Common Knowledge and Disparate Priors: When it is O.K. to Agree to Disagree

Hellman, Ziv

15 May 2007

Common Knowledge and Disparate Priors – When it is O.K. to Agree to Disagree

Ziv Hellman
zivyahel @ gmail.com[1]

Abstract. Abandoning the oft-presumed common prior assumption, partitioned type spaces with disparate priors are studied. It is shown that in the two-player case, a unique fundamental pair of priors (p_1, p_2) can be identified in each type space, from whose properties boundaries on the possible ranges of expected values under common knowledge can be derived. In the limit as p_1 and p_2 approach each other, $p = p_1 = p_2$ is a common prior, and standard results stemming from the common prior assumption are recapitulated. It is further shown that this two-player fundamental pair of priors is a special case of the n-player situation, where a representative n-tuple of fundamentally associated priors $(p_0, ..., p_{n-1})$ can be selected, out of at most $n-1$ such n-tuples, to play an analogous role.
JEL Classification Numbers: C70, D82, D84.

1. Introduction

1.1 The Common Prior Assumption

The common prior assumption (CPA) is, in a sense, one (mathematically rigorous) answer to the age-old philosophical question 'how can reasonable and honest individuals come to disagree?' The CPA, as widely adopted in much of economics, game theory and decision theory literature, responds to this question via what has come to be called 'the Harsányi Doctrine', namely the position that all women and men are 'created equal' with respect to probability assessments in the absence of information – hence the term common prior – and all differences in probabilities should therefore, in principle, be traced to asymmetries in information received over time.

It is difficult to over-state the pervasiveness of the common prior assumption. It suffices in this regard to quote the words of Aumann (1987), which still hold true despite the years that have passed since they were written:

> 'Common priors are explicit or implicit in the vast majority of the differential information literature in economics and game theory… The assumption is pervasive in the enormous literature on rational expectations, trading in securities, bargaining under incomplete information, auctions, repeated games, signalling, discrimination, insurance, principal-agent, moral hazard, search, entry and exit, bankruptcy, what have you. Citing the relevant papers would

[1] I would like to thank Oliver Board for helpful comments and advice.

> make our references longer than our text. Occasionally the definitions do pay lip-service to the possibility of distinct priors p^i; but usually this is quickly abandoned, and in the theorems and examples, one returns to common priors'.

The CPA is also a crucial assumption under-pinning the celebrated 'agreeing to disagree'[2] paper of Aumann (1976), which proves a surprising theorem showing that it is impossible for private information to lead to divergent beliefs under conditions of common knowledge. Numerous authors have since extended this result and applied it to interactions between agents in various situations. The typical result is a 'no-bet' or 'no-trade' theorem (cf. Milgrom and Stokey (1982), Sebenius and Geanakopolos (1983)) – agents who start with common prior distributions will *never* agree to engage in speculative trade based on differences in private information that they subsequently receive. As soon as it becomes common knowledge that they wish to trade, their expectations for the value of assets in question become identical.

As Nau (1995) points out 'these results are perceived to be a problem for the theory of speculative markets: asymmetric information alone cannot be responsible for the existence of large stock exchanges… It is a point which is crucial for the understanding of the very complex speculative markets we see nowadays'.

The disconnect between no bet/no trade theoretical models and empirical reality may be due to several factors, among them risk-aversion issues, bounded rationality, lack of common knowledge, the cost of information exchange, and errors in information transmission and/or reception. In this paper, however, attention will be focussed solely on the common prior assumption.

The CPA may, and has, been challenged (cf. Gul (1998) and Nau (1995) for only two of many examples) for being an assumption that is far too strong to be believed to apply in reality as much as one might suspect given how often it is assumed in models in the literature. The CPA is often accompanied by a story that postulates that the current probabilistic beliefs of players all stem from temporal Bayesian updating conducted under conditions of asymmetric information – a story that may be fictional and/or irrelevant to the model being studied. Furthermore, it supposes that if one goes back sufficiently far in this historical story, there was a point in time when all the players were in possession of the same information *and* were in full agreement on a prior probability distribution. Philosophically, one may object there never was a primeval moment in time when all individuals were in exactly the same state of information – everyone receives different sensory data and filters it through his or her own cognitive model of the physical and social environment from the moment of birth.

Partly in response to this chorus of objections, efforts were conducted in the 1990s to seek out a full characterisation of when the CPA may and may not hold in a model. These efforts were crowned with success, resulting in at least two different characterisations.

[2] It should be noted that although the result is nearly always called 'agreeing-to-disagree', in actual fact it states precisely the opposite – under common knowledge and common priors the players can *not* agree to disagree.

Independently, Morris (1995), Samet (1998b) and Feinberg (2000), proved one characterisation, in which the presence or absence of a common prior depends on the absence or presence of at least one bet which seems, from each player's private perspective, to guarantee him or her positive expected value. In addition to this, Samet (1998a) defined and proved a characterisation based on the convergence of 'infinite iterated expectations'. Both of these characterisations will play major roles in this paper.

Since then, however, there has been very little written on the next obvious question: given an apparent need for the systematic study of models that do not presume common priors, and the existence to hand of full characterisations of the common prior situation, what non-trivial results can be attained in models in which the common prior assumption is removed as an axiom, and instead disparate priors[3] are taken into account? This paper represents an attempt to begin answering this question.

We are especially inspired by Morris (1995) (and Bernheim (1986) before him) in seeking to illustrate that it is not the case, despite what is sometimes claimed, that 'anything can happen' if the common prior assumption is relaxed. To the contrary, we strive here to show that even with disparate priors it is possible to derive interesting bounding theorems under conditions of common knowledge, and indeed to place common priors inside a broader context, so that – in line with a 'correspondence principle' – the standard results stemming from the common prior assumption re-emerge in the limit as disparate priors approach common priors.

An important debt is also acknowledged to Samet (1998a) and Nehring (2001), for theorems appealed to, methods of proofs, and ideas and inspiration in general.

1.2 Summary of Results

The formal treatment of the beliefs of players is usually conducted by representing those beliefs by use of a *partitioned type space*. In this model, players are assumed not to know everything about the world, and instead to consider a set of possible states, only one of which is the true state. The players are not perfectly informed, and are unsure which is the true state. Each player's knowledge is represented by an information partition, which divides the states into a number of mutually disjoint and exhaustive subsets. If two states are in the same partition, the player cannot tell them apart; instead, the player has a probability distribution on each partition, giving the likelihood of which state in the partition is the true state, under the assumption that the true state is located somewhere within that partition.

It is convenient to summarise player i's knowledge by way of a 'type matrix' M_i from which the player's partitions and probability distribution can easily be read. These matrices also have the desirable properties that player i's expectation of a random

[3] The term 'common prior' is universally used to describe the case in which players share at least one prior between them. There is no such uniformity of terminology in the literature to describe the converse situation. That case is termed here 'disparate priors', but the same concept has been called 'heterogeneous priors', 'distinct priors', 'unequal priors' 'non-common priors' or 'uncommon priors' in other papers.

variable *f* is simply given by $M_i f$, and a prior probability distribution for player *i* by $p_i M_i = p_i$. Such a distribution is called a common prior if a single p satisfies $pM_i = p$ simultaneously for all players.

The Morris-Feinberg characterisation of common priors (Morris (1995), Feinberg (2000)), establishes the existence of common priors via the absence of mutually profitable bets. Samet's characterisation (Samet (1998a)), in contrast, shows that the existence of a common prior can be interpreted more directly in terms of the players' beliefs as encoded in the type space and type matrices.

Samet himself provides an intuitive explanation of his result in this way: imagine that Adam and Eve – who have both excelled in their studies at the same school of economics – are asked what return they expect on IBM stock. Having been exposed to different sources of information, we oughtn't be surprised if the two provide different answers. But we can then go on to ask Eve what she thinks Adam's answer was. Being a good Bayesian, she can compute the expectation of various answers and come up with Adam's expected answer. Likewise, Adam can provide us with what he expects was Eve's answer to that question. This process can continue, moving back and forth between Eve and Adam, theoretically forever. There are, in this example, two possible infinite sequences of alternating expectations, one that starts with Eve and one that starts with Adam.

Samet calls this process 'obtaining an iterated expectation', and shows that there exists a common prior if and only if both of these sequences converge to the same limit. He achieves this result by representing Adam's beliefs[4] by a type matrix M_1 and Eve's beliefs by type matrix M_2. These then form two 'permutation matrices', $M_{\sigma_1} = M_2 M_1$, which is intended to be used for the process of obtaining iterated expectations starting with Adam, and $M_{\sigma_2} = M_1 M_2$, which does the same for the iterated expectations starting with Eve. It then turns out to be the case that both M_{σ_1} and M_{σ_2} are ergodic Markov matrices, and hence by standard results in Markov chain theory, each of them has a unique invariant probability measure, which may be labelled respectively p_1 and p_2. We can call these 'Samet probability measures'. It is then shown in Samet (1998a) that if $p_1 \neq p_2$, Adam and Eve cannot share a common prior. On the other hand, if $p_1 = p_2$, then not only is there a common prior, it has positively been identified – it is precisely $p := p_1 = p_2$.

This is a remarkable result, made all the more remarkable by the fact that it applies results developed in Markov theory for the study of stochastic processes to answer a question that seems not to be even remotely related. But it still leaves some remaining questions. For one thing, as Nehring (2001) points out, the condition is 'epistemically somewhat opaque. In particular, by looking at the limit, it is no longer transparent who does the expecting, and even what the direct object of expectation is; only some "ultimate

[4] For the sake of simplicity here, we will make the mild technical assumption that the entire relevant state space is the meet of the type spaces of Adam and Eve.

object of expectation" is given'. Infinite objects of contemplation are certainly not unknown in game theory, which pioneered the study of such infinitary statements as 'she knows that I know that she knows that I know that …'. Never the less, a finitary characterisation of the same concept can be expected to add insight. Secondly, one might also ask if this characterisation can be expressed in a way that is intrinsic to the subject to hand, without directly referring to Markov chain concepts. Finally, Samet (1998a) shows that if the limits of the iterated expectations do not converge to one and the same limiting vector, there is no common prior, but is silent about what those limits *do* tell us in the absence of a common prior.

Proposition 1 of this paper answers this last question by showing that, in the 2-player case, whether or not the iterated expectations limits converge to the same vector, the Samet probability measures are, under all conditions, priors. In fact, maintaining the notation of p_1 and p_2 as above, p_1 is a prior for player 1, and p_2 is a prior for player 2.

Furthermore, letting p' be a prior for player 1 and p'' be a prior for player 2 (so that $p'M_1 = p'$ and $p''M_2 = p''$) we define the pair (p', p'') as being *balanced* if they satisfy the equations $p'M_2 = p''$ and $p''M_1 = p'$ – in a sense, Adam's prior from this pair is 'fundamentally paired' to Eve's prior when Adam plugs in Eve's prior into his equation defining a prior, and *vice versa*. This simply-stated definition gives us the finitary characterisation sought in the previous paragraph: the proof of Proposition 1 shows that the Samet probability measures are always a balanced pair of priors and Proposition 2 shows that in each single-meet two-player type space, a balanced pair of probability measures is also a pair of Samet measures. The conclusion is that a unique balanced pair is guaranteed always to exist. A common prior can then be characterised as a *self-balancing* pair, meaning a balanced pair (p_1, p_2) such that $p_1 = p_2$.

The balanced pair (p_1, p_2), however, contains more information than just the answer to the question 'does a common prior exist?'. When the priors are disparate, the vector $p_1 - p_2$ serves as a measure of 'how far' the type space is from having a common prior, and under conditions of common knowledge it encodes implications regarding bounds on the range of possible expected values. This is expressed formally in propositions 3, 4 and 5, along with the definition that players 1 and 2 having ε-*separated priors* with respect to a random variable *f* if their balanced pair (p_1, p_2) satisfies $|(p_1 - p_2)f| = \varepsilon$.

In particular, writing player *i*'s expectation of *f* as $E_i f$, proposition 3 shows that under the condition of ε-separated priors, if it is common knowledge at a particular state that $E_1 f = \alpha_1$ and $E_2 f = \alpha_2$, then $|\alpha_1 - \alpha_2| = \varepsilon$ – which is a generalisation of Aumann's agreeing-to-disagree theorem, as that theorem is recapitulated in the limit as $|(p_1 - p_2)|$ approaches zero. Proposition 4 goes further, showing that if the players have ε-separated priors with respect to *f* and it is common knowledge at a particular state that $E_1 f \geq E_2 f$, then it cannot also be the case that it is common knowledge that $E_1 f - E_2 f > \varepsilon$. This is a generalisation of the main ‘no-bet’ result of Sebenius and Geanakoplos (1983) as again that theorem is recaptured in the limit as $|(p_1 - p_2)|$ approaches zero.

In fact, the conclusions of these celebrated theorems can hold true under certain conditions even when $|(p_1 - p_2)| > 0$. What counts is the vector-space geometric relationship of a random variable f with respect to the vector difference $p_1 - p_2$ – if f is perpendicular to $p_1 - p_2$, then $(p_1 - p_2)f = 0$, and the players cannot agree to disagree under common knowledge. What happens when there is a common prior is that each and every vector f is perpendicular to $p_1 - p_2 = 0$. Otherwise, the non-zero projection of f on $p_1 - p_2$ is crucial. From this perspective, the literature on ‘agreeing-to-disagree’ type results stemming from the CPA can be considered the study of the special limit case of type spaces in which $p_1 - p_2 = 0$.

This leads naturally to the question: how ubiquitous is the case of common priors within the general set of type spaces of two players? This is explored in Proposition 6. Somewhat surprisingly at first glance, the answer depends on the cardinalities of the information partitions the players. A type space is defined to be *complementarily-partitioned* if the sum of the cardinalities of the partitions of player 1 and of player 2 equals one plus the cardinality of the state space. Within the space of type spaces that are not complementarily-partitioned, the set of type spaces with common priors is nowhere dense, indicating that common priors should be extremely rare. However, amongst the complementarily-partitioned type spaces, this no longer holds, as examples show, and common priors can be much more common.

The geometric relationship between random variables and the vector $p_1 - p_2$ also has implications regarding how ubiquitous 'mutually profitable bets' are between players. As shown in Proposition 7 and the discussion following it, it turns out that under disparate priors, there is a large cardinality of such bets, which might indicate that the various 'no trade/no betting' theorems that exist fail – spectacularly – to describe reality as we see it, simply because they assume a common prior, when in fact common priors are rare.

So far in this summary we have dealt solely with the two-player case. Many of the results extend to the n-player case, as shown in section 4. The results on balanced priors extend in a particularly elegant way: Samet (1998a) shows that for each permutation of the n players, one can associate an invariant measure, and then derives a simple test: a common prior exists if and only if all the $n!$ measures coincide, in which case the common prior has been identified. In Propositions 8 and 9 in this paper, we show that the full set of $n!$ vectors is not necessary, as it breaks down naturally into $n-1$ n-tuples of priors, where each such n-tuple satisfies the conditions that it is an orbit of a certain group of permutations and each element of the n-tuple is a prior of a unique player. It then turns out to be the case that a common prior exists if and only if the elements of any single such n-tuple completely coincide, in which case the common prior has been identified.

This last point ties into a subject that has received little attention in the literature: the efficient calculation of a common prior, given a type space. Samet's results indicate that by using numerical methods developed for calculating the invariant measures of Markov matrices, it is possible to calculate explicitly a common prior when such exists, but in the n-player case this might require as many as $n!$ separate calculations, a significant calculation burden. The results here show that this burden can be reduced appreciably, by cutting the number of Markov-type invariant measure calculations to one, followed by at most $n-1$ straightforward vector-matrix multiplications.

1.3 Outline of Paper

The broad outline of the paper includes preliminary definitions and results in section 2; the two-player case is explored in section 3, and section 4 is devoted to the n-player case. Proofs appear in the body of the paper, except for a couple of observations and lemma, whose proofs are relegated to the appendix, when it was felt that the full details of the proofs would do more to hinder than help the flow of ideas.

2. Preliminary Definitions and Results

Formally, a *type space* for a set of players $I = \{1,\ldots,n\}$ is a tuple $\langle I, \Omega, (\Pi_i, t_i)_{i\in I}\rangle$, where Ω is a finite set, whose elements are called *states*. Π_i is a partition of Ω for each $i \in I$. Subsets of Ω are *events*. For each $\omega \in \Omega$, $\Pi_i(\omega)$ is the element of the partition Π_i which contains ω. For each *i*, t_i is a function $t_i : \Omega \to \Delta^\Omega$, which associates with each state ω the *type* of *i* at ω, a probability distribution over Ω. The type function t_i for each i must satisfy two conditions: for each $\omega \in \Omega$, $t_i(\omega)(\Pi_i(\omega)) = 1$; and for each $\omega' \in \Pi_i(\omega)$, $t_i(\omega') = t_i(\omega)$.

It will be assumed here, without further comment, that for each *i* and ω, $t_i(\omega)(\{\omega\}) > 0$ – mainly because then the results of Samet (1998a) can be freely adduced. An enquiry along similar lines to that conducted here without this assumption is possible in principle, but doing so would require replacing the notion of common knowledge with that of common 1-belief (as defined in Moderer and Samet (1989)) and the notion of the meet by events *E* which are minimal non-empty events for which *E* is the common 1-belief in *E*.

Given a state space Ω and a set of players *I*, $T(\Omega, I)$ will stand for the set of all possible type spaces $\langle I, \Omega, (\Pi_i, t_i)_{i\in I}\rangle$. When the cardinality of *I* is fixed and understood, $T(\Omega)$ can stand for $T(\Omega, I)$.

In the sequel a fixed type-space $\langle I, \Omega, (\Pi_i, t_i)_{i\in I}\rangle \in T(\Omega, I)$ will often be assumed as given. When varying considerations, such as whether *I* contains two players or *n* players, are relevant, they will be explicitly stated .

The *meet* of $(\Pi_i)_{i\in I}$ is the partition Π of Ω which is the finest among all partitions that are coarser than Π_i for each *i*.

Given a state $\omega \in \Omega$, an event *A* is *common knowledge at* ω if and only if *A* contains the member of the meet Π that contains ω. Equivalently, *A* is common knowledge if *A* is the union of all the elements of Π contained in *A*.

The vector space R^{Ω} will play a prominent role in this research. Given a vector $v \in R^{\Omega}$ and a real number α, $H(v,\alpha)$ will denote the hyperplane defined by $vx = \alpha$, and similarly $H^{+}(v,\alpha)$ denotes the open half-space defined by $vx > \alpha$ and $H^{-}(v,\alpha)$ the open half-space defined by $vx < \alpha$.

Following standard conventions (see for example Gale (1960, page 43)), depending on the context, 0 will sometimes be understood to mean the origin in R^{Ω}, i.e. the vector $(0,0,0,\ldots,0)$. Given a vector $v = (\xi_j)$, $v \geqq 0$ means $\xi_j \geq 0$ for all *j*, $v > 0$ means $\xi_j > 0$ for all *j*, and $v \geq 0$ means $\xi_j \geq 0$ for all *j* but $v \neq 0$. If $v, w \in R^{\Omega}$, $v \geqq w$, $v > w$ and $v \geq w$ respectively stand for $v - w \geqq 0$, $v - w > 0$ and $v - w \geq 0$.

Probability measures on Ω will be considered row vectors in R^{Ω}. Random variables (i.e. real-valued functions on Ω) and the column vectors in R^{Ω} corresponding to them will frequently be used interchangeably here. For a probability measure *p* and random variable *f*, the expectation of *f* with respect to *p* is the vector dot product $pf = \sum_{\omega} p(\omega) f(\omega)$. The special random variable 1_A for an event *A* is the characteristic function getting the value 1 if $\omega \in A$ and 0, and given a probability measure *p*, the probability of event *A* under the measure *p* is the expectation of 1_A with respect to *p*, i.e. $p1_A = \sum_{\omega} p(\omega) 1_A(\omega) = \sum_{\omega \in A} p(\omega) = p(A)$.

For a given type space define for each player *i* the *type matrix* M_i in R^{Ω^2}, by $M_i(\omega,\omega') = t_i(\omega)(\{\omega'\})$, which is a Markov matrix representing t_i as if were a Markovian transition function.

For each random variable *f* on Ω, the expectation of player *i* of that random variable, when viewed as a function of the state, is again a random variable $E_i f$ given by $E_i f(\omega) = \sum_{\omega' \in \Omega} f(\omega') t_i(\omega)(\{\omega'\})$, which can more simply be written as a vector dot-product for each $\omega \in \Omega$: $E_i f(\omega) = t_i(\omega) \cdot f$. Given the definition of M_i, $M_i f = E_i f$, and in this paper the notation $M_i f$ and $E_i f$ will therefore often be used interchangeably to mean the same thing.

The conditional expectation of a random variable *f*, conditional on an event *A*, is also definable as just another random variable in this context, as follows:

$$E_i(f \mid A)(\omega) = \begin{cases} \frac{1}{t_i(\omega)(A)} \sum_{\omega' \in \Omega} 1_A(\omega') f(\omega') t_i(\omega)(\{\omega'\}) & t_i(\omega)(A) \neq 0 \\ 0 & t_i(\omega)(A) = 0 \end{cases}$$

If it is the case that for a given event Y, random variable f and real-number α, $E_i f(\omega) = \alpha$ uniformly for all $\omega \in Y$, then $E_i(f \mid Y) = \alpha$ can be written unambiguously, without the necessity of specifying the particular states. From the definitions, it follows that for each $Y \in \Pi_i$ $E_i(f \mid Y)$ is uniform over all states.

Observation. Given any $k \geq 1$ and type matrix M_i, $M_i^k = M_i$.

The proof of the observation appears in the appendix.

Corollary. For any $k \geq 1$, random variable *f*, and $i \in I$, $E_i^k f = E_i f$. □

Given an event *A* and regarding 1_A as a column vector, $M_i 1_A$ is another column vector which can be regarded as the random variable such that

$$M_i 1_A(\omega) = \sum_{\omega' \in A} t_i(\omega)(\{\omega'\}) = t_i(\omega)(A)$$

Given a type space, one can ask whether the space might have come to exist, in its current state, from a space with no information at all, by the players acquiring new information over time and updating their beliefs in a Bayesian manner. Each player's possible initial belief on the no-information primeval space is called a prior. In general, given player *i*'s current type, there will not be a single prior from which the player could have arrived at the current state from the (hypothetical) primeval past – there will be a set of possible priors. A main question is then whether or not the agents have a *common prior*, meaning a possible initial identical belief that implies the differences in probability assessments currently seen amongst the players can be attributed solely to asymmetric information received over time.

More formally, a *prior* for $i \in I$ at state ω is a probability measure $p \in \Delta^{\Omega}$ such that for each event A $p(A \mid \Pi_i(\omega)) = t_i(\omega)(A)$, whenever the conditional probability measure is defined. A probability measure is a prior for *i*, without the local specification, if it is a prior for *i* at each and every state ω.

Given a particular player *i*, each type of that player, $t_i(\omega)$, is a prior at ω. In fact, the set of all priors for player *i* can be identified with the convex hull of all of *i*'s types (cf. Samet (1998b)).

The vector dot-product equation $p(M_i 1_A) = \sum_{\omega\in\Omega} p(\omega) t_i(\omega)(A)$ establishes that a probability measure *p* on Ω is a prior for *i* iff it is an invariant probability measure for M_i, i.e. $pM_i = p$. A common prior, therefore, is a single *p* such that simultaneously for all players *i*, $pM_i = p$.

For a fixed $Q \in \Pi$ and for each *i*, the restriction of the type-matrix M_i to *Q*, denoted by M_i^Q, is defined by

$$M_i^Q(\omega,\omega') = \begin{cases} t_i(\omega)(\{\omega'\}) & \omega \in Q \\ 0 & \omega \notin Q \end{cases}$$

For any random variable random variable *f* on Ω, $E_i^Q f$ is defined as $M_i^Q f$ regarded as a random variable.

Given the *n* type matrices defined by a type space, for any permutation σ of *I* write

$$M_\sigma = M_{\sigma(1)} \cdots M_{\sigma(n)}$$

and for any random variable *f*,

$$E_\sigma f = E_{\sigma(1)} \cdots E_{\sigma(n)} f$$

The *iterated expectation* of *f* with respect to σ is the sequence $((E_\sigma)^k f)_{k=1}^\infty$.

The definition of M_σ^Q is the obvious one

$$M_\sigma^Q = M_{\sigma(1)}^Q \cdots M_{\sigma(n)}^Q$$

Samet (1998a) includes the following result, which will be crucial for the sequel:

Theorem (Samet). For each $Q \in \Pi$ and $i \in I$, $M_{\sigma_i}^Q$ is ergodic and therefore has a unique invariant probability measure $p_{\sigma_i}^Q$. The ergodicity of this matrix then further implies that the iterated expectation of any random variable f with respect to σ_i, given by $\lim_{k\to\infty}(M_{\sigma_i}^Q)^k f$, converges at every state to $p_{\sigma_i}^Q f$ within each *Q* – in words, the iterated expectations are common knowledge and uniform in each state. On each $Q \in \Pi$, the players have a common prior if and only if for all $i, j \in I$, $p^Q_{\sigma_i} = p^Q_{\sigma_j}$ – hence there exists at most one common prior on Q.

The following theorem, which appears as Lemma 3 in Nehring (2001), will also be needed:

Theorem (Nehring). Define $[f]$, for any random variable f, to be the smallest linear subspace L of R^{Ω} containing f with the property that $E_i g \in L$ whenever $g \in L$, for any player i and random variable g. Then given any finite sequence $(i_1, i_2, ..., i_K)$ of elements in I, with $K \geq 2$, and any random variable f, there exist random variables $\{g_i\}_{i \in I}$ in $[f]$ such that

$$E_{i_K} E_{i_{K-1}} ... E_{i_2}(f - E_1 f) = \sum_{i \in I} E_{i_K}(g_i - E_i g_i)$$

3. Two Players

3.1 Identifying the Priors

Throughout this section, the fixed type space $\langle I, \Omega, (\Pi_i, t_i)_{i \in I} \rangle$ will be assumed to satisfy the constraint that the cardinality of I is equal to 2, and the players will be labelled player 1 and player 2. Because this labelling is arbitrary, some of the results will be worded as applying to player 1 with respect to player 2 in certain symmetrical situations, with the understanding that the symmetry immediately implies that they apply just as well to player 2 with respect to player 1.

When there are two players, there are only two possible permutations of the set of players – which will be labelled here $\sigma_1 = (2,1)$ and $\sigma_2 =$ identity – and hence two permutation matrices $M_{\sigma_1} = M_2 M_1$ and $M_{\sigma_2} = M_1 M_2$. For each $Q \in \Pi$, $M^Q_{\sigma_1}$ and $M^Q_{\sigma_2}$ each have a unique invariant probability measure, respectively $p^Q_{\sigma_1}$ and $p^Q_{\sigma_2}$. We will call $p^Q_{\sigma_1}$ and $p^Q_{\sigma_2}$ the *Samet probability measures* of the type space with respect to Q.

It will be assumed, temporarily, that $\Pi = \{\Omega\}$, so that there is no need to specify $Q \in \Pi$, and we can write p_{σ_i} in place of $p^Q_{\sigma_i}$, etc., easing the notational burden in formulae and proofs. The more general case of multiple elements of the meet will be returned to later.

Note that in what follows there is no assumption that $p_{\sigma_1} = p_{\sigma_2}$ – in other words, we are explicitly permitting the possibility of disparate priors.

Proposition 1. p_{σ_1} is a prior for player 1, and p_{σ_2} is a prior for player 2.

Proof. By definition of invariant probability measure, $p_{\sigma_1} M_{\sigma_1} = p_{\sigma_1} M_2 M_1 = p_{\sigma_1}$. Multiplying on the right by M_2, this leads to $p_{\sigma_1} M_2 M_1 M_2 = p_{\sigma_1} M_2$. Rewriting this as $(p_{\sigma_1} M_2)(M_1 M_2) = p_{\sigma_1} M_2$ makes clear that $p_{\sigma_1} M_2$ is an invariant probability measure of $M_1 M_2 = M_{\sigma_2}$. But as M_{σ_2} is ergodic, it has a unique such invariant probability measure, which we already labelled as p_{σ_2}, leading to the conclusion that $p_{\sigma_1} M_2 = p_{\sigma_2}$.

We can now run the following series of calculations. First multiply on the right by M_2:

$$p_{\sigma_1} M_2 M_2 = p_{\sigma_2} M_2$$

But from an earlier observation, $M_2{}^2 = M_2$, so

$$p_{\sigma_1} M_2 = p_{\sigma_2} M_2$$

We started this chain of calculations with $p_{\sigma_1} M_2 = p_{\sigma_2}$, so we conclude that

$$p_{\sigma_2} = p_{\sigma_2} M_2$$

In other words, the unique Samet probability measure of M_{σ_2}, p_{σ_2}, is also an invariant measure of M_2, hence a prior for player 2. By entirely symmetric considerations, we can just as readily conclude that the unique Samet probability measure of M_{σ_1}, p_{σ_1}, is an invariant measure of M_1, i.e. $p_{\sigma_1} = p_{\sigma_1} M_1$, hence a prior for player 1. ▯

Note: The previous result can also be understood within the context of the derivation of $p_{\sigma_i} f$ through the infinite process $\lim_{k\to\infty} (M_{\sigma_i})^k f$, for an arbitrary random variable *f*, as follows. Set $g = M_2 f$. Consider the column vector which is uniformly $p_{\sigma_2} g$. This is equal to $\lim_{k\to\infty} (M_{\sigma_2})^k g$. But the sequence

$$s = M_{\sigma_2} g, M_{\sigma_2} M_{\sigma_2} g, \ldots$$

can be re-written as

$$s = (M_1 M_2) M_2 f, (M_1 M_2 M_1 M_2) M_2 f, \ldots$$

and again as

$$s = M_1 (M_2 M_2) f, M_1 M_2 M_1 (M_2 M_2) f, \ldots$$

or

$$s = M_1 M_2 f, M_1 M_2 M_1 M_2 f, \ldots$$

so that $\lim_{k\to\infty} s = \lim_{k\to\infty} (M_{\sigma_2})^k g = \lim_{k\to\infty} (M_{\sigma_2})^k f$. But this in turn means that $p_{\sigma_2} g = p_{\sigma_2} f$, or $p_{\sigma_2} M_2 f = p_{\sigma_2} f$. As *f* was selected arbitrarily, it may be concluded that $p_{\sigma_2} M_2 = p_{\sigma_2}$. A symmetric argument yields $p_{\sigma_1} M_1 = p_{\sigma_1}$.

Notation. We will henceforth label the Samet probability measures of the permutation matrices, p_{σ_1} and p_{σ_2}, more simply as p_1 and p_2, for ease of notation. As just shown, p_1 and p_2 are guaranteed to be priors, every type space has a unique pair of such priors, and they satisfy the equations $p_1 M_1 = p_1$, $p_2 M_2 = p_2$, $p_1 M_2 = p_2$ and $p_2 M_1 = p_1$.

3.2 Balanced Pairing

The previous proposition showed that the Samet probability measures of the permutation matrices satisfy certain equations involving the type matrices. These equations, it turns out, characterise these measures, so that we can use them for a definition more intrinsic to the study of type-spaces that avoids the appeal to concepts from Markov chain theory (even though we will still lean on results from that theory for existence and uniqueness) .

Definition. Given a pair of type matrices M_1 and M_2, a pair of probability measures (p_1, p_2) are a *balanced pair* if they satisfy the equations $p_1 M_2 = p_2$ and $p_2 M_1 = p_1$.

Proposition 2. For each single-meet two-player type space, there exists a unique balanced pair of probability measures (p_1, p_2), and this pair is a pair of priors, so that in addition $p_1 M_1 = p_1$ and $p_2 M_2 = p_2$. The existence of a common prior is equivalent to the existence of a self-balanced prior, meaning a balanced pair (p_1, p_2) such that $p_1 = p_2$.

Proof. Suppose that a balanced pair (p_1, p_2) exists. Then $p_2 M_1 M_2 = p_1 M_2 = p_2$. Hence p_2 is an invariant probability measure of the ergodic Markov matrix M_{σ_2}, and therefore unique and equal to one of the Samet probability measures.

The respective conclusion for p_1 is arrived at by entirely symmetric argumentation. Such a balanced pairing must always exist, by a previous proposition, because the Markov matrices M_{σ_1} and M_{σ_2} are guaranteed to have invariant probability measures.

If there is a common prior *p*, then by definition simultaneously $pM_2 = p$ and $pM_1 = p$, and we have trivially identified the unique balanced pair, (p, p), hence p is self-balanced. On the other hand, if the two elements of the balanced pair coincide, a common prior has been identified, simply because the elements of the pair *are* priors. □

One perspective on balanced pairs is the following: start with the set of priors of player 1, i.e. the set $P_1 = \{q_1 \mid q_1 M_1 = q_1\}$, and the corresponding set of priors of player 2, $P_2 = \{q_2 \mid q_2 M_2 = q_2\}$. Define a mapping $\xi : P_1 \to P_2$ by $\xi(q_1) = q_1 M_2$ – to see that this is well-defined on the range, simply note that $(q_1 M_2) M_2 = q_1 M_2$ because M_2 is idempotent. Similarly define $\eta : P_2 \to P_1$ by $\eta(q_2) = q_2 M_1$.

So every prior q_1 of player 1 has a 'ξ-mate' $\xi(q_1)$, and every prior q_2 of player 2 has an 'η-mate' $\eta(q_2)$. The question is: when is a prior the η-mate of its own ξ-mate? To answer this, define $\eta\xi : P_1 \to P_1$ and $\xi\eta : P_2 \to P_2$. The last proposition implies that the mapping $\eta\xi$ has a unique fixed point p_1, so that $p_2 = \xi(p_1)$ satisfies $p_2 M_1 = p_1$. This also identifies p_2 as the unique fixed point of $\xi\eta$, and the balanced pair is then (p_1, p_2).

As M_1 and M_2 are varied, the corresponding balanced pair (p_1, p_2) varies as well. The vector $p_1 - p_2$, as a function of M_1 and M_2, can serve as a rough measure of 'how far' the type space is from having a common prior, given that a common prior exists if and only if $p_1 - p_2 = 0$. Under conditions of common knowledge, the identity of the vector $p_1 - p_2$ has implications regarding bounds on the range of possible expected values, as discussed after considerations of a couple of examples.

3.3 Examples

Example. In this example, $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, $I = \{1,2\}$, $\Pi_1 = \{\{\omega_1, \omega_2, \omega_3\}, \{\omega_4\}\}$, $\Pi_2 = \{\{\omega_1, \omega_2\}\{\omega_3, \omega_4\}\}$, $\Pi = \{\{\omega_1, \omega_2, \omega_3, \omega_4\}\}$.

The type matrices are

$$M_1 = \begin{bmatrix} 1/3 & 1/3 & 1/3 & 0 \\ 1/3 & 1/3 & 1/3 & 0 \\ 1/3 & 1/3 & 1/3 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$M_2 = \begin{bmatrix} 1/2 & 1/2 & 0 & 0 \\ 1/2 & 1/2 & 0 & 0 \\ 0 & 0 & 1/2 & 1/2 \\ 0 & 0 & 1/2 & 1/2 \end{bmatrix}$$

The permutation matrices are

$$M_{\sigma_1} = M_2 M_1 = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 \\ \frac{1}{6} & \frac{1}{6} & \frac{1}{6} & \frac{1}{2} \\ \frac{1}{6} & \frac{1}{6} & \frac{1}{6} & \frac{1}{2} \end{bmatrix}$$

$$M_{\sigma_2} = M_1 M_2 = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{6} & \frac{1}{6} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{6} & \frac{1}{6} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{6} & \frac{1}{6} \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix}$$

The balanced priors are the self-balanced $p := p_1 = p_2 = [0.25,\ 0.25,\ 0.25,\ 0.25]$, hence this vector is also the unique common prior in this example. A quick calculation indicates that indeed $pM_1 = p$ and $pM_2 = p$.

Example. In this example, $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, $I = \{1,2\}$, $\Pi_1 = \{\{\omega_1, \omega_2\}\{\omega_3, \omega_4\}\}$, $\Pi_2 = \{\{\omega_1, \omega_3\}\{\omega_2, \omega_4\}\}$, $\Pi = \{\{\omega_1, \omega_2, \omega_3, \omega_4\}\}$.

The type matrices are

$$M_1 = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 \\ \frac{1}{2} & \frac{1}{2} & 0 & 0 \\ 0 & 0 & \frac{3}{4} & \frac{1}{4} \\ 0 & 0 & \frac{3}{4} & \frac{1}{4} \end{bmatrix}$$

$$M_2 = \begin{bmatrix} \frac{1}{3} & 0 & \frac{2}{3} & 0 \\ 0 & \frac{2}{3} & 0 & \frac{1}{3} \\ \frac{1}{3} & 0 & \frac{2}{3} & 0 \\ 0 & \frac{2}{3} & 0 & \frac{1}{3} \end{bmatrix}$$

The permutation matrices are

$$M_{\sigma_1} = M_2 M_1 = \begin{bmatrix} 1/6 & 1/6 & 1/2 & 1/6 \\ 1/3 & 1/3 & 1/4 & 1/12 \\ 1/6 & 1/6 & 1/2 & 1/6 \\ 1/3 & 1/3 & 1/4 & 1/12 \end{bmatrix}$$

$$M_{\sigma_1} = M_2 M_1 = \begin{bmatrix} 1/6 & 1/3 & 1/3 & 1/6 \\ 1/6 & 1/3 & 1/3 & 1/6 \\ 1/4 & 1/6 & 1/2 & 1/12 \\ 1/4 & 1/6 & 1/2 & 1/12 \end{bmatrix}$$

The balanced priors are

$$p_1 = \left[\frac{1125}{4950}, \frac{1125}{4950}, \frac{243}{495}, \frac{135}{990}\right]$$

$$p_2 = \left[\frac{7}{33}, \frac{8}{33}, \frac{14}{33}, \frac{4}{33}\right]$$

Again, straight-forward vector calculations show that the equations $p_1 M_1 = p_1$, $p_2 M_2 = p_2$, $p_1 M_2 = p_2$ and $p_2 M_1 = p_1$ are satisfied, as expected.

3.4 Common Knowledge and Disparate Priors

We can now relax the assumption that $\Pi = \{\Omega\}$, and consider the general case in which the meet contains several elements. There are now (at most) two balanced priors on each element Q of the meet, p_1^Q and p_2^Q. We can also trivially assign balanced priors to each state, in the sense that given a state $\omega \in \Omega$ the associated balanced priors can be defined by $p_i^{\omega} := p_i^{Q(\omega)}$, where $Q(\omega)$ is the element of the meet containing ω.

Proposition 3. In a 2-player type space, given a random variable f and a state ω^*, if $\left|(p_i^{\omega^*} - p_j^{\omega^*})f\right| = \varepsilon$, then if it is common knowledge at ω^* that player 1's expectation of f is α_1 and player 2's expectation of f is α_2, then $|\alpha_1 - \alpha_2| = \varepsilon$.

Proof. As defined above, $E_i f(\omega) = t_i(\omega) f$. Under the assumption of mutual common knowledge, the valuations $E_1 f(\omega) = \alpha_1$ and $E_2 f(\omega) = \alpha_2$ hold uniformly for all $\omega \in \Pi(\omega^*)$.

Consider the expression $p_i^{\omega^*}(E_i f - E_j f)$. As it was assumed that $E_1 f(\omega) = \alpha_1$ and $E_2 f(\omega) = \alpha_2$ uniformly, the vector $|E_i f - E_j f|$ is uniformly $|\alpha_1 - \alpha_2|$. By definition, every prior is a probability measure, and hence $\sum_\omega p_i^{\omega^*}(\omega) = 1$, so that $\left|p_i^{\omega^*}(E_i f - E_j f)\right|$ $= |\alpha_1 - \alpha_2| \sum_\omega p_i^{\omega^*}(\omega) = |\alpha_1 - \alpha_2|$.

But $p_i^{\omega^*}(E_i f - E_j f)$ is equal to $p_i^{\omega^*} M_i f - p_i^{\omega^*} M_j f$, which by above results is given by $p_i^{\omega^*} f - p_j^{\omega^*} f = (p_i^{\omega^*} - p_j^{\omega^*}) f$. □

This result can be understood from elementary considerations of the balanced priors as Samet invariant probability measures on the permutation matrices, as calculated in the infinite limit of iterated expectations. If it is common knowledge that player 1's expectation of f is α_1, then the infinite sequence $M_1 f, M_2 M_1 f, M_1 M_2 M_1 f, \ldots$ is constantly uniformly α_1, hence trivially is uniformly α_1 in the limit. A similar statement holds if player 2's expectation of f is α_2, with respect to the sequence $M_2 f, M_1 M_2 f, M_2 M_1 M_2 f, \ldots$. It is therefore not surprising that under full common knowledge, $(p_1^{\omega^*} - p_2^{\omega^*}) f$ turns out to be the same as $\alpha_1 - \alpha_2$.

Note that this implies that the possible spread of expected values under common knowledge depends on the vector geometry of the random variable *f* with respect to the vector $p_1^{\omega^*} - p_2^{\omega^*}$. If *f* is perpendicular to $p_1^{\omega^*} - p_2^{\omega^*}$, then $(p_1^{\omega^*} - p_2^{\omega^*}) f = 0$, and the players cannot agree to disagree under common knowledge – which is exactly what happens when there is a common prior, because then $p_1^{\omega^*} - p_2^{\omega^*} = 0$ and each and every vector is perpendicular to 0. In other cases, the non-zero projection of *f* on $p_1^{\omega^*} - p_2^{\omega^*}$ is crucial. (Note that the vector $p_1^{\omega^*} - p_2^{\omega^*}$ itself has constraints on its possible values: $\sum_{\omega \in \Omega} p_1^{\omega^*}(\omega) - p_2^{\omega^*}(\omega) = 0$, because $\sum_{\omega \in \Omega} p_1^{\omega^*}(\omega) = 1$ and $\sum_{\omega \in \Omega} p_2^{\omega^*}(\omega) = 1$.)

Proposition 3 thus naturally motivates the following definition:

Definition: In a 2-player type space with balanced priors p_1^{ω} and p_2^{ω} at a state ω, for each random variable f, players 1 and 2 will be termed to have ε*-separated priors with respect to f at* ω if $\left|(p_i^{\omega} - p_j^{\omega}) f\right| = \varepsilon$.

As a corollary of the proof of the last proposition we can get a yet stronger result, attained by weakening the insistence that the common-knowledge expectations of the two players be given by precise values α_1 and α_2, and assuming only common knowledge of the fact that one player has greater expectations than the other player.

Proposition 4. In a 2-player type space, if the players have ε-separated priors with respect to random variable f at state ω^*, and if it is common knowledge at ω^* that $E_1 f \geq E_2 f$, then it cannot also be the case that it is common knowledge that $E_1 f - E_2 f > \varepsilon$. Similarly, it cannot be the case that it is common knowledge that $E_1 f - E_2 f < \varepsilon$. Thus, either $E_1 f(\omega) - E_2 f(\omega) = \varepsilon$ for all $\omega \in \Pi(\omega^*)$, or there is at least one $\omega_1 \in \Pi(\omega^*)$ with $E_1 f(\omega_1) - E_2 f(\omega_1) > \varepsilon$ and at least one $\omega_2 \in \Pi(\omega^*)$ such that $E_1 f(\omega_2) - E_2 f(\omega_2) < \varepsilon$.

Proof. If at ω^* it is common knowledge that $E_1 f \geq E_2 f$ and $E_1 f - E_2 f > \varepsilon$, then $p_1^{\omega^*}(E_1 f - E_2 f) = \sum_\omega p_1^{\omega^*}(\omega)[E_1 f(\omega) - E_2 f(\omega)] > \varepsilon$, because the latter expression is a weighted average of elements, each of which is strictly greater than ε. But as in the above proof, $p_1^{\omega^*}(E_1 f - E_2 f)$ – which must be greater than or equal to zero by the assumption of common knowledge that $E_1 f \geq E_2 f$ – is by previous results equal to $p_1^{\omega^*} M_1 f - p_1^{\omega^*} M_2 f$, which is $p_1^{\omega^*} f - p_2^{\omega^*} f = (p_1^{\omega^*} - p_2^{\omega^*}) f = \varepsilon$, the last equality following from the assumption of ε-separated priors with respect to *f*. This is a contradiction.

Similarly, if it is common knowledge that $E_1 f \geq E_2 f$ and $E_1 f - E_2 f < \varepsilon$, we derive a contradiction to the assumption that $(p_1^{\omega^*} - p_2^{\omega^*}) f = \varepsilon$. □

It can readily be seen that this proposition implies the previous one – if it is common knowledge that player 1's expectation of *f* is uniformly α_1 and player 2's expectation of *f* is uniformly α_2, where $\alpha_1 > \alpha_2$ without loss of generality, then the assumption of ε-separation with respect to *f* implies that $\alpha_1 - \alpha_2$ can be neither less than or greater than ε – hence it is precisely ε.

In the special case that *f* is perpendicular to $(p_1^{\omega^*} - p_2^{\omega^*})$, the proposition states that if it is common knowledge that $E_1 f \geq E_2 f$, then it cannot also be the case that it is common knowledge that $E_1 f - E_2 f > 0$, hence $E_1 f = E_2 f$ and in this case there can be no agreement on disagreement.

We can also consider a case intermediate between the two previous propositions, in which the expectation of only one player is by common knowledge uniformly a precise value, and ask what implications that has on the values of the expectation of the other player.

Proposition 5. In a 2-player type space, if the players have ε-separated priors with respect to random variable f at state ω^*, and if $p_1^{\omega^*} E_i f = \alpha$ for one of the indexes i and some number α, then it must be the case that for $j \neq i$, $E_j f|_{\Pi(\omega^*)}$, as a vector, is located either in the hyperplane $H(p_1^{\omega^*}, \alpha + \varepsilon)$ or the hyperplane $H(p_1^{\omega^*}, \alpha - \varepsilon)$. In particular, if it is common knowledge at ω^* that $E_i f = \beta$ for some number β, and $E_j f \geq E_i f$, then $E_j f \in H(p_1^{\omega^*}, p_1^{\omega^*} \beta + \varepsilon)$.

Proof. Again, we work with $p_1^{\omega^*}(E_j f - E_i f)$. By assumption of ε-separated priors, this is equal to $\pm\varepsilon$. But it is also equal to $p_1^{\omega^*} E_j f - p_1^{\omega^*} E_i f$. The assumption that $p_1^{\omega^*} E_i f = \alpha$ then implies that $p_1^{\omega^*} E_j f = \alpha \pm \varepsilon$, which is the same as saying $E_j f|_{\Pi(\omega^*)}$ is located either in $H(p_1^{\omega^*}, \alpha + \varepsilon)$ or $H(p_1^{\omega^*}, \alpha - \varepsilon)$. The special case of common knowledge that $E_i f = \beta$ implies that $p_1^{\omega^*} E_i f = p_1^{\omega^*} \beta$, and hence $E_j f \geq E_i f$ implies $E_j f \in H(p_1^{\omega^*} \beta + \varepsilon)$. □

3.5 Common Priors as the Limiting Case of Disparate Priors

With respect to the above propositions, the sharpest results are obtained in the special case in which $p_1^{\omega^*} = p_2^{\omega^*}$:

- If $p_1^{\omega^*} = p_2^{\omega^*}$ and $f = 1_A$, $(p_i^{\omega^*} - p_j^{\omega^*})1_A = 0$ for all A, hence by Proposition 3 under common knowledge the players can never agree to disagree on the probability of the occurrence of an event, and we recapitulate the theorem of Aumann (1976).

- If $p_1^{\omega^*} = p_2^{\omega^*}$ and it is common knowledge at ω^* that $E_1 f \geq \alpha$ and $E_2 f \leq \alpha$, then since $(p_i^{\omega^*} - p_j^{\omega^*}) f = 0$, Proposition 4 implies that $E_1 f - E_2 f = 0$, hence $E_1 f = E_2 f = \alpha$, recapitulating the main 'no-bet' result of Sebenius and Geanakoplos (1983).

But the condition $p_1^{\omega^*} = p_2^{\omega^*}$ is equivalent to the existence of a common prior over $\Pi(\omega^*)$, hence the propositions may be considered generalisations of these well-known CPA agreeing-to-disagree results.

As stated in the introduction to this paper, the CPA has often been criticised in the past, especially when the CPA leans on a supposed 'dynamic story' – the view that players assessing differing expectations of events do so solely because of differences in the private information they possess respectively, because in some hypothetical past they shared a common prior, with their current beliefs posterior to a, perhaps distant, past of shared probabilities. Gul (1998) argues that 'since there never was a prior stage, the prior distribution is meaningless'.

Aumann (1998)[5], in reply, essentially restates the position of Aumann (1987), which includes the assertion that 'people with different information may legitimately entertain different probabilities, but there is no rational basis for people who have always been fed precisely the same information to do so'. In the zeal to highlight 'differences in information' as the sole bearer of distinction, Aumann (1998) postulates that 'if … the beliefs at an "actual" prior stage are different and *not* commonly known, then there must be differential information already at that stage', and then argues that analysis must proceed to a further earlier stage until all differences in information have been purged and a common primeval prior can be identified. He is even willing to go so far as to say 'if one sets forth all relevant information in sufficient detail, then in principle, there should be no room for differing probabilities. When we say all relevant information, we mean *all*: the schools the players attended, their childhood experiences, even their genes (which indirectly reflect the experience of previous generations).'

The results of this paper shed further light on matters at the heart of the Gul-Aumann debate. The players – or any observers for that matter – need no more information than knowledge of the type space itself – i.e. the tuple $\langle \Omega, (\Pi_1, t_1), (\Pi_2, t_2) \rangle$ – in the present, in order to identify all the elements of the meet $\{Q \mid Q \in \Pi\}$ and the corresponding set of type matrices $\{M_i^Q \mid i \in I, Q \in \Pi\}$. Simple matrix multiplication then yields the permutation matrices $M_{\sigma_1}^Q$ and $M_{\sigma_2}^Q$.

Obtaining the balanced priors $\{p_i^Q \mid i \in I, Q \in \Pi\}$ is then a matter of calculating the invariant probability measures of the permutation matrices. Numerical methods for doing so, either in some cases using direct methods with exact results, or in others using iterative methods converging up to a 'reasonable' tolerance, are the subject of active research (see for example Stewart (1994)). (The computational burden becomes even lighter when one considers the fact that it necessary at each Q to calculate only one of the Samet probability measures – say, p_1^Q – and then the other can be obtained by the computationally simpler method of direct matrix multiplication, given that $p_2 = p_1 M_2$.)

In principle, therefore, there is a computationally efficient and well-defined algorithmic procedure for going from the type space to the full set of balanced priors. With these latter to hand, the analysis locally at any state $\omega \in \Omega$ can proceed in one of two ways.

[5] Aumann (1998) also includes a formalisation of an argument in favour of the CPA that runs essentially along the following lines: Beliefs are based on information. If all information is removed, all that is left is an empty shell. Since there is no reason to distinguish between empty shells, individuals must start with common priors. Bernheim (1986) terms such an argument for common priors 'assuming the conclusion'.

i) If $p_1^{\omega} = p_2^{\omega}$ there is a common prior. One may in this case embrace the Harsányi doctrine and assert that 'differences in probabilities solely express differences in information'. A (possibly fictional) historical account of prior stages, in which differences in information led to disparate information partition refinements and differential probability assessments based on Bayesian updating against the common prior, may be adduced. Even if one chooses not to resort to conjuring the past, the existence of a common prior justifies quoting any of the large number of agreeing-to-disagree type results under common knowledge that have been proved since the seminal work of Aumann (1976).

ii) If $p_1^{\omega} \neq p_2^{\omega}$, there can be no common prior. One may again suppose a (possibly fictional) historical account of prior stages but in this case, in a kind of reversal of the Harsányi doctrine, in the primeval past the players begin unequal, with a fundamental disagreement regarding the 'true' prior, one player believing p_1^{ω} and the other p_2^{ω}. As asymmetric information is obtained over time by the players, their information partitions diverge, along with their respective probability assessments under Bayesian updating from their different priors, so that both differential information *and* subjective probability differences contribute to the divergences. Even if one chooses not to resort to conjuring the past, in this case the players are fully justified in agreeing-to-disagree. Consideration of $(p_i^{\omega} - p_j^{\omega})f$, for any random variable, indicates how far apart the players can be when agreeing-to-disagree under common knowledge with respect to *f*, as proved in the above propositions.

In summary, it is the set of vectors $\{p_i^Q - p_j^Q \mid Q \in \Pi\}$ that contains the information for the above-derived bounds on expected values under conditions of common knowledge. Since the values of $\{p_i^Q - p_j^Q \mid Q \in \Pi\}$ can be derived from the type space, it follows that one needs no more than knowledge of the tuple $\langle \Omega, (\Pi_1, t_1), (\Pi_2, t_2) \rangle$ for these results. From this perspective, the entire corpus of literature on agreeing-to-disagree type results, such as 'no-bet', 'no-trade', etc., stemming from the CPA is the study of the special 'limit case' of a particular subset of the set of all type spaces $\langle \Omega, (\Pi_1, t_1), (\Pi_2, t_2) \rangle$ – namely, the set of type spaces from which it can be deduced that the vectors $\{p_i^Q - p_j^Q \mid Q \in \Pi\}$ are uniformly zero.

3.6 The Rareness and Ubiquity of Common Priors

The next obvious question is: how 'special' is a common prior situation within the space of all type spaces? If one were to select a random sampling of type spaces, should one expect common priors to be ubiquitous or rare?

Fixing the state space Ω, let $T(\Omega)$ denote the set of all type spaces of two players over Ω sharing a single-element meet. (The loss of generality for the sake of simplicity is tolerable as extending the proofs in this section to the case of multiple-element meets is straight-forward). Because I and Ω are fixed here, an element $\tau \in T(\Omega)$ is completely determined by its associated partitions, which we can label $\Pi_{1,\tau}$ and $\Pi_{2,\tau}$, along with the associated type functions, $t_{1,\tau}$ and $t_{2,\tau}$.

Next, let $\alpha_1 : T(\Omega) \to M_1(\Omega)$ be the mapping of $T(\Omega)$ to $M_1(\Omega)$, the set of all type matrices of player 1, with $\alpha_2 : T(\Omega) \to M_2(\Omega)$ playing the same role for player 2. Let $\alpha : T(\Omega) \to M_1(\Omega) \times M_2(\Omega)$ stand for the bijective mapping taking each element $\tau \in T(\Omega)$ to $(\alpha_1(\tau), \alpha_2(\tau))$. The space $M_1(\Omega) \times M_2(\Omega)$ can be considered a sub-space of $R^{|\Omega|^2} \times R^{|\Omega|^2}$, with the latter endowed with the standard topology of vector spaces over the reals, hence the space $M_1(\Omega) \times M_2(\Omega)$ naturally inherits a sub-space topology. As α is a bijection, we can give $T(\Omega)$ the topology that makes α a topological isomorphism.

Let $\beta : M_1(\Omega) \times M_2(\Omega) \to \Delta^{\Omega} \times \Delta^{\Omega} \subset R^{\Omega} \times R^{\Omega}$ be the mapping that takes each element of the space $M_1(\Omega) \times M_2(\Omega)$ to the unique corresponding pair of balanced priors. Let the mapping γ be further defined by $\gamma : (p_1, p_2) \mapsto p_1 - p_2 \in R^{\Omega}$ for each pair of balanced priors (p_1, p_2).

Lemma. The mapping $\xi = \gamma \circ \beta \circ \alpha$ is continuous.

The proof of this lemma appears in the appendix.

Definition. A type space τ in $T(\Omega)$ will be said to be *complementarily-partitioned* if the cardinalities of its associated partitions, $\Pi_{1,\tau}$ and $\Pi_{2,\tau}$, satisfy $|\Pi_{1,\tau}| + |\Pi_{1,\tau}| = |\Omega| + 1$.

Notation. The set of non-complementarily-partitioned type spaces within $T(\Omega)$ will be labelled $\overline{T}(\Omega)$. Denote further the sub-space of $\overline{T}(\Omega)$ consisting of type spaces that share a common-prior between them by $\overline{C}(\Omega)$.

Proposition 6. $\overline{C}(\Omega)$ is nowhere dense in $\overline{T}(\Omega)$.

Proof. This is proved in two steps.

1. $\overline{C}(\Omega)$ has empty interior: Let τ be an arbitrary element in $\overline{C}(\Omega)$, with common prior p. Let $m = |\Pi_{1,\tau}|$ and $n = |\Pi_{2,\tau}|$. Using the earlier defined mapping α, $\alpha(\tau) = (M_1, M_2)$, a pair of type matrices. Associated with M_1 is $p_1^1, p_1^2, ..., p_1^m \in \Delta^{\Omega}$, where each p_1^j is a distinct row in M_1 corresponding to one of the partitions in $\Pi_{1,\tau}$, and similarly associated with M_2 is $p_2^1, p_2^2, ..., p_2^n \in \Delta^{\Omega}$. We can now form the convex hulls $X(p_1^1, p_1^2, ..., p_1^m)$ and

$X(p_2^1, p_2^2, ..., p_2^n)$, of dimensions $m-1$ and $n-1$, respectively, such that p is the unique point of intersection of these two convex polytopes, which are constrained to be within Δ^{Ω}, a polytope of dimension $|\Omega|-1$.

The assumption that $\tau \in \overline{T}(\Omega)$, and therefore not complementarily-partitioned, means that $m-1+n-1<|\Omega|$. This, plus the fact that the convex hulls intersect solely at a single point, implies that an arbitrarily small deformation of one or the other can pull them apart from the point of intersection – in the terminology of differential topology, all such intersections are non-transversal (cf. Guillemin and Pollack (1974)).

In the specific context here, this translates into the possibility of finding a type matrix M_1' with associated $p_1'^1, p_1'^2, ..., p_1'^m$ such that M_1' is within an arbitrarily small ε-ball of M_1, and such that the convex hull $X(p_1'^1, p_1'^2, ..., p_1'^m)$ has no intersection with $X(p_2^1, p_2^2, ..., p_2^n)$. Then $\tau' = \alpha^{-1}(M_1', M_2)$ is a type-space within ε of τ, and $\tau' \notin \overline{C}(\Omega)$.

2. $\overline{C}(\Omega)$ is closed: $\overline{C}(\Omega)$ can be defined as $\xi^{-1}(0)|_{\overline{T}(\Omega)}$, where $\xi = \gamma \circ \beta \circ \alpha$ is continuous, by the above lemma. It is therefore a closed set.

This suffices to show that $\overline{C}(\Omega)$ is nowhere dense in $\overline{T}(\Omega)$. □

The restriction to non-complementarily-partitioned type spaces is necessary for the above proof to work. The heart of the proof is essentially the claim that given a pair of type matrices M_1, M_2 that share a common prior, one of them can be 'perturbed' by an arbitrarily small ε into another type matrix such that the new pair does not have a common prior. But it is a theorem of differential topology that transverse sub-manifolds of complementary dimension intersect in 0-manifolds – i.e. isolated points. The following examples illustrate the implications this has for the question of the ubiquity of common priors.

Example. Let $\Omega = \{1,2,3\}$, and let the type space τ be defined by $\Pi_{1,\tau} = \{1,2\},\{3\}$, $\Pi_{2,\tau} = \{1\},\{2,3\}$, with $t_{1,\tau}(1) = 1/2$, $t_{1,\tau}(2) = 1/2$, $t_{1,\tau}(3) = 1$ and $t_{2,\tau}(1) = 1$, $t_{2,\tau}(2) = 1/3$, $t_{2,\tau}(3) = 2/3$. The corresponding type matrices are

$$M_1 = \begin{bmatrix} 1/2 & 1/2 & 0 \\ 1/2 & 1/2 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

$$M_2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{3} & \frac{2}{3} \\ 0 & \frac{1}{3} & \frac{2}{3} \end{bmatrix}$$

which share the common prior (1/4,1/4,1/2).

But there is no perturbation of the matrices which will lead to a situation of disparate priors. In fact, given any arbitrary type functions t_1 and t_2, such that

$$M_1 = \begin{bmatrix} t_1(1) & t_1(2) & 0 \\ t_1(1) & t_1(2) & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

$$M_2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & t_2(2) & t_3(3) \\ 0 & t_2(2) & t_3(3) \end{bmatrix}$$

are type matrices, M_1 and M_2 will have a common prior between them. Geometrically, this is an example of transversally intersecting one-dimensional lines in a two-dimensional space.

Example. Let $\Omega = \{1,2,3,4,5\}$, and let the type space τ be defined by $\Pi_{1,\tau} = \{1,2\},\{3,4,5\}$, $\Pi_{2,\tau} = \{1\},\{4\},\{5\},\{2,3\}$, with $t_{1,\tau}(1) = 1/2$, $t_{1,\tau}(2) = 1/2$, $t_{1,\tau}(3) = 1/2$, $t_{1,\tau}(4) = 1/4$, $t_{1,\tau}(5) = 1/4$ and $t_{2,\tau}(1) = 1$, $t_{2,\tau}(4) = 1$, $t_{2,\tau}(5) = 1$, $t_{2,\tau}(2) = 1/4$, $t_{2,\tau}(3) = 3/4$. The corresponding type matrices

$$M_1 = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 \\ \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} \\ 0 & 0 & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} \\ 0 & 0 & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} \end{bmatrix}$$

$$M_2 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{4} & \frac{3}{4} & 0 & 0 \\ 0 & \frac{1}{4} & \frac{3}{4} & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

share the common prior (0.125, 0.125, 0.375, 0.1875, 0.1875).

There is no perturbation of these matrices that will lead to a situation of disparate priors. Geometrically, this is an example of a 1-dimensional polytope transversally intersecting a 3-dimensional polytope inside a 4-dimensional space.

It is easy to conjure up examples of such complementarily-partitioned type-spaces in any dimension. A trivial but instructive example in n-dimensions is

$$M_1 = \begin{bmatrix} \frac{1}{n} & \frac{1}{n} & \cdots & \frac{1}{n} \\ \frac{1}{n} & \frac{1}{n} & \cdots & \frac{1}{n} \\ . & & & \\ . & & & \\ . & & & \\ \frac{1}{n} & \frac{1}{n} & \cdots & \frac{1}{n} \end{bmatrix}$$

$$M_2 = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ . & & & \\ . & & & \\ . & & & \\ 0 & 0 & \cdots & 1 \end{bmatrix}$$

Here, M_1 is associated with a single point, whilst M_2 is associated with the entire $n-1$ dimensional polytope Δ^n. An intersection – meaning a common prior – is inevitable in such a situation.

3.7 Betting and Disparate Priors

As previously noted, the main characterisation of common priors in the literature is the Morris-Feinberg theorem that states there is no common prior in a type space if and only if there is at least one random variable t with respect to which the players can agree to take opposite sides of a bet.

More formally, the players *fail to bet* on t at some state ω if it is not common knowledge at ω that player 1's expectation of t is positive, $E_1 t > 0$, and player 2's expectation of $-t$ is also positive, $E_2(-t) > 0$. If the players do not fail to bet on t, then t is a *mutually acceptable bet*. There is *no betting* amongst the players at ω if they fail to bet on any random variable t. Sebenius and Geanakoplos (1983) established that if there is a common prior, the players will never bet under common knowledge. Morris (1995) and Feinberg (2000), independently, proved the converse – if there is no common prior, the players can always identify a random variable t which is mutually acceptable.

Given the amount of information that the balanced priors p_1 and p_2 of a type space bear with respect to common knowledge and common priors, it is natural, in light of the Morris-Feinberg theorem, to enquire whether knowledge of the balanced priors can also provide information about mutually acceptable bets between the players. In order to study this matter, we lean on Nehring's Theorem, mentioned above, which in its two-dimensional version reduces to the following: recalling the definition of the permutation matrices $M_{\sigma_1} = M_2 M_1$ and $M_{\sigma_2} = M_1 M_2$, given M_{σ_1}, any integer j and any random variable f on Ω, there exist random variables g_1 and g_2 in the linear space $[f]$ such that $(M_{\sigma_2})^j (f - M_1 f) = M_1(g_1 - M_1 g_1) + M_1(g_2 - M_2 g_2)$, and a similar statement, *mutatis mutandis*, holds for M_{σ_2}.

Proposition 7. In a 2-player type space, given a random variable f and a state ω^*, if $(p_i^{\omega^*} - p_j^{\omega^*}) f \neq 0$, there exists a mutually acceptable bet λ such that $\lambda \in [f]$.

Proof. Suppose without loss of generality that $(p_2^{\omega^*} - p_1^{\omega^*}) f > 0$ and $(p_1^{\omega^*} - p_2^{\omega^*}) f < 0$. Consider the expression $(M_{\sigma_2})^j (f - M_1 f)$. Letting the integer j grow without bound, $\lim_{j\to\infty} (M_{\sigma_2})^j (f - M_1 f)$ is equal to a vector whose elements are uniformly equal to $p_2^{\omega^*}(f - M_1 f)$. But $p_2^{\omega^*}(f - M_1 f) = p_2^{\omega^*} f - p_1^{\omega^*} f$, and it was already assumed that this last expression is greater than 0 – in other words, $\lim_{j\to\infty} (M_{\sigma_2})^j (f - M_1 f) > 0$, as a vector inequality.

This in turn implies, because the state space Ω is finite, that for some finite k, $(M_{\sigma_2})^k (f - M_1 f)$ is uniformly greater than 0. By Nehring's Theorem, then, there exist random variables g_1 and g_2 in $[f]$ such that

$$(M_{\sigma_2})^k (f - M_1 f) = M_1(g_1 - M_1 g_1) + M_1(g_2 - M_2 g_2) > 0$$

But for either i, $M_i(g_i - M_i g_i) = M_i g_i - M_i g_i = 0$, hence $M_1(g_2 - M_2 g_2) > 0$. Therefore, by setting $\lambda_1 = (g_2 - M_2 g_2)$, $M_1 \lambda_1 > 0$ and $M_2 \lambda_1 = 0$.

Next, since $(p_1^{\omega^*} - p_2^{\omega^*})f < 0$, $(p_1^{\omega^*} - p_2^{\omega^*})(-f) > 0$, and by reasoning similar to the above we can arrive at random variables h_1 and h_2 in $[f]$ such that for some finite n,

$$(M_{\sigma_1})^n (f - M_2 f) = M_2(h_1 - M_2 h_1) + M_2(h_2 - M_1 h_2) > 0$$

from which it follows that, setting $\lambda_2 = (h_2 - M_1 h_2)$, $M_2 \lambda_2 > 0$ and $M_1 \lambda_2 = 0$.

Finally, setting $\lambda = \lambda_1 - \lambda_2$, λ by construction satisfies the condition that player 1's expectation of it is positive, $E_1 \lambda > 0$ while player 2's expectation of $-\lambda$ is also positive, $E_2(-\lambda) > 0$. □

This last proposition has at least two interesting implications.

Firstly, imagine two individuals who insist on finding a bet they can conduct between themselves. They can write down their type matrices and calculate their balanced priors. If the balanced priors are equal, they have a common prior and by the Sebenius-Geanakoplos theorem they can stop right there – they will not be able to find any random variable on which to bet. If the priors are disparate, at state ω^*, all they need do is identify a function f such that $(p_i^{\omega^*} - p_j^{\omega^*})f \neq 0$, and then follow the iterative steps appearing in Nehring (2001) and the procedure in the above proof to calculate (admittedly not necessarily in a computationally efficient manner) a mutually acceptable bet λ.

Secondly, from this we see that in situations of lack of common priors, the Morris-Feinberg result holds rather strongly – there are a very large cardinality of mutually acceptable bets. To be more precise, start with the observation that the players needn't work terribly hard to find a function f such that $(p_i^{\omega^*} - p_j^{\omega^*})f \neq 0$. Denote by Z the set of random variables g such that $(p_i^{\omega^*} - p_j^{\omega^*})g = 0$. As $p_i^{\omega^*} - p_j^{\omega^*} \neq 0$, Z is a hyper-plane in R^{Ω} – and so is a set of dimension less than $|\Omega|$ and therefore of Lebesgue measure zero in R^{Ω}. In other words, chances are that by selecting a random f in R^{Ω}, the players can apply the above procedure to find a mutually acceptable bet. Even if not, suppose they have selected an arbitrary $g \in Z$. For each $0 \leq \alpha < 1$, the vector given by $h = \alpha g + (1-\alpha)(p_i^{\omega^*} - p_j^{\omega^*})$ satisfies $(p_i^{\omega^*} - p_j^{\omega^*})h \neq 0$ – and now again the procedure can be followed to find a mutually acceptable bet.

Given this, it might not be surprising to discover a great deal of bets being concluded under conditions of disparate priors.

4. *N* Players

4.1 Balanced Priors

In this section, the cardinality of *I*, the set of players in the type space $\langle I, \Omega, (\Pi_i, t_i)_{i \in I} \rangle$, will be any finite number *n*. The number of type matrices is obviously also *n*, labelled M_1 through M_n.

With *n* players, there are $n!$ different permutations. For each permutation σ, there is an associated permutation matrix $M_\sigma = M_{\sigma(1)} M_{\sigma(2)} \ldots M_{\sigma(n)}$. Given $Q \in \Pi$, each M_σ^Q is ergodic and has a unique Samet invariant probability measure labelled p_σ^Q. Permutations will be assumed to operate on these invariant probability measures by way of $\mu p_\sigma^Q = p_{\mu\sigma}^Q$ for any permutation μ.

For notational simplicity, we will again assume temporarily that $\Pi = \{\Omega\}$.

Notation: Certain subsets of Σ_n, the set of all permutations on *n* objects, will be of special interest. For each $j \in I$, the set of all permutations σ such that $\sigma(n) = j$ will be labelled $\Sigma_{\rightarrow \cdots j}$, and a typical element in it will be written $\sigma_{\rightarrow \cdots j} \in \Sigma_{\rightarrow \cdots j}$. Similarly, the set of all permutations σ such that $\sigma(1) = j$ will be labelled $\Sigma_{j \cdots \rightarrow}$, and a typical element in it will be written $\sigma_{j \cdots \rightarrow} \in \Sigma_{j \cdots \rightarrow}$. The set of invariant probability measures $\{p_{\sigma_{\rightarrow \cdots j}}\}$ associated with the permutation matrices $\{M_{\sigma_{\rightarrow \cdots j}}\}$ such that $\sigma_{\rightarrow \cdots j} \in \Sigma_{\rightarrow \cdots j}$ will be denoted $\Psi_{\rightarrow \cdots j}$, and $\Psi_{j \cdots \rightarrow}$ denotes the obvious equivalent for elements of $\Sigma_{j \cdots \rightarrow}$. Note the following cardinalities: $|\Sigma_{\rightarrow \cdots j}| = |\Sigma_{j \cdots \rightarrow}| = (n-1)!$, whilst for the probability measures we have an upper bound on distinct cardinalities, $|\Psi_{\rightarrow \cdots j}| \geq (n-1)!$ and $|\Psi_{j \cdots \rightarrow}| \geq (n-1)!$.

One particular permutation will be important enough here to be singled out: define η to be the permutation defined by:

$$\begin{pmatrix} 1 & 2 & 3 & \ldots & i & \ldots & n \\ n & 1 & 2 & \ldots & i-1 & \ldots & n-1 \end{pmatrix}$$

Clearly, for any $\sigma_{j \cdots \rightarrow} \in \Sigma_{j \cdots \rightarrow}$, $\eta \sigma_{j \cdots \rightarrow} \in \Sigma_{\rightarrow \cdots j}$.

Proposition 8. For each permutation σ, $p_\sigma M_{\sigma(n)} = p_\sigma$ and $p_\sigma M_{\sigma(1)} = p_{\eta\sigma} \in \Psi_{\rightarrow \cdots \sigma(1)}$.

Proof. Select an arbitrary *j* and an arbitrary $p_{\sigma_{j\ldots\to}} \in \Psi_{j\ldots\to}$. By definition, $\sigma_{j\ldots\to}(1) = j$ and $p_{\sigma_{j\ldots\to}} M_{\sigma_{j\ldots\to}} = p_{\sigma_{j\ldots\to}}$.

Write out $M_{\sigma_{j\ldots\to}}$ as $M_j \ldots M_k$, where $k = \sigma_{j\ldots\to}(n)$. Follow this by multiplying both sides of $p_{\sigma_{j\ldots\to}} M_{\sigma_{j\ldots\to}} = p_{\sigma_{j\ldots\to}}$ on the right by M_j. We are thus lead to the equation $p_{\sigma_{j\ldots\to}} M_j \ldots M_k M_j = p_{\sigma_{j\ldots\to}} M_j$. Rewriting this as $(p_{\sigma_{j\ldots\to}} M_j)(\ldots M_k M_j) = p_{\sigma_{j\ldots\to}} M_j$, or equally well as $(p_{\sigma_{j\ldots\to}} M_j) M_{\eta\sigma_{j\ldots\to}} = p_{\sigma_{j\ldots\to}} M_j$, indicates that $p_{\sigma_{j\ldots\to}} M_j$ is the unique Samet invariant probability measure of $M_{\eta\sigma_{j\ldots\to}}$. But the Samet probability measure of $M_{\eta\sigma_{j\ldots\to}}$ already has a label, $p_{\eta\sigma_{j\ldots\to}}$, so in particular

$$p_{\sigma_{j\ldots\to}} M_j = p_{\eta\sigma_{j\ldots\to}}$$

Now, $\eta\sigma_{j\ldots\to} \in \Sigma_{\to\ldots j}$, so that $p_{\eta\sigma_{j\ldots\to}} \in \Psi_{\to\ldots j}$, which in other words (recalling that $\sigma_{j\ldots\to}(1) = j$) means

$$p_{\sigma_{j\ldots\to}} M_{\sigma_{j\ldots\to}(1)} \in \Psi_{\to\ldots j}$$

We can now run the following series of calculations, based on $p_{\sigma_{j\ldots\to}} M_j = p_{\eta\sigma_{j\ldots\to}}$. First multiply on the right by M_j:

$$p_{\sigma_{j\ldots\to}} M_j M_j = p_{\eta\sigma_{j\ldots\to}} M_j$$

But $M_j^{\,2} = M_j$, so

$$p_{\sigma_{j\ldots\to}} M_j = p_{\eta\sigma_{j\ldots\to}} M_j$$

We started this chain of calculations with $p_{\sigma_{j\ldots\to}} M_j = p_{\eta\sigma_{j\ldots\to}}$, so we conclude that

$$p_{\eta\sigma_{j\ldots\to}} = p_{\eta\sigma_{j\ldots\to}} M_j$$

In other words, the Samet invariant probability measure of $M_{\eta\sigma_{j\ldots\to}}$, $p_{\eta\sigma_{j\ldots\to}}$, is also an invariant measure of M_j, hence a prior for player *j*.

Since by definition $\eta\sigma(n) = j$, the set $\{\eta p_{\sigma_{j\ldots\to}}\}$ as $\sigma_{j\ldots\to}$ ranges over all elements of $\Sigma_{j\ldots\to}$, is just $\Psi_{\to\ldots j}$, and hence $p_{\sigma_{\to\ldots j}} M_j = p_{\sigma_{\to\ldots j}}$ for all $p_{\sigma_{\to\ldots j}} \in \Psi_{\to\ldots j}$. As *j* was selected arbitrarily, it follows that for any permutation σ, $p_\sigma M_{\sigma(n)} = p_\sigma$, and the proof is complete. ∎

Corollary. Each element of the set $\Psi_{\to\cdots j}$ is a prior for player *j*. □

Given the corollary, we can call, for each *j*, the set of $\Psi_{\to\cdots j}$ the *set of balanced priors* for player *j*, and we might as well label it $\Psi_j \equiv \Psi_{\to\cdots j}$. There are n such sets, Ψ_1 through Ψ_n, each of which is of cardinality at most $(n-1)!$, for a total of $n!$ balanced priors. We can label the totality of the balanced priors, $\bigcup_{i\in I} \Psi_i$, by Ψ.

As the proof shows, the type matrices $\{M_i\}_{i\in I}$ play the following roles here: for each *j*, M_j acts, by way of the action pM_j, as a mapping of all of $\Psi_{j\cdots\to}$ to all of $\Psi_{\to\cdots j} \equiv \Psi_j$; in addition, for each $i \neq j$, it maps one element in each Ψ_i to an element in Ψ_j; and it maps each element in Ψ_j to itself.

The n-player version of Proposition 5 of Samet (1998a) follows readily: if there is a common prior *p*, then for any two permutations σ and σ', $pM_\sigma = pM_{\sigma'} = p$, and all the balanced priors are equal by definition. In the other direction, if all the balanced priors coincide, a common prior has been identified, simply because the balanced priors *are* priors.

4.2 Orbits of Priors

It would seem from the previous section that *a priori* one may need to calculate all the $n!$ Samet probability measures in order to answer the question 'is there a common prior' for a given n-player type space. This is quite a calculational burden, given the n-fold matrix multiplication needed for working out each matrix M_σ, and then the effort required for working out the invariant probabilities of these Markov matrices. Fortunately, it is possible to prove a theorem that indicates an easier way.

Definition. For any $p_\sigma \in \Psi$, the orbit of p_σ under the action of the permutation η – i.e., the n-element set $\{p_0, p_1, \ldots, p_{n-1}\}$, where for each $i \in \{0,\ldots,n-1\}$, $p_i = \eta^i p_\sigma$ – will be termed an *orbit of balanced priors*, or just an orbit of priors for short.

Lemma. The set of balanced priors Ψ can be partitioned into $n-1$ distinct orbits of priors. Each such orbit contains exactly one representative from each element of $\{\Psi_j\}_{j\in I}$.

Proof. That the orbits of priors partition the space Ψ into $n-1$ distinct subsets follows from standard results in the theory of group actions and orbits.

Next, select an arbitrary p_σ and consider its orbit $\{p_0, p_1, \ldots, p_{n-1}\}$. From the previous result that $p_\mu M_{\mu(1)} = p_{\eta\mu}$ for any permutation μ, we can immediately conclude that for each i, $p_{i+1} = p_i M_{\sigma(i)}$. One the other hand, as $\eta^i \sigma$ is the permutation associated with p_i, we can also write $p_{i+1} = p_{\eta^i \sigma} M_{\eta^i \sigma(1)} \in \Psi_{\eta^i \sigma(1)}$, and therefore as i goes from 0 to $n-1$, each p_i is located in a different set of balanced priors. But the set $\{\Psi_{\eta^i \sigma(1)}\}_{i \in \{0,\ldots,n-1\}}$ is nothing other than $\{\Psi_j\}_{j \in I}$, and we conclude that each orbit of priors contains exactly one representative from each element of $\{\Psi_j\}_{j \in I}$. □

Proposition 9. A type space $\langle I, \Omega, (\prod_i, t_i)_{i \in I} \rangle$ has a common prior if and only if each orbit of priors $\{p_0, p_1, \ldots, p_{n-1}\}$ satisfies $p_0 = p_2 = \ldots = p_{n-1}$.

Proof. If there is a common prior, then all the elements of Ψ coincide, hence trivially all the elements of each orbit of priors coincide.

In the other direction, suppose that an arbitrary orbit of priors $O = \{p_0, p_1, \ldots, p_{n-1}\}$ satisfies $p_0 = p_2 = \ldots = p_{n-1}$. By the lemma, each $p_i \in O$ is located within a distinct element of the set $\{\Psi_j\}_{j \in I}$. But each such Ψ_j is a set of priors for player *j*. We have therefore identified a common prior for the type space. □

This proposition considerably reduces the informational and calculational burden for establishing the existence of a common prior for a given type space. For one thing, the set of all balanced priors Ψ contains *n*! elements; each orbit of priors consists of *n* elements. Secondly, the algorithm for identifying a common prior is now reduced to selecting an arbitrary permutation σ, forming the permutation matrix M_σ, calculating its invariant Samet probability measure p_σ, and then iteratively forming the orbit of priors $\{p_0, p_1, \ldots, p_{n-1}\}$ by setting $p_0 = p_\sigma$ and $p_{i+1} = p_i M_{\sigma(i)}$. If for some *i*, $p_{i+1} \neq p_i$, it can be concluded that the type space does not have a common prior. If on the other hand $p_0 = p_2 = \ldots = p_{n-1}$, not only have we established that there is a common prior *p*, we have precisely identified it. It follows that the study of type spaces with common priors is the study of type spaces all of whose orbits of priors are uniformly equal.

Since any orbit of priors contains all the information needed to ascertain whether or not there is a common prior, we can select one arbitrarily to serve as the representative orbit of priors $\{p_0, p_1, \ldots, p_{n-1}\}$ for a particular type space.

Note that though the orbits of balanced priors contain information regarding the existence or non-existence of common priors shared between all *n* players in a type space, they do not tell us anything about common priors amongst proper subsets of the set of all players. There are well-know examples of type spaces that have no common priors but in which every pair of players share a common prior between them.

Finally, note that when the type space consists of only two players, there is only one possible orbit of priors, which is the entire 2-element space of balanced priors $\{p_{\sigma_1}, p_{\sigma_2}\}$, and we recapitulate the results of the previous section in the case of two players.

4.3 Characterising Balanced Priors

Just as in the 2-player case it is possible to provide a definition of balanced priors that is intrinsic to type spaces and does not make reference to Markov chain concepts.

Definition. Given an n-player type space, the elements of an n-tuple of probability measures $\{p_0, p_1, ..., p_{n-1}\}$ will be termed *balanced* if there exists a mapping $\mu : \{0,...,n-1\} \to I$ such that for $j < n-1$, $p_{j+1} = p_j M_{\mu(j)}$, and $p_0 = p_{n-1} M_{\mu(n-1)}$.

It is clear that given an n-tuple of balanced probability measures $\{p_0, p_1, ..., p_{n-1}\}$, for each j, $p_j M_{\mu(j)} M_{\mu(j+1)} ... M_{\mu(n-1)} M_{\mu(0)} ... M_{\mu(j-1)} = p_j$, hence they are all Samet probability measures. This insight leads to a straight-forward proof, which we omit here, that any n-tuple of balanced probability measures is an orbit of balanced priors, and hence there are between one to $n-1$ distinct such tuples in any n-player type space.

4.4 Mutually Acceptable Bets

In the n-player case, a set of random variables $\{t_i\}_{i \in I}$ is a *feasible bet* if $\sum_{i \in I} t_i(\omega) = 0$ for all $\omega \in \Omega$. The players *fail to bet* on a feasible bet $\{t_i\}_{i \in I}$ at some state ω if it is not common knowledge at ω that each player's expectation $E_i t_i$ of his own bet t is positive, i.e. $E_i t_i > 0$. A *mutually acceptable bet* is a feasible bet $\{t_i\}_{i \in I}$ which does satisfy the condition that $E_i t_i > 0$ for all players. There is *no betting* amongst the players at ω if they fail to bet on any feasible $\{t_i\}_{i \in I}$.

The n-player Morris-Feinberg theorem then states there is no common prior in a type space if and only if there is a mutually acceptable bet $\{t_i\}_{i \in I}$.

Given Proposition 9, it is natural to enquire whether it is possible to derive the conclusion of the Morris-Feinberg theorem directly from consideration of orbits of priors. With the assistance of ideas from Nehring (2001), it turns out that this is true.

Proposition 10. If the elements of any orbit of priors at a state ω^* in a type space $\langle I, \Omega, (\Pi_i, t_i)_{i \in I} \rangle$ fail to be uniformly equal, then there exists a mutually acceptable bet.

Proof. Let $\{p_0^{\omega^*}, p_1^{\omega^*}, ..., p_{n-1}^{\omega^*}\}$ be an arbitrary orbit of priors such that for some l, $p_l^{\omega^*} \neq p_{l+1}^{\omega^*}$. Let σ^l be the permutation associated with $p_l^{\omega^*}$, with $j = \sigma^l(1)$, so that

$p^{\omega^*}_{\sigma^l} \equiv p^{\omega^*}_l$, and label the permutation associated with $p^{\omega^*}_{l+1}$ by $\sigma^{l+1} = \eta\sigma^l$, so that $p^{\omega^*}_{\sigma^{l+1}} \equiv p^{\omega^*}_{l+1}$. By previous results, $p^{\omega^*}_{l+1} = p^{\omega^*}_l M_j \equiv p^{\omega^*}_{\sigma^l} M_{\sigma^l(1)}$. Next, select a random variable f such that $p^{\omega^*}_{\sigma^l} f - p^{\omega^*}_{\sigma^{l+1}} f > 0$ – that such can be found is guaranteed by the fact that $p^{\omega^*}_l \neq p^{\omega^*}_{l+1}$.

Consider the expression $(M_{\sigma^l})^k (f - M_j f)$. Letting the integer k grow without bound, $\lim_{k\to\infty} (M_{\sigma^l})^k (f - M_j f)$ is equal to a vector whose elements are uniformly equal to $p^{\omega^*}_{\sigma^l}(f - M_j f)$. But $p^{\omega^*}_{\sigma^l}(f - M_j f) = p^{\omega^*}_{\sigma^l} f - p^{\omega^*}_{\sigma^{l+1}} f$, and it was already assumed that this last expression is greater than 0 – in other words, $\lim_{k\to\infty} (M_{\sigma^l})^k (f - M_j f) > 0$, as a vector inequality.

This in turn implies, because the state space Ω is finite, that for some finite k, $(M_{\sigma^l})^k (f - M_j f) > 0$ is uniformly greater than 0. Because the k-fold concatenation of $(\sigma^l(1),\ldots,\sigma^l(n))$ is a finite sequence of elements of I, whose initial element is $\sigma^l(1) = j$, by Nehring's Theorem there exist random variables $\{g_i^1\}_{i\in I}$, in $[f]$ such that

$$(M_{\sigma^l})^k (f - M_j f) = \sum_{i\in I} M_j (g_i^1 - M_i g_i^1) > 0$$

Setting, for each $i \neq j$, $\lambda_i^1 = -(g_i^1 - M_i g_i^1)$, and $\lambda_j^1 = \sum_{i\in I} (g_i^1 - M_i g_i^1)$, it is clear that $\sum_{i\in I} \lambda_i^1 = 0$, and that for $i \neq j$, $E_i \lambda_i^1 = 0$, but $E_j \lambda_j^1 > 0$.

Now, since $(M_{\sigma^l})^k (f - M_j f) > 0$, it follows from the properties of type matrices that for each $t > 1$, $M_{\sigma^l(t)} (M_{\sigma^l})^k (f - M_j f) > 0$. But then again we can apply Nehring's Theorem, each time for a finite sequent of elements of I whose initial element is $\sigma^l(t)$, to obtain $\{g_i^t\}_{i\in I}$, in $[f]$ such that, following the same recipe as above, we can define for each $i \neq \sigma^l(t)$, $\lambda_i^t = -(g_i^t - M_i g_i^t)$, and $\lambda_{\sigma^l(t)}^t = \sum_{i\in I} (g_i^t - M_i g_i^t)$. Clearly, $\sum_{i\in I} \lambda_i^t = 0$, and for $i \neq j$, $E_i \lambda_i^t = 0$, but $E_{\sigma^l(t)} \lambda_{\sigma^l(t)}^t > 0$.

Finally, setting $\lambda_i = \sum_{t\in I} \lambda_i^t$, we have $\sum_{i\in I} \lambda_i = 0$, and $E_i \lambda_i > 0$ for all i. □

As with the analogous proposition presented in the previous section, this result indicates there is no lack of mutually acceptable bets in situations of disparate priors. Individuals who wish to engage in betting should have no problem identifying an endless number of feasible bets of the form $\{t_i\}_{i\in I}$ they would all accept – even though they are aware that $\sum_{i\in I} t_i(\omega) = 0$ in all states. As pointed out in Feinberg (2000), each such mutually acceptable bet is analogous to the existence of n securities such that if the players are risk-neutral and have the same utility functions then at every state of the world the sum of what they are willing to pay for the securities is always greater than the total worth of the securities.

Example. In this example, $\Omega = \{\omega_1, \omega_2, \omega_3\}$, $I = \{1,2,3\}$, $\Pi_1 = \{\{\omega_1, \omega_2\}\{\omega_3\}\}$, $\Pi_2 = \{\{\omega_1, \omega_3\}\{\omega_2\}\}$, $\Pi_3 = \{\{\omega_1\}, \{\omega_2, \omega_3\}\}$, $\Pi = \{\{\omega_1, \omega_2, \omega_3\}\}$.

The type matrices are

$$M_1 = \begin{bmatrix} 1/2 & 1/2 & 0 \\ 1/2 & 1/2 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

$$M_2 = \begin{bmatrix} 1/2 & 0 & 1/2 \\ 0 & 1 & 0 \\ 1/2 & 0 & 1/2 \end{bmatrix}$$

$$M_3 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1/3 & 2/3 \\ 0 & 1/3 & 2/3 \end{bmatrix}$$

There are two orbits of priors, given (to eight decimal places) by:

[0.30769231 , 0.23076923 , 0.46153846]
[0.38461538 , 0.23076923 , 0.38461538]
[0.30769231 , 0.30769231 , 0.38461538]

and

[0.28571429 , 0.28571429 , 0.42857143]
[0.35714286 , 0.28571429 , 0.35714286]
[0.35714286 , 0.21428571 , 0.42857143]

and it can be concluded from cursory inspection of either orbit that there is no common prior in this type space.

5. Appendix – Remaining Proofs

Observation. Given any $k \geq 1$ and type matrix M_i, $M_i^k = M_i$.

Proof. Recall that by assumption, for each $\omega \in \Omega$, $t_i(\omega)(\Pi_i(\omega)) = 1$, and for each $\omega' \in \Pi_i(\omega)$, $t_i(\omega') = t_i(\omega)$.

Label the states of Ω as $\omega_1, \ldots, \omega_s$, and select arbitrarily ω_l, ω_m. Then $M_i(\omega_l, \omega_m) = t_i(\omega_l)(\{\omega_m\})$ and by the definition of matrix multiplication

$$M_i^2(\omega_l, \omega_m) = \sum_{j=1}^{s} [t_i(\omega_l)(\{\omega_j\})][t_i(\omega_j)(\{\omega_m\})]$$

We can immediately note that for all *j* such that $\omega_j \notin \Pi_i(\omega_l)$, $t_i(\omega_l)(\{\omega_j\}) = 0$, so that for the sake of working out the sum in the above equation we may restrict attention only to those *j* such that $\omega_j \in \Pi_i(\omega_l)$. We will accordingly write

$$M_i^2(\omega_l, \omega_m) = \sum_{j=1}^{s'} [t_i(\omega_l)(\{\omega_j\})][t_i(\omega_j)(\{\omega_m\})]$$

with s' indicating that the sum is only over those *j* such that $\omega_j \in \Pi_i(\omega_l)$.

Suppose first that $\omega_m \in \Pi_i(\omega_l)$. Then by assumption, as *j* varies, $t_i(\omega_j) = t_i(\omega_l)$ so that $t_i(\omega_j)(\{\omega_m\})$ is a fixed value, equal to $t_i(\omega_l)(\{\omega_m\})$. This fixed value can be pulled out of the sum, leading to the equation

$$M_i^2(\omega_l, \omega_m) = t_i(\omega_l)(\{\omega_m\}) \sum_{j=1}^{s'} t_i(\omega_l)(\{\omega_j\})$$

But because by assumption $t_i(\omega)(\Pi_i(\omega)) = 1$, we can write $\sum_{j=1}^{s'} t_i(\omega_l)(\{\omega_j\}) = 1$, hence $M_i^2(\omega_l, \omega_m) = t_i(\omega_l)(\{\omega_m\}) = M_i(\omega_l, \omega_m)$.

Next suppose that $\omega_m \notin \Pi_i(\omega_l)$ so that $M_i(\omega_l, \omega_m) = 0$. As *j* varies over those *j* such that $\omega_j \in \Pi_i(\omega_l)$, $t_i(\omega_j)(\{\omega_m\}) = 0$. This fixed value can again be pulled out of the sum, leading immediately to the conclusion that $M_i^2(\omega_l, \omega_m) = 0 = M_i(\omega_l, \omega_m)$.

The general result, for M_i^k, follows by straightforward *k*-fold iteration of this result. □

Lemma (to Proposition 6). The mapping $\xi = \gamma \circ \beta \circ \alpha$ is continuous.

Proof. The mapping β maps elements of a real vector space to elements of a real vector space. Hence to establish its continuity we can rely on the Heine definition of continuity and consider an arbitrary sequence $(M_1^1, M_2^1), (M_1^2, M_2^2), ..., (M_1^j, M_2^j), ...,$ of elements of $M_1(\Omega) \times M_2(\Omega)$, and the associated sequence $\beta(M_1^1, M_2^1), \beta(M_1^2, M_2^2), ..., \beta(M_1^j, M_2^j), ...$ Suppose that $\lim_{j \to \infty}(M_1^j, M_2^j) = (M_1^0, M_2^0)$.

For each $k \geq 1$, define the infinite sequence $(M_{\sigma_1}^1)^k, (M_{\sigma_1}^2)^k, ..., (M_{\sigma_1}^j)^k, ...$ and the sequence $(M_{\sigma_2}^1)^k, (M_{\sigma_2}^2)^k, ..., (M_{\sigma_2}^j)^k, ...,$ where $M_{\sigma_1}^j = M_2^j M_1^j$ and $M_{\sigma_2}^j = M_1^j M_2^j$. Then $\lim_{j \to \infty}(M_{\sigma_1}^j)^k = (M_{\sigma_1}^0)^k$ and $\lim_{j \to \infty}(M_{\sigma_2}^j)^k = (M_{\sigma_2}^0)^k$, where $M_{\sigma_1}^0 = M_2^0 M_1^0$ and $M_{\sigma_2}^0 = M_1^0 M_2^0$. But for each j, $\lim_{k \to \infty}(M_{\sigma_1}^j)^k$ approaches a probability matrix A_j each of whose rows is a probability vector $p_{\sigma_1}^j$ that is one of the balanced priors associated with the pair (M_1^j, M_2^j), and $\lim_{k \to \infty}(M_{\sigma_2}^j)^k$ gives a matrix each of whose rows is the other balanced prior. Then $\lim_{j \to \infty} \lim_{k \to \infty}(M_{\sigma_1}^j)^k$ approaches a probability matrix each of whose rows is one of the balanced priors associated with the matrices (M_1^0, M_2^0), and the same holds for $\lim_{j \to \infty} \lim_{k \to \infty}(M_{\sigma_2}^j)^k$. Hence $\lim_{j \to \infty} \beta(M_1^j, M_2^j) = \beta(M_1^0, M_2^0)$, and we have proved that β is continuous.

The mapping $\gamma : (p_1, p_2) \mapsto p_1 - p_2$ is clearly continuous, as is α, being by definition a topological isomorphism. Hence $\xi = \gamma \circ \beta \circ \alpha$ is continuous. ∎

References

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

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

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

Bernheim, B. Axiomatic Characterizations of Rational Choice in Strategic Environments, Scand. J. of Economics 88 (3) (1986) 473-488

Feinberg, Y. Characterizing Common Priors in the Form of Posteriors, Journal of Economic Theory 91 (2) (2000), 127-179

Gale, D. The Theory of Linear Economic Models, New York, McGraw-Hill, 1960.

Geanakoplos, J. and Polemarchakis, H. We Can't Disagree Forever, Journal of Economic Theory 28 (1) (1982), 192-200

Guillemin, V. and Pollack, A. Differential Topology, Prentice-Hall, New Jersey, 1974.

Gul, F. A Comment on Aumann's Bayesian View, Econometrica 66 (4) (1998), 923 -927.

Milgrom, P. and Stokey, N., Information, Trade and Common Knowledge, Journal of Economic Theory, 26 (1982), 17 – 27

Monderer, D. and Samet, D. (1989). Approximating Common Knowledge with Common Belief, Games and Economic Behavior 1, 170-190

Morris, S. Trade with Heterogeneous Prior Beliefs and Asymmetric Information, Econometrica, 62 (1995), 1327-1347

Nau, R. The Incoherence of Agreeing to Disagree. Theory and Decision 39 (1995) 219-239

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

Samet, D. Iterated Expectations and Common Priors, Games and Economic Behavior, 24 (1998), 131-141

Samet, D. Common Priors and Separation of Convex Sets, Games and Economic Behavior, 24 (1998), 172-173

Sebenius, J., and Geanakoplos, J. Don't bet on it: Contingent agreements with asymmetric information, Journal of the American Statistical Association 78 (1983), 424–426.

Stewart, W. Introduction to the Numerical Solutions of Markov Chains, 1994, Princeton University Press, Princeton, New Jersey