Common Priors For Like-Minded Agents

Klaus Nehring

Department of Economics

University of California, Davis*

August 2003

*E-mail address: kdnehring@ucdavis.edu; Address for manuscript correspondence: Institute for Advanced Study, Einstein Drive, Princeton, New Jersey 08540, Main Tel: 609-734-8000; Main Fax: 609-924-8399.

ABSTRACT

Two agents are like-minded when their beliefs are equal once conditioned on knowledge of both of their types. Assuming the existence of an outside observer that is commonly known to be like-minded and uninformative about the insiders, we derive the existence of a common prior among the insiders, with the outsiders beliefs (appropriately conditioned) serving as the common prior. A key advantage of like-mindedness is its fully local definition, which allows to distinction between consistency of agent's actual beliefs and of beliefs they merely view as possible.

By later including agents' "epistemic attitudes" among the primitives, we derive like-mindedness from reasonableness judgments about each others attitudes. In this richer framework, one can model alternative conceptions of intersubjective rationality as constraints on such reasonableness judgements.

Keywords: Common Prior Assumption, like-mindedness, incomplete information, intersubjective rationality, pluralism, relativism.

JEL classification: C70, D80.

1. INTRODUCTION

The Common Prior Assumption (CPA) plays a central role in information economics and game theory . This key role is due in part to a set of substantial core concepts and results such as the notion of correlated equilibrium (Aumann 1974, 1987) and the No Trade Theorem (Milgrom-Stokey 1982). Moreover, and probably more importantly, the CPA is central for a methodological reason, especially under incomplete information: without it, it seems to hard to restrict the beliefs of different agents and their types in a transparent and controlled way. This methodological rationale presumably explains why almost the entire literature continues to assume common priors, even though this assumption is widely viewed as empirically *false,* and frequently even as normatively *unwarranted* (see f.i. Morris (1995)).

If the continued appeal of the CPA was entirely due to mere convenience (or intellectual laziness), comparable for example to the assumption of a representative agent in macroeconomics, it would not merit a detailed foundational exploration. But this not be the case: while differences in information are not the only reason for empirical differences in beliefs, they appear to be the dominant ones. A world in which beliefs were as unconstrained across individuals as preferences are would probably look very different from the world we inhabit; in particular, it would exhibit financial trades based on differences in beliefs on a scale far greater than is observed in practice.

Thus, at an intuitive level at least, the CPA appears valid as an important approximation or benchmark.[1] However, in situations of incomplete information, that is: in situations in which agents are mutually uncertain about each others' beliefs, and without a preceding stage in which beliefs were commonly known, there is a significant gap between the formal statement of the CPA and its underlying intuitive content. Indeed, Gul (1998) has even questioned whether the CPA can be transparently interpreted at all in this context (see also Dekel-Gul (1997) and Lipman (1997)). Paraphrasing Samet (1998a), Gul's question asks "how can one tell, from agents' actual belief hierarchies alone, whether their beliefs are consistent with the CPA". This *meaningfulness* question has by now been successfully addressed in a number of papers in the literature (Bonanno-Nehring (1999), Feinberg (2000), Halpern (2002), Nehring (2001), Samet (1998a,1998b)).

In the present paper, we want to go further by asking the *explanatory* question as to what substantive empirical and/or normative assumptions *underly* the CPA, what facts in the widest sense (including possibly facts about agents' knowledge and rationality) *bring about* the CPA. We shall

[1] Whether it is an appropriate one will depend, of course, on the context and the modeler's judgment.

argue below that none of the existing characterizations is satisfactory for this purpose.

Like-minded Agents

To address this explanatory issue, we will derive the CPA from an underlying notion of "like-mindedness" under appropriate auxiliary assumptions. Intuitively, two agents are like-minded if they assess uncertainties in fundamentally in the same way, if they attribute differences in their beliefs fully to differences in their information. We define like-mindedness of two agents formally as equality of their beliefs, *conditional* on knowledge of both agents' entire belief hierarchies; the conditioning ensures that the agents' beliefs are compared on the basis of the same (hypothetical) information. Empirically, agents may fail to be like-minded for example due to differences in temperament, cognitive strategy, professional training, Weltanschauung.

The definition of like-mindedness is most easily illustrated in the special case of one-sided incomplete information among two agents, call them Ego and Alter. For example, Ego, a patient, is uncertain about the beliefs of Alter, a medical doctor consulted for an imminent surgery; for the sake of the argument, Ego's beliefs are commonly known. In this context, a "state of the world" is fully described by specifying the "state of nature" (success or failure of the surgery) and Alter's probability distribution over states of nature, together with Ego's joint probability distribution over states of nature and Alter's first-order beliefs. Ego and Alter are *like-minded* at state α, if Ego's probability distribution over states of nature *conditional on Alter's being what it is at* α agrees with Alter's.

Example 1: Ego certain that Alter like-minded

State of the world	τ	ω_2	ω_3	ω_4
State of nature	success	failure	success	failure
Ego	0.4	0.4	0.04	0.16
Alter	0.5	0.5	0.2	0.8

In example 1, in the true state τ, Alter gives the surgery a chance of 50 % to succeed, in contrast to Ego, whose probability is 44%. Nonetheless, at all states Alter and Ego are like-minded; for example, at τ, *conditional* on Alter's estimate of success being equal to 50 %, Ego's is 50 % as well . Thus, Ego *knows* that she is like-minded to Alter. From Ego's point of view, Alter's beliefs can be viewed as the outcome of conditioning on Ego's "prior" based on his private information: Ego can view Alter has her other, better-informed self, her *alter ego.* Note that Ego's prior can be viewed as a common prior for the two.

Contrast this with the following example 2, in which Ego is unsure about Alter's like-mindedness. Specifically, Ego is unsure whether Alter is "balanced" or "overconfident"; to model this, we include Alter's psychology in the state of nature.

Example 2: Ego unsure whether Alter like-minded

State of the world	τ	ω_2	ω_3	ω_4
State of nature	(success,balance)	(failure,balance)	(success,overconfidence)	(failure,overconfidence)
Ego	0.25	0.25	0.15	0.35
Alter	0.5	0.5	0.5	0.5

Here, Ego and Alter are in fact like-minded (at τ), but they nonetheless *agree to disagree* about the probability of the surgery's success, as it is common knowledge that Ego's probability is 40% while Alter's is 50%. This explained by Ego's fear that Alter may be overconfident, conditional on which she reduces her estimate of success to 30%. Obviously, the two agents' interactive beliefs cannot be represented by a common prior.

Common Priors with a Like-Minded, Uninformed Outsider

It is clear from examples 1 and 2 that the CPA is satisfied under one-sided incomplete information if and only if the agent whose beliefs are commonly known knows that she is like-minded with all others (in which case this is also common knowledge).[2] However, in the general case of many-sided incomplete information, common knowledge of like-mindedness turns out to be insufficient for the

[2] Under one-sided incomplete information, the notion of like-mindedness as presented here is not new; it has come up before in dynamic one-person settings in which Ego and Alter represent to the same person at an earlier respectively later date, and plays a central role there in the justification of Bayesian updating . Like-mindedness is non-trivial by possibly failing to hold even in this *intra*-personal context: for example, sober real-world Egos often fail to be like-minded with their future drunk Alters.

The relevant literature to date is largely philosophical; "Like-mindedness" appears there as "reflection principle".See, for example, Goldstein (1983), van Fraassen (1984), Maher (1993).

existence of a common prior, as shown in section 3.2. Additional assumptions are needed. In this paper, we will close the gap by assuming the existence of a like-minded, uninformed outsider. By "uninformed", we will mean "uninformed from the point of view of the insiders", that is "*uninformative*" for the insiders.[3] In line with the incomplete information setting, uninformativeness does not require that insiders know anything about the outsider's beliefs; uninformativeness seems plausible in many situations in which the outsider is "sufficiently removed" from the scene of inside action.

The main result of the paper, Theorem 1 shows that common knowledge of like-mindedness with the outsider together with common knowledge of his uninformativeness yields a common prior among insiders that is given by the outsiders' appropriately conditioned beliefs; we dub this "personalized" version of the CPA "External Harsanyi Consistency". By equating the common prior with the conditional probabilities of a particular individual, External Harsanyi Consistency renders the common prior an ordinary personal probability. This makes it possible to transparently impose specific assumptions on the common prior directly; by contrast, Samet's (1998) elegant internal characterization of the common prior as an infinite limit of higher-order expectations does not, by itself, render the content of such assumptions similarly transparent.[4] In the converse direction, Theorem 1 also shows that common knowledge of External Harsanyi Consistency plus a regularity condition implies common knowledge of like-mindedness and uninformativeness; the latter is therefore not merely an ad-hoc additional assumption that happens to yield the desired conclusion, but is part and parcel of the very notion of External Harsanyi Consistency.

Like-mindedness and the Intersubjective Rationality of Beliefs

The analysis in the core of the paper is situated in a standard type-space framework in which all assumptions are formulated as conditions on the agents' probabilistic belief hierarchies. In section 6, we develop a richer framework in which agents' "epistemic attitudes" are introduced as independent primitives, and like-mindedness of beliefs is derived from the recognition of other agents as "equally rational" (we will say "co-rational"). This framework allows one to formulate competing normative positions on the content of intersubjective rationality in terms of alternative restrictions on these equivalence relations. We distinguish three types of positions: a *rationalist* position that can be

[3]More precisely, an outside agent is "*uninformative*" if any insider's beliefs about the state of nature and insiders' internal belief hierarchies are indepedent of the outsider's beliefs about these.

[4]While it tells one when agents' belief hierarchies are consistent with particular assumptions, it is not clear how Samet's definition will, in general, allow one to judge their adequacy in particular epistemic situations.

viewed as capturing the Harsanyi doctrine (cf. Aumann 1987), a *pluralist* position that allows for some intrinsic differences of beliefs that cannot be attributed to differences of information, and, finally, a *relativist* position that rejects the normativity of any restrictions of beliefs across agents.

Comparison to the Literature

Our main result, Theorem 1, derives the CPA from common knowledge of events (like-mindedness and uninformativeness) that are not necessarily commonly known. By their logical structure, these "fully local" properties can distinguish between what is true of agents' beliefs de facto from what is true of beliefs that are merely viewed as possible by some agents, or that only "commonly possible", that is: not impossible on the basis of what is commonly known. In section 3.4, we provide a formal definition of what makes a property "fully local", and show that like-mindedness is the strongest, fully local property entailed by the existence of a common prior. We take full locality of properties involved in a derivation/characterization of the CPA as the decisive criterion that distinguishes a genuinely *explanatory* derivation.

By contrast, all of the contributions to the above-mentioned literature establishing the meaningfulness of the CPA under incomplete information characterize the CPA in terms of a events that, if true, must necessarily be commonly known, and therefore are not fully local.[5] Most of these characterizations are based on the absence of any generalized "agreement to disagree" in the sense of Aumann (1976). The difference between agreement- and like-mindedness based derivations of the CPA is illustrated by Example 2 in which there is no common prior, and in which Ego and Alter agree to disagree, and this is common knowledge. Both of these facts are *explained* by Ego's uncertainty about Alter's like-mindedness reflected in the non-likemindedness of Alter and Ego at the counterfactual states ω_3 and ω_4; clearly, like-mindedness is more primitive a notion than agreement.[6]

This difference in logical structure shows up starkly in the implied versions of the No Trade theorem under incomplete information. While agreement-based characterizations of the CPA render

[5] Such events can be called "intrinsically public". In the notation of section 2 below, an event E is intrinsically public if $E = K_I^* E$. The existence of a common prior, formulated as an event (later called "Harsanyi consistency") is itself an intrinsically public event.

Our discussion assumes the Truth axiom; the general case without the Truth axiom is studied in Bonanno-Nehring (1999).

[6] One may be tempted to motivate these Agreement-based characterizations as capturing "*common knowledge* of like-mindedness", but this involves an evident fudge as it implicitly appeals to a fully local notion of like-mindedness.

this result essentially tautological, the like-mindedness based derivation via Theorem 1 combined with Aumann's (1976) original Agreement Theorem preserves its striking character; see section 5 for more details. In the concluding section 7, we also point out that the present fully local foundation of common priors suggests a natural generalization to a derivation of "almost common priors" based on (almost) common knowledge of almost-like-mindedness.

Organization

The remainder of the paper is organized as follows. After introducing the general framework in Section 2, a formal definition of like-mindedness is proposed and discussed in section 3. It is shown that like-mindedness is the strongest fully local property entailed by the existence of a common prior. Its motivation is also elaborated in a dynamic setting in which the agents mutually reveal their belief-hierarchies. Section 4 introduces the notions of an uninformative outsider and of External Harsanyi Consistency, and establishes the main result of the paper, Theorem 1. The Theorem, and in particular the key uninformativeness assumption, are then discussed in Section 5. Section 6 derives like-mindedness from co-rationality judgments regarding other agents' epistemic attitudes, and defines alternative types of rationality norms governing such judgments. Section 7 concludes. All proofs are collected in the appendix.

2. BAYESIAN TYPE SPACES

Definition 1 *A* **rooted Bayesian Type Space** *is a tuple*
$\mathcal{B} = \langle N, \Omega, \tau, \Theta, \theta, \{p_i\}_{i\in N}\rangle$ *, where*

- N *is a finite set of* agents.
- Ω *is a finite set of* states *(or possible worlds). The subsets of* Ω *are called* events.
- τ *is the true state.*
- Θ *is the set of* "states of nature".
- $\theta : \Omega \to \Theta$ *specifies, for each* $\alpha \in \Omega$*, the state of nature* θ^α *obtaining at* α.
- *for every agent* $i \in N$, $p_i : \Omega \to \Delta(\Omega)$ *(where* $\Delta(\Omega)$ *denotes the set of probability distributions over* Ω*) is a function that specifies, for each* $\alpha \in \Omega$*, his probabilistic beliefs* p_i^α *at* α.

A type space is simply a state space in which at any state α, the agents' beliefs at that state p_i^α are specified. As a result, an agents' belief at a state describes not only his beliefs about facts of nature, but also his beliefs about other agents' (first-order) beliefs about states of nature, hence also his beliefs about agents' higher-order beliefs about states of nature, thus in effect: an entire belief hierarchy. For example, $p_i^\alpha(\{\omega | p_j^\omega(rain) \geq 0.7\})$ denotes agent $i's$ probability at state α that agent j believes that it will rain with at least 70% probability. A state in a type space can be thus be thought of as a notational device for describing the belief hierarchies of each agent.[7] Fixing a particular state τ as the "root" fixes a particular profile of belief hierarchies.[8]

We will maintain the following two assumptions.

Assumption 1 *(Introspection)* For all $\alpha \in \Omega$ and all $i \in N$: $p_i^\alpha(\{\omega \in \Omega \mid p_i^\omega = p_i^\alpha\}) = 1$.

Assumption 2 *(Truth)* For all $\alpha \in \Omega$ and all $i \in N$: $p_i^\alpha(\{\alpha\}) > 0$.

Introspection says that agents are always (at any state α) certain of own belief p_i^α. Truth states that, for any state that may occur, agents will have put positive probability on that state if it occurs; thus Truth assumes that agents are never wrong in their probability-one beliefs. While standard, this assumption is not unrestrictive.[9]

Let $\|p_i = p_i^\alpha\|$ denote the event $\{\omega \in \Omega \mid p_i^\omega = p_i^\alpha\}$. An agent "knows" an event E at α ("$\alpha \in K_i E$") if he is certain of it, i.e. if $p_i^\alpha(E) = 1$. This endows the interactive Bayesian model with knowledge operators $K_i : 2^\Omega \to 2^\Omega$, for $i \in N$. For the associated possibility correspondences $\mathcal{P}_i : \alpha \mapsto \{\omega \in \Omega \mid \alpha \notin K_i(\Omega \backslash \{\omega\})\}$, one has $\mathcal{P}_i(\alpha) = \text{supp } p_i^\alpha = \|p_i = p_i^\alpha\|$ by Introspection and Truth; in particular, the family $\mathcal{P}_i(\Omega) := \{\mathcal{P}_i(\omega) \mid \omega \in \Omega\}$ is i's *type partition* . For a set of agents $J \subseteq N$, "common knowledge among the agents in J" is given by an operator $K_J^* : 2^\Omega \to 2^\Omega$ with associated possibility operator $\mathcal{P}_J^*$. First, define an auxiliary operator "everybody in J knows" $K_J : 2^\Omega \to 2^\Omega$ by setting $K_J(E) := (\cap_{i \in J} K_i E)$. E is *common knowledge among the agents in* J if everybody in J knows that E, and if everybody in J knows that everybody in J knows that E, and so forth. Formally,

$$K_J^*(E) := K_J(E) \cap K_J\left(K_J(E)\right) \cap K_J\left(K_J\left(K_J(E)\right)\right) \cap ... \quad .$$

[7] By results due to Armbruster-Boege (1979) and Mertens-Zamir (1985), any profile of probabilistic belief hierarchies has a type-space representation; the assumption that the state space Ω is finite is restrictive but entirely standard. Infinite state-spaces are considered in Feinberg (2000) and Halpern (2002).

[8] Rooted type spaces have been introduced in Bonanno-Nehring (1999).

[9] See Bonanno-Nehring (1999) for a detailed study of its relaxation.

3. LIKE-MINDEDNESS

3.1 Definition

Intuitively, it is clear that the CPA can only be expected to hold if "all individuals look at the world in fundamentally the same way"; otherwise, their beliefs may differ even in situations of complete information, and thereby violate the CPA. The task of the present section is to formalize this informal background assumption in terms of appropriate conditions on individuals' belief hierarchies; in section 6, we will explicitly introduce agents' "mind-sets" or "epistemic attitudes" as additional primitives of the model and derive like-mindedness from them.

Under complete information, the formal content of like-mindedness is equality of beliefs. Under incomplete information, one needs to "control for" potential asymmetries in information. This motivates the following condition.

Definition 2 *i and j are* **like-minded** at state α *("$\alpha \in \mathbf{LM}_{ij}$") if, for all $E \subseteq \Omega$,*

$$p_i^\alpha(E/\|p_j = p_j^\alpha\|) = p_j^\alpha(E/\|p_i = p_i^\alpha\|) \quad .$$

Interpretation: i and j are like-minded at state α if their subjective probabilities on any event E agree, conditional on their being informed about each other's *entire* belief hierarchies.

3.2 Like-mindedness and Common Priors

The goal of this paper is to derive under an appropriate set of assumptions the existence of a common prior among a set of agents I from like-mindedness of the agents. Following Bonanno-Nehring (1999), the following is an appropriate local definition of a common prior μ as an event in the presence of the Truth axiom[10].

Definition 3 (Harsanyi Consistency) $\alpha \in \mathbf{HC}_\mu$ *if $\mu(\mathcal{P}_I^*(\alpha)) = 1$ and, for all $\beta \in \mathcal{P}_I^*(\alpha)$ and all $i \in I : \mu(\|p_i = p_i^\beta\|) > 0$ and $\quad p_i^\beta = \mu(./\|p_i = p_i^\beta\|)$; moreover, let $\mathbf{HC} := \cup_{\mu \in \Delta(\Omega)} \mathbf{HC}_\mu$.*

As observed there, the local common prior μ is unique, has support $\mathcal{P}_I^*(\alpha)$, and is commonly known (that is: $\mathbf{HC}_\mu = K_I^*(\mathbf{HC}_\mu)$ as well as $\mathbf{HC} = K_I^*(\mathbf{HC})$). Moreover, it is easily seen that Harsanyi

[10] Bonanno-Nehring (1999) show that, in the absence of the Truth axiom, there are multiple reasonable local formulations of the common prior assumption.

Consistency implies common knowledge of like-mindedness among the agents. This follows from simply observing that if $\alpha \in \mathbf{HC}_\mu$, then for any $i, j \in I$,

$$p_i^\alpha(E/\|p_j = p_j^\alpha\|) = \mu(E/\|p_j = p_j^\alpha\| \cap \|p_i = p_i^\alpha\|) = p_j^\alpha(E/\|p_i = p_i^\alpha\|).$$

Ideally, one would like the converse to hold as well, but this fails in general. To see this, consider a Bayesian type space with two agents $I = \{1, 2\}$ and $\Omega = \{\omega_1, .., \omega_4\}$ such that $\mathcal{P}_1(\Omega) = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$ and $\mathcal{P}_2(\Omega) = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$. Since the sets $\|p_1 = p_1^\alpha\| \cap \|p_2 = p_2^\alpha\|$ are singletons, like-mindedness is trivially satisfied, whether or not there exists a common prior.

Generalizing this example, it is clear that common knowledge of like-mindedness among two agents implies a common prior between them only in degenerate situations. We will show in the next section how this impasse can be overcome in the presence of a "like-minded, uninformed outsider". Before doing so, we will present a dynamic interpretation of like-mindedness (section 3.3) and show that like-mindedness is strongest "fully local" property entailed by Harsanyi Consistency (section 3.4).

3.3. A Dynamic Interpretation

3.3.1 Interactive Bayesian Updating.—

The dynamic interpretation requires a 2-period extension of the above model, in which agents revise their beliefs at date 2 via Bayesian updating on a received information signal, and in which the signal-generating process is commonly known. Formally, a dynamic version is obtained by treating instances of the same individual $i \in J$ at different dates t as different "agents" $i_t \in N := J \times \{1, 2\}$. A "state" α describes now, besides a possible state of nature, a possible history of individuals' beliefs over time.

The commonly known signal-generating process can be described by a family of information partitions $\{\mathcal{F}_i\}_{i \in J}$. With $\mathcal{F}_i(\alpha)$ denoting the cell of the partition $\mathcal{F}_i$ containing the state $\alpha \in \Omega$, the signal received by i in α can be identified with $\mathcal{F}_i(\alpha)$ which, literally, is the set of states at which i receives the same signal as at α.

Definition 4 *A* **Bayesian Type Space with Updating** *is a tuple* $(J, \mathcal{B}, \{\mathcal{F}_i\}_{i \in J})$, *where*

- $\mathcal{B}$ *is a rooted Bayesian type space with* $N = J \times \{1, 2\}$
- *For each* $i \in J$, $\mathcal{F}_i$ *is a partition of* Ω.

- *For all $i \in J$, all $\alpha \in \Omega$ and all $E \subseteq \Omega$, $p_{i_2}^{\alpha}(E) = p_{i_1}^{\alpha}(E/\mathcal{F}_i(\alpha))$.*

Note that we have allowed the information partitions $\mathcal{F}_i$ to be arbitrary partitions of Ω; in particular, information may be obtained not only about the natural state, but also about other agents' beliefs. The following fact summarizes some elementary consequences of the definition; in particular, individuals know which signal they have received (part i)) and remember their past beliefs (part ii)). Part iii) states a simple characterization of the type-partition of date-2 agents.

Fact 1 *For all $i \in J$ and all $\alpha \in \Omega$:*

i) $\mathcal{F}_i(\alpha) = K_{i_2}(\mathcal{F}_i(\alpha))$,

ii) $\|p_{i_1} = p_{i_1}^{\alpha}\| = K_{i_2}(\|p_{i_1} = p_{i_1}^{\alpha}\|)$,

iii) $\|p_{i_2} = p_{i_2}^{\alpha}\| = \|p_{i_1} = p_{i_1}^{\alpha}\| \cap \mathcal{F}_i(\alpha)$.

3.3.2 A Dynamic Thought Experiment.—

In this set-up, it is straightforward to model the following simple thought experiment. Suppose that the two individuals were to reveal their entire belief hierarchies truthfully to the other at date 1, and update their beliefs on this information at date 2 according to Bayes' rule. Then the agents' beliefs will agree at date 2 if and only if the agents are like-minded at date 1.

Say that two agents are minimally like-minded if their that if their entire belief hierarchies happen to be common knowledge between the two, these must coincide.

Definition 5 *The agents i and j are* **minimally like-minded** at state α *("$\alpha \in$ $\mathbf{MLM}_{ij}$") if $\alpha \in K^*_{\{i,j\}}(\|p_i = p_i^{\alpha}\| \cap \|p_j = p_j^{\alpha}\|)$ implies $p_i^{\alpha} = p_j^{\alpha}$.*

Consider now a Bayesian Type Space with Updating in which two individuals are informed of each other's date-1 belief hierarchy. Then at date 2, each individual will know the others' date-1 belief hierarchy. Using this knowledge, he can infer how the other will revise her beliefs on the basis of her new information, and can thus infer her date 2 posterior beliefs. Likewise, he knows that the other agent must have figured out his posterior beliefs as well; thus, revelation of agents' entire belief hierarchies leads to common knowledge of posteriors within one iteration.[11] As a result, minimal like-mindedness at date 2 is equivalent to like-mindedness at date 1. This is summarized in the following proposition.

[11] By contrast, if only beliefs about particular events are revealed as in Geanakoplos/Polemarchakis (1982), common knowledge of posteriors is achieved only after sufficiently many iterations.

Proposition 1 *In a Bayesian Type Space with Updating* $(J, \mathcal{B}, \{\mathcal{F}_k\}_{k \in J})$ *with* $\mathcal{F}_i = \mathcal{P}_j(\Omega)$ *and* $\mathcal{F}_j = \mathcal{P}_i(\Omega)$ *for given* $i, j \in J$, $\mathbf{LM}_{\{i_1, j_1\}} = \mathbf{MLM}_{\{i_2, j_2\}}$.

Proposition 1 can be viewed as deriving like-mindedness from the "rock-bottom" concept of minimal like-mindedness, by appealing to a (possibly counterfactual) situation in which agents will have identical information. The beauty of this counterfactual is its conceptual well-definedness in terms of a well-specified process of information generation. By contrast, the counterfactuals that have been commonly used to justify the common prior assumption from minimal like-mindedness lack such operational well-definedness. This applies to the postulate of an original situation of "no information" as in Aumann (1987), as well as to the elimination of informational asymmetries by some (ill-specified) process of forgetting as in Aumann (1998). Many have been struck by these counterfactuals as extravagant and inadmissible (see Binmore-Brandenburger (1990), Dekel-Gul (1997), and Gul (1998)).

We conclude this section with a discussion of a non-trivial feature of the definition of like-mindedness, namely the conditioning on the entire belief-hierarchy of the other. To see more clearly why this is necessary, it is instructive to compare **LM** to the following simpler and superficially perhaps more attractive criterion $\widetilde{\mathbf{LM}}$:

$$\alpha \in \widetilde{\mathbf{LM}} \text{ if, for all } E \subseteq \Omega,\ p_i^\alpha(E/\|p_j(E) = p_j^\alpha(E)\|) = p_j^\alpha(E/\|p_i(E) = p_i^\alpha(E)\|).$$

In contrast to **LM**, $\widetilde{\mathbf{LM}}$ fails to follow from the existence of a common prior. Indeed, note that $\alpha \in \widetilde{\mathbf{LM}}$ implies that the agents' subjective probabilities on an event E must be equal whenever they are *mutually* known;[12] it was Aumann's (1976) seminal insight, however, that an implication of this kind holds also for probabilities that are *commonly* known.

Likewise, the process of mutual belief revelation of section 3.2 yields a backward justification only **LM**, and not of $\widetilde{\mathbf{LM}}$, .since if only the other's original estimate of the probability on E, her posterior on E may again be uncertain; in this case, **MLM** will be satisfied vacuously whether or not **LM** is.

While $\widetilde{\mathbf{LM}}$ is a restriction on agents' second-order beliefs (on what they believe about each other's beliefs), **LM** entails no finite-order restriction on beliefs at all. This is the way it must be: Lipman (1997) has shown that even the CPA entails *no* finite-order restrictions on beliefs (beyond those implied by the truth axiom).

[12]This is a restriction on agents' second-order beliefs (on what they believe about each other's beliefs); by contrast, LM entails no finite-order restriction on beliefs at all. This is the way it must be: Lipman (1995) has shown that even the CPA entails *no* finite-order restrictions on beliefs (beyond those implied by the truth axiom).

3.4 Like-Mindedness as a Fully Local Property

We will now provide a rigorous definition of the notion of like-mindedness as a Fully Local Property. Formally, this requires considering a universe of Bayesian type spaces which, for notational simplification, we take in this section to be triples $\langle N, \Omega, \{p_i\}_{i \in N} \rangle$. Let Ω^* be an infinite "alphabet of states". Let $\mathcal{T}$ denote the set of all Bayesian type spaces with $\Omega \subseteq \Omega^*$ satisfying Introspection and Truth. A *property* (of interactive beliefs) is a mapping $\Phi : \mathcal{B} \to \Phi(\mathcal{B}) \subseteq \Omega$; the interpretation is that $\Phi(\mathcal{B})$ is the set of states at which this property is satisfied. Two type spaces $\mathcal{B} = \langle N, \Omega, \{p_i\}_{i \in N} \rangle$ and $\mathcal{B}' = \langle N, \Omega', \{p'_i\}_{i \in N} \rangle$ are *locally equivalent at* $\alpha \in \Omega \cap \Omega'$ if $\mathcal{P}_i(\alpha) = \mathcal{P}'_i(\alpha)$ and $p^{\alpha}_{i|\mathcal{P}_i(\alpha)} = p'^{\alpha}_{i|\mathcal{P}_i(\alpha)}$ for all $i \in N$. A property is **fully local** if, for all $\mathcal{B}, \mathcal{B}' \in \mathcal{T}$ and all $\alpha \in \Omega \cap \Omega'$ such that $\mathcal{B}$ and $\mathcal{B}'$ are locally equivalent at $\alpha, \alpha \in \Phi(\mathcal{B})$ if and only if $\alpha \in \Phi(\mathcal{B}')$. Note that like-mindedness, viewed as a mapping from type spaces to events, is a fully local property, since it only depends on agents' beliefs at α in the type space representation[13]; by contrast, Harsanyi Consistency, for example, is not. Note also that the definition of a fully local property can be meaningfully translated into one in which the type spaces $\mathcal{B}$ are viewed as elements of the universal type space a la Mertens and Zamir (1985).[14] Intuitively, a fully local property is determined by only the direct relations between agents' beliefs, disregarding their content as belief hierarchies that is obtained from unpacking their type-space representation. It is this focus on the direct relations between agents' beliefs that allows fully local properties to distinguish between the actual satisfaction of the property, and its satisfaction at possible, or even merely "commonly possible" states (elements of $\mathcal{P}^*_I(\alpha)$).

The key observation is that any fully local property that is entailed by Harsanyi Consistency is also entailed by like-mindedness; like-mindedness is therefore the *strongest fully local property* that is entailed by Harsanyi Consistency! In view of the discussion of section 3.2, this implies that there is no fully local property whose being commonly known entails Harsanyi Consistency for all Bayesian type spaces. Formally, one has the following result, with set-notation for properties to be read point-wise.

---

[13] The same holds for uninformativeness as defined in section 4.

[14] We omit the technical details. The basic idea is the following: If states are identified as profiles of belief hierarchies and a state of nature, the equality of the set of locally possible states $\cup_{i \in N} \mathcal{P}_i(\alpha) = \cup_{i \in N} \mathcal{P}'_i(\alpha)$ would need to be formulated as a bijection, and the equality of belief maps "$p^{\alpha}_i = p'^{\alpha}_i$ for all $i \in N$" simply as isomorphism under this projection. This isomorphism ignores by construction the identity (in terms of their associated belief hierarchies) of the states, and thus only takes into account the structure of agents' beliefs as "first-order" beliefs about the profile of belief hierarchies and the state of nature.

Proposition 2 *Let Φ be any fully local property such that* $\mathbf{HC} \subseteq \Phi$. *Then* $\cap_{i,j \in N} \mathbf{LM}_{i,j} \subseteq \Phi$.

4. OBSERVATION BY A LIKE-MINDED AND UNINFORMED OUTSIDER

To derive the existence of a common prior from like-mindedness conditions, we shall now include among the set of agents a like-minded "outsider"; thus, the set of individuals J will be made up of a set of "insiders" $i \in I$ and an outsider o, i.e. $J = I \cup \{o\}$. We will show that if the outsider is uninformed in an appropriate sense, the insiders' beliefs are Harsanyi consistent.

What should it mean for an outsider to be "uninformed"? A simple, though overly restrictive definition would require that the insiders know *everything* known by the outsider. Since the outsider knows his belief hierarchy, so must every insider. Thus, if this is common knowledge, the outsiders' belief hierarchy is commonly known as well. If, in addition, the outsider knows himself to be like-minded with every insider, then it is easily seen that the outsider's beliefs serve as a common prior among the insiders and himself, generalizing in straightforward manner the examples with one-sided incomplete information in the introduction. This, essentially, can be viewed as an interpersonal version of the gist of Aumann's (1998) argument (although, formulating his argument intertemporally, Aumann does not appeal to the notion of like-mindedness). However, assuming the outsider's beliefs to be commonly known seems unreasonably strong in most contexts, and indeed conflicts with the incomplete information picture of the world that motivates this work in the first place. A much more plausible and conceptually more satisfactory definition envisions the outsider as *uninformative,* that is, roughly speaking, as lacking any private information about the state of nature (and about the insiders beliefs about that state) that the insiders do not posses already.

To capture this formally, one needs to focus on the belief-closed event-subspace $\mathcal{A}(I)$ describing the state of nature as well as the hierarchies of the beliefs of the insiders about each others beliefs about the state of nature; specifically, let $\mathcal{A}(I)$ be the smallest algebra $\mathcal{A}$ such that $\mathcal{A}$ contains all events of the form $\theta^{-1}(T)$ for $T \subseteq \Theta$ and such that, , for all $E \in \mathcal{A}, i \in S, c \in [0,1] : \|p_i(E) = c\| \in \mathcal{A}$.[15] We will refer to the events $A \in \mathcal{A}$ as *internal* events. For each $i \in J$, let q_i^α denote the restriction of p_i^α to this event-subspace $\mathcal{A}(I)$; likewise, $q_i^\alpha(./E)$ denotes the restriction of $p_i^\alpha(./E)$ to $\mathcal{A}(I)$ for $E \subseteq \Omega$. We will write q_I for $(q_i)_{i \in I}$.

The above considerations motivate the following definitions.

[15] In the following, one could replace $\mathcal{A}(I)$ by any algebra $\mathcal{A}$ that is *belief-closed for I* in that, for all $E \in \mathcal{A}, i \in S, c \in [0,1] : \|p_i(E) = c\| \in \mathcal{A}$.

Definition 6 *i) The outsider is* **like-minded** *at* α *("$\alpha \in$ **LMO**"), if, for all* $i \in I$,

$$q_i^\alpha(./\|p_o = p_o^\alpha\|) = q_o^\alpha(./\|p_i = p_i^\alpha\|).$$

ii) The outsider is **uninformative** *at* α *("$\alpha \in$ **UIO**"), if, for all* $i \in I$,

$$q_i^\alpha = q_i^\alpha(./\|p_o = p_o^\alpha\|).$$

Thus, the outsider is "uninformative" about internal events if getting to know the outsider's beliefs does not change any insiders' beliefs about these. Clearly, if the insiders know the outsiders' beliefs, i.e. $p_i^\alpha(\|p_o = p_o^\alpha\|) = 1$ for all i, the outsider is uninformative at α. Note that, for the purpose of getting an exact characterization result, we have restricted like-mindedness between the insiders and the outsiders here to internal events. Also, we do not require like-mindedness among the insiders about these events, as this property will be derived from the others.

We now want to show that common knowledge of like-mindedness and uninformativeness imply consistency of the insiders' beliefs with the common prior assumption, with the outsiders beliefs functioning as "common prior". This is formalized as follows. The first part is a local definition of "Harsanyi consistency among insiders"; the second states that the outsiders' beliefs over $\mathcal{A}(I)$, conditioned on the insiders' common knowledge component, serves as their internal common prior.

Definition 7 *i)* **(Internal Harsanyi Consistency)** $\alpha \in \mathbf{IHC}_\mu$ *if for all* $\beta \in \mathcal{P}_I^*(\alpha)$ *and all* $i \in I : \mu(\|q_i = q_i^\beta\|) > 0$ *and* $q_i^\beta = \mu(./\|q_i = q_i^\beta\|)$.

ii) **(External Harsanyi Consistency)** $\alpha \in \mathbf{EHC}$ *if* $\alpha \in \mathbf{IHC}_{q_o^\alpha(./\mathcal{P}_I^*(\alpha))}$.

To illustrate the interplay of the various definitions, consider the following example with a single "inside" agent (for maximum simplicity) and an outside observer. Note that while the notion of External Harsanyi Consistency is conceptually uninteresting in this case, it is mathematically non-vacuous.

Example 3.

State of the world	$\omega_1 = \tau$	ω_2	ω_3	ω_4
State of nature	success	failure	success	failure
Ego	0.4	0.4	0.1	0.1
Alter	0.5	0.5	0.7	0.3

Let Alter denote the single insider i, and Ego the outsider o. $\mathcal{A}(I)$ is simply the partition {success,fail}= $\{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$. Hence **LMO** $=\{\omega_1, \omega_2\}$; thus at τ, Alter knows Ego to be like-minded, while Ego is unsure of this. Since Ego's beliefs are commonly known, **UIO** $= \Omega$. In the single-insider case, **IHC** $= \Omega$ by Introspection. On the other hand, while there obviously does not exist a common prior among the two agents (at any state), **EHC** $=\{\omega_1, \omega_2\}$, since $q_o^\omega(./\|p_i = p_i^\omega\|) = (0.5, 0.5)$ for all ω. Thus, at τ, Alter knows that his beliefs are EHC with Ego's, while Ego does not.[16]

It is also instructive to let Ego take the role of the insider, and Alter that of the outsider. Again $\mathcal{A}(I = \{Ego\}) = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$. In this case, again **UIO** $= \Omega$, this time non-trivially in that Ego does not know Alter's beliefs, but does not care to know, either. Again, we have **LMO** $=$ **EHC** $=\{\omega_1, \omega_2\}$. Note that, at τ, the insiders' beliefs are EHC with the outsider, but the insider is not sure of this.

This example reveals a couple of general features of the notion of External Harsanyi Consistency. First, EHC does not imply that insiders know the outsider's beliefs on $\mathcal{A}(I)$, nor does it imply that the outsider knows the insiders to be consistent, or that he knows their internal common prior if they are consistent. Second, EHC neither implies nor is implied by Harsanyi Consistency among the insiders together with the outsider. Third, while **IHC** $= \mathbf{K}_I^*(\mathbf{IHC})$, **EHC** can be strictly contained in $\mathbf{K}_I^*(\mathbf{EHC})$.[17]

Common knowledge of like-mindedness and uninformativeness implies common knowledge of External Harsanyi Consistency and a bit more, namely "Regularity" of the outsider's beliefs defined as follows.

Definition 8 ("Regularity") $\alpha \in$ **REG** *if, for any* $\beta, \gamma \in \mathcal{P}_o(\alpha)$ *and any* $i \in I$*:* $q_o^\alpha(./\|p_i = p_i^\beta\|) = q_o^\alpha(./\|p_i = p_i^\gamma\|)$ *whenever* $q_i^\beta = q_i^\gamma$ *and* $\mathcal{P}_I^*(\beta) = \mathcal{P}_I^*(\gamma)$.

Intuitively, Regularity says that, in learning from some insider i about internal events, the outsider cares only about $i's$ own beliefs about these events, as well as what is commonly known among the insiders. This seems very weak; roughly speaking, it only excludes a *further* role of the insider's beliefs about (other's beliefs about) the outsider.

The following is the main result of the paper; it says in particular that when like-mindedness and uninformativeness are common knowledge, the insiders' beliefs are externally Harsanyi consistent.

Theorem 1 $K_I^*(\mathbf{LMO}) \cap K_I^*(\mathbf{UIO}) = K_I^*(\mathbf{EHC}) \cap K_I^*(\mathbf{REG})$.

[16] In the degenerate case of single insider, it is easily verified that **EHC** $=$ **LMO** $\cap$ **UIO**.

[17] Indeed, $\alpha \in \mathbf{K}_I^*(\mathbf{EHC})$ if and only if $\alpha \in \mathbf{EHC} \cap \mathbf{K}_I^*(\|q_o^\omega(./\mathcal{P}_I^*(\omega)) = q_o^\alpha(./\mathcal{P}_I^*(\alpha))\|)$.

The "converse" direction of Theorem 1, i.e. the entailment of like-mindedness and uninformativeness by common knowledge of External Harsanyi Consistency plus Regularity is valuable for two reasons. Firstly, it ensures that common knowledge of like-mindedness and uninformativeness entails *only* the Harsanyi consistency of the belief-hierarchies among insiders, without further restrictions. Secondly, it shows that "being uninformative" is part and parcel of thinking of the outsider as bearer of the common prior, rather than merely an ad-hoc auxiliary assumption.

As to the proof of Theorem 1 (which can be found in the appendix), the common knowledge of like-mindedness and uninformativeness at some state α immediately imply common knowledge of a rudimentary version of **EHC**, namely that, for all $\beta \in \mathcal{P}_I^*(\alpha)$,

$$q_i^\beta = q_i^\beta(./\|p_o = p_0^\beta\|) = q_o^\beta(./\|p_i = p_i^\beta\|).$$

The difficulty of the proof is to show that the putative prior at α $q_o^\alpha(./\mathcal{P}_I^*(\alpha))$ has global reach, i.e. that $q_o^\beta(./\|p_i = p_i^\beta\|)$, which turns out to equal $q_o^\beta(./\|q_i = q_i^\beta\| \cap \mathcal{P}_I^*(\alpha))$ by regularity, in turn is identical to $q_o^\alpha(./\|q_i = q_i^\beta\| \cap \mathcal{P}_I^*(\alpha))$. The key to this is a lemma that shows that the putative prior $q_o^\alpha(./\mathcal{P}_I^*(\alpha))$ has full support, i.e. that, for all $\beta \in \mathcal{P}_I^*(\alpha)$, $q_o^\alpha(\|q_I = q_I^\beta\|/\mathcal{P}_I^*(\alpha)) > 0$.[18]

An interesting special case of Theorem 1 arises in situations in which the outsider knows what is common knowledge among insiders. This is likely to happen when the group of insiders is large and/or heterogeneous, simply because in such situations little will be common knowledge among insiders. Note that since in this case by definition $p_o^\alpha(\mathcal{P}_I^*(\alpha)) = 1$, the putative prior $q_o^\alpha(./\mathcal{P}_I^*(\alpha))$ is given by the observer's unconditional beliefs q_o^α. If this is common knowledge among all together with like-mindedness and uninformativeness, then q_o^α and indeed p_o^α will be commonly known; moreover, by like-mindedness, the latter will amount to a common prior over all agents. Formally, let **WIO** ("weakly informed outsider") denote the event that the outsider knows all that is common knowledge among insiders, $\mathbf{WIO} := \cap_{E \in 2^\Omega}[(K_I^* E)^c \cup K_o E]$.[19] The following is a corollary to Theorem 1. The assumption $\mathcal{A}(J) = 2^\Omega$ is merely technical and means that states are fully specified as belief-hierarchies about Θ among J, and contain no extra information.

Proposition 3 *Suppose $\alpha \in K_J^*(\mathbf{WIO})$ and assume $\mathcal{A}(J) = 2^\Omega$. Then $\alpha \in \mathbf{HC}_{p_o^\alpha}$ if and only if $\alpha \in K_J^*(\mathbf{LMO}) \cap K_J^*(\mathbf{UIO})$.*

[18] This is a non-trivial consequence of uninformativeness; it is for example in general not true that for all $\beta \in \mathcal{P}_I^*(\alpha)$, $p_o^\alpha(\{\beta\}/\mathcal{P}_I^*(\alpha)) > 0$. The latter will fail to be the case if at β some agent is certain that the outsider cannot have the beliefs p_o^α.

[19] Note that within $K_J^*(\mathbf{WIO})$, common knowledge with respect to I and J are the same. Formally, one can show that $K_J^*(\mathbf{WIO}) \cap K_I^* E = K_J^*(\mathbf{WIO}) \cap K_J^* E$, for all $E \in 2^\Omega$.

Thus Proposition 3 *derives* the existence of an individual whose beliefs are public and serve as a common prior. This elimination of uncertainty about the "inside observer" o is a highly non-trivial consequence of the *conjunction* of common knowledge of weak informedness, uninformativeness and like-mindedness, and would in general be far from true in the absence of like-mindedness, as illustrated for instance by example 3 with Alter as the outsider.

5. DISCUSSION

Theorem 1 goes beyond the results available in the literature in two ways: by fully localizing the intersubjective consistency condition, relying on like-mindedness rather than Agreement, and by equating the common prior among insiders with the (conditional) probability of a particular agent. These advances come at the price of referring to the beliefs of an outsider who, moreover, is assumed to be uninformative. We will now discuss these three aspects of Theorem 1 in turn.

5.1 Uninformativeness

Common Knowledge of uninformativeness can be understood as representing, within the static confines of a Bayesian type space in which there is no dynamic notion of receiving information, the notion that insiders treat the outsider as if he had no relevant information beyond information that is publicly available to the insiders. This is certainly a strong (if negative) assumption, but far from an unreasonable one in many circumstances; for example, it is a plausible and standard view of the average private financial investor.[20]

The appeal of the uninformativeness assumption is substantially strengthened by the fact that it must hold merely for *some* suitably chosen individual.[21] Uninformativeness seems plausible especially for individuals sufficiently "far from the scene", that is: with little knowledge about the variables of interest (Θ) and the actors involved. Note that in this case, insiders will typically be uncertain about the outsiders' beliefs.

[20] Note that this is consistent with investors viewing each others as overconfident (and hence not like-minded). However, in this case an investor's beliefs , while uninformative about the fundamentals, may be highly informative about the beliefs of other investors with a similar psychology.

[21] To strengthen the latter point, one can show that in fact the mere possibility (with positive probability) of the existence of such an individual suffices; see the earlier working paper version (Nehring 1998) for details.

5.2. External Harsanyi Consistency

External Harsanyi consistency equates the common prior among insiders with the (conditional) probability of a particular individual. This renders the common prior an ordinary subjective probability, and thereby makes it possible to impose specific assumptions directly on the common prior, in order to specify agents' interactive beliefs by conditioning on their type partitions. This is common modeling practice in information economics and game theory; in fact, we are not aware of a single case in which a model has been defined directly in terms of assumptions on the belief hierarchies.

This practice is commonly justified by assuming a prior stage in which agents beliefs *were* common knowledge, and are remembered ex interim. The standard practice is a fine way of telling a story, but pays the price of simplifying the nature of interactive uncertainty dramatically and by fairly brute force. Clearly, any agent at the prior stage can be viewed as an uninformed outsider whose beliefs are commonly known. Theorem 1 improves on this story by allowing to dispense with the assumptions of common knowledge of agents' belief ex ante and perfect memory of them, thereby preserving a genuine incomplete-information quality. The crux, however, is that at least one ex-ante agent i must be commonly known to be uninformative ex interim. In many situations, this will not be the case, since typically learning about agent i's ex-ante beliefs will be informative about his ex-post beliefs for some other agents, as typically agent i's private information will cut across time. So while sensible in particular situations, we believe the intertemporal interpretation overall to be far less attractive than the interpersonal one.

In contrast to the external characterization of a common prior in Theorem 1, Samet (1998a) has given an internal one in terms of the agents higher-order iterated expectations about random variables. Samet's characterization elegantly does the job it is supposed to do, but does not on its own allow one to transparently interpret assumptions on the common prior itself.

5.3 A Fully Local Characterization

The second and, in our view, the main advance of Theorem 1 over existing characterizations of the CPA is its fully local character, being derived from common knowledge of underlying properties, rather than as equivalent to a property that is intrinsically common known, as is Agreement (the absence of any agreement to disagree) and its variants. As explained in the next section, this makes it possible to derive the CPA from facts about particular agents at particular states, and shows that a *foundation* of the CPA requires positive common knowledge assumptions.

One of the classic implications of the CPA under asymmetric information is the impossibility of mutually profitable equilibrium (common knowledge) trades (Milgrom-Stokey 1982). Originally, this was viewed as a striking consequence of the combination of common priors and common knowledge of rationality. It should cause not a little queasiness, therefore, if in the more general incomplete information setting, the CPA is simply *defined* in terms of absence of disagreement, and thus, effectively, absence of mutually profitable trade.[22] Clearly, the entire *explanatory content* of the No Trade theorem is lost in this approach.

The queasiness is overcome in the present approach. For here, Agreement respectively No Trade become again highly non-trivial consequences of the underlying epistemic premises, and, in particular, continue to rely on substantive common knowledge assumptions; indeed, it would appear that the most direct way to derive a no trade result from common knowledge of like-mindedness and uninformativeness would simply be to take the route via the derived common prior by combining Theorem 1 with Milgrom-Stokey's (1982) original No Trade Theorem.

6. LIKE-MINDEDNESS AND RATIONALITY

So far, to ensure "operationality in principle", we have defined like-mindedness directly in terms of agents' belief hierarchies. This prevents one, however, from understanding like-mindedness rigorously as an expression of agents' rationality, for rationality norms apply directly only to single agents, not to sets of agents. Consider, for example, a two agent situation with complete information in which it is common knowledge that Ego is sober and Alter is drunk. Does rationality require that Ego's beliefs about Alter's accident risk be identical to that of Alter's own beliefs? If so, of whom does rationality require this? Presumably not of Ego: given knowledge of Alter's beliefs, Ego may well beg to differ, on grounds of Alter's inability to form adequate beliefs in his drunken state.[23] On the other hand, does rationality demand of Alter to equate his beliefs to Ego's? Here opinions will differ. Many subjectivists will deny this, deeming the assessment of probabilities an irreducibly

[22] If anything, the no trade theorem is even less trivial under incomplete information, and should a priori be harder to establish, not "easier".

[23] Similar arguments have been used in the single-person intertemporal context against the "Reflection principle" (a special case of Likemindedness), and, by consequence, against the unconditional rationality of belief revision by Bayesian updating.

personal matter[24]; we will refer to this position as *"relativist"*.[25] On the other hand, adherents of the "Harsanyi doctrine" (cf. Aumann (1987)) will argue that, since Alter and Ego share the same information by assumption, Alter should equate his beliefs to Ego's; we will refer to this position as *"rationalist"*. Finally, an intermediate *"pluralist"* view would accept that Alter and Ego may rationally agree to disagree, provided that Alter and Ego differ in "epistemically legitimate" aspects. Being drunk versus being sober may not qualify, but being an "optimist" versus being a "pessimist" (or being a Hegel versus being a Schopenhauer) may, on a pluralist view. The goal of this section is not to argue for one of these positions, nor even to explicate them with any degree of thoroughness. Rather, we want to introduce a framework in which these positions can be captured formally, and to flesh out the content and significance of the central notion of like-mindedness more fully in the process.

The first step is to add to a rooted Bayesian type space in the sense of section 2 two families of mappings $\{y_i : \Omega \to Y\}_{i \in N}$ and $\{\Gamma_{ij} : \Omega \to 2^Y\}_{i,j \in N}$. Here, Y describes the universe of possible "epistemic attitudes" such as the pair of characteristics (sober,optimist); $y_i(\alpha)$ specifies agents i's epistemic attitude at state α; note that it is w.l.o.g. to take this set to be the same for all agents. To sidestep some conceptual issues, we will assume that agents always know their epistemic attitude. The set $\Gamma_{ij}(\alpha) \subseteq Y$ describes the set of epistemic attitudes $y \in Y$ of agent j that agent i (with type $y_i(\alpha)$) views "on par" with himself at state α, that he views as "reasonable", or "**co-rational**", as we shall say to emphasize the relational character of this judgement. Co-rationality of another's attitude will form the basis for an agent to accept like-mindedness as a constraint on his interactive beliefs. $\Gamma_{ij}(\alpha)$ may depend on α, to allow for others' uncertainty about i's co-rationality judgments. Note that Γ_{ij} may well differ from Γ_{ji}. For example, if i is a rationalist while j is a relativist, $\Gamma_{ij}(\alpha)$ may be large while $\Gamma_{ji}(\alpha)$ may be essentially empty.

The enriched framework can be summarized by the following definition.

Definition 9 *An* **extended Bayesian type space** *is a tuple*
$\mathcal{E} = \langle N, \Omega, \tau, \Theta, \theta, \{p_i\}_{i \in N}, Y, \{y_i\}_{i \in N}, \{\Gamma_{ij}\}_{i,j \in N} \rangle$ *, where*

- $\langle N, \Omega, \tau, \Theta, \theta, \{p_i\}_{i \in N} \rangle$ *is a rooted Bayesian type space,*
- Y *is a universe of possible epistemic attitudes*

[24] For an explicit rejection of the normative validity of the common prior assumption, see Morris (1995).

[25] Not as "subjectivist" or "personalist", since these terms apply to single-person decision making, and by themselves do not imply a view on the intersubjective issues at stake here.

- $y_i : \Omega \to Y$ *specifies, for each* $\alpha \in \Omega$, *agent* i'*s epistemic attitude obtaining at* α, *such that* $\alpha \in K_i \| y_i = y_i^\alpha \|$.

- $\Gamma_{ij} : \Omega \to 2^Y$ *, for each* $\alpha \in \Omega$, *the set of types of agent* j *judged by agent* i *to be co-rational to him.*

We are now ready to define "interactive rationality" of an agents' beliefs.

Definition 10 (Interactive Rationality) *The agent* i *is rational at* α *("*$\alpha \in \mathbf{RAT}_i$*) if, for all* $j \in N$, $y_j(\alpha) \in \Gamma_{ij}^\alpha$ *implies* $\alpha \in \mathbf{LM}_{ij}$.

Define events Λ_{ij} by setting $\alpha \in \Lambda_{ij}$ iff $y_j(\alpha) \in \Gamma_{ij}^\alpha$; one can think of Λ_{ij} as "*judged* like-mindedness", in contrast to the *effective* like-mindedness events $\mathbf{LM}_{ij}$ of sections 3 through 5.[26] The link to the central result of the paper, Theorem 1, is established by the following simple observation.

Observation 1 $\mathbf{RAT}_i \cap \mathbf{RAT}_j \cap (\Lambda_{ij} \cup \Lambda_{ji}) \subseteq \mathbf{LM}_{ij}$.
In particular, $K_I^* \left(\cap_{i \in I} \mathbf{RAT}_i\right) \cap K_I^* \left(\cap_{i \in I} \Lambda_{io}\right) \subseteq K_I^*(\mathbf{LMO})$.

Interactive rationality is merely agent-relative, relating given individual co-rationality judgments to entailed like-mindedness restrictions on beliefs. A richer notion of rationality invokes agent-non-relative constraints on the co-rationality conditions. Formally, these can be captured by a **intersubjective rationality norm** $\Gamma \subseteq Y \times Y$; say that an agent is **intersubjectively rational** at α if he is interactively rational at α and if $\Gamma_{ij}^\alpha \supseteq \Gamma$. Then common knowledge of like-mindedness is entailed by common knowledge of intersubjective rationality whenever Γ is sufficiently rich.

As mentioned above, one can expect a wide range of positions on the nature of the correct co-rationality norm Γ. On the one extreme, a *relativist* position is represented by an essentially empty co-rationality norm, leaving everything to unmoored individual judged. On the other extreme, one can imagine the relation to be universal, asserting that any two Bayesian agents are co-rational to each other. But this seems implausibly strong, as the sober-versus-drunk example indicates. A weaker and more sensible view would assert that any agent of sufficient "epistemic competence" is co-rational to any other. This gives rise to a co-rationality relation Γ of the form $\Gamma = Y \times C$, where C is the set of "epistemically competent" types. Epistemic competence can be a matter of

[26]Note that for the purpose of defining interactive rationality, one could have taken a more abstract approach by simply taking the Λ_{ij} as primitives.

brain chemistry, emotional state, intellectual capacity. Intersubjective rationality norms with this structure will be called "*rationalist*". Intermediate between these two positions is the *pluralist* view on which co-rationality is genuinely relational (mathematically: not a product set): for example, a neoclassical economist may be required to respect the belief of any other (competent) neoclassical economist, whether or not the latter has left-wing or right-wing political views. He may not be required to respect the beliefs of a deconstructionist, nor may the latter be required to respect his beliefs. (Would deconstructionists be required to respect other deconstructionist views ??). Different pluralist positions are naturally distinguished by the inclusiveness of Γ. Qualitatively, a natural distinction would between those who put the burden of proof on denying co-rationality (the "*rationalist pluralists*"), and those put that burden on imposing co-rationality (the "*relativist pluralists*").

7. CONCLUSION

The basic goal of this paper was to derive the existence of a common prior among a set of agents from like-mindedness among the agents; it was achieved with the help of a like-minded, uninformed outsider. Essential to the explanatory character of our derivation was the definition of like-mindedness and uninformativeness as fully local properties. In sections 3 through 5, these properties were defined in terms of agents actual beliefs; in section 6, we derived like-mindedness itself from agents' equivalence judgments concerning others epistemic attitudes.

The fully local perspective invites – and should make possible – future generalizations in which agents' beliefs are "almost" consistent with common priors. As indicated in the introduction, such generalizations would be desirable as improvements of the stylized description of empirical reality; on the other hand, they would also demonstrate the conceptual robustness of the CPA and thereby solidify its appeal as an approximation. In a fully local approach, there are two distinct sources of approximation: a merely almost-common knowledge[27] of the underlying condition (like-mindedness), and a continuous weakening of that condition to "almost like-mindedness". In the enriched framework of section 6, almost like-mindedness in turn could be derived from a generalization of the subjective co-rationality relations to subjective similarity metrics over epistemic attitudes, together with an interactive rationality principle according to which similarity of epistemic attitudes across agents implies almost like-mindedness. We leave the development of these ideas to future research.

[27] Presumably in the sense of Monderer-Samet (1989).

APPENDIX: PROOFS

Proof of Fact 1.

i) By Bayesian updating, for any $i \in J$ and any $\beta \in \Omega$, supp $p_{i_2}^{\beta} \subseteq \mathcal{F}_i(\beta)$. In particular, for any $\beta \in \mathcal{F}_i(\alpha)$, supp $p_{i_2}^{\beta} \subseteq \mathcal{F}_i(\alpha)$; i.e. $\mathcal{F}_i(\alpha) \subseteq K_{i_2}(\mathcal{F}_i(\alpha))$. The converse follows from the Truth axiom.

ii) Bayesian updating preserves certain beliefs; by Introspection, these include the initial knowledge of the prior $p_{i_1}^{\alpha}$. This shows $\|p_{i_1} = p_{i_1}^{\alpha}\| \subseteq K_{i_2}(\|p_{i_1} = p_{i_1}^{\alpha}\|)$. The converse follows from Truth or Introspection.

iii) $\underline{\|p_{i_2} = p_{i_2}^{\alpha}\| \supseteq \|p_{i_1} = p_{i_1}^{\alpha}\| \cap \mathcal{F}_i(\alpha)}$.

Take $\alpha \in \Omega$ and $\beta \in \|p_{i_1} = p_{i_1}^{\alpha}\| \cap \mathcal{F}_i(\alpha)$. By Bayesian updating, for any $E \subseteq \Omega$, $p_{i_2}^{\alpha}(E) = p_{i_1}^{\alpha}(E/\mathcal{F}_i(\alpha))$ and $p_{i_2}^{\beta}(E) = p_{i_1}^{\beta}(E/\mathcal{F}_i(\beta))$. Since by assumption $\mathcal{F}_i(\beta) = \mathcal{F}_i(\alpha)$ and $p_{i_1}^{\beta} = p_{i_1}^{\alpha}$, it follows that $p_{i_2}^{\beta} = p_{i_2}^{\alpha}$ as desired.

$\underline{\|p_{i_2} = p_{i_2}^{\alpha}\| \subseteq \|p_{i_1} = p_{i_1}^{\alpha}\| \cap \mathcal{F}_i(\alpha)}$.

Take $\alpha \in \Omega$ and $\beta \in \|p_{i_2} = p_{i_2}^{\alpha}\|$, i.e. such that $p_{i_2}^{\beta} = p_{i_2}^{\alpha}$. By i), $\alpha \in K_{i_2}(\mathcal{F}_i(\alpha))$, hence by the assumption on β also $\beta \in K_{i_2}(\mathcal{F}_i(\alpha))$, which implies $\beta \in \mathcal{F}_i(\alpha)$. Analogously, one obtains from ii) the implication $\beta \in \|p_{i_1} = p_{i_1}^{\alpha}\|$, as needed. $\square$

Proof of Proposition 1.

If the agents reveal their types at date 1, these are always commonly known between them at date 2.

Lemma 1 $K_{\{i_2,j_2\}}^{*}(\|p_{i_2} = p_{i_2}^{\alpha}\| \cap \|p_{j_2} = p_{j_2}^{\alpha}\|) = \Omega.$

Proof of lemma.

By fact 1, for any $\alpha \in \Omega$, $\|p_{i_2} = p_{i_2}^{\alpha}\| = \|p_{i_1} = p_{i_1}^{\alpha}\| \cap \|p_{j_1} = p_{j_1}^{\alpha}\| = \|p_{j_2} = p_{j_2}^{\alpha}\|$. Writing E_α for$\|p_{i_2} = p_{i_2}^{\alpha}\| = \|p_{j_2} = p_{j_2}^{\alpha}\|$, one has from Introspection $\alpha \in E_\alpha = K_{i_2}E_\alpha = K_{j_2}E_\alpha = K_{\{i_2,j_2\}}^{*}E_\alpha$. $\square$.

Proof of the Proposition.

Part 1: $\mathbf{LM}_{\{i_1,j_1\}} \supseteq \mathbf{MLM}_{\{i_2,j_2\}}$.

Take $\alpha \in \mathbf{MLM}_{\{i_2,j_2\}}$. By lemma 1, one has $p_{i_2}^{\alpha} = p_{j_2}^{\alpha}$. By Bayesian updating at α, $p_{i_2}^{\alpha} = p_{i_1}^{\alpha}(./\|p_{j_1} = p_{j_1}^{\alpha}\|)$ as well as $p_{j_2}^{\alpha} = p_{j_1}^{\alpha}(./\|p_{i_1} = p_{i_1}^{\alpha}\|)$. It follows that $\alpha \in \mathbf{LM}_{\{i_1,j_1\}}$.

Part 2: $\mathbf{LM}_{\{i_1,j_1\}} \subseteq \mathbf{MLM}_{\{i_2,j_2\}}$.

Straightforward. $\blacksquare$

Proof of Proposition 2.

We shall prove the modus tollens. Assume thus that $\cap_{i,j\in N}\mathbf{LM}_{i,j} \nsubseteq \Phi$, i.e. that, for some $\mathcal{B} \in \mathcal{T}$ and some $\alpha \in \Omega$, $\alpha \in \cap_{i,j\in N}\mathbf{LM}_{i,j}(\mathcal{B})$, but $\alpha \notin \Phi(\mathcal{B})$. It is easily verified that there exist a type-space $\mathcal{B}'$ that is locally equivalent to $\mathcal{B}$ at α such that $\alpha \in \mathbf{HC}(\mathcal{B}')$. On the other hand, since Φ is fully local, $\alpha \notin \Phi(\mathcal{B}')$. $\square$

Proof of Theorem 1.

$K_I^*(\mathbf{LMO} \cap \mathbf{UIO}) \subseteq K_I^*(\mathbf{EHC})$.

Consider $\alpha \in K_I^*(\mathbf{LMO} \cap \mathbf{UIO})$, $\beta \in \mathcal{P}_I^*(\alpha)$ and any $i \in I$. Take $\gamma \in \mathcal{P}_I^*(\alpha) \cap \mathcal{P}_o(\alpha)$ from lemma 2 below such that $q_I^\gamma = q_I^\beta$. Then by lemma 3 below,

$$q_i^\beta = q_i^\gamma = q_o^\alpha(./\|q_i = q_i^\gamma\| \cap \mathcal{P}_I^*(\alpha)) = q_o^\alpha(./\|q_i = q_i^\beta\| \cap \mathcal{P}_I^*(\alpha)).$$

Lemma 2 *If $\beta \in \mathcal{P}_I^*(\alpha)$ and $\alpha \in K_I^*(\mathbf{UIO})$ then there exists $\gamma \in \mathcal{P}_I^*(\alpha) \cap \mathcal{P}_o(\alpha)$ such that $q_I^\gamma = q_I^\beta$.*

By assumption, using a standard characterization of $\mathcal{P}_I^*(\alpha)$, there exist sequences $\{\omega_k\}_{k=1,...,n}$ in $\mathcal{P}_I^*(\alpha)$ and $\{i_k\}_{k=1,...,n-1}$ in I such that $\omega_1 = \alpha$, $\omega_n = \beta$ and such that $\omega_{k+1} \in \mathcal{P}_{i_k}(\omega_k)$ for $k = 1, ..., n-1$. The proof is by induction on the length ℓ of a sequence connecting α and β. The claim of the lemma holds trivially for ℓ equal to 1. Suppose it to hold for $\ell = n-1$, i.e. that there exists $\delta \in \mathcal{P}_o(\alpha) \cap \mathcal{P}_I^*(\alpha)$ such that $q_I^\delta = q_I^{\omega_{n-1}}$.

First, $\beta \in \mathcal{P}_{i_{n-1}}(\omega_{n-1})$ implies $p_{i_{n-1}}^{\omega_{n-1}}(\|q_I = q_I^\beta\|) > 0$. Since $\|q_I = q_I^\beta\| \in \mathcal{A}(I)$ and $q_{i_{n-1}}^{\omega_{n-1}} = q_{i_{n-1}}^\delta$,

$$p_{i_{n-1}}^\delta(\|q_I = q_I^\beta\|) > 0. \tag{1}$$

We need to show that $p_o^\alpha(\mathcal{P}_I^*(\alpha) \cap \|q_I = q_I^\beta\|) > 0$. By way of contradiction, assume this to be false. I.e., since $p_o^\delta = p_o^\alpha$, assume that

$$p_o^\delta(\mathcal{P}_I^*(\alpha) \cap \|q_I = q_I^\beta\|) = 0.$$

By Truth (of the observer's beliefs), this implies

$$\|p_o = p_o^\delta\| \subseteq (\mathcal{P}_I^*(\alpha) \cap \|q_I = q_I^\beta\|)^c.$$

Since $p_{i_{n-1}}^\delta(\mathcal{P}_I^*(\alpha)) = 1$, therefore

$$p_{i_{n-1}}^\delta(\|q_I = q_I^\beta\|/\|p_o = p_o^\delta\|) = p_{i_{n-1}}^\delta(\mathcal{P}_I^*(\alpha) \cap \|q_I = q_I^\beta\|/\|p_o = p_o^\delta\|) = 0.$$

Since $\delta \in \mathcal{P}_I^*(\alpha) \subseteq \mathbf{UIO}$ and $\|q_I = q_I^\beta\| \in \mathcal{A}(I)$,

$$p_{i_{n-1}}^\delta(\|q_I = q_I^\beta\|) = p_{i_{n-1}}^\delta(\|q_I = q_I^\beta\|/\|p_o = p_o^\delta\|).$$

Combining the last two equations yields

$$p_{i_{n-1}}^\delta(\|q_I = q_I^\beta\|) = 0,$$

in contradiction to (1), as desired. $\square$

Lemma 3 *If $\gamma \in \mathcal{P}_I^*(\alpha) \cap \mathcal{P}_o(\alpha)$ and $\alpha \in K_I^*(\mathbf{LMO} \cap \mathbf{UIO})$, then*

$$q_i^\gamma = q_o^\alpha(./\|q_i = q_i^\gamma\| \cap \mathcal{P}_I^*(\alpha)).$$

Take any $\omega \in \mathcal{P}_I^*(\alpha) \cap \mathcal{P}_o(\alpha)$ such that $q_i^\omega = q_i^\gamma$. Since $\omega \in \mathcal{P}_I^*(\alpha) \subseteq \mathbf{LMO} \cap \mathbf{UIO}$, once can infer that

$$q_i^\omega = q_i^\omega(./\|p_o = p_0^\omega\|) = q_o^\omega(./\|p_i = p_i^\omega\|).$$

Moreover, $q_o^\omega = q_o^\alpha$ since $\omega \in \mathcal{P}_o(\alpha)$. One thus obtains

$$q_i^\gamma = q_i^\omega = q_o^\omega(./\|p_i = p_i^\omega\|) = \ q_o^\alpha(./\|p_i = p_i^\omega\|). \tag{2}$$

Equation (2) implies that

$$\begin{aligned} q_o^\alpha(./\|q_i &= q_i^\gamma\| \cap \mathcal{P}_I^*(\alpha)) = \sum_{\omega \in \|q_i = q_i^\gamma\| \cap \mathcal{P}_I^*(\alpha)} q_o^\alpha(./\|p_i = p_i^\omega\|) \cdot p_o^\alpha(\|p_i = p_i^\omega\|/\|q_i = q_i^\gamma\| \cap \mathcal{P}_I^*(\alpha)) = \\ &= \sum_{\omega \in \|q_i = q_i^\gamma\| \cap \mathcal{P}_I^*(\alpha)} q_i^\gamma \cdot p_o^\alpha(\|p_i = p_i^\omega\|/\|q_i = q_i^\gamma\| \cap \mathcal{P}_I^*(\alpha)) = q_i^\gamma, \end{aligned}$$

which establishes the desired conclusion. $\square$

$\underline{K_I^*(\mathbf{LMO} \cap \mathbf{UIO}) \subseteq K_I^*(\mathbf{REG})}$.

Follows immediately from lemma 3.

$\underline{K_I^*(\mathbf{EHC} \cap \mathbf{REG}) \subseteq K_I^*(\mathbf{LMO} \cap \mathbf{UIO})}$.

Take $\alpha \in K_I^*(\mathbf{EHC} \cap \mathbf{REG})$, $\beta \in \mathcal{P}_I^*(\alpha)$ and any $i \in I$.

Since $\beta \in \mathbf{EHC}$, $q_i^\beta = q_o^\beta(./\|q_i = q_i^\beta\| \cap \mathcal{P}_I^*(\beta))$.

Since also $\beta \in \mathcal{P}_i(\beta)$ by Truth and $\beta \in \mathbf{REG}$, $q_o^\beta(./\|q_i = q_i^\beta\| \cap \mathcal{P}_I^*(\beta)) = q_o^\beta(./\|p_i = p_i^\beta\|)$, hence also

$$q_i^\beta = q_o^\beta(./\|p_i = p_i^\beta\|). \tag{3}$$

It follows that

$$q_i^\beta(./\|p_0 = p_0^\beta\|) = q_o^\beta(./\|p_i = p_i^\beta\| \cap \|p_0 = p_0^\beta\|) = q_o^\beta(./\|p_i = p_i^\beta\|) = q_i^\beta, \tag{4}$$

Equation 4 yields $\beta \in$ **UIO** immediately.

Similarly, the conjunction of equations 3 and 4 yields $\beta \in$ **LMO**. ■

Proof of Proposition 3.

Necessity is straightforward. To prove sufficiency, take $\alpha \in K_J^*(\mathbf{WIO}) \cap K_J^*(\mathbf{LMO}) \cap K_J^*(\mathbf{UIO})$. Note that at $\alpha \in K_J^*(\mathbf{WIO})$,common knowledge with respect to I and J are the same, i.e. $\mathcal{P}_J^*(\alpha) = \mathcal{P}_I^*(\alpha)$; assume w.l.o.g. that $\mathcal{P}_J^*(\alpha) = \Omega$. Hence $q_o^\alpha(./\mathcal{P}_I^*(\alpha)) = q_o^\alpha$ and thus, by Theorem 1, for all $\beta \in \Omega$, $\beta \in \mathbf{EHC}_{q_o^\alpha(./\mathcal{P}_I^*(\alpha))} = \mathbf{EHC}_{q_o^\alpha}$, and $q_o^\beta = q_o^\alpha$; thus q_o^α is constant on Ω. It follows that for all $E \in \mathcal{A}(I)$ and all $c \in [0,1] : \|q_o(E) = c\|$ equals Ω or $\emptyset$. Thus $\mathcal{A}(I)$ is in fact belief-closed for J, not just I, which implies that $\mathcal{A}(I) = \mathcal{A}(J) = 2^\Omega$. But this means that $q_o^\alpha = p_o^\alpha$, from which the claim follows. ■

REFERENCES

[1] Armbruster, W. and W. Boege (1979), Bayesian Game Theory, in *Game Theory and Related Topics,* ed. by O. Moeschlin and D. Pallaschke, Amsterdam, North-Holland.

[2] Aumann, R. (1976), Agreeing to Disagree, *Annals of Statistics,* 4, 1236-1239.

[3] Aumann, R. (1987), Correlated Equilibrium as an Expression of Bayesian Rationality, *Econometrica,* 55. 1-18.

[4] Aumann, R. (1998), Common Priors: A Reply to Gul, *Econometrica,* 66, 929-938.

[5] Binmore, K. and A. Brandenburger (1990), Common Knowledge and Game Theory, in: K. Binmore, *Essays on the Foundations of Game Theory,* Oxford, Basil Blackwell.

[6] Blume, L., Brandenburger, A. and E. Dekel (1986), Lexicographic Probabilities and Choice under Uncertainty, *Econometrica,* 59, 61-79.

[7] Boege, W. and T. Eisele (1979), On Solutions of Bayesian Games, *International Journal of Game Theory* 8, 193-215.

[8] Bonanno, G. and K. Nehring (1999), How to Make Sense of the Common Prior Assumption under Incomplete Information, *International Journal of Game Theory,* 28, 409-434.

[9] Bonanno, G. and K. Nehring (1997), Assessing the Truth Axiom under Incomplete Information, Working Paper, University of California, Davis, forthcoming *Mathematical Social Sciences.*

[10] Bonanno, G. and K. Nehring (1998), On the Logic and Role of Negative Introspection of Common Belief, *Mathematical Social Sciences,* 35, 17-36..

[11] Dekel, E. and F. Gul (1997), Rationality and Knowledge in Game Theory, in: Kreps D. M. and K.F. Wallis (eds.), *Advances in Economics and Econometrics*, vol. 1, Cambridge, Cambridge UP, 87-172.

[12] Feinberg, Y. (2000), Characterizing the Existence of a Common Prior in the Form of Posteriors, *Journal of Economic Theory,* 91, 127-179.

[13] van Fraassen, B. (1984), Belief and the Will, *Journal of Philosophy,* 81, 235-256.

[14] Geanakoplos, J. and H. Polemarchakis (1982), We can't disagree forever, *Journal of Economic Theory,* 26, 363-390.

[15] Gul, F. (1998), A Comment on Aumann's View, *Econometrica,* 66, 923-927.

[16] Halpern, J. (2002), Characterizing the Common Prior Assumption, *Journal of Economic Theory,* 106, 316-355.

[17] Harsanyi, J. (1967-86), Games with Incomplete Information Played by Bayesian Players, Parts I-III, *Management Science,* 8, 159-182, 320-334, 486-502.

[18] Lipman, B. (1997), Approximate Common Priors, mimeo, University of Western Ontario.

[19] Maher, P. (1993), *Betting on Theories,* Cambridge University Press.

[20] Mertens, J.-F. and S. Zamir (1985), Formulation of Bayesian Analysis for Games with Incomplete Information, *International Journal of Game Theory,* 14, 1-29.

[21] Milgrom, P. and N. Stokey (1982), Information, Trade and Common Knowledge, *Journal of Economic Theory,* 26, 12-27.

[22] Morris, S. (1994), Trade with Heterogeneous Prior Beliefs and Asymmetric Information, *Econometrica,* 62, 1327-1347.

[23] Morris, S. (1995), The Common Prior Assumption in Economic Theory, *Economics and Philosophy,* 11, 227-253.

[24] Nehring, K. (2001), Common Priors under Incomplete Information: A Unification, *Economic Theory,* 18, 535-553.

[25] Samet, D. (1998a), Iterated Expectations and Common Priors, *Games and Economic Behavior* 24, 131-141.

[26] Samet, D. (1998b), Common Priors and the Separation of Convex Sets, *Games and Economic Behavior* 24, 173-175.