HOW COMMON ARE COMMON PRIORS?

ZIV HELLMAN AND DOV SAMET

Abstract. To answer the question in the title we consider a fixed state space with a profile of partitions that have a single common knowledge component and let the the types, that is, the posterior probabilities, vary. How common is it for types to have a common prior? The answer depends on the partition profile. We say that it is maximal if any refinement of its elements results in a profile that has more than one common knowledge component. If the partition profile is maximal, then every profile of types has a common prior. If it is not maximal, then the set of types that have a common prior is topologically small: it is nowhere dense in the set of all types.

1. Introduction

Ever since the introduction of games with incomplete information by Harsanyi (1967, 1968a, b), the use of common priors in models of differential information has been ubiquitous. The assumption that players posterior beliefs are derived from a prior common to all of them plays an essential role in the No Disagreements theorem of Aumann (1976) and in the No Trade theorems that followed. It is also a basic building block of correlated equilibrium as the ultimate expression of common knowledge of rationality (Aumann (1987)). Moreover, as pointed out in that paper, the assumption of a common prior is pervasive and explicit or implicit in the vast majority of the differential information literature in economics and game theory. The assumption that players share a common prior, sometimes known as the Harsnyi doctrine, has been debated, sometimes vehemently. Aumann‘s 1987 paper sparked a debate between him and Gul (Aumann (1998) and Gul (1998)) in which the latter argued that an over-reliance on assuming common priors is unrealistic.

The debate over the extent to which common priors should or should not be automatically assumed leads naturally to the question in the title of this paper: how common *are* common priors? In other words, given a (finite) state space Ω, n players and a partition profile $\Pi = (\Pi_1, \ldots, \Pi_n)$ associating a partition Π_i of the state space with each player, if one assigns an arbitrarily selected type function to each player, how likely will the resulting type space have a common prior?

It turns out that the answer to the question depends very much on the partition profile. If the partition profile Π is what we term tight, as defined in this paper, then no matter what type functions are associated with the players, a common prior exists. In contrast, if Π is not tight, the set of type profiles that have a common prior is nowhere dense in the set of all type profiles for Π.

This leads to a surprising corollary. As we show in the paper, to determine whether or not a partition profile $\Pi = (\Pi_1, \ldots, \Pi_n)$ is tight all one needs to know is how to count, as it depends solely on the total number of partition elements: a partition profile whose meet is a singleton is tight if and only if $\sum_{i=1}^{n} |\Pi_i| = (n-1)|\Omega| + 1$, and it is always true that $\sum_{i=1}^{n} |\Pi_i| \leq (n-1)|\Omega| + 1$. This means there is a phase change that occurs when the maximal value of the total number of partition elements is attained i.e. as long as the sum total of partition elements is less than $(n-1)|\Omega|+1$ 1, it is almost always the case that an arbitrary type profile will *not* have a common prior, but if the sum of partition elements *does* equal $(n-1)|\Omega|+1$, a common prior *always* exists.

2. Preliminaries

Partitions of state spaces. Let Ω be a finite set called a *state space*, and $I = \{1,\ldots,n\}$ a set of *agents* with $n \geq 2$. A *partition profile* for Ω is a vector $\Pi = (\Pi_1,\ldots,\Pi_n)$ of partitions of Ω. We write $\Pi'_i \succeq \Pi_i$ when the partition Π'_i refines Π_i, and $\Pi'_i \succ \Pi_i$ when $\Pi'_i \succeq \Pi_i$ and $\Pi'_i \neq \Pi_i$. For partition profiles Π and Π', we write $\Pi \succeq \Pi'$ when for each i, $\Pi'_i \succeq \Pi_i$, and $\Pi \succ \Pi'$ when $\Pi \succeq \Pi'$ and $\Pi' \neq \Pi$. The *meet* of Π is the partition Π_c which is the finest among all partitions that are coarser than Π_i for each i.[1]

Types. Denote by $\Delta(\Omega)$ the set of all probability functions on Ω. A *type function* for Π_i is a function $t_i\colon \Pi_i \to \Delta(\Omega)$ that satisfies for each $\pi \in \Pi_i$, $t_i(\pi)(\pi) = 1$. The probability function $t_i(\pi)$ is the *type* of i in each state of π. A *type profile* for Π is a vector of type functions $t = (t_1,\ldots,t_n)$. Denote by $T(\Pi)$ the set of all type profiles for Π. We consider $T(\Pi)$ as a subset of $\times_{i=1}^{n} R^{\Omega\times\Pi_i}$.

Priors and common priors. A *prior* for t_i is a probability function $p \in \Delta(\Omega)$, such that for each $\pi \in \Pi_i$ with $p(\pi) > 0$, $t_i(\pi)(\cdot) = p(\cdot|\pi)$. Contrasting a prior for t_i with the probability functions $t_i(\pi)$, the latter are referred to as the *posterior* probabilities of i. It is easy to see that p is a prior for t_i if an only if it is a convex combination of i's types, $\{t_i(\pi) \mid \pi \in \Pi_i\}$. A type profile t has a *common prior* if there is a probability function in $\Delta(\Omega)$ which is a prior for t_i for each i.

3. Main results

Definition 1. *A partition profile Π is* **tight** *if for each partition profile Π' that satisfies $\Pi' \succ \Pi$, $\Pi'_c \succ \Pi_c$.*

The Main Theorem. *Let Π be a partition profile, for Ω, then*

- *If Π is tight then each type profiles in $T(\Pi)$ has a common prior.*
- *If Π is not tight then the set of type profiles in $T(\Pi)$ that have a common prior is nowhere dense in $T(\Pi)$.*

The next proposition characterizes tight partition profiles the meet of which is a singleton by their size.

Proposition 1. *For each partition profile Π the meet of which is a singleton, $\sum_{i=1}^{n} |\Pi_i| \leq (n-1)|\Omega|+1$. The partition profile Π is tight if and only if $\sum_{i=1}^{n} |\Pi_i| = (n-1)|\Omega|+1$.*

Note that the dimension of the set of types $T(\Pi)$ is $\sum_{i=1}^{n} |\Omega| - |\Pi_i|$ and the dimension of set of priors $\Delta(\Omega)$ is $|\Omega|-1$. The inequality in the proposition says that for each Π with $|\Pi_c| = 1$, the dimension of the set of types is at least as that of the set of priors, and equality holds for the tight partition profiles.

[1] Following Aumann's (1976) definition of common knowledge in terms of the meet, elements of the meet are referred sometimes as common knowledge components.

4. Proofs

Proof of the Main theorem: We state first some well known facts about meets. Denote by $\cup\Pi$ the set of the elements in all the partitions in Π. We say that two elements π and π' in $\cup\Pi$ *connect* if there is a sequence $\pi_1, \ldots, \pi_c$ in $\cup\Pi$ such that $\pi = \pi_1$, $\pi' = \pi_c$, and for $k = 1, \ldots, m-1$, $\pi_k \cap \pi_{k+1} \neq \emptyset$. The connectivity relation is an equivalence relation on $\cup\Pi$. Each element of the meet Π_c is a union of the elements of an equivalence class in $\cup\Pi$.

We say that a partition profile Π' is a *minimal* refinement of Π, if Π' is obtained form Π by refining only one partition Π_i, and this refinement is obtained by splitting a single element $\pi \in \Pi_i$ into two sets. We record the following simple result.

Lemma 1. *A partition profile* Π *is tight if and only if for any minimal refinement* Π' *of* Π*,* $\Pi'_c \succ \Pi_c$*.*

Proof: If Π is tight, then by definition the condition in the lemma holds. Conversely, suppose that for every minimal refinement Π' of Π, $\Pi'_c \succ \Pi_c$, and let Π'' be a refinement of Π. Then there exists a minimal refinement of Π, Π', such that $\Pi'' \succeq \Pi' \succ \Pi$. The claim follows since $\Pi''_c \succeq \Pi'_c \succ \Pi'_c$. ∎

We explain now why it is enough to prove the theorem for partition profiles the meet of which is a singleton. For each element π of the meet Π_c of Π, and for each i, the set $\Pi_i(\pi)$ of the elements of Π_i contained in π, partition it. Thus, $\Pi(\pi) = (\Pi_1(\pi), \ldots, \Pi_n(\pi))$ is a partition profile of π, and π is the only element of its meet. Obviously, a partition profile Π is tight if and only if for each element π in the meet Π_c, $\Pi(\pi)$ is tight. Thus, to prove the first part of the theorem it is enough to show it for partition profiles the meet of which is a singleton.

Similarly, if $\pi \in \Pi_c$, and the set of types for $\Pi(\pi)$ that have a common prior is nowhere dense in the set all types for $\Pi(\pi)$, then the set of types for Π that have a common prior is nowhere dense in the set all types for Π. Thus, it is enough to prove the second part of the theorem for partition profiles the meet of which is a singleton. Therefore, we assume henceforth that Π_m is a singleton.

We prove the first part of the theorem by induction on the size of the state space Ω. The structure of the proof is simple. Starting with a tight partition profile Π, we refine it to Π' where the meet of Π' has two elements. We prove that these two elements are tight. Starting with a type profile t for Π, we construct in a obvious way a type profile t' for Π'. Using the induction hypothesis we have a common prior for each of the elements of Π'_c. Using these common priors we construct a common prior for Π.

For $|\Omega| = 1$ the claim is trivial. Suppose that $|\Omega| \geq 2$ and the claim holds for all state spaces of size smaller than $|\Omega|$. Let Π be tight. Since Π_c is a singleton there must be an i and $\pi \in \Pi_i$ such that $|\pi| > 1$. Let π be such an element, and consider the minimal refinement of Π, Π', obtained by splitting π into π_1 and π_2. Since Π is tight, the meet Π'_c has at least two elements. Since every element in $\cup\Pi$ connects with π, it follows that every element in $\cup\Pi'$ connects with either π_1 or π_2. Thus, Π'_c has exactly two elements: Ω_1 which contains all elements of Π' that connect to π_1 and contains, of course, π_1 and Ω_2 which contains all elements of Π' that connect to π_2 and in particular π_2 itself.

For $k = 1, 2$ and each i denote by $\Pi'_i(\Omega_k)$ the set of elements of Π'_i that are contained in Ω_k. Then $\Pi'(\Omega_k) = (\Pi'_1(\Omega_k), \ldots, \Pi'_n(\Omega_k))$ is a partition profile for Ω_k. We claim that $\Pi'(\Omega_k)$ is tight. Suppose it is not, then by Lemma 1 there exists a minimal refinement of $\Pi'(\Omega_k)$, denoted $\hat{\Pi}'(\Omega_k)$, the meet of which is a singleton. Suppose that this minimal refinement is obtained by splitting $\hat{\pi}$ into $\hat{\pi}'$ and $\hat{\pi}''$.

Examine first the case that $\hat{\pi} \neq \pi_k$. Note, that all the elements in $\cup\hat{\Pi}'(\Omega_k)$ connect with π_k. Consider now the minimal refinement of Π, denoted $\hat{\Pi}$, obtained by splitting $\hat{\pi}$ as in $\hat{\Pi}'(\Omega_k)$. Then, all the elements of $\cup\hat{\Pi}$ that are contained in Ω_k connect to π (which is the union of π_1 and π_2) per our assumption, and all the elements of $\cup\hat{\Pi}$ that are contained in Ω_{3-k} also connect to π. Thus the meet of $\hat{\Pi}$ is a singleton contrary to the maximality of Π.

Now, examine the case that $\hat{\pi} = \pi_k$. Then, $\pi_k = \hat{\pi}' \cup \hat{\pi}''$. All the elements in $\cup\hat{\Pi}'(\Omega_k)$ connect with $\hat{\pi}'$. Consider the minimal refinement of Π, $\hat{\Pi}$, obtained by splitting π into $\hat{\pi}''$ and $\hat{\pi}' \cup \pi_{3-k}$. All the elements in $\cup\hat{\Pi}$ that are contained in Ω_k connect to $\hat{\pi}' \cup \pi_{3-k}$ through $\hat{\pi}'$, per our assumption, and all the elements of $\cup\hat{\Pi}$ that are contained in Ω_{3-k} also connect to $\hat{\pi}' \cup \pi_{3-k}$ through π_{3-k}. Thus the meet of $\hat{\Pi}$ is a singleton contrary to the maximality of Π.

Consider a type profile $t = (t_1, \dots, t_n)$ for Π. We define a type profile $t' = (t'_1, \dots, t'_n)$ for Π' as follows: t agrees with t' on all the elements of Π that were not split. For the two parts of π, π_1 and π_2, define t'_i as follows. If $t_i(\pi_k) > 0$ then $t'_i(\pi_k)(\cdot) = t_i(\pi)(\cdot \mid \pi_k)$. If $t_i(\pi_k) = 0$ then $t'_i(\pi_k)$ is defined arbitrarily. The restriction on t' to elements in $\Pi'_i(\Omega_k)$, denoted $t'(\Omega_k)$, is a type profile for this partition profile. By the induction hypothesis $t'(\Omega_k)$ has a common prior p_k. We think of p_k as being a probability function on Ω that vanishes on Ω_{3-k}. Using p_1 and p_2 we construct a common prior for t.

Assume first that either $t_i(\pi)(\pi_k) = 0$ or $p_k(\pi_k) = 0$. In either case p_{3-k} is a common prior for t. Assume, then, that for $k = 1, 2$, $t_i(\pi)(\pi_k) > 0$ and $p(\pi_k) > 0$, and consider the equation $t_i(\pi)(\pi_1)/t_i(\pi)(\pi_2) = ap_1(\pi_1)/(1-a)p_2(\pi_2)$. By the positivity assumption, it has a solution $a \in (0, 1)$. The probability function $ap_1 + (1-a)p_2$ is a prior for Π.

To prove the second part of the theorem let C be the set of type profiles in $T(\Pi)$ that have a common prior. We show first that C is closed. Indeed, suppose that (t^n) is a sequence in C, and $t^n \to t$. For each n let p^n be a common prior of t^n. Because of the compactness of $\Delta(\Omega)$, we can assume without loss of generality that for some $p \in \Delta(\Omega)$, $p^n \to p$. Suppose now that for $\pi \in \Pi_i$, $p(\pi) > 0$. Then for n large enough $p^n(\pi) > 0$. For such n, $t_i^n(\pi)(\cdot) = p^n(\cdot)/p^n(\pi)$. In the limit we get $t_i(\pi)(\cdot) = p(\cdot)/p(\pi)$, which shows that p is a common prior of t, and thus $t \in C$.

Define a type function t_i to be *positive* if for each $\pi \in \Pi_i$ and $\omega \in \pi$, $t_i(\pi)(\omega) > 0$. Denote by P the set of positive type profiles. Clearly, P is the relative interior of $T(\Pi)$. Since C is closed, the set $P \cap C^c$ is open. Also, its complement contains C. Therefore to show that C is nowhere dense in $T(\Pi)$ it is enough to show that the closure of $P \cap C^c$ is $T(\Pi)$. To show this we prove that the closure of $P \cap C^c$ contains $P \cap C$. Thus the closure of $P \cap C^c$ contains also P and hence it is indeed $T(\Pi)$.

Assume that Π is not tight. By Lemma 1 there exists a minimal refinement, Π', of Π the meet of which is a singleton. Suppose that Π' is obtained by splitting $\pi \in \Pi_i$ into the two sets π_1 and π_2.

Let t be a type profile in $P \cap C$ with a common prior p. We show that we can find a type profile $\hat{t} \in P \cap C^c$ arbitrarily close to t. Define first a type profile t' for Π' as follows: t' agrees with t on all elements of Π that were not split. For $k = 1, 2$, $t'_i(\pi_k)(\cdot) = t_i(\pi)(\cdot \mid \pi_k)$. It is obvious that p is also a common prior for t'. Define now a type profile $\hat{t}$ for Π as follows: $\hat{t}$ agrees with t on all elements of all partitions other than $\pi \in \Pi_i$. For π define $\hat{t}(\pi)(\cdot) = at_i(\pi)(\cdot \mid \pi_1) + (1-a)t_i(\pi)(\cdot \mid \pi_2)$, where $a = (1-\varepsilon)t_i(\pi)(\pi_1)$ for small enough positive ε. Thus, on π_1 the probabilities assigned by $\hat{t}_i(\pi)$ are those of $t_i(\pi)$ multiplied by $1-\varepsilon$. On π_2 the probabilities assigned by $\hat{t}_i(\pi)$ are those of $t_i(\pi)$ multiplied by $1+\varepsilon$. We define the type profile $\hat{t}'$ for Π' in the same way t' was defined. It is clear that $\hat{t}' = t'$.

Suppose now that $\hat{t}$ has a common prior q, then q is also a prior for $\hat{t}'$. But t' is positive and Π'_c is a singleton. Therefore, by Samet (1998), t' can have only one prior, and thus $p = q$. But p cannot be a prior for $\hat{t}$ as $p(\cdot \mid \pi) = t_i(\pi)(\cdot) \neq \hat{t}_i(\pi)(\cdot)$. Therefore, $\hat{t}$ does not have a common prior, i.e., $\hat{t} \in P \cap C^c$. ∎

Proof of Proposition 1: Let Π be a tight partition profile the meet of which is a singleton. The proof is by induction on the size on Ω. If Ω has a single element then the equality in the proposition is obvious. Suppose the equality is proved for all state spaces smaller than Ω, and consider the minimal refinement, Π' of Π described in the proof of the first part of the main theorem. Then, by the induction hypothesis, for $k = 1, 2$, $\sum_{i=1}^n |\Pi_i(\Omega_k)| = (n-1)|\Omega_k| + 1$. Adding the the two equations and noting that $\sum_{i=1}^n |\Pi_i(\Omega_1)| + \sum_{i=1}^n |\Pi_i(\Omega_2)| = \sum_{i=1}^n |\Pi_i| + 1$ completes the proof of the equality.

If Π is not tight, then it must have a refinement which is tight, and therefore it satisfies the inequality in of the proposition. ∎

References

Aumann, R. J. (1976), Agreeing to disagree, *Ann. Statist.*, 4(6), 1236–1239.

Aumann, R. J. (1999), Interactive epistemology I: Knowledge, *Internat. J. Game Theory*, 28(3), 263–300.

Fagin, R., J. Y. Halpern, Y. Moses, and M. Y. Vardi (1995), *Reasoning about knowledge.* MIT Press, Cambridge, MA.

Nyarko, Y. (1990), Most games violate the Harsanyi doctrine, Unpublished manuscript, `//ideas.repec.org/p/cvs/starer/91-39.html`.

Samet, D. (1990), Ignoring ignorance and agreeing to disagree, *J. of Economic Theory*, 52, 190–207.

Samet, D. (1998), Iterated expectations and common priors, *Games and Economic Behavior*, (24), 131-141.

The Faculty of Management, Tel Aviv University. Currently on leave at ILLC, University of Amsterdam and CWI, Amsterdam, The Netherlands.