Almost Common Priors

Ziv Hellman

The Department of Mathematics and the Centre for the Study of Rationality,
The Hebrew University of Jerusalem.

Abstract. What happens when priors are not common? We show that for each type profile τ over a knowledge space $(\Omega, \mathbf{\Pi})$, where the state space Ω is connected with respect to the partition profile $\mathbf{\Pi}$, we can associate a value $0 \leq \varepsilon \leq 1$ that we term the *prior distance* of τ. If τ has ε prior distance, then for any bet f amongst the players, it cannot be common knowledge that each player expects a positive gain of $\varepsilon \|f\|_\infty$, thus extending no betting results under common priors. Furthermore, as more information is obtained and partitions are refined, the prior distance, and thus the extent of common knowledge disagreement, decreases.

1 Introduction

What happens if priors are not common? Can one measure 'how far' a belief space is from a common prior, and use that to approximate standard results that apply under the common prior assumption?

The assumption that players' posterior beliefs in models of differential information are derived from a common prior has been ubiquitous in the literature since (Ha1) introduced the concept in his groundbreaking work on games with incomplete information. Indeed, as pointed out in (Au2), the assumption of a common prior 'is pervasively explicit or implicit in the vast majority of the differential information literature in economics and game theory'. Although more than a score of years have passed since those lines were published, they retain their full force.

But despite its pervasiveness, the common prior assumption was, and still is, debated and challenged (see (Gul) and (Au3)). In the traditional framework, players are assumed to have a common prior distribution over the set of possible states of the world at some *ex ante* time, before they have any information at all. The receipt of information by the players leads to partitions of the states of the world, and each player updates his beliefs, using Bayes' Rule, to derive the posteriors. But the priors are primitives of this model, and there is no *a priori* reason, given the construction of the model, to suppose the players should agree on a prior. Furthermore, as (Gul) notes, in many cases of interest all that observers have are profiles of posteriors, not priors, and there are examples of posteriors that could not possibly have been derived from common priors. He also questions the assumption that there is a relevant hypothetical 'prior stage' at all. (HS1) further show that the set of profiles of posteriors that admit common priors is a topologically small set within the space of all profiles of posteriors, which means that a generic profile of posteriors will not have a common prior.

Surprisingly, there has been very little accomplished to date in the systematic study of situations of non-common priors. So, what if we pick up the gauntlet, and do *not* impose the assumption of a common prior? One of the justifications for the common prior assumption that is often raised is the claim that, once we begin to relax the common priors assumption, 'anything is possible', in the sense that heterogeneous priors allow "sufficient freedom as to be capable of generating virtually any outcome". (The quote is from (Sa4). A similar argument appears in (Mo2)).

It is the intention of this paper to show that this is not so, at least with regards to no disagreements results, initiated by the seminal paper of (Au1). We show that to the contrary, the players' profile of posteriors always establishes an upper bound on the extent of disagreements under common knowledge. In fact, following this line of inquiry leads to the conclusion that, in a sense, the prior stage, if such exists, can be ignored; the posteriors alone suffice for establishing bounds on disagreements, and often those bounds will be tighter than the bounds we would have been led to believe exist from consideration of the 'historical' priors alone.

The main results of the paper are, in brief, as follows. Let τ be a type profile for a two-player knowledge space $(\Omega, \mathbf{\Pi})$, where Ω is a finite state space that is connected with respect to the partition profile $\mathbf{\Pi}$.[1] Then we show that we can always associate τ with a value

[1] The assumption that Ω is connected is adopted in the introduction for simplicity. As detailed in the body of the paper, the natural units for considering the subject of nearness of priors are the elements of the meet of a partition profile. This arises even

ε, normalised to $0 \leq \varepsilon \leq 1$, which we term the prior distance of τ, and find a pair of priors that are ε-almost common priors. This is the intuitive measure of 'how far' the space is from a common prior.

The significance of this definition is as follows. Let $E_i f(\omega)$ denote player i's expected value of a random variable f at the state ω. If τ has ε prior distance, and if it is common knowledge at any state ω^* that $E_1 f(\omega^*) \geq \eta_1$, and $E_2 f(\omega^*) \leq \eta_2$, then $|\eta_1 - \eta_2| \leq 2\varepsilon \|f\|_\infty$ (the factor of 2 is there because f can take negative values; if f is restricted to non-negative values, that doubling factor disappears).

This result can be reformulated in a way that generalises the standard 'no betting' characterisation of the existence of common priors, which states that τ has a common prior if and only if there is no bet f, in which player 1 takes the opposite side of the bet to player 2 in each state, such that it is common knowledge that both agents expect a positive gain from the bet (see (Mo1), (Sa2), and (Fe1)). However, suppose that there is no common prior, and that we have found a bet f at which it is common knowledge that both players expect a positive gain, but that we charge the players a transaction cost for undertaking the bet that is higher than the most each of them can expect to gain — then the players will refrain from taking the bet. The ε prior distance of τ implies a systematic way to do this. We can show that if τ has ε prior distance, then for every random variable f, it cannot be common knowledge that in every state ω, $E_1 f(\omega) > \varepsilon \|f\|_\infty$, and $E_2 f(\omega) < -\varepsilon \|f\|_\infty$. This can be interpreted as saying that for any f, it cannot be the case that player 1 takes the opposite side to player 2 of the bet f, and it is common knowledge that both expect to gain more than $\varepsilon \|f\|_\infty$.

In the n-player case, we again associate each type profile with a prior distance. If the prior distance is ε, then there is no n-tuple of random variables $f = (f_1, \ldots, f_n)$, such that $\sum_i f_i = 0$ and it is common knowledge that in every state, $E_i f(\omega) > \varepsilon \|f\|_\infty$. Again, this can be interpreted as establishing a bound on common knowledge of gains from bets.

---

if we limit consideration to common priors, as there are examples of type profiles that have a common prior whose support lies only in some of the meet elements. If we focus on connected state spaces, this technical matter can be ignored.

2 Preliminaries

2.1 Distributions

For a set Ω, denote by $\Delta^\Omega \subset \mathbb{R}^\Omega$ the simplex of probability distributions over Ω. An *event* is a subset of Ω, with the set of events denoted by Σ. A *random variable* f over Ω is any element of $\mathbb{R}^\Omega$. Given a probability measure $\mu \in \Delta^\Omega$ and a random variable f, the *expected value of* f with respect to μ, $E_\mu f := \sum_{\omega \in S} f(\omega)\mu(\omega)$. For an event A such that $\mu(A) > 0$, $E_\mu(f|A) := \sum_{\omega \in A} f(\omega)\mu(\omega)/\sum_{\omega \in A} \mu(\omega)$. Relative to a given probability distribution, the expected value of an event H is the expected value of the standard characteristic function 1^H which is defined as:

$$1^H(\omega) = \begin{cases} 1 & \text{if } \omega \in H \\ 0 & \text{if } \omega \notin H \end{cases}$$

2.2 Knowledge and Belief

A *knowledge space* for a nonempty, finite set of *players* I, is a pair $(\Omega, \mathbf{\Pi})$. In this context, Ω is a nonempty set called a *state space*, and $\mathbf{\Pi} = (\Pi_i)_{i \in I}$ is a *partition profile*, where for each $i \in I$, Π_i is a partition of Ω into measurable sets with positive measure. When working with a knowledge space $(\Omega, \mathbf{\Pi})$, an element $\omega \in \Omega$ is typically termed a *state*. For each $\omega \in \Omega$, we denote by $\Pi_i(\omega)$ the element of Π_i containing ω. Π_i is interpreted as the information available to player i; $\Pi_i(\omega)$ is the set of all states that are indistinguishable to i when ω occurs.

Player i is said to *know* an event E at ω if $\Pi_i(\omega) \subseteq E$. We define for each i a *knowledge operator* $K_i \colon 2^\Omega \to 2^\Omega$, by $K_i(E) = \{\omega \mid \Pi_i(\omega) \subseteq E\}$. Thus, $K_i(E)$ is the event that i knows E. A partition Π' is a *refinement* of Π if every element of Π' is a subset of an element of Π. Refinement intuitively describes an increase of knowledge.

The *meet* of $\mathbf{\Pi}$, denoted $\wedge\mathbf{\Pi}$, is the partition that is the finest among the partitions that are simultaneously coarser than all the partitions Π_i. We will denote by $\mathbf{\Pi}(\omega)$ the element of $\wedge\mathbf{\Pi}$ containing ω. $\mathbf{\Pi}$ is called *connected* when $\wedge\mathbf{\Pi} = \{\Omega\}$.

A *type function* for Π_i is a function $t_i : \Omega \to \Delta^\Omega$ that associates with each state ω a distribution in Δ^Ω, in which case the latter is termed the *type* of i at ω. Each type function t_i further satisfies the following two conditions:

(a) $t_i(\omega)(\Pi_i(\omega)) = 1$, for each $\omega \in \Omega$;
(b) t_i is constant over each element of Π_i.

A *type profile* for $\mathbf{\Pi}$ is a vector of type functions, $\tau = (t_i)_{i \in I}$, where for each i, t_i is a type function for Π_i, which intuitively represents the player's beliefs. For $f \in \mathbb{R}^\Omega$, denote by $E_i f$ the element of $\mathbb{R}^\Omega$ defined by $E_i f(\omega) = t_i(\omega) f$.

2.3 Common Priors and Common Knowledge

A *prior* for a type function t_i is a probability distribution $p \in \Delta^\Omega$, such that for each $\pi \in \Pi_i$, if $p(\pi) > 0$, and $\omega \in \pi$, then $t_i(\omega)(\cdot) = p(\cdot|\pi)$. If we start with a probability distribution p over Ω, and then consider a partition Π_i, we can always form a type function t_i for which p is a prior, by applying Bayes' Rule, relative to $\Pi_i(\omega)$, at each state ω. Relative to a type profile τ, denote the set of all priors of player i by $P_i(\tau)$, or simply by P_i when τ is understood.[2] In general, P_i is a set of probability distributions, not a single element; as pointed out by (Sa2), P_i is the convex hull of all of i's types.

A *common prior* for the type profile τ is a probability function $p \in \Delta^\Omega$ which is a prior for each player i.[3] A type profile τ is called *consistent* (the term is due to Harsányi) when it has a common prior.

An event $E \subseteq \Omega$ is *self-evident* if for all $\omega \in E$ and each $i \in I$

(1) $$\Pi_i(\omega) \in E.$$

In particular, every element of the meet, $M \in \wedge\mathbf{\Pi}$, is self-evident.

[2] Strictly speaking, the set of priors of a player i depends solely on i's type function t_i, not on the full type profile τ. However, since we are studying connections between sets of priors of different players, we will find it more convenient to write $P_i(\tau)$, as if P_i is a function of τ.

[3] Contrasting a prior for t_i with the types $t_i(\omega, \cdot)$, the latter are referred to as the posterior probabilities of i.

An event E is *common knowledge at* $\omega \in \Omega$ iff there exists a self-evident event $F \ni \omega$ such that for all $i \in I$

$$F \subseteq K_i(E). \tag{2}$$

In fact, the element of the meet containing ω is also known as the *common knowledge component* of ω, because it is the smallest self-evident set containing ω.

2.4 Characterisation of the Existence of Common Priors

We adopt the standard notation that for vectors $x_1, x_2 \in \mathbb{R}^\Omega$, $x_1 > x_2$ means that $x_1(\omega) > x_2(\omega)$ for all $\omega \in \Omega$, and $x_1 > 0$ means that $x_1(\omega) > 0$ for all ω. $x_1 \geq x_2$ means that $x_1(\omega) \geq x_2(\omega)$ for all $\omega \in \Omega$, and there is at least one ω^* such that $x_1(\omega^*) = x_2(\omega^*)$.

In the two-player case, i.e. $|I| = 2$, the existence of common priors is characterised by the statement:

> A type space τ has a common prior iff there is no $f \in \mathbb{R}^\Omega$ such that
> $$E_1 f > 0 > E_2 f.$$

When $|I| = n$, the characterisation of the existence of common priors is accomplished by:

> A type space τ has a common prior iff there are no $f_1, \ldots, f_n \in \mathbb{R}^\Omega$, such that $\sum_{i=1}^n f_i = 0$, and $E_i f_i > 0$ for all $i \in I$.

The two-player characterisation is a special case of the n-player case, slightly reworded.

The functions f_i, which sum to zero, can be interpreted as a bet between the players. The condition $E_i f_i(\omega) > 0$, for each state ω, amounts to saying that the positivity of $E_i f_i$ is always common knowledge amongst the players. Thus, a necessary and sufficient condition for the existence of a common prior is that there is no bet for which it is always common knowledge that all players expect a positive gain. This establishes a fundamental, and remarkable, two-way connection between posteriors and priors.

The most accessible proof of this result is in (Sa2). It was proved by (Mo1) for finite type spaces and independently by (Fe1) for compact type spaces. (BN1) proved it for finite type spaces with two agents.

2.5 Background Assumptions

Through the main body of this paper, the state space Ω will be assumed to satisfy $|\Omega| = m$, where m is a positive integer. All partition profiles will further be assumed to be connected, i.e., all meets will all be assumed to be singletons. The extension of results to non-connected spaces is straightforward.[4] Working with a connected space Ω has the advantage that a statement such as 'it is common knowledge that $E_i f > 0$' reduces to '$E_i f(\omega) > 0$ for all $\omega \in \Omega$.

3 The Two Player Case

3.1 Prior Distance

As noted in Section 2.4, the shortest and most direct proof of the Common Prior Characterisation Theorem, appears in (Sa2). It is ultimately based on the observation that the players' sets of priors, $\{P_i\}_{i \in N}$, are compact and convex sets. In the two-player case, the proof runs as follows:

Since a common prior exists if and only if $P_1 \cap P_2 \neq \emptyset$, this immediately implies that there does not exist a common prior if and only if P_1 and P_2 can be strongly separated, i.e., if and only if there is a random variable $g \in \mathbb{R}^\Omega$ and $c \in \mathbb{R}$, such that $x_1 g > c > x_2 g$, for every $x_1 \in P_1$ and $x_2 \in P_2$.

[4] We mention here in passing that even when relating to common priors it is natural to restrict consideration to meet elements, rather than entire state spaces. Consider a simple example: let τ be defined by $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, $\Pi_1 = \Pi_2 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$, $t_1 = \{\{\frac{1}{3}, \frac{2}{3}\}, \{\frac{1}{2}, \frac{1}{2}\}\}$, $t_2 = \{\{\frac{1}{3}, \frac{2}{3}\}, \{\frac{1}{4}, \frac{3}{4}\}\}$. τ is a positive type profile, and it is in fact consistent: it has a single common prior, $p = \{\frac{1}{3}, \frac{2}{3}, 0, 0\}$. But while the players cannot agree to disagree under common knowledge at states ω_1 and ω_2, they will have disagreements at states ω_3 and ω_4. If we break down $\mathbf{\Pi}$ into its two meet elements, $M_1 = \{\omega_1, \omega_2\}$ and $M_2 = \{\omega_3, \omega_4\}$, it is clear that the source of this is that τ_{M_1} has a common prior, but τ_{M_2} does not.

The proof then proceeds, by appropriately subtracting terms from the coordinates of g to yield f such that $x_1 f > 0 > x_2 f$, and then concludes, since this must hold true for the extreme points of P_1 and P_2, that a common prior fails to exist if and only if $E_1 f > 0 > E_2 f$, which is equivalent to common knowledge amongst the players that one player ascribes positive expectation at every state to f, and that the other player ascribes negative expectation to f at every state. Hence we have found a bet with respect to which there is common knowledge that both players believe they have positive expectation of winning.

The starting point here is the further observation that one can draw even more information from the separation between the sets of priors. To begin with, there is a standard measure of distance between two compact sets. Let $\|\cdot\|$ be any norm on $\mathbb{R}^{\Omega}$. Then define the $d_{\|\cdot\|}$-distance between the common prior sets, $d_{\|\cdot\|}(P_1, P_2)$, by

$$d_{\|\cdot\|}(P_1, P_2) := \inf_{\rho_1 \in P_1, \rho_2 \in P_2} \|\rho_1 - \rho_2\| . \tag{3}$$

Furthermore, by compactness, there are distributions $\mu_1 \in P_1$ and $\mu_2 \in P_2$, at which the minimal distance $d_{\|\cdot\|}(P_1, P_2)$ is attained. We choose to measure the distance between the sets of priors using the L^1 norm, $\|x\|_1 := \sum_{i=1}^{m} |x_i|$. This gives us, for each type profile τ, the unique ε we seek:

Definition 3.11. Given a type profile τ, the *prior distance* of τ is

$$\mathrm{pd}(\tau) := \frac{d_{\|\cdot\|_1}(P_1, P_2)}{2}$$

♦

The division by two in Definition 3.11 is for the sake of normalisation: $0 \le \mathrm{pd}(\tau) \le 1$ for any τ. Since τ has a common prior if and only if the intersection of P_1 and P_2 is non-empty, it has a common prior if and only if $\mathrm{pd}(\tau) = 0$. The prior distance in a sense measures 'how far' the type space is from having a common prior.

Definition 3.12. Given a type profile τ such that $\mathrm{pd}(\tau) = \varepsilon$, a pair of priors $\mu_1 \in P_1$ and $\mu_2 \in P_2$ that are of L^1 distance 2ε from each other and satisfy the property that there exist a pair of parallel supporting hyperplanes of P_1 and P_2 at μ_1 and μ_2, respectively, are *nearest priors of* τ. Given a pair of nearest priors μ_1 and μ_2, a point μ that is equidistant between μ_1 and μ_2 is an *almost common prior* of τ. We will also sometimes call a pair of nearest priors ε*-nearest priors*, and an almost common prior an ε*-almost common prior*, we when wish to stress that $\mathrm{pd}(\tau) = \varepsilon$. ♦

3.2 Bounded Common Knowledge Disagreement

Looking again at the proof of the characterisation of the existence of common priors in (Sa2), notice that it hinges on the existence of a random variable $g \in \mathbb{R}^{\Omega}$, and $\alpha \in \mathbb{R}$, such that

$$x_1 g > \alpha > x_2 g$$

for all $x_1 \in P_1$ and $x_2 \in P_2$. This can be re-written as the existence of $\beta, \gamma \in \mathbb{R}$ such that

$$(4) \qquad x_1 g \geq \beta > \alpha > \gamma \geq x_2 g \text{ for all } x_1 \in P_1 \text{ and } x_2 \in P_2.$$

Intuitively, $|\beta - \gamma|$ represents the 'extent of common knowledge disagreement' between the players, with respect to g. The question that will concern us in this section is: can we identify a g, and β, γ, such that (4) obtains, *and* $|\beta - \gamma|$, the extent of common knowledge disagreement, is maximal, among all random variables g in some bounded subset of $\mathbb{R}^{\Omega}$?

To make the connection between almost common priors and common knowledge disagreement, we make use of a fundamental property of duality. Let $\mathrm{pd}(\tau) = \varepsilon$, and let μ_1 and μ_2 be ε-nearest priors, with μ the corresponding ε-almost common prior. Denote the support functional of P_i by

$$h_i(f) = \sup_{\phi \in P_i} (\phi \cdot f),$$

and then denote

$$(5) \qquad \delta_\tau = \max_{\{f \in \mathbb{R}^{\Omega} : \|f\|_\infty \leq 1\}} [\mu \cdot f - h_1(f)].$$

Applying the Minimum Norm Duality Theorem (see Theorem 1 on page 136 of (Lu1)), $\delta_\tau = \varepsilon$, and there is a random variable f^* such that $\|f^*\|_\infty \leq 1$ and δ_τ is attained. Furthermore, f^* is aligned with $\mu - \mu_1$, and hence (by Definition 3.12), is also aligned with $\mu_1 - \mu_2$, and $-f^*$ is the random variable of L^∞-norm less than or equal to one at which $\mu \cdot f - h_2(f))$ is maximised.

Geometrically, identifying f^* is tantamount to finding the parallel supporting hyperplanes of P_1 and P_2 of greatest distance.[5] By Equation 5, there is a β such that

$$x_1 f^* \geq \beta$$

for all $x_1 \in P_1$ (thus defining a hyperplane), and similarly there is a γ such that

$$\gamma \geq x_2 f^*$$

for all $x_2 \in P_2$. As a consequence of the Minimum Norm Duality Theorem,

$$|\beta - \gamma| = 2\varepsilon,$$

the L^1 distance between these two hyperplanes (and between P_1 and P_2), and furthermore, there are no $g \in \mathbb{R}^\Omega$ with $\|g\|_\infty \leq 1$, and $b, d \in \mathbb{R}$, such that $|b - d| > 2\varepsilon$ and $x_1 g \geq b > d \geq x_2 g$ for every $x_1 \in P_1$ and $x_2 \in P_2$. We have thus accomplished the goal we set in the paragraph immediately after Equation (4).

Proposition 3.21. A two-player type profile τ has ε prior distance if and only if

1. there exists a random variable $\|f^*\|_\infty \leq 1$, and a pair $(\mu_1, \mu_2) \in P_1 \times P_2$, such that $E_{\mu_1} f^* - E_{\mu_2} f^* = 2\varepsilon$; and
2. there does not exist any random variable $\|g\|_\infty \leq 1$ such that $E_{\varphi_1} g - E_{\varphi_2} g > 2\varepsilon$ for every pair $(\varphi_1, \varphi_2) \in P_1 \times P_2$.

Proof. Suppose that $\mathrm{pd}(\tau) = \varepsilon$, with ε-nearest priors (μ_1, μ_2) and ε-almost common prior μ. As shown in the previous paragraphs, that

[5] The study of the maximal distance between parallel supporting hyperplanes of convex sets has recently been a focus of efforts in statistical learning theory, where it goes under the name of seeking the 'maximal margin hyperplane', or 'optimal hyperplane'.

implies the existence of a random variable f^*, with $\|f^*\|_\infty \leq 1$, such that $\mu_1 f^* - \mu f^* + \mu f^* - \mu_2 f^* = E_{\mu_1} f^* - E_{\mu_2} f^* = 2\varepsilon$. At the same time, there does not exist g, with $\|g\|_\infty \leq 1$, such that for every pair $\varphi_1 \in P_1$ and $\varphi_2 \in P_2$, there exist $\beta, \gamma \in \mathbb{R}$, such that $\varphi_i g \geq \beta$ and $\varphi_j g \leq \gamma$ and $\beta - \gamma > 2\varepsilon$, i.e., $E_{\varphi_1} g - E_{\varphi_2} g > 2\varepsilon$.

In the other direction, if there exists an f^* as in the statement of the proposition, then f^* defines a pair of parallel supporting hyperplanes H_1 of P_1 and H_2 of P_2 of maximal distance apart 2ε. Choosing $\mu_1 \in H_1 \cap P_1$ and $\mu_2 \in H_2 \cap P_2$ gives a pair of ε-nearest priors. ∎

Proposition 3.22. A two-player type profile τ has ε prior distance if and only if

1. there does not exist any random variable $\|f\|_\infty \leq 1$ such that $E_1 f - E_2 f > 2\varepsilon$ for all $\omega \in \Omega$; and
2. there exists $\|f^*\|_\infty \leq 1$ such that $E_1 f^* - E_2 f^* \geq 2\varepsilon$.

Proof. There exists a random variable f, with $\|f\|_\infty \leq 1$, such that for all $\omega \in \Omega$, $E_1 f(\omega) - E_2 f(\omega) > 2\varepsilon$, if and only if for every ω, and type t_1 of player 1 at ω and type t_2 of player 2 at ω, $t_1 f - t_2 f > 2\varepsilon$. But every prior $p_i \in P_i$ is a convex combination of the player i's types. Hence this holds if and only if for every pair of priors $(\varphi_1, \varphi_2) \in P_1 \times P_2$, $\varphi_1 f - \varphi_2 f > 2\varepsilon$.

Noting that $\varphi_1 f - \varphi_2 f \leq \|\varphi_1 - \varphi_2\|_1 \|f\|_\infty$, we then immediately have that $\mathrm{pd}(\tau) = \varepsilon$ implies that there does not exist a random variable $\|f\|_\infty \leq 1$ such that for all $\omega \in \Omega$, $E_1 f(\omega) - E_2 f(\omega) > 2\varepsilon$. In addition, $\mathrm{pd}(\tau) = \varepsilon$ implies the existence of $\|f^*\|_\infty \leq 1$ such that $\varphi_1 f^* - \varphi_2 f^* \geq 2\varepsilon$ for all priors, while $\mu_1 f^* - \mu_2 f^* = \varepsilon$ for a pair of nearest priors μ_1, μ_2. But this can only hold true if $E_1 f^* - E_2 f^* \geq 2\varepsilon$.

In the other direction, the statement of the proposition implies that there does not exist any random variable $\|g\|_\infty \leq 1$ such that $E_{\varphi_1} g - E_{\varphi_2} g > 2\varepsilon$ for every pair of priors $(\varphi_1, \varphi_2) \in P_1 \times P_2$, and there exists $\|f^*\|_\infty \leq 1$ such that $E_{\mu_1} f^* - E_{\mu_2} f^* = 2\varepsilon$ (obtained by convex combinations giving positive weight only to types that are 2ε apart), so that by Proposition 3.21, $\mathrm{pd}(\tau) = \varepsilon$. ∎

Corollary 3.21. Let τ be a two-player type profile of ε prior distance, and let $\omega^* \in \Omega$. Let $f \in \mathbb{R}^\Omega$ be a random variable, and let

$\eta_1, \eta_2 \in \mathbb{R}$. If it is common knowledge at ω^* that player 1's expected value of f is greater than or equal to η_1, and player 2's expected value of f is less than or equal to η_2, then

$$|\eta_1 - \eta_2| \leq 2\varepsilon \, \|f\|_\infty .$$

Proof. Suppose that $|\eta_1 - \eta_2| > 2\varepsilon \, \|f\|_\infty$. Let $g := \frac{f}{\|f\|_\infty}$. Then g satisfies $\|g\|_\infty \leq 1$, yet $|E_1 g(\omega) - E_2 g(\omega)| > 2\varepsilon$ for all ω, contradicting the assumption that $\mathrm{pd}(\tau) = \varepsilon$. ∎

This also leads to a generalisation of the No Disagreements Theorem of (Au1), which reduces to it when $\varepsilon = 0$.

Corollary 3.22. Let τ be a two-player type profile of ε-prior distance, and let $\omega^* \in \Omega$. Let H be an event. If it is common knowledge at $\omega^* \in \Omega$ that $E_1(H) = \eta_1$ and $E_2(H) = \eta_2$, then $|\eta_1 - \eta_2| \leq \varepsilon$.

Proof. Let $f \in \mathbb{R}^\Omega$ satisfy $0 \leq f(\omega) \leq 1$ for all $\omega \in \Omega$. Then it cannot be the case that $|E_1 f(\omega) - E_2 f(\omega)| > \varepsilon$ for all ω. Suppose by contradiction that this statement holds. Let $g = 2f - 1$. Then $g \in \mathcal{D}^\infty$, and $|E_1 g(\omega) - E_2 g(\omega)| > 2\varepsilon$ for all ω, contradicting the assumption that $\mathrm{pd}(\tau) = \varepsilon$.

Consider the standard characteristic function 1^H. Since $0 \leq 1^H(\omega) \leq 1$ for all ω, and the expected value of 1^H at every state is the expected value of the event H at that state, the conclusion follows. ∎

Finally, we have our generalisation of the No Betting characterisation:

Theorem 1. *A two-player type profile* τ *satisfies* $\mathrm{pd}(\tau) = \varepsilon$ *if and only if*

1. *there does not exist a random variable* f *such that* $\max_\omega f(\omega) - \min_\omega f(\omega) \leq 2$, *and* $E_1 f > \varepsilon$ *and* $E_2 f < -\varepsilon$*; and*
2. *there exists* f^* *such that* $\max_\omega f^*(\omega) - \min_\omega f^*(\omega) \leq 2$, *and* $E_1 f^* \geq \varepsilon$ *and* $E_2 f^*(\omega) \leq -\varepsilon$.

Proof. By Proposition 3.22, $\mathrm{pd}(\tau) = \varepsilon$ iff there does not exist any random variable f such that $\|f\|_\infty \leq 1$ and $E_1 f - E_2 f > 2\varepsilon$, and there exists an f^* such that $\|f^*\|_\infty \leq 1$ and $E_1 f^* - E_2 f^* \geq 2\varepsilon$. But the existence of a random variable f such that $\|f\|_\infty \leq 1$ and $E_1 f - E_2 f > 2\varepsilon$ for all $\omega \in \Omega$ is equivalent to the statement that there exist $b, c, d \in \mathbb{R}$ such that $E_1 f(\omega) \geq b > c > d \geq E_2 f(\omega)$ for all ω, and $b - c = c - d = \varepsilon$. Defining $g := f - c$ yields $E_1(g)(\omega) > \varepsilon$ and $E_2(g)(\omega) < -\varepsilon$ for all ω, and clearly $\max_\omega g(\omega) - \min_\omega g(\omega) \leq 2$. Similar reasoning applies to f^*. ∎

We can interpret this in the following way: suppose, in the context of player in a type space of ε prior distance, that a book-maker seeks to make a 'Dutch book'-type profit, by exploiting the belief differences of the players with respect to a normalised random variable f of common knowledge disagreement between the players. He therefore proposes that player 1 take a long position on f, and that player 2 take the opposite short position, so that if the true state is revealed to be ω^*, player 1 gains/loses $f(\omega^*)$, and player 2 gains/loses $-f(\omega^*)$. If set up correctly, both players believe that they cannot lose, but in fact it is the book-maker who cannot lose – whatever he owes one player when the bet is to be paid is covered by what the other player pays him, while he profits from brokerage fees paid by both players. The proposition shows that, no matter which L^∞ normalised f is used, it will never be the case that it is common knowledge that both players expect to gain more than ε. Scaling $\|f\|_\infty$ by a scalar α scales the bound by the same α, so that in general, if $\mathrm{pd}(\tau) = \varepsilon$, it cannot be the case that it is common knowledge that both players expect to gain more than $\varepsilon \|f\|_\infty$ by taking opposites sides of a bet f. When ε is very small, this may imply significant limitations on betting, in cases in which there is no common prior.

4 The n-Player Case

Now let $|I| = n$, so that τ is an n player type profile over the state space Ω, where $|\Omega| = m$. There are n sets of priors $P_i(\tau)$, one for each player, and a common prior exists if and only if $\bigcap_{i \in I} P_i \neq \emptyset$.

Following an idea in (Sa2), consider the bounded, closed, and convex subsets of $\mathbb{R}^{mn}$

$$X := P_1 \times P_2 \cdots \times P_n, \tag{6}$$

and

$$Y := \{(p, p, \ldots, p) \in \mathbb{R}^{nm} | p \in \Delta^m\}. \tag{7}$$

Clearly, $\bigcap_{i=1}^n P_i = \emptyset$ iff X and Y are disjoint. When they are disjoint, we can, as in the previous section, seek points in X and Y, respectively, that are of minimal L^1 distance apart.

Definition 4.01. Given an n-player type space τ, and the corresponding spaces X and Y, as defined in Equations (6) and (7), if the L^1 distance between X and Y is δ, define the *prior distance* of τ, denoted $\mathrm{pd}(\tau)$, to be $\frac{\delta}{n}$. If n-tuples $x = (\mu_1, \ldots, \mu_n) \in X$ and $y = (p, \ldots, p) \in Y$ are of minimal distance δ apart, such that there exist a pair of parallel supporting hyperplanes of X and Y at x and y, respectively, then x is a tuple of *nearest priors*, and p is an *almost common prior* of τ.

As in the two-player case, division by n ensures normalisation. To see this, let $\|f\|_\infty \leq 1$ define parallel supporting hyperplanes of X and Y of L^1 distance δ apart, as in Definition 4.01, and x be nearest priors with p a corresponding almost common prior. Then $|xf - pf| = \delta$, but since p and x_i for all $1 \leq i \leq n$ are elements of Δ^m, it follows that $\delta \leq n$, hence $0 \leq \mathrm{pd}(\tau) \leq 1$.

The generalisation of the No Betting characterisation is then as follows.

Theorem 2. *An n-player type profile τ satisfies $\mathrm{pd}(\tau) = \varepsilon$ iff*

1. *there does not exist an n-tuple of random variables $(f_1, \ldots, f_n) \in \mathbb{R}^{nm}$ such that $\sum_i f_i = 0$, and for all i, $\max_\omega f_i(\omega) - \min_\omega f_i(\omega) \leq 2$ and $E_i f > \varepsilon$; and*
2. *there exists an n-tuple of random variables f^* such that $\sum_i f_i = 0$, and for all i, $\max_\omega f_i(\omega) - \min_\omega f_i(\omega) \leq 2$ and $E_i f \geq \varepsilon$.*

Proof. Denote $\delta = n\varepsilon$. Using reasoning similar to that applied several times in proofs above in the two-player case, $\mathrm{pd}(\tau) = \varepsilon$ if and only if there does not exist $f = (f_1, \ldots, f_n) \in \mathbb{R}^{nm}$ such that $\|f_i\|_\infty \leq 1$ for all i, and for all $x = (x_1, \ldots, x_n)$ in X and $y = (p, \ldots, p)$ in Y, $xf - yf > \delta$, and there exists $f^* = (f_1^*, \ldots, f_n^*) \in$

$\mathbb{R}^{nm}$ with $\|f_i^*\|_\infty \leq 1$, and nearest priors $x^* = (\mu_1, \ldots, \mu_n) \in X$ and a corresponding almost common prior p with $y^* = (p, \ldots, p)$, such that $x^* f^* - y^* f^* = \delta$.

The condition $xf - yf > \delta$, for all $x \in X$, and all $y \in Y$, with $f \in \mathbb{R}^{nm}$, $\|f_i\|_\infty \leq 1$, is equivalent to the existence of $b, c, d \in \mathbb{R}$ such that $xf \geq b > d \geq yf$ and $b - d > n\varepsilon$. We may then subtract d/n from all the components of f, yielding a new function $g = (g_1, \ldots, g_n)$ satisfying $\max_\omega g_i(\omega) - \min_\omega g_i(\omega) \leq 1$ for each i, such that for all $x \in X$, and all $p \in \Delta^m$

$$\sum_{i=1}^{n} x_i g_i > n\varepsilon,$$

and

$$\sum_{i=1}^{n} p g_i < 0.$$

Since the last inequality holds for all $p \in \Delta^m$, it is equivalent to $\sum_i g_i < 0$. Furthermore, since the coordinates of x_i are non-negative, uniformly increasing the coordinates of the g_i leaves $\sum_i x_i g_i > \delta$, so that we may assume that $\sum_i g_i = 0$.

The fact that $\sum_i x_i g_i > \delta = n\varepsilon$ still leaves the possibility that $x_i g_i < \varepsilon$ for some i. But let x_i^* be the point that minimizes $x_i g_i$ over P_i, for each i. Since $\sum_{i=1}^n x_i^* g_i > n\varepsilon$, there are constants c_i guaranteeing $x_i^* g_i + c_i > \varepsilon$ for each i, satisfying $\sum_i c_i = 0$. Denoting by e the vector in $\mathbb{R}^m$ whose every coordinate is 1, define $h_i = g_i + c_i e$. Then $\max_\omega h_i(\omega) - \min_\omega h_i(\omega) \leq 2$ for each i, $\sum_i h_i = \sum_i g_i = 0$, $\sum_i x_i h_i > n\varepsilon$, and for each $x_i \in P_i$, $x_i f_i > \varepsilon$. The conclusion of the theorem follows. ∎

We can interpret this in the following way: suppose that a book-maker seeks to make a 'Dutch book'-type profit, by exploiting the belief differences of the players, given a type profile with prior distance ε. He offers to sell the n players a portfolio of securities, one per player. Each player i buys the security f_i, and his gain/loss depends on the value of that security when the true state ω is revealed. These securities are carefully crafted in such a way that the book-maker cannot lose, by arranging that $\sum_i f_i = 0$, but are also crafted so that each player believes he or she has a sure win, since $E_i f_i(\omega) > 0$ at each state ω. Then it cannot be the case that it is common knowledge that every player expects to gain at least $\varepsilon \left\| f \right\|_{\infty}$.

5 Conclusion

In the traditional framework of Bayesian theory, one assumes that each player starts with a prior distribution over the set of possible states of the world. As (differential) information is obtained by the players, a partition profile is formed, which in turn leads to a type profile (i.e., posteriors) as Bayes' Rule is applied to each player's prior at each of his partition elements.

The dictionary entry for the word 'prior' includes two definitions: preceding in time/order; or preceding in importance. A prior distribution is called such because according to the traditional story it precedes in time the information leading to the posterior, but there is also a sense in which it is often perceived as preceding the posterior in importance, as being the 'prime factor' from which the posterior is passively derived.

The results of the previous sections indicate that, at least with regards to disagreements under common knowledge, the posteriors play a role that is more important than the priors. To see what is meant by this, for simplicitly restrict attention to the 2-player case (the argument applies just as well to the n-player case), let Ω be finite, and suppose that the players begin with priors φ_0^1 and φ_0^2, where the total variation distance between φ_0^1 and φ_0^2 is $\varepsilon_0 > 0$. As information is obtained, a partition profile $\mathbf{\Pi}_1$ (which we will assume for simplicity has a singleton meet) is formed, and a corresponding type profile τ_1 is derived from the prior. Even though we have explicitly started with the priors φ_0^1 and φ_0^2, the set of priors of player 1 in this type profile, $P_1(\tau_1)$, will generally contain more than one point (in addition to φ_0^1), and a similar statement holds true for $P_2(\tau_1)$. It follows that $\varepsilon_1 = \mathrm{pd}(\tau_1)$, the prior distance of τ_1, satisfies $\varepsilon_1 \leq \varepsilon_0$. That is, it cannot be greater than ε_0 – we can summarise this insight as 'increasing information can never increase (common knowledge) disagreements' – but it may well be smaller. Let φ_1^1 and φ_1^2 be the ε_1-almost common priors of τ_1. If $\varepsilon_1 < \varepsilon_0$, we may disregard the 'historical' derivation of τ_1 from φ_0^1 and φ_0^2, and for all intents and purposes deal with τ_1 'as if' it was derived instead from φ_1^1 and φ_1^2.

This process can be continued. Let $\Omega \succ \mathbf{\Pi}_1 \succ \mathbf{\Pi}_2 \succ \ldots$ be a sequence of proper partition refinements (again, for simplicity, suppose that they all have singleton meets; alternatively, we can fix a 'true state' ω^* and focus on the sequence of connected components $\mathbf{\Pi}_1(\omega^*) \succ \mathbf{\Pi}_2(\omega^*) \succ \ldots$). Form a corresponding sequence of type profiles $\tau_1, \tau_2, \ldots$, by deriving, as above, τ_1 from $\mathbf{\Pi}_1$ and the original priors φ_0^1 and φ_0^2, deriving τ_2 from $\mathbf{\Pi}_2$ and the ε_1-almost common priors φ_1^1 and φ_1^2 of τ_1, and so forth. We then have a corresponding sequence $\varepsilon_0 \geq \varepsilon_1 \geq \varepsilon_2 \geq \ldots$, where for each j, ε_j is the prior distance of τ_j.

By results appearing in (HS1), the sequence $\varepsilon_0 \geq \varepsilon_1 \geq \varepsilon_2 \geq \ldots$ must end,[6] at some $n \leq |\Omega|-1$, with $\varepsilon_n = 0$, at which point we are as far as possible from the original separation of the priors, arriving at a situation in which we may as well have started with a common prior. It is in this sense that the posterior, i.e. the structure of the partition profile and the type profile, plays a greater role in determining limits of disagreements than the prior.

In the context of common priors, this remark is not new. (Ge1), for example, presents a version of the well-known envelopes problem in which the players refrain from betting, not because their posteriors are derived from common priors, but because they know that the posteriors could have been derived from a common prior, and hence they know they cannot disagree. What we have here is an extension of this principle to all type profiles – what count for bounding disagreements are not the historical priors, but the fictional almost common priors from which the posteriors could have been derived.

The question of the need to assume a dynamic framework, with an explicit prior stage followed by a 'current', or posterior stage, plays a central role in a debate over the common prior assumption between (Gul) and (Au3). (Gul) argues against always adopting such a dynamic view. In his reply, (Au3), while defending the dynamic

---

[6] It is shown in (HS1) that a partition profile that is tight, as defined there, always has a common prior. In the case of a finite state space and two players, a partition profile is tight if $|\Pi_1|+|\Pi_2| = |\Omega|+1$. It follows that a process of proper refinements, which always add partition elements, inevitably brings about a tight partition profile.

framework as ‘perfectly legitimate and intuitive’, also agrees that it would be desirable to deal ‘directly in terms of the “current”, posterior probabilities, without any reference, either implicit or explicit, to any prior stage.’ As the discussion here shows, all the data needed for bounds on disagreements can be known from the posterior probabilities, without reference to a prior stage. Indeed, even if there was, historically, a prior stage, one is better off ignoring the ‘true’ priors and considering instead the almost common priors.

Bibliography

[Au1] Aumann, R. J.: Agreeing to Disagree, Ann. Statist., **4(6)** (1976) 1236–1239.

[Au2] Aumann, R. J.: Correlated Equilibrium as an Expression of Bayesian Rationality, Econometrica, **55** (1987) 1–18.

[Au3] Aumann, R. J.: Common Priors: A Reply to Gul, Econometrica, **66** (1998) 929–938.

[BN1] Bonanno, G., Nehring K.: How to make sense of the common prior assumption under incomplete information, Int. J. Game Theory, **28** (1999) 409–434.

[Fe1] Feinberg, Y.: , Characterizing Common Priors in the Form of Posteriors, J. Econ. Theory, **91** (2000) 127–179.

[Ge1] Geanakoplos, J.: , Common Knowledge, Handbook of Game Theory, Vol. 2 (1994) 1439-1496, Elsevier Science B.V., the Netherlands.

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

[Ha1] Harsányi, J. C.: Games with Incomplete Information Played by Bayesian Players, Management Science, **14** (1967-8) 159–182, 320–334, 486–502.

[He1] Heifetz, A.: The Positive Foundation of the Common Prior Assumption, Games and Economic Behavior, **56** (2006) 105–120.

[HS1] Hellman, Z., Samet, D.: How Common are Common Priors, Discussion Paper 532, (2010) The Centre for the Study of Rationality, Hebrew University.

[Lu1] Luenberger, D. G.: Optimization by Vector Space Methods, (1969) John Wiley and Sons, New York, NY.

[Mo1] Morris, S.: Trade with Heterogenous Prior Beliefs and Asymmetric Information, Econometrica, **62** (1995) 1327–1347.

[Mo2] Morris, S. (1995b): The Common Prior Assumption in Economic Theory, Economics and Philosophy, **11** (1995) 227–253.

[Sa1] Samet, D.: Iterated Expectations and Common Priors, Games and Economic Behavior, **24** (1998) 131–141.

[Sa2] Samet, D.: Common Priors and Separation of Convex Sets, Games and Economic Behavior, **24** (1998) 172–174.

[Sa3] Samet, D.: Quantified Beliefs and Believed Quantities, J. Econ. Theory, **95** (2000) 169–185.

[Sa4] Samuelson, L.: Modeling Knowledge in Economic Analysis, Journal of Economic Literature, **42** (2004) 367–403.

[SG1] Sebenius, J., Geanakoplos, J.: Don't Bet on it: Contingent Agreements with Asymmetric Information, Journal of the American Statistics Association, **78** (1983) 424–426.