Check for updates

Econometrica, Vol. 71, No. 4 (July, 2003), 1255–1267

FINITE ORDER IMPLICATIONS OF COMMON PRIORS

BY BARTON L. LIPMAN[1]

I characterize the implications of the common prior assumption for finite orders of beliefs about beliefs at a state and show that in finite models, the only such implications are those stemming from the weaker assumption of a common support. More precisely, given any finite N and any finite partitions model where priors have the same support, there is another finite partitions model with common priors that has the same nth order beliefs and knowledge for all $n \leq N$.

KEYWORDS: Common prior assumption, universal belief space, common knowledge, common belief.

1. INTRODUCTION

THE COMMON PRIOR ASSUMPTION (CPA) is used in virtually all of game theory and information economics. I give a precise definition below, but for now, I will say that the CPA is the assumption that the beliefs of every agent are generated by updating a single prior over the state space. It is crucial to many important results such as Aumann's (1976) agreeing-to-disagree result and the no-trade theorem (see, e.g., Milgrom and Stokey (1982)). Morris (1995) summarizes the intense debate regarding the assumption.

I characterize the finite order implications of the common prior assumption—that is, what restrictions it imposes for beliefs about the beliefs about . . . the beliefs of others, where "beliefs about" is repeated only a finite number of times and analogously for knowledge. I show that in finite models, the only finite order implications of the CPA are those stemming from the weaker assumption that priors have a common support. This fact has three implications. First, in practical terms, the distinction between common priors and common knowledge is smaller than one might think. Second, the primary force of the common prior assumption is on the entire infinite hierarchy of beliefs, not any finite portion of it. Finally, there is a sense in which "almost all" priors are common. However, the topology in which this is true has drawbacks.

A precise statement of the result requires some background. Imagine we have a complete description of the world including the true value of some parameter that is unknown to the agents, the beliefs of the agents (their *first-order beliefs*) about this parameter, their beliefs about the first-order beliefs of others (*second-order beliefs*), their beliefs about the second-order beliefs of others (*third-order beliefs*), etc. I refer to a specification of a parameter value together with an infinite hierarchy of beliefs for each player as a *world*, reserving *state* for a related notion. As Mertens and Zamir (1985) showed, such a description of the *actual* world generates a collection of *possible* worlds, one of which is the

[1] I wish to thank Jim Bergin, Eddie Dekel, Joe Halpern, Stephen Morris, Klaus Nehring, Shmuel Zamir, seminar audiences at NYU, Columbia, and Pittsburgh, and especially an editor and four anonymous referees for helpful comments, SSHRCC and the NSF for financial support, and the University of Pennsylvania for an enjoyable visit during which some of this research was carried out. Former title: "Approximately Common Priors."

actual world. These other worlds are fictitious constructs, used to clarify our understanding of the actual world.[2] The actual world also identifies a prior over the set of possible worlds for each agent. This prior is not unique since any prior yielding the same conditional probabilities on the appropriate events would suffice. However, the actual world does determine whether there is a single prior we could attribute to all agents.

In other words, the actual world—the infinite hierarchy of beliefs about beliefs for each player—completely determines whether or not the beliefs satisfy the CPA. This may seem surprising since we usually describe the common prior assumption as an *ex ante* statement, one regarding the agent's beliefs before getting information. But Mertens and Zamir's theorem tells us that whether the CPA holds is determined at the interim level—that is, in the actual world or after receipt of information.

To state my question more precisely, then, what are the implications of the common prior assumption for finite order beliefs at a world? The answer depends on two issues concerning the definition of the CPA. First, does having a common prior require agents to give positive probability to the true world? While most models assuming common priors also assume that the only possible worlds are those in the support of the common prior, an assumption often called *nondelusion*, these are conceptually distinct assumptions. Nondelusion is a statement about the relationship between beliefs and the external world, not the relationship between the beliefs of different agents. Consequently, I define the CPA without requiring nondelusion.[3]

Second, is a common prior supposed to be *consistent with* the beliefs of all agents or to *generate* these beliefs? If the common prior gives zero probability to an agent receiving certain information, then it does not generate that agent's beliefs in response to that information in the sense that the beliefs do not derive from the prior by Bayes' rule. On the other hand, if each agent's beliefs are generated by updating the common prior whenever his information does have positive prior probability, then beliefs are *consistent* with the common prior. I define the *common prior assumption* to be the stronger requirement: that there is a single prior with the property that for every agent, every information event has positive prior probability and the agent's beliefs are generated by Bayesian updating.[4]

For an example inconsistent with common priors, suppose there are two agents and two parameter values, *high* and *low*. Suppose agent 1's second-order beliefs give probability 1 to the pair of statements: (1) the parameter value is *low* and (2) agent 2 puts probability 0 on the parameter value being *low*. Clearly, agents cannot update from a common prior to such beliefs. More generally, one implication of the CPA for finite levels of beliefs is that it is common knowledge that each agent gives probability 0 to the occurrence of an event while some agent gives that event 0 probability. I refer to this property as *weak consistency*. It is not hard to see that weak consistency is equivalent to the priors of the agents having the same support.

My main result is that this is the only restriction the CPA has for finite order beliefs at a world generated by a finite partitions model. (I explain the use of partitions models in the next section. See Remark 1 in the following section on the role of finiteness.) That is, given any world ω from a finite partitions model satisfying weak consistency and any finite

[2] See Dekel and Gul (1997) for a similar perspective.

[3] It is easy to use my results to show that adding nondelusion to the definition of the CPA changes my characterization only by adding nondelusion to the list of finite order implications.

[4] It is not difficult to show that the weaker version of the CPA has virtually no finite order implications if we treat nondelusion as separate from the CPA. For further discussion, see Bonnano-Nehring (1999) and Halpern (2002).

N, there is another world, ω^N, which is consistent with the common prior assumption and where for all $n \leq N$, nth order beliefs and knowledge at ω^N are the same as those at world ω. This implies that dropping the CPA is almost the same as dropping common knowledge assumptions. More precisely, given a world where certain facts are common knowledge but priors are different, we can find a world with common priors in which nothing changes except that these facts are arbitrarily high mutual knowledge but *not* common knowledge.

The theorem shows that the CPA is primarily an infinite order condition. Since weak consistency is the only implication of the CPA for finite orders of beliefs, the main implications of common priors are on the full infinite hierarchy of beliefs, not any finite portion. This view is supported by Bannano and Nehring (1999) and Feinberg (2000). These authors show that the common prior assumption is equivalent to no common knowledge of disagreement[5] and hence to a statement about infinitely many levels of beliefs. My result holds because no common knowledge of disagreement has minimal implications for any finite set of levels of beliefs.[6]

The theorem has an unexpected corollary: the set of worlds consistent with common priors is dense in the set of all worlds. On the other hand, I discuss below some concerns with the topology used for this result.

The next section contains basic definitions and an example. The results are in Section 3. Section 4 contains concluding remarks. Except as noted, proofs are in the Appendix.

2. THE MODEL AND AN ILLUSTRATIVE EXAMPLE

2.1. *Model*

The simplest way to construct worlds is by use of a *partitions model*. Given a *parameter space* Θ, and a finite set of players $\mathcal{I} = \{1, \ldots, I\}$, I define a partitions model to consist of the following objects:

1. A finite set of *states*, S.
2. A function $f : S \to \Theta$ where $f(s)$ is the value of the unknown parameter in state s.
3. For each player $i \in \mathcal{I}$, a partition Π_i of S, where $\pi_i(s)$ denotes the event of Π_i containing s.
4. For each player $i \in \mathcal{I}$ and each event $\pi \in \Pi_i$, a conditional belief $\mu_i(\cdot \mid \pi)$ over π.

I use $\mathcal{M}$ to denote a typical partitions model $(S, f, \{(\Pi_i, \mu_i)\}_{i\in\mathcal{I}})$.

A *state* is an implicit representation of a world, where the hierarchy of beliefs is defined through the partitions and conditional beliefs. Given a state s, we can determine the world to which s corresponds recursively. First, $f(s)$ is the parameter value for the world to which s corresponds. Next we must identify each player's first-order beliefs at s. To do so, note that μ_i together with f induces a probability distribution on Θ for each $\pi \in \Pi_i$. This gives i's first-order beliefs. Formally, the first-order beliefs of i at state s, say $\delta_i^1[s]$, are defined by

$$\delta_i^1[s](B) = \mu_i(f^{-1}(B) \mid \pi_i(s))$$

for each measurable $B \subseteq \Theta$. Player i's second-order beliefs are determined by the partitions for the other players. In the two agent case, player 1's second-order beliefs at s, say

[5] A similar result is contained in Morris (1994). See Samet (1996) for a very elegant restatement.

[6] I thank Klaus Nehring for pointing out this interpretation.

$\delta_1^2[s]$, are a probability distribution over $\Theta \times \Delta(\Theta)$, where $\Delta(\Theta)$ is the set of first-order beliefs for player 2.[7] Formally,

$$\delta_1^2[s](B) = \mu_1(\{s' \in S \mid (f(s'), \delta_2^1(s')) \in B\} \mid \pi_1(s))$$

for each measurable $B \subseteq \Theta \times \Delta(\Theta)$. Continuing recursively, we can define $\delta_i^n[s]$ for every finite n, every agent i, and every state s. Let $\delta_i[s] = (\delta_i^1[s], \delta_i^2[s], \dots)$. The world corresponding to state s, denoted $\omega(s)$, is $(f(s), \delta_1[s], \dots, \delta_I[s])$.

In short, given a partitions model and a state in that model, we can construct a world; that is, a point in the Mertens-Zamir universal beliefs space.[8] However, a given world in the universal beliefs space corresponds to many states in many different partitions models partly because we can separate *belief* and *knowledge* in a partitions model. The conditional probability distributions define probability beliefs at each state as discussed above. In addition, the partitions define nonprobabilistic knowledge. For example, suppose $\{s_1, s_2\}$ is an event in player i's partition but $\mu_i(s_1 \mid \{s_1, s_2\}) = 0$. If we think of knowledge as probability 1 belief, then we could interpret this as saying that the agent knows the state is s_2. However, contrast this partitions model with one where i's partition event is $\{s_2\}$. In the latter case, there is no ambiguity that the agent knows s_2 is true. In the former, we could think of the agent giving s_1 some infinitesimal probability. While I do not claim that this difference is critical, it is a formal difference that could be important in some contexts. However, in either case, s_2 corresponds to the same world in the universal beliefs space since this depends on beliefs only. My results on matching beliefs do include probability 1 beliefs; hence if one prefers this definition of knowledge, one can simply ignore my results on matching knowledge as well.[9]

A partitions model satisfies the common prior assumption if there exists a probability distribution μ on S with $\mu(\pi) > 0$ for all $\pi \in \cup_i \Pi_i$ such that for all i, all $\pi \in \Pi_i$, and all $E \subseteq S$, $\mu_i(E \mid \pi) = \mu(E \cap \pi)/\mu(\pi)$.

2.2. *An Example*

Suppose we have two players, two possible values of the unknown parameter, $\Theta = \{\theta_1, \theta_2\}$, and two states, $S = \{s_1, s_2\}$. The function f is given by $f(s_j) = \theta_j$, $j = 1, 2$. Both players' partitions are trivial: $\Pi_i = \{S\}$ for $i = 1, 2$. Because the partitions are trivial, the conditional beliefs are priors. Assume player 1's prior is $\mu_1(s_1) = 2/3$, while player 2's prior has $\mu_2(s_1) = 1/3$. Let $\mathscr{M}$ denote this partitions model.

[7] We could equivalently define second-order beliefs to be a distribution over $\Theta \times \Delta(\Theta) \times \Delta(\Theta)$—that is, over Θ and the first-order beliefs of both players. Naturally, we would then impose the requirement that each player knows his own beliefs. The approach in the text follows Brandenburger and Dekel (1993), while the approach described in this footnote follows Mertens–Zamir.

[8] While not directly relevant, it is true that any point in the universal beliefs space can be constructed in this fashion from some partitions model, though this converse requires a less restrictive class of partitions models than I use. See Tan and Werlang (1988) or Brandenburger and Dekel (1993). This relationship between partitions models and the universal beliefs space makes it possible to relate certain theorems in one setting to theorems in the other. However, this statement does not imply that *all* theorems in one setting can be immediately translated into the other.

[9] To be more precise, one can change the definition of weak consistency (see Section 3) to $\tau = S$ and thus eliminate the use of knowledge there. If one does so, my proof that weak consistency of a state implies that is consistent with the common prior assumption to every finite level goes through unchanged.

For concreteness, interpret the parameter as the value of an asset owned by player 1 where $\theta_1 = 1$ and $\theta_2 = 2$ and assume the players are risk neutral. At both states, player 1's expected value is 4/3, while player 2's expected value is 5/3. Hence it is common knowledge that both agents would strictly gain if 2 bought the asset from 1 for, say, 3/2. The no-trade theorem implies that if risk neutral agents have common priors, it can never be common knowledge that there are strict gains from trade. Thus this example illustrates the well-known fact that the conclusion of the no-trade theorem can fail without common priors.

I now show that we can match beliefs at each state to an arbitrarily high order in a partitions model with common priors. I do so by constructing a sequence of partitions models where each model has a state s^N that matches the beliefs at s in $\mathcal{M}$ up to order N. In each of these models, the difference in beliefs is replaced by a difference in information. However, this replacement is not without cost: we *cannot* avoid changing beliefs at very high orders.

So fix an integer N and construct the Nth model, $\overline{\mathcal{M}}^N$, as follows. The set of states is $\overline{S}^N = S \times \{1, \dots, 2^N\}$. Intuitively, state (s, k) is the kth copy of state s from our original partitions model. Accordingly, the new function relating states and the unknown parameter, $\bar{f}^N$, is given by $\bar{f}^N(s, k) = f(s)$. The common prior is that all states in $\overline{S}^N$ are equally likely.

To generate 1's beliefs, his information must reflect a greater likelihood of state s_1 than s_2. Hence his partition, $\overline{\Pi}_1^N$, includes all events of the form $\{(s_1, 2k-1), (s_1, 2k), (s_2, k)\}$ for $k = 1, \dots, 2^{N-1}$. Since player 1 "uses up" his s_1 copies faster than his s_2 copies, the partition also includes $\{(s_2, 2^{N-1}+1), \dots, (s_2, 2^N)\}$. Note that at *every* event except the last, player 1's beliefs give probability 2/3 to θ_1 and 1/3 to θ_2. The partition for player 2, $\overline{\Pi}_2^N$, is analogous: it contains every event of the form $\{(s_1, k), (s_2, 2k-1), (s_2, 2k)\}$ for $k = 1, \dots, 2^{N-1}$ plus the event $\{(s_1, 2^{N-1}+1), \dots, (s_1, 2^N)\}$. Again, at all events except this last one, player 2's beliefs over Θ match his beliefs from the original model.

It is not hard to show that nth order beliefs at $(s, 1)$ in $\mathcal{M}^N$ are the same as nth order beliefs at s in $\mathcal{M}$ for all $n \leq N$ for $s = s_1, s_2$. As noted above, we have matched first-order beliefs: in state $(s, 1)$, player 1's first-order beliefs put probability 2/3 on θ_1 and 1/3 on θ_2 or, in the obvious vector notation, (2/3, 1/3), the same as his first-order beliefs at s_1 in the original partitions model. Similarly for player 2. To get second-order beliefs for player 1 at $(s, 1)$, note that the events of 2's partition that 1 considers possible are $\{(s_1, 1), (s_2, 1), (s_2, 2)\}$ and $\{(s_1, 2), (s_2, 3), (s_2, 4)\}$. The set of first-order beliefs for 2 at these events form the support of player 1's second-order beliefs. 2's first-order beliefs at each of these events are (1/3, 2/3), so player 1 puts probability 1 on these being player 2's first-order beliefs, the same as 1's second-order beliefs at s_1 in the original model. Again, player 2 is analogous.

In general, player 1's nth order beliefs are determined by the partition events that are $n-1$ "steps" from $(s, 1)$. The event in 1's partition containing $(s, 1)$ is one step away from $(s, 1)$. An event in i's partition is n steps away from $(s, 1)$ if it intersects an event in j's partition, $j \neq i$, which is $n-1$ steps from $(s, 1)$. It is not hard to see that all events less than N steps away from $(s, 1)$ have the same first-order beliefs as in the original model. Hence player 1 attaches probability 1 to player 2 attaching probability 1 to... to player 1's beliefs being (2/3, 1/3) and player 2's beliefs being (1/3, 2/3) up to order N.

The partitions model $\overline{\mathcal{M}}^N$ is "like" $\mathcal{M}$ in the sense that beliefs at certain states in $\overline{\mathcal{M}}^N$ are the same as beliefs in $\mathcal{M}$ up to a high but finite order. On the other hand, the ex ante models are quite different. In the original partitions model, there is zero probability that player 1 puts probability 1 on θ_2 but this event has positive probability in the new partitions model. Put differently, "closeness" is an interim notion, focusing on closeness of actual worlds, not ex ante models.

If we define knowledge to mean probability 1 belief, then all statements regarding knowledge are preserved up to level N since these are special cases of statements about beliefs. Even if we define knowledge in terms of partitions, however, such statements are preserved. A formal statement of this is notationally heavy and so is deferred.

While the new partitions model matches statements about mutual knowledge to a very high degree, it does so only up to a *finite* degree. Hence statements about common knowledge are not matched. In the original model, beliefs over Θ are common knowledge. In the new model, both players know that both know that . . . that both know the beliefs over Θ for a very large number of "know that" 's. But the beliefs are *not* common knowledge.

The importance of this point is underscored when we interpret the parameter as the value of an asset held by player 1 where $\theta_1 = 1$ and $\theta_2 = 2$. Note that at $(s, 1)$ in $\mathscr{M}^N$, player 1's expected value of the asset is 4/3 and player 2's is 5/3, as in the original model. However, these beliefs are not common knowledge. Instead, for all $n \leq N$, it is mutual knowledge of order n at $(s, 1)$ that these are the expected values. Hence it is mutual knowledge of every order $n \leq N$ that there are gains from trade. However, at $(s, 1)$ it is *not* mutual knowledge of order $N + 1$ that there are gains from trade. Hence this example illustrates that the conclusion of the no-trade theorem does not hold if we replace common knowledge with mutual knowledge of high but finite order. The fact that the no-trade theorem no longer holds if we drop the common prior assumption or if we replace common knowledge with mutual knowledge of finite order is well-known. The new observation is that these two results are the same result.

A natural question to ask is whether we can match beliefs at *all* orders instead of only up to point. It is easy to see that extending the construction above does not work. Replacing N with ∞ gives a prior that is uniform on a countable set—an impossibility. In fact, Mertens–Zamir's results imply that no such construction is possible. They showed that a specification of all finite orders of beliefs determines whether the beliefs are consistent with common priors. Hence if they are not, there is no way to change the model to one with a common prior while matching *all* orders of beliefs.

3. RESULTS

I first require further definitions. Certain events in the state space for a partitions model $\mathscr{M}$ correspond to situations we can describe in terms of Θ and $\mathscr{I}$ only. Since these two sets are used for all partitions models, we can relate such events across models. These events will refer to levels of beliefs and knowledge, so it is useful to first define

$$K_i(E) = \{s \in S \mid \pi_i(s) \subseteq E\},$$

$$B_i^p(E) = \{s \in S \mid \mu_i(E \mid \pi_i(s)) = p\}.$$

The former is standard: $K_i(E)$ is the set of states where every state that i believes possible is contained in E or, more colloquially, where i knows E. The latter is the set of states where i gives probability p to the event E.[10] Both operators depend on the partitions model $\mathscr{M}$ but this dependence is omitted from the notation. Instead, if I need to refer to

[10] While the notation I use is the same as theirs, this should not be confused with Monderer and Samet's (1989) notion of p-belief.

operators or events from different models, say $\mathcal{M}$ and $\overline{\mathcal{M}}$, I will use a $^-$ on top of a variable to refer to the object from $\overline{\mathcal{M}}$.

Let

$$A_1 = \{E \subseteq S \mid E = f^{-1}(\widehat{\Theta}) \text{ for some } \widehat{\Theta} \subseteq \Theta\}.$$

For any collection of subsets $\mathcal{E}$ of S, let $\sigma(\mathcal{E})$ denote the σ-algebra generated by these sets. Let $\mathcal{A}_1 = \sigma(A_1)$. I refer to these events as *1 measurable*. Recall that first-order beliefs are probability distributions over Θ, so we can think of them equivalently as probability distributions over the 1 measurable events. In general, *N measurable events* are those on which Nth order beliefs are defined. We can define these events recursively as follows. For any collection of events $\mathcal{E}$, let

$$\widehat{A}(\mathcal{E}) = \{E \subseteq S \mid E \in \mathcal{E} \text{ or } \exists F \in \mathcal{E}, i \in \mathcal{I}, p \in [0,1] \\ \text{with } E = B_i^p(F) \text{ or } E = K_i(F)\}.$$

Let $A_{n+1} = \widehat{A}(\mathcal{A}_n)$ and $\mathcal{A}_{n+1} = \sigma(A_{n+1})$. The events in $\mathcal{A}_N$ are *N measurable events.*

A state in one partitions model corresponds to a set of beliefs and knowledge that could be held at a state in another model, making these states "equivalent." To make this precise, I define two sets of equivalence relations. Suppose we have two models $\mathcal{M}$ and $\overline{\mathcal{M}}$ (where these could be the same model). Recall that a 1 measurable event corresponds simply to a subset of Θ. Hence for 1 measurable events E and $\overline{E}$, it seems natural to call these events "equivalent" if they correspond to the same subset of Θ. That is, I define $E =_0 \overline{E}$ if $f(E) = \bar{f}(\overline{E})$.[11] Similarly, I define $s \sim_0^* \bar{s}$ if $f(s) = \bar{f}(\bar{s})$. Note that I could equivalently start by defining $\sim_0^*$ and then define $=_0$ by $E =_0 E$ iff for all $s \in E$ and all $\bar{s}$ with $s \sim_0^* \bar{s}$, we have $\bar{s} \in \overline{E}$ and conversely.

The definition is more complex at higher levels because there are more ways two states could differ. For two states to be equivalent at level n, the same statements about knowledge and beliefs up to level n should hold at each. More precisely, I define $s \sim_n^* \bar{s}$ if $s \sim_{n-1}^* \bar{s}$ and for all n measurable E and $\overline{E}$ with $E =_{n-1} \overline{E}$, $s \in K_i(E)$ iff $\bar{s} \in \overline{K}_i(E)$ and for all p, $s \in B_i^p(\overline{E})$ iff $\bar{s} \in \overline{B}_i^p(\overline{E})$. To understand this, recall that the new events that come in at level n are those that involve statements of knowledge or belief regarding events at level $n-1$. Hence if E and $\overline{E}$ represent the same events at $n-1$, equivalence of s and $\bar{s}$ at level n requires knowledge and beliefs regarding these events to be the same. We can then define $=_n$ over $n+1$ measurable events E and $\overline{E}$ by $E =_n \overline{E}$ iff for all $s \in E$ and all $\bar{s} \sim_n^* s$, we have $\bar{s} \in \overline{E}$ and analogously for all $\bar{s} \in \overline{E}$.

I say that a state s in a partitions model $\mathcal{M}$ satisfies the CPA to level N if there is a model $\overline{\mathcal{M}}$ satisfying the CPA and a state $\bar{s}$ in that model such that $s \sim_N^* \bar{s}$.

Finally, I need to define common knowledge. For any $E \subseteq S$, let $K_G(E)$ denote the event that "everyone knows" E—that is, that E is group knowledge in the sense that

$$K_G(E) = \bigcap_{i=1}^{I} K_i(E).$$

Recursively define order n group knowledge, K_G^n, by $K_G^{n+1}(E) = K_G(K_G^n(E))$ where $K_G^1 = K_G$. Finally, $CK(E)$ is the event that E is common knowledge:

$$CK(E) = \bigcap_{n=1}^{\infty} K_G^n(E).$$

[11] Note that this relation is defined only on 1 measurable events. Hence E is the set of *all* states with parameter values in $f(E)$ and likewise for $\overline{E}$.

Given a partitions model $\mathscr{M}$, let τ denote the event that each agent gives zero probability to the occurrence of any event while another agent gives it zero probability. That is,

$$\tau = \{s' \in S \mid \mu_i(E \cap B_j^0(E) \mid \pi_i(s')) = 0, \forall i, j, E\}.$$

DEFINITION 1: State s in partitions model $\mathscr{M}$ satisfies *weak consistency* if $s \in CK(\tau)$.

Weak consistency of s is equivalent to the statement that for all s' reachable from s—that is, all s' in a partition event n steps from s for some finite n—$\mu_i(s' \mid \pi_i(s')) = 0$ for some i implies $\mu_j(s' \mid \pi_j(s')) = 0$ for all j. Hence if we construct a prior for each agent by taking any strictly positive distribution over the events in i's partition, say β_i, and defining $\mu_i(s) = \mu_i(s \mid \pi_i(s))\beta_i(\pi_i(s))$, the priors will all have the same support. In this sense, weak consistency is equivalent to the assumption that the priors of the agents have the same support.

My main result is Theorem 1.

THEOREM 1: *A state satisfies the common prior assumption to level N for all finite N if and only if it satisfies weak consistency.*

The necessity part of this theorem is easy to see. If s does not satisfy weak consistency, then there is some n for which τ is not group knowledge of order n at s. Suppose $\bar{s} \sim_N^* s$ in partitions model $\overline{\mathscr{M}}$ where N is much larger than n and $\overline{\mathscr{M}}$ satisfies the CPA. Since statements of nth order knowledge at s and $\bar{s}$ must be the same, $\bar{\tau}$ is not order n mutual knowledge at $\bar{s}$ where $\bar{\tau} =_n \tau$. That is, there must be a state $\bar{s}'$ in $\overline{\mathscr{M}}$ not ruled out by order n group knowledge such that $\bar{\mu}(\overline{E} \cap \overline{B}_j^0(\overline{E}) \mid \bar{\pi}_i(\bar{s}')) > 0$, where $\bar{\mu}$ is the common prior. But then the common prior gives positive probability to a state where j's updating from this prior gives it zero probability, a contradiction.

The sufficiency proof is in the Appendix because some of the details are rather tedious. To see the intuition, consider the following example. Suppose there are two players, $\Pi_1 = \Pi_2 = \{S\}$, $\Theta = S$, and $f(s) = s$. Finally, let μ_1 and μ_2 denote the beliefs of the two players where $\mu_1 \neq \mu_2$. I construct $\overline{\mathscr{M}}$ similarly to the example in Section 2.2 by letting $\overline{S} = S \times \{1, \dots, K\}$ for some K and $\bar{f}(s, k) = f(s)$. I will refer to state $(s, k) \in \overline{S}$ as the kth copy of s and let $\overline{C}_k$ denote the set of all kth copies. I complete the construction so that $s \sim_N^* (s, 1)$ for every s. To do so, let one event in player 1's partition be the event that consists of all the first copies, that is, $\overline{C}_1$. To match first order beliefs, we need $\bar{\mu}(s, 1)/\bar{\mu}(\overline{C}_1) = \mu_1(s)$. Once we choose $\bar{\mu}(\overline{C}_1)$, then, this equation determines $\bar{\mu}$ for all first copies.

Given this, how will we make player 2's first-order beliefs work out properly at $(s, 1)$? The simplest way is to construct this event of player 2's partition by setting it equal to the set of first *and* second copies. That is, 2 has a partition event equal to $\overline{C}_1 \cup \overline{C}_2$. The presence of the second copies gives us a degree of freedom to get 2's beliefs right. More specifically, we need $[\bar{\mu}(s, 1) + \bar{\mu}(s, 2)]/[\bar{\mu}(\overline{C}_1) + \bar{\mu}(\overline{C}_2)] = \mu_2(s)$. Given $\bar{\mu}(C_1)$, we already determined $\bar{\mu}(s, 1)$ for each s. Given these values plus $\bar{\mu}(C_2)$, this equation will determine $\bar{\mu}(s, 2)$ for every s.

From here, the pattern is clear: Player 1 has a partition event with all the second and third copies, one with all the fourth and fifth copies, etc., while 2 has a partition event with all the first and second copies, one with all the third and fourth, etc. This will completely determine $\bar{\mu}$ as a function of the prior probabilities assigned to the $\overline{C}_k$ sets. The only thing remaining is to ensure that we can choose these probabilities so that $\bar{\mu}$ is a legitimate probability distribution. Weak consistency ensures that this can be done.

REMARK 1: If we allow S to be infinite, the characterization becomes more complex. In Lipman (2002), I show that if S is countable, then a *tail consistency* condition is also required. Intuitively, this condition says that the tails of the distributions for different players cannot go to zero at different rates.

A corollary to this result is that every world in the Mertens-Zamir universal beliefs space is arbitrarily close to a world that satisfies the CPA. That is, the set of worlds satisfying the common prior assumption is dense in the universal beliefs space. The proof of the corollary is very simple but a careful statement requires many details that have not been relevant to this point. For brevity, I give only a sketch of the basic ideas. More details and a proof of Corollary 1 below are available in notes on my web page, http://people.bu.edu/blipman.

The universal beliefs space, denoted Ω, is the set of all "coherent"[12] worlds. Recall that a world is a specification of a parameter value plus the infinite hierarchies of beliefs of each player. Hence Ω is a subset of an infinite product space where the component spaces are Θ and spaces of probability distributions on various sets. E.g., first order beliefs are probability distributions on Θ, so one component is the set of such distributions. As usual, I use the weak topology for spaces of probability distributions. Given the product structure, it seems natural to use the (relativized) product topology to extend the component topologies to Ω. This is the topology used by Mertens–Zamir. I follow them in assuming that Θ is compact.

As discussed in Section 2.1, any partitions model and state s in that model uniquely identifies a particular world denoted $\omega(s)$ in the universal beliefs space. Mertens–Zamir prove that every world in the universal beliefs space is arbitrarily close to one generated by a partitions model with a finite S. In other words, while finiteness of S is a restriction in general, for considering closeness of worlds, it is without loss of generality.

I define a world $\omega \in \Omega$ to be weakly consistent if it is $\omega(s)$ for some weakly consistent s in a partitions model. I define ω to be consistent with common priors if it is $\omega(s)$ for some s in a partitions model that satisfies the CPA.

COROLLARY 1: *The closure of the set of worlds consistent with common priors is* Ω.

To see this, recall that every world is arbitrarily close to one that corresponds to a state in a (finite) partitions model. Hence we can restrict attention to such worlds. All such worlds are arbitrarily close to worlds that are weakly consistent because we can move an arbitrarily small amount of probability to turn a world that is not weakly consistent into one that is.[13] Hence we can restrict attention to weakly consistent worlds. Theorem 1 shows that every such world is arbitrarily close in the product topology to one consistent with common priors.

[12] Loosely, a world is coherent if it is common belief that each player's beliefs are consistent across orders in the sense that his order n beliefs are the marginal over the appropriate subspace of his order $(n+1)$ beliefs.

[13] To see this, consider a world that does not satisfy weak consistency but does generate a finite belief-closed set. Then it must be true that in the partitions model to which this belief-closed set corresponds, the supports of the priors differ. We can change this model to one with the same partitions and f but where the supports of the priors are the same across agents by changing the priors by an arbitrarily small amount. It is not hard to see that the belief-closed set this new partitions model generates is arbitrarily close to the one with which we began.

It is important to remember that strategic behavior is *not* continuous with respect to this topology.[14] In other words, while a world that violates common priors is close in the product topology to one that satisfies common priors, the kind of behavior that can occur in these two worlds is very different. Consequently, the denseness result should not be interpreted as implying that worlds violating common priors are "uninteresting."

4. CONCLUSION

One implication of Theorem 1 is that a local result (a result about a state, not the model as a whole) that requires the common prior assumption must depend on infinitely many orders of beliefs. More intuitively, if the common prior assumption is needed, one cannot drop common knowledge assumptions. For example, the fact that the agreeing to disagree result and no-trade theorems don't hold without common priors indicates that, as is well-known, these results rely on their common knowledge assumptions. Intuitively, common knowledge assumptions translate a local statement to a global one. The new partitions model I construct to match beliefs at a state cannot match the original model globally. Hence it should not be surprising that results, local or not, whose proofs rely on global statements will not carry over to noncommon priors.

Dept. of Economics, Boston University, 270 Bay State Rd., Boston, MA 02215, U.S.A, and Dept. of Economics, University of Wisconsin, Madison, WI, U.S.A.; blipman@bu.edu; http://people.bu.edu/blipman/

Manuscript received January, 1997; final revision received June, 2002.

APPENDIX: PROOF OF THEOREM 1

Necessity is discussed in the text. To show sufficiency, suppose we have a state, say s^*, that satisfies weak consistency. Without loss of generality, assume every state $s \in S$ is reachable from s^*. Let S_+ be the set of states such that $\mu_i(s \mid \pi_i(s)) > 0$ for some i and $S_0 = S \backslash S_+$. By weak consistency, for all $s \in S_+$, $\mu_i(s \mid \pi_i(s)) > 0$ for *all* i. For each i, fix a probability distribution β_i on Π_i such that $\beta_i(\pi_i) > 0$ for all $\pi_i \in \Pi_i$. For each i, define $\mu_i(s) = \beta_i(\pi_i(s))\mu_i(s \mid \pi_i(s))$. This can be thought of as i's prior. By construction, the support of every μ_i is S_+. Let $\mathcal{L}$ denote any constant such that $\mathcal{L} > \mu_j(s)/\mu_i(s)$ for all i, j, and $s \in S_+$. By finiteness of S, such an $\mathcal{L}$ exists.

I define a new partitions model $\overline{\mathcal{M}}$ as follows. Let $\overline{S} = S \times \{1, \ldots, NI\}$. Let $\bar{f}(s,k) = f(s)$ for all k. For each player i and state s, I define $N+1$ sets corresponding to different "copies" of state s for i. Specifically, for $n = 0, \ldots, N$,

$$C_n^i(s) = \{(s,k) \in \overline{S} \mid (n-1)I + i + 1 \leq k \leq nI + i\}.$$

(Because $(s,k) \in \overline{S}$ implies $k \geq 1$, for $n = 0$, the k's are effectively just $1, \ldots, i$. Similarly, $k \leq NI$ is the effective upper bound for $n = N$.) I will refer to $C_n^i(s)$ as i's nth copies of s. Note for later use that each of these sets has at most I elements.

Player i's partition, $\overline{\Pi}_i$, consists of the sets

$$\bigcup_{s \in \pi_i} C_n^i(s), \quad n = 0, \ldots, N, \quad \pi_i \in \Pi_i.$$

Intuitively, i has $N+1$ copies of each of his partition events in the original model.

[14] See Kajii and Morris (1998) on topologies on beliefs with respect to which the equilibrium correspondence is upper semicontinuous.

I now define $\bar{\mu}$. For an $\alpha > 0$ to be determined below, for $i = 1, \ldots, I$ and $k = 0, \ldots, N-1$, recursively define

$$\bar{\mu}(s, 1) = \alpha\mu_1(s)$$

and

$$\bar{\mu}(s, kI+i) = \mathscr{L}^{kI+i-1}\alpha\mu_i(s) - \sum_{j=1}^{kI+i-1} \bar{\mu}(s, j).$$

Note that $\bar{\mu}(s, kI+i) = 0$ for all $s \in S_0$. For $s \in S_+$, $\bar{\mu}(s, kI+i) > 0$ iff

$$\mathscr{L}^{kI+i-1}\alpha\mu_i(s) > \sum_{j=1}^{kI+i-1} \bar{\mu}(s, j).$$

For $i \geq 2$, the right-hand side is $\mathscr{L}^{kI+i-2}\alpha\mu_{i-1}(s)$. Hence for $s \in S_+$, the inequality holds iff $\mathscr{L} > \mu_{i-1}(s)/\mu_i(s)$, which holds by definition of $\mathscr{L}$. A similar argument applies for $i = 1$. Hence $\bar{\mu}(s, k) > 0$ iff $s \in S_+$. Also,

$$\sum_{s\in S}\sum_{k=1}^{NI} \bar{\mu}(s, k) = \sum_{s\in S} \mathscr{L}^{NI-1}\alpha\mu_I(s) = \mathscr{L}^{NI-1}\alpha.$$

Hence by setting $\alpha = (1/\mathscr{L})^{NI-1}$, we ensure that this is one.

The following lemma shows that this construction yields a partitions model with the appropriate properties.

LEMMA 1: *For all* $s, (s, k) \in \overline{S}$, *and* $n = 1, \ldots, N$, $(s, k) \sim_n^* s$ *for* $k = 1, \ldots, (N-n)I + n$.

PROOF: The proof is by induction on n. Since $\bar{f}(s, k) = f(s)$, we have $(s, k) \sim_0^* s$ for all k, proving the claim for $n = 0$.

Suppose we have shown the result up to some n. Fix n measurable events E and $\overline{E}$ with $E =_{n-1} \overline{E}$. By definition, then, $s \in E$ and $\bar{s} \sim_{n-1}^* s$ implies $\bar{s} \in \overline{E}$ and conversely. Hence by the induction hypothesis, $s \in E$ iff $(s, k) \in E$ for $k = 1, \ldots, (N-n+1)I + n - 1$.

For knowledge, recall that $s \in K_i(E)$ iff $\pi_i(s) \subseteq E$. By the above, $\pi_i(s) \subseteq E$ iff for all $s' \in \pi_i(s)$, $(s', k) \in \overline{E}$ for $k = 1, \ldots, (N-n+1)I + n - 1$. By construction, $\bar{\pi}_i(s, k)$ is a collection of copies of states $s' \in \pi_i(s)$. That is, it is the set of (s', k') with $s' \in \pi_i(s)$ and k' in a certain range. Since each partition event contains at most I "versions" of each s, then if $k + I - 1 \leq (N-n+1)I + n - 1$, we must have $\bar{\pi}_i(s, k) \subseteq \overline{E}$. This holds iff $k \leq (N-n)I + n$. Hence $s \in K_i(E)$ iff $(s, k) \in \overline{K}_i(\overline{E})$ for $k = 1, \ldots, (N-n)I + n$.

Turning to beliefs, let

$$\overline{E}(s) = \{(s', k') \in \overline{E} \mid s' = s\}.$$

I now show that for all $s' \in \pi_i(s) \cap E$,

$$\frac{\mu_i(s \mid \pi_i(s))}{\mu_i(s' \mid \pi_i(s))} = \frac{\bar{\mu}(\overline{E}(s) \mid \bar{\pi}_i(s, k))}{\bar{\mu}(\overline{E}(s') \mid \bar{\pi}_i(s, k))}$$

for all i and for $k = 1, \ldots, (N-n)I + n$. This will imply $\mu_i(E \mid \pi_i(s)) = \bar{\mu}(\overline{E} \mid \bar{\pi}_i(s, k))$ and hence $s \in B_i^p(E)$ iff $(s, k) \in B_i^p(E)$ for k in the relevant range.

To show the claim, fix any i, any k between 1 and $(N-n)I + n$, any s, and any $s' \in \pi_i(s) \cap E$. Note that

$$\frac{\mu(\overline{E}(s) \mid \bar{\pi}_i(s, k))}{\bar{\mu}(\overline{E}(s') \mid \bar{\pi}_i(s, k))} = \frac{\bar{\mu}(\overline{E}(s) \cap \bar{\pi}_i(s, k))}{\bar{\mu}(\overline{E}(s') \cap \bar{\pi}_i(s, k))}.$$

By construction, $\bar{\pi}_i(s,k)$ is a union over $s' \in \pi_i(s)$ of some level of copies of s'. Let m denote this level—that is, $\bar{\pi}_i(s,k) = \cup_{s' \in \pi_i(s)} C^i_m(s')$. Since $k \leq (N-n)I+1$ and since each collection of copies includes no more that I elements, we know that the largest k' with $(s',k') \in \bar{\pi}_i(s,k)$ satisfies

$$k' \leq k+I-1 \leq (N-n)I+1+I-1 = (N-n+1)I \leq (N-n+1)I+n-1.$$

By the argument above, we know that for all k' in this range, $(s,k') \in \overline{E}$ and, since $s' \in E$, $(s',k') \in \overline{E}$. Hence $\overline{E}(s) \cap \bar{\pi}_i(s,k) = C^i_m(s)$ and likewise for s'. So

$$\frac{\bar{\mu}(\overline{E}(s) \mid \bar{\pi}_i(s,k))}{\bar{\mu}(\overline{E}(s') \mid \bar{\pi}_i(s,k))} = \frac{\bar{\mu}(C^i_m(s))}{\bar{\mu}(C^i_m(s'))}.$$

For $m = 1, \ldots, N-1$, this is

$$\frac{\sum_{j=(m-1)I+i+1}^{mI+i} \bar{\mu}(s,j)}{\sum_{j=(m-1)I+i+1}^{mI+i} \bar{\mu}(s',j)} = \frac{\sum_{j=1}^{mI+i} \bar{\mu}(s,j) - \sum_{j=1}^{(m-1)I+i} \bar{\mu}(s,j)}{\sum_{j=1}^{mI+i} \bar{\mu}(s',j) - \sum_{j=1}^{(m-1)I+i} \bar{\mu}(s',j)}.$$

But $\sum_{j=1}^{\ell I+i} \bar{\mu}(s,j) = \mathscr{L}^{\ell I+i-1}\mu_i(s)$. Substituting,

$$= \frac{[\mathscr{L}^{mI+i-1} - \mathscr{L}^{(m-1)I+i-1}]\mu_i(s)}{[\mathscr{L}^{mI+i-1} - \mathscr{L}^{(m-1)I+i-1}]\mu_i(s')} = \frac{\mu_i(s)}{\mu_i(s')} = \frac{\mu_i(s \mid \pi_i(s))}{\mu_i(s' \mid \pi_i(s'))}$$

as was to be shown. The cases of $m=0$ and $m=N$ are analogous.

Hence for every n measurable E and $\overline{E}$ with $E =_{n-1} \overline{E}$, we have $s \in B^p_i(E)$ iff $(s,k) \in \overline{B}^p_i(\overline{E})$ for $k = 1, \ldots, (N-n)I+n$. *Q.E.D.*

In light of the lemma, we see that $(s,k) \sim^*_N s$ for $k = 1, \ldots, N$, so s is consistent with the CPA to level N. Since N is arbitrary, the conclusion follows. *Q.E.D.*

REFERENCES

AUMANN, R. (1976): "Agreeing to Disagree," *Annals of Statistics*, 4, 1236–1239.

BONNANO, 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.

BRANDENBURGER, A., AND E. DEKEL (1993): "Hierarchies of Beliefs and Common Knowledge," *Journal of Economic Theory*, 59, 189–198.

DEKEL, E., AND F. GUL (1997): "Rationality and Knowledge in Game Theory," in *Advances in Economic Theory, Seventh World Congress*, ed. by D. M. Kreps and K. F. Wallis. Cambridge: Cambridge University Press, 87-172.

FEINBERG, Y. (2000): "Characterizing Common Priors in the Form of Posteriors," *Journal of Economic Theory*, 91, 127–179.

HALPERN, J. (2002): "Characterizing the Common Prior Assumption," *Journal of Economic Theory*, 106, 316–355.

KAJII, A., AND S. MORRIS (1998): "Payoff Continuity in Incomplete Information Games," *Journal of Economic Theory*, 82, 267–276.

LIPMAN, B. (2002): "Finite Order Implications of Common Priors in Infinite Models," Working Paper, Boston University.

MERTENS, J.-F., AND S. ZAMIR (1985): "Formalization of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 14, 1–29.

MILGROM, P., AND N. STOKEY (1982): "Information, Trade, and Common Knowledge," *Journal of Economic Theory*, 26, 17–27.

MONDERER, D., AND D. SAMET (1989): "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior*, 1, 170–190.

MORRIS, S. (1994): "Trade with Heterogeneous Prior Beliefs and Asymmetric Information," *Econometrica*, 62, 1327–1347.
——— (1995): "The Common Prior Assumption in Economic Theory," *Economics and Philosophy*, 11, 227–283.
SAMET, D. (1996): "Common Priors and Separation of Convex Sets," Tel Aviv University Working Paper.
TAN, T., AND S. WERLANG (1988): "On Aumann's Notion of Common Knowledge—An Alternative Approach," Working Paper.