, the common prior assumption is shown to hold, and equilibrium behavior results.[39]

Such results do not explore the consequences of selecting a non-Nash rationalizable outcome in preference to a Nash outcome, and thus violating the common prior assumption. Roth (1992) has considered the phenomenon of predatory pricing. The empirical puzzle is that firms enter markets, encounter predatory pricing, and then withdraw from the market. In an equilibrium model (i.e., under the common prior assumption), entrants would anticipate the incumbent's predatory strategy and not enter in the first place. Roth characterizes the relationship between rationalizable strategies and the conjectures which support them, and shows how further restrictions on beliefs (beyond the assumption of common knowledge of rationality, but not assuming common priors) serve to further characterize possible patterns of play. The empirical puzzle of entrants withdrawing in the face of predatory pricing is intuitively explained as a consequence of strategic uncertainty.

Contrast this explanation with standard equilibrium models. In the light of the profession's distaste for differences in prior beliefs, and taste for asymmetric information, there are a plethora of asymmetric information stories (reviewed by Roth) accounting for the empirical puzzle. But what is it about assuming, say, uncertainty about other firms' costs that is less *ad hoc* than assuming new entrants are uncertain of how the incumbent will behave (consistent with common knowledge of rationality)?[40]

Repeated games in general represent a setting where equilibrium solution concepts, and the perfect foresight common prior they imply, seem to strain credulity in an especially stark way. One response has been to consider evolutionary models of play in repeated games.[41] But evolutionary models typically assume extreme myopia on the part of

[39] Moulin (1979), Milgrom and Roberts (1990), Guesnerie (1992), Carlsson and van Damme (1993).

[40] See Brandenburger (1993) for a general discussion of 'strategic' versus 'structural' uncertainty in games.

[41] Foster and Young (1990), Kandori, Mailath and Rob (1992).

players. A compromise approach assumes that while each player's action is a best response to some rationalizable conjecture over future play, his conjecture, and so also his actions, change according to an evolutionary process.[42] Thus, allowing for possibly heterogeneous beliefs over the strategies of other players (consistent with common knowledge of rationality), the set of outcomes is refined not by imposing the common prior assumption (and thus equilibrium play) but by imposing other restrictions, motivated by evolutionary stories, on how beliefs and actions evolve through time.

5.2 Asset Pricing

There has been extensive work on asset pricing allowing for heterogeneous beliefs.[43] 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 individuals' 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 price of an asset is always greater than or equal to each individual's expected value of future payments from that 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's (1936) notion of speculation in the *General Theory*: speculation is buying an asset for its short term capital gain, at a price higher than the value of the discounted stream of future returns.

Morris (1993b) showed how this argument can be used to explain a paradox concerning the share prices of initial public offerings. When shares in a firm are traded publicly for the first time, we have a perfect example of an environment where there have not been learning opportunities, and thus there is scope for heterogeneous prior beliefs. Through time, the share price converges to the fundamental value, as all

[42] This line of research has been pursued by Matsui and Rob (1991), Lagunoff and Matsui (1994) and Matsui (1993).

[43] Lintner (1969). See Varian (1993) for a survey and Harris and Raviv (1993) for an important recent contribution.

traders learn the true data generating process. But the initial heterogeneity leads to a finitely lived price bubble. This is consistent with the 'hot issue' phenomenon in the pricing of initial public offerings (Ritter, 1991): initial public offerings underperform the market in their first few years of public trading, so that the opening market price must have been too high.[44]

In the above model, beliefs and thus prices converge asymptotically. But if, in addition, 'regime shifts' introduce non-stationarities into the observed data, then traders may never learn the true data generating process (see Kurz, 1994a, for a frequentist version of this argument).

Pursuing a different approach, Varian (1985) showed that in an Arrow-Debreu single good economy with common preferences but heterogeneity of prior beliefs, increasing heterogeneity of beliefs (for a given average belief) decreases equilibrium asset prices if risk aversion does not decline too rapidly with increasing income.[45]

5.3 Trading Volume

Conventional accounts of day to day movements of the stock market take it for granted that the arrival of new information explains both price changes and the volume of trade. This is assumed to be true both for private (asymmetrically distributed) information and for publicly released information – where trading volume presumably reflects a lack of consensus in the interpretation of the information.[46] Both elements have been criticized in the economics literature in the light of the common prior assumption and the no trade results it implies: consensus in the interpretation of new public information is assumed, and asymmetric information, in the absence of other reasons for trade, cannot account for the trading volume.[47]

Since the arrival of new information surely does lead to trade, various

[44] Allen, Morris and Postlewaite (1993) extend this line of argument, in the presence of asymmetric information, to show the existence of a price bubble in a rational expectations equilibrium of a finite horizon economy, where each individual knows an asset is worthless, but nonetheless buys it at a strictly positive price in the hope of re-selling it to someone else. This result does not depend on heterogeneous prior beliefs (differences in risk aversion and endowments will suffice), but the most intuitive motivation of the bubble relies on heterogeneous priors.

[45] Racetrack betting is a simple financial market where the implications of heterogeneous beliefs can be explored. A well-known anomaly is that 'long shots' – i.e., low probability winners – are a particularly bad investment at equilibrium prices. Quandt (1986) gives an explanation based on risk-loving behavior while Shin (1992) gives an explanation based on heterogeneous beliefs and asymmetric information.

[46] Beaver (1968), discussing the relation between accounting releases and trading volume, notes 'the economist's notion that volume reflects a lack of consensus regarding the price'.

[47] This is true even if differences in (posterior) beliefs are so complex that it would take a long time for individuals to converge to a common posterior. As long as there is initially common knowledge that there are no *ex ante* gains from trade, rational traders would not trade during the learning process.

tactics have been employed to bypass the no trade results in explaining financial markets. One explanation is that, if markets are incomplete, the arrival of new information will lead to trades that could not be made, contingent on that information, *ex ante*.[48] But Ross has argued that

> 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.

The most common modelling approach to bypass the no trade results is to simply assume that there exists some unmodelled population of traders who are prepared to trade. Such 'noise traders' may be thought to be irrational, or their behavior may be determined by factors exogenous to the model. The existence of noise traders is sufficient to get around no trade results, since other 'rational' traders believe with positive probability that they are trading with such noise traders. Allowing small numbers of noise traders accounts for certain qualitative features of financial markets and provides a framework for empirical analysis of trading volume.[49] But, while noise traders may represent an econometrically tractable route around no trade results, we cannot be satisfied with models incorporating noise traders until an adequate account of such behavior is given.

Heterogeneous prior beliefs among traders can certainly provide one such account. However, an interesting explanation along these lines must identify situations where heterogeneity of prior beliefs will be more or less important. We can make some safe generalizations about where individuals 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). We expect larger trading volumes to be associated with the arrival of new information which is open to a wide variety of interpretations.

Morris (1994) has identified which types of differences in prior beliefs will lead to trade (in the presence of asymmetric information). Varian (1989) extended the single asset mean variance framework in which Grossman (1976) studied rational expectations equilibria, 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 individuals' different interpretation of new information, while price changes reflect an average of individuals' different beliefs.[50]

[48] Such re-trading can serve to 'complete' the market. See Duffie and Huang (1985).
[49] Kyle (1985), Black (1986), De Long, Shleifer, Summers and Waldman (1990).
[50] See also Biais and Boessarts (1993).

5.4 Non-Convergent Beliefs

There are a number of reasons why heterogeneous beliefs might fail to converge – to each other or to the truth (see Section 4.2). One reason was because individuals' beliefs might fail to put positive probability on the true state. Nyarko (1991b) considers a model of a price-setting monopolist facing an unknown linear demand curve, that is, the slope and intersect are unknown. If the monopolist does not put positive weight on the true slope and intersect, then the monopolist's actions and beliefs may cycle forever, and the monopolist will never observe anything which contradicts his prior beliefs. If Nyarko's model was generalized to a oligopolistic setting, heterogeneous beliefs would persist.

Another setting where heterogeneous prior beliefs might persist is where individuals put positive probability on the truth, but learning is endogenous, so it is costly to discover the truth. Thus an employer might believe (with high probability) that a certain class of workers is of low ability. If the probability was sufficiently high, it would not even be worthwhile to hire a worker in order to discover if his belief was correct.[51]

Similarly, Basu (1992) argues that agricultural input subsidies (which he believes are damaging) are sustained in India because a majority of the population believe that alternative policies would be disastrous. Such beliefs are not proved right or wrong because the policy is not pursued. It is only natural to assume that they may be heterogeneous among the population. 'Ideology' in beliefs about economic policy may represent differences in beliefs uncontradicted by experience. Harrington (1993) argues forcefully that voters have heterogeneous beliefs about the efficacy of alternative public policies.

6. CONCLUSION

Economists are increasingly questioning a taboo on relaxing the common prior assumption. This paper has attempted to provide a framework for thinking about the issue. It was argued that 'rationality' arguments in favor of the common prior assumption are rather weak. But the commitment to the common prior assumption has been more of a methodological one: as Aumann (1987) puts it, 'the common prior assumption enables one to zero in on purely informational issues in analyzing economic (and other interactive) models of uncertainty'. But not all economic issues are informational, and there are some cases

[51] Thus Coate and Loury (1993) consider equilibria with labor market discrimination where beliefs which are incorrect are in some sense self-fulfilling, because out-of-equilibrium outcomes are not observed. Heterogeneity of prior beliefs is implicit in Coate and Loury (1993); Farmer and Terrell (1994) make the heterogeneity explicit in a different model of discrimination.

where differences in prior beliefs are essential to understanding economic phenomena. It is hard to characterize exactly what those cases are. But in Section 5, I gave examples of work where explicit modelling of heterogeneous prior beliefs seems to help our understanding and offered a few suggestions about where such modelling might be most useful.

REFERENCES

Akerlof, G. and Dickens, W. 1982. 'The economic consequences of cognitive dissonance'. *American Economic Review*, 72:307–19

Allais, M. 1953. 'Le comportement de l'homme rationnel devant le risque. Critique des postulats de l'école américaine'. *Econometrica*, 21:503–46

Allen, F., S. Morris and A. Postlewaite. 1993. 'Finite bubbles with short sales constraints and asymmetric information'. *Journal of Economic Theory*, 61:206–29

Anscombe, F. and R. Aumann. 1963. 'A definition of subjective probability'. *Annals of Mathematical Statistics*, 34:199–305

Aumann, R. 1976. 'Agreeing to Disagree', *The Annals of Statistics*, 4:1236–9

Aumann, R. 1987. 'Correlated equilibrium as an expression of Bayesian rationality'. *Econometrica*, 55:1–18

Basu, K. 1992. 'Comment' in *Proceedings of the World Bank Annual Conference on Development Economics*. L. Summers and S. Shah, (eds.)

Bayes, T. 1763. 'An essay towards solving a problem in the doctrine of chances'. Reprinted in *Biometrika*, 1956, 45:293–315

Beaver, W. 1968. 'The information content of annual earnings announcements'. *Empirical Research in Accounting: Selected Studies*, supplement to *Journal of Accounting Research*, 6:67–92

Becker, G. and G. Stigler. 1977. 'De gustibus non est disputandum'. *American Economic Review*, 67:76–90

Bernheim, D. 1984. 'Rationalizable strategic behavior'. *Econometrica*, 52:1007–28

Bernheim, D. 1986. 'Axiomatic characterizations of rational choice in strategic environments'. *Scandinavian Journal of Economics*, 88:473–88

Bewley, T. 1986. 'Knightian decision theory: Part I'. Cowles Foundation Discussion Paper No. 807, Yale University

Biais, B. and P. Boessarts. 1993. 'Asset prices and volume in a beauty contest'. Unpublished

Black, F. 1986. 'Noise'. *Journal of Finance*, 41:529–43

Blume, L. and D. Easley. 1993. 'Rational expectations and rational learning'. Unpublished

Böge, W. and T. Eisele. 1979. 'On the solutions of Bayesian games'. *International Journal of Game Theory*, 8:193–215

Borel, E. 1924. 'Apropos of a treatise on probability'. In *Studies in Subjective Probability*, Kyburg and Smokler (eds.), 1st edn. John Wiley and Sons, 1964

Brandenburger, A. 1993. 'Strategic and structural uncertainty in games'. Unpublished

Brandenburger, A. and E. Dekel. 1987. 'Rationalizability and correlated equilibria'. *Econometrica*, 55:1391–402

Brandenburger, A. and E. Dekel. 1993. 'Hierarchies of beliefs and common knowledge'. *Journal of Economic Theory*, 59:189–98

Brandenburger, A., E. Dekel and J. Geanakoplos. 1992. 'Correlated equilibrium with generalized information structures'. *Games and Economic Behavior*, 4:182–201

Broome, J. 1989. 'Should social preferences be consistent?' *Economics and Philosophy*, 5:7–18

Carlsson, H. and E. van Damme. 1993. 'Global games and equilibrium selection'. *Econometrica*, 61:989–1018

Carnap, R. 1950. *Logical Foundations of Probability*. University of Chicago Press

Coate, S. and G. Loury. 1993. 'Will affirmative action policies eliminate negative stereotypes?' *American Economic Review*, 83:1220-40

De Finetti, B. 1937. 'La prevision: ses lois logiques, ses sources subjectives'. *Annales de l'Institute Henri Poincaré*, 7:1–68. English trans. reprinted in *Studies in Subjective Probability*, Kyburg and Smokler (eds.), 2nd edn. 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*, 98:703–38

Debreu, G. 1974. 'Excess Demand Functions'. *Journal of Mathematical Economics*, 1:15–21

Duffie, D. and C. Huang. 1985. 'Implementing Arrow–Debreu equilibria by continuous trading of few long-lived securities'. *Econometrica*, 53:1337–56

Farmer, A. and M. Terrell. 1994. 'Discrimination, Bayesian updating of employer beliefs and human capital accumulation'. Unpublished

Foster, D. and P. Young. 1990. 'Stochastic evolutionary game dynamics'. *Theoretical Population Biology*, 38:219–32

Fudenberg, D. and D. Levine. 1993. 'Self-confirming equilibrium'. *Econometrica*, 61:523–45

Geanakoplos, J. 1989. 'Game theory without partitions, and applications to speculation and consensus'. Cowles Foundation Discussion Paper No. 914, Yale University

Geanakoplos, J. and H. Polemarchakis. 1982. 'We can't disagree forever'. *Journal of Economic Theory*, 28:192–200

Grossman, S. 1976. 'On the efficiency of competitive stock markets where agents have diverse expectations'. *Journal of Finance*, 31:573–85

Guesnerie, R. 1992. 'An exploration of eductive justifications of the rational-expectations hypothesis'. *American Economic Review*, 82:1254–78

Gul, F. 1993. 'A note on the Bayesian view in game theory and economics'. Stanford University Research Paper No. 1191

Hahn, F. 1977. 'Exercises in conjectural equilibrium analysis'. *Scandinavian Journal of Economics*, 79:210-26

Hammond, P. 1983. 'Ex-post optimality as a dynamically consistent objective for collective choice under uncertainty'. In *Social Choice and Welfare*, P. Pattanaik and M. Salles (eds.). North-Holland

Harrington, J. 1993. 'Economic policy, economic performance, and elections'. *American Economic Review*, 83:27–42

Harris, M. and A. Raviv. 1993. 'Differences in opinion make a horse race'. *Review of Financial Studies*, 6:473–506

Harrison, M. and D. Kreps. 1978. 'Speculative investor behavior in a stock market with heterogeneous expectation'. *Quarterly Journal of Economics*, 92:323–36

Harsanyi, J. 1967. 'Games with incomplete information played by "Bayesian" players. Part I: the basic model'. *Management Science*, 14:159–82

Harsanyi, J. 1968a. 'Games with incomplete information played by "Bayesian" players. Part II: Bayesian equilibrium points'. *Management Science*, 14:320-34

Harsanyi, J. 1968b. 'Games with incomplete information played by "Bayesian" players. Part III: the basic probability distribution of the game'. *Management Science*, 14:486–502

Heckman, J. and T. McCurdy. 1988. 'Empirical tests of labor-market equilibrium: an evaluation'. *Carnegie-Rochester Conference Series on Public Policy*, 28:231–58

Kahneman, D., P. Slovic and A. Tversky. 1982. *Judgements under Uncertainty: Heuristics and Biases*. Cambridge University Press

Kalai, E. and E. Lehrer. 1993a. 'Rational learning leads to Nash equilibrium'. *Econometrica*, 61:1019–45

Kalai, E. and E. Lehrer. 1993b. 'Subjective equilibrium in repeated games,' *Econometrica*, 61:1231–40

Kandori, M., G. Mailath and R. Rob. 1992. 'Learning, mutation and long run equilibria in games: a general theory and applications'. *Econometrica*, 61:29–56
Keynes, J. 1921. *A Treatise on Probability*. Macmillan & Co
Keynes, J. 1936. *The General Theory of Employment, Interest and Money*, MacMillan & Co
Knight, F. 1921. *Risk, Uncertainty and Profit*. Hougton-Miller
Kurz, M. 1994a. 'Asset prices with rational beliefs'. Unpublished
Kurz, M. 1994b. 'On the structure and diversity of rational beliefs'. *Economic Theory*, 4:1–24
Kyle, A. 1985. 'Continuous auctions and insider trading'. *Econometrica*, 53:1315–35
Lagunoff, R. and A. Matsui. 1994. 'Agitators and free riders on the path to cooperation: the evolution of cooperation in mechanisms for public projects'. CARESS Working Paper No. 94–05, University of Pennsylvania
Levi, I. 1980. *The Enterprise of Knowledge*. MIT Press
Lintner, J. 1969. 'The aggregation of investors' diverse judgements and preferences in purely competitive markets'. *Journal of Financial and Quantitative Analysis*, 4:347–400
Matsui, A. 1993. 'Evolution and rationality'. CARESS Working Paper No. 93–19, University of Pennsylvania
Matsui, A. and R. Rob. 1991. 'Evolution, rationality and equilibrium selection in societal games'. CARESS Working Paper No. 91–31, University of Pennsylvania
Mertens, J.-F. and S. Zamir. 1985. 'Formulation of Bayesian analysis for games with incomplete information'. *International Journal of Game Theory*, 14:1–29
Milgrom, P. and J. Roberts. 1990. 'Rationalizability, learning and equilibrium in games with strategic complementarities'. *Econometrica*, 58:1255–78
Milgrom, P. and N. Stokey. 1982. 'Information, trade and common knowledge'. *Journal of Economic Theory*, 26:17–27
Mises, R. von. 1957. *Probability, Statistics and Truth*. George Allen and Unwin Ltd
Mongin, P. 1993. 'Consistent Bayesian aggregation'. CORE Discussion Paper No. 9319, forthcoming in *Journal of Economic Theory*
Morris, S. 1991. 'The role of beliefs in economic theory'. Ph.D. Dissertation, Yale University
Morris, S. 1993a. 'Risk, uncertainty and hidden information'. CARESS Working Paper No. 93–10, University of Pennsylvania. Forthcoming in *Theory and Decision*
Morris, S. 1993b. 'Price bubbles and learning'. University of Cambridge Economic Theory Discussion Paper No. 191
Morris, S. 1994. 'Trade with heterogeneous prior beliefs and asymmetric information'. *Econometrica*, 62:1327–48
Moulin, H. 1979. 'Dominance solvable voting schemes'. *Econometrica*, 58:1255–78
Myerson, R. 1991. *Game Theory: Analysis of Conflict*. Harvard University Press
Nau, R. and K. McCardle. 1990. 'Coherent behavior in noncooperative games'. *Journal of Economic Theory*, 50:424–44
Nau, R. and K. McCardle. 1991. 'Arbitrage, rationality and equilibrium,' *Theory and Decision*, 31:199–240
Nyarko, Y. 1991a. 'Most games violate the Harsanyi doctrine'. Working Paper No. 91–39, C.V. Starr Center for Applied Economics, New York University
Nyarko, Y. 1991b. 'Learning in mis-specified models and the possibility of cycles'. *Journal of Economic Theory*, 55:416–27
Pearce, D. 1984. 'Rationalizable strategic behavior and the problem of perfection,' *Econometrica*, 47:655–78
Quandt, R. 1986. 'Betting and equilibrium'. *Quarterly Journal of Economics*, 101:201–7
Ramsey, F. 1926. 'Truth and probability'. In *The Foundations of Mathematics and other Logical Essays*. Kegan Paul, 1930
Ritter, J. 1991. 'The long-run performance of initial public offerings'. *Journal of Finance*, 46:3–27
Ross, S. 1989. 'Intertemporal asset pricing'. In *Theory of Valuation: Frontiers in Modern Financial Theory*. S. Bhattacharya and G. Constantinides, (eds.). Rowman and Littlefield

Roth, D. 1992. 'Rationalizable predatory pricing'. Forthcoming in *Journal of Economic Theory*

Rothschild, M. 1974. 'A two-armed bandit theory of market pricing'. *Journal of Economic Theory*, 9:185–202

Rubinstein, A. 1993. 'On price recognition and computational complexity in a monopolistic model'. *Journal of Political Economy*, 101:473–84

Savage, L. 1954. *The Foundations of Statistics*. John Wiley and Sons

Schick, F. 1986. 'Dutch bookies and money pumps'. *Journal of Philosophy*, 83:112–19

Sebenius, J. and J. Geanakoplos. 1983. 'Don't bet on it: contingent agreements with asymmetric information'. *Journal of the American Statistical Association*, 78:382, 424–26

Shin, H. 1992. 'Prices of state contingent claims with insider traders, and the favorite-longshot bias'. *Economic Journal*, 102:411, 426–35

Solow, R. 1976. 'Economics. Is something missing?' In *Economics History and the Modern Economist*. W. Parker, (ed.). Basil Blackwell

Sonnenschein, H. 1973. 'Do Walras's identity and continuity characterize the class of community excess demand functions?' *Journal of Economic Theory*, 6:345–54

Varian, H. 1985. 'Divergence of opinion in financial markets: a note'. *Journal of Finance*, 40:309–17

Varian, H. 1989. 'Differences of opinion in financial markets'. In *Financial Risk: Theory, Evidence and Implications*. C. Stone, (ed.). Kluwer Academic Publications

Varian, H. 1992. 'Differences of opinion'. In *The New Palgrave Dictionary of Money and Finance*, pp. 668–70. J. Eatwell, M. Milgate and P. Newman (eds.). Macmillan/Stockton Press