Common Priors under Incomplete Information: A Unification

Klaus Nehring
Department of Economics
University of California, Davis
email: kdnehring@ucdavis.edu

First version: June 1999
This version: March 2000

ABSTRACT

While the meaningfulness of the common prior assumption (CPA) under incomplete information has been established recently by various authors, its epistemic rationale has not yet been adequately clarified. To do so, we provide a characterization of the CPA in terms of a new condition called "Mutual Calibration", and argue that it constitutes a more transparent and more primitive formalization of the Harsanyi Doctrine under incomplete information than the existing characterizations. Our analysis unifies the understanding of the CPA under incomplete information, clarifying the role of higher-order expectations and of the difference between situations with only two and those with at least three agents.

In the concluding section, the analysis is applied to the problem of defining Bayesian consistency of the intertemporal beliefs of a single-agent with imperfect memory. The CPA yields a notion of "Bayesian updating without a prior".

1. INTRODUCTION

The Common Prior Assumption (CPA) is a fundamental and pervasive assumption in much of game theory and the economics of information. It is often motivated as expressing the principle "that differences in probability estimates of distinct individuals should be explained by differences in information"; Aumann (1987, p.7) refers to this as the *Harsanyi Doctrine.* The Harsanyi Doctrine itself can be justified either normatively as a requirement of intersubjective rationality, or pragmatically, to enable one to "zero in on purely informational issues in analyzing economic models with uncertainty"[1]. Under *asymmetric information*, when there is a real prior stage with commonly known beliefs, the CPA is a compelling formalization of the Harsanyi Doctrine. Under genuinely *incomplete information* without a prior stage, the CPA is not a transparent assumption on the primitives of the model, i.e. the agents' belief hierarchies (see especially Gul (1998) and Lipman (1997)); there is thus an issue whether the CPA is *meaningful* under incomplete information. Moreover, it is not clear that the Harsanyi Doctrine itself can be formulated in a meaningful way, since the model no longer contains a substantive notion of "receiving information". As emphasized by Dekel and Gul (1997), the CPA stands in danger of losing its *epistemic rationale* as a result. We consider these two challenges to the interpretation of the CPA to be major advances in the understanding of the notion of incomplete information.[2]

Recently, a small literature has emerged that establishes the CPA's meaningfulness via representation theorems. Most contributions characterize the CPA as essentially equivalent to the absence of mutually profitable bets or variants thereof ("No Betting"); see Feinberg (2000) and Bonanno-Nehring (1999)[3], as well as Samet (1998b) and Halpern (1998)[4]). All of these characterizations are incomplete in the sense of establishing only the existence of a common prior, without giving content to the common prior itself. The gap has been closed by Samet's (1998a) work using infinitely deep expectations about expectations ... about expectations of random variables; we shall refer to his conditions as *Samet Consistency.* This is an important achievement, since modelling specific economic situations typically requires assumptions on the "artificial" common prior; Samet's char-

[1] Aumann (1987, p.13); we concur with Morris' (1995) assessment that the methodological motivation is the more important and most widely shared one.

[2] The distinction of incomplete and asymmetric information is of fundamental importance; in particular, as Dekel-Gul (1997) and Gul (1998) argue forcefully, it is without loss of generality to assume common knowledge of the model only in a hierarchy framework representing incomplete information.

[3] Bonanno-Nehring (1999) emphasize in addition the role of truth conditions on probability-one beliefs.

[4] These were preceded by Morris's (1994) work under asymmetric information.

acterization shows that these can be interpreted in terms of agents' belief hierarchies, at least in principle. On the other hand, Samet Consistency seems less than fully transparent epistemically (cf. the discussion in section 6 where it is formally defined)[5]. To overcome its deficiencies, in particular: to obtain a satisfactory formal expression of the Harsanyi doctrine under incomplete information, while preserving the advantages of Samet's ingenious use of higher-order expectations, we provide in this paper yet a third characterization in terms of a condition called "Mutual Calibration"[6]. Mutual Calibration can be viewed as a common core of Samet Consistency and No Betting; its analysis leads to a unified understanding of the CPA under incomplete information.

Mutual Calibration says the following, roughly. Under incomplete information, it is perfectly natural that some agent, Bob, expects another agent, Ann, to be negatively surprised about the realization of some random variable f , $E_{Bob}(f - E_{Ann}f) < 0$ due to "implicit"[7] private information on Bob's part. This intuitive explanation ceases to be viable, however, if Bob's expectation of Ann's negative surprise is common knowledge, in other words: evident to both of them. Then Bob's expectation of Ann's surprise would have to be attributed at least in part to a purely "attitudinal" difference between them, specifically an assessment of Ann by Bob as "optimistically biased" about f. This contradicts the spirit of the Harsanyi doctrine, and is accordingly ruled out by Mutual Calibration. – Similarly, a third agent, Yang, may expect Bob to expect Ann to be negatively surprised about f due to differences of implicit information, i.e. $E_{Yang}E_{Bob}(f - E_{Ann}f) < 0$. Common knowledge of events of this kind, however, cannot be explained in this way, and is also ruled out by Mutual Calibration. The main result of the paper, Theorem 1, shows Mutual Calibration to be equivalent to the existence of a common prior. While our debt to Samet (1998a) will be evident, the differences between Mutual Calibration and Samet Consistency are substantial. Conceptually, Mutual Calibration is naturally viewed as a multi-agent generalization of "Comprehensive Agreement" (the absence of any "*agreement to disagree*" in the sense of Aumann (1976)), while Samet Consistency is not. This is reflected in the mathematics: in Mutual Calibration higher-order expectations play no independent role in the case of only two agents; their role is to generate jointly and not merely pairwise common priors when their are at least three agents. No analogous difference between two- and

[5]For example, it is not clear how Samet consistency could be defined syntactically.

[6]Mutual Calibration as a relation between epistemic agents has to be distinguished, of course, of "well-calibratedness" in the sense of Dawid (1982) which describes a relation between one agent and the world.

[7]By "implicit information" we mean aspects of agents' beliefs that are taken by others to be informative, i.e. that, if revealed, would lead the latter to revise their beliefs.

many-agent situations exists for Samet Consistency.[8] Moreover, Mutual Calibration has a natural dynamic extension which states that any repeated announcement of (higher-order) expectations of another's surprise is self-defeating ("I cannot predict forever that you will be surprised", section 5).

We shall argue that Mutual Calibration constitutes a more satisfactory reformulation of the Harsanyi doctrine under incomplete information than the existing conditions, in particular than No Betting. This is a live issue only for the many-agent case. With two agents, the existence of a mutually profitable pair of bets is equivalent to an "agreement to disagree" about the expected profitability of either agent's bet, which arguably is the most clear-cut violation of the Harsanyi doctrine under incomplete information available[9]; see section 3. In the many-agent case, this equivalence no longer holds and needs to be replaced by the statement that the existence of a mutually profitable vector of bets entails mutual miscalibration among the agents in their assessment of the expected profitability of *some* agent's bet; in other words, the mutual non-profitability of some bet can be derived from Mutual Calibration concerning the expected profitability of each agents' bet (Proposition 2). Mutual Calibration can thus be viewed as *more primitive* a condition than No Betting, and, we shall argue, by consequence as a more satisfactory primitive[10] formalization of the Harsanyi doctrine; see section 7 for details.

In dual manner, mutual miscalibration among agents concerning some random variable gives rise to a mutually beneficial bet based on agents' higher-order expectations about that random variable (Proposition 3). This leads to a new proof of the equivalence of the CPA to No Betting as a corollary to our main result, Theorem 1, which shows in an epistemically transparent manner *how* the non-existence of a common prior generates mutually profitable betting opportunities. Perhaps contrary to initial appearances, No Betting, Mutual Calibration, and Samet Consistency are cut from rather similar cloth. Indeed, in view of the inherent simplicity of the CPA, a different state of affairs would have been rather surprising.[11]

In the final section 8, we apply the foregoing analysis to the problem of characterizing of dynamic coherence of a single individual with imperfect memory whose belief dynamics are modeled in an intertemporal belief-hierarchy / type-space framework. To the best of our knowledge, this problem

[8] Indeed, Samet motivates his contribution entirely in the context of two agents.

[9] Bonanno-Nehring (1999).

[10] Of course, in a *derived* sense all characterizations of the CPA are legitimate formalizations of the Harsanyi doctrine if any characterization is.

[11] By contrast, Aumann's (1998) reply to Gul proposes a qualitatively different justification of the CPA, which goes however beyond the "static" incomplete information model by considering alternative information structures.

has not been addressed before in either the economics or the philosophy literature[12]. We invoke Theorem 1 to establish the meaningfulness and epistemic rationale of the intertemporal CPA, i.e. of "*Bayesian updating without a* (historical) *prior*".

We discuss the application to single-agent belief dynamics rather extensively for two reasons. First, it is of significant interest in its own right, conceptually as well as for game-theoretic and economic applications. In particular, in the context of modelling learning in Bayesian spirit, the standard assumption of perfect memory is very strong and drives many results. For example, it implies that agents who are exposed to an infinite number of draws from a fixed distribution over their infinite lifetime will eventually learn the true distribution. Yet for real-world agents learning is not the automatic product of ample exposure. If they do not pay attention, or do not sufficiently remember their past experiences, they may learn very little or not at all; just recall trying to learn the vocabulary of a foreign language!

There is also the hope of some pedagogical cross-fertilization. In particular, while the Harsanyi Doctrine is highly controversial in an interpersonal context[13], it commands almost universal consent in an intertemporal setting in the form of Bayesian updating. On the other hand, generalizing Bayesian updating to situations with imperfect memory is a conceptually non-trivial task, since such situations are characterized as essentially by genuinely incomplete information as are situations of interpersonal uncertainty. As a result, the CPA faces the same foundational issues across time as across individuals. Indeed, these issues go to the heart of imperfect memory: after all, how seriously would one be taking these imperfections if one were to postulate a prior that existed in historical time at epistemic birth and that is remembered and at hand for updating at all times?

The plan of the paper is as follows. After setting up the formal framework in section 2, we introduce the Comprehensive Agreement among two agents as a benchmark for formalizing the Harsanyi doctrine under incomplete information in section 3. Section 4 defines "Mutual Calibration" among many agents and demonstrates its equivalence to the existence of a common prior. A natural dynamic extension is sketched in section 5; it expresses the intuitive notion that agents treat other agents' beliefs as "fully conveying information". Mutual Calibration is compared to Samet Consistency in section 6 and related to No Betting in section 7. The concluding section 8 considers the CPA problematic in an intertemporal setting. The remaining proofs are collected in an appendix.

---

[12] With the partial exception of Bonanno-Nehring (1999) who point out the existence of such an isomorphism.

[13] Cf. e.g. Morris (1995).

2. FORMAL FRAMEWORK

Definition 1 *A* **type space** *is a tuple* $\langle I, \Omega, \{p_i\}_{i \in I} \rangle$, *where*

- I *is a finite set of* agents.
- Ω *is a finite set of* states*; the subsets of* Ω *are called* events.
- *for every agent* $i \in I$, p_i *is a function that specifies, for each state* $\alpha \in \Omega$, *his probabilistic beliefs* $p_i^\alpha : 2^\Omega \to \mathbf{R}$ *at* α.

A state in a type space can be thought of as a notational device for describing the belief hierarchies of each agent.[14] Throughout, agents are assumed to be certain of and correct in their beliefs ("Introspection" and "Truth"). The second assumption is made for simplicity but is not unrestrictive; its bite lies mainly in assuming that agents are always certain of the correctness of each others' beliefs.[15]

> **Introspection)** *For all* $\alpha \in \Omega$ *and all* $i \in I$: $p_i^\alpha(\{\omega \in \Omega \mid p_i^\omega = p_i^\alpha\}) = 1$.
>
> **Truth)** *For all* $\alpha \in \Omega$ *and all* $i \in I$: $p_i^\alpha(\{\alpha\}) > 0$.

For all $\alpha \in \Omega$, let $\Pi_i(\alpha) := \text{supp } p_i^\alpha$. By Introspection and Truth, the family $\Pi_i := \{\Pi_i(\omega) \mid \omega \in \Omega\}$ is a partition of Ω, i's *type partition.*[16] An agent "knows" an event E at α if he is certain of it, i.e. if $E \supseteq \Pi_i(\alpha)$. For a set of agents J, let Π_J denote the meet (finest common coarsening) of the partitions $\{\Pi_i\}_{i \in J}$, with $\Pi_J(\alpha)$ denoting the cell of the meet containing state α. E is *common knowledge among the agents in* J if everybody in J knows that E, and if everybody in J knows that everybody in J knows that E, and so forth. Formally, E is "common knowledge among J" at α if $E \supseteq \Pi_J(\alpha)$.

Definition 2 *A probability measure* μ *is a* **common prior among J** *at* α, *if* $\mu(\alpha) > 0$ *and if, for all* $i \in J$, $\omega \in \Pi_J(\alpha)$ *and* $E \subseteq \Omega$, $p_i^\omega(E) = \mu(E/\Pi_i(\omega))$ *whenever* $\mu(\omega) > 0$.

[14]By results due to Armbruster-Boege (1979), Mertens-Zamir (1985), Brandenburger-Dekel (1993), any profile of probabilistic belief hierarchies has a type-space representation; the assumption that the state space Ω is finite is restrictive but entirely standard.

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

[16]We have not included the type partitions among the primitives to emphasize their derived status, especially to prevent a reading of types as "signals".

In the following, “common knowledge” and “common prior” will be short for “common knowledge / common prior among I”. Note the restrictions “$\mu(\alpha) > 0$” and “for all $\omega \in \Pi_J(\alpha)$”; these are necessary to obtain a *local* definition of a common prior that is well-defined in terms of the belief-hierarchies of the agents in J at state α.[17] In view of the partitional structure of the Π_i, it is easily verified that, for any two common priors μ and μ' at α, their posteriors on the common-knowledge event $\Pi_J(\alpha)$ coincide and have full support, i.e. that $\mu(./\Pi_J(\alpha)) = \mu'(./\Pi_J(\alpha))$ and supp $\mu(./\Pi_J(\alpha)) = \Pi_J(\alpha)$. In the sequel, without any loss from the relevant local point of view, we will assume throughout that $\Pi_I = \{\Omega\}$. This entails uniqueness of any common prior among I; moreover, the clause “$\mu(\alpha) > 0$” in its definition becomes redundant.

A random variable f is a real-valued function on Ω. For any random variable f, agent i's expectation of that random variable, when viewed as a function of the state, is again a random variable $E_i f$ given by $(E_i f)(\alpha) = \sum_{\omega\in\Omega} p_i^\alpha(\omega) f(\omega)$. We will make frequent use of the fact that $E_i E_i f = E_i f$ due to Introspection. For a probability measure μ on Ω, let $E_\mu f$ denote the expectation of f with respect to μ, $E_\mu f = \sum_{\omega\in\Omega} \mu(\omega) f(\omega)$. Finally, the indicator function associated with the event A is denoted by 1_A.

A common prior μ can be viewed as a “neutral” probability measure in terms of which every agent is “calibrated” in the sense that under μ no agent is expected to be surprised in a particular direction once the true state is revealed; an agent's surprise is given by the random variable $f - E_i f$.[18]

Observation 1 *The probability measure μ is a common prior for I at α if and only if at α, for any $i \in I$ and any random variable f on Ω,*

$$E_\mu(f - E_i f) = 0. \tag{1}$$

This view of the CPA in terms of expectations of random variables rather than conditioning on events has the advantage of tying in naturally with the proposed characterizing conditions of the CPA, all of which are based on expectations of random variables.

[17] Alternative definitions are possible, but they all coincide when Truth is assumed; see Bonanno-Nehring (1999).

[18] Various authors (Mertens-Zamir (1985), Feinberg (1996), Samet (1998a)) have used equation (1) in the form $E_\mu f = E_\mu E_i f$ for $f = 1_A$, $A \subseteq \Omega$, to define Bayesian priors.

3. COMPREHENSIVE AGREEMENT AMONG TWO AGENTS AS BENCHMARK

Under complete information (characterized by the fact that each agent's beliefs are commonly known), the formal content of the Harsanyi Doctrine is trivial: agent's beliefs must be identical. Under incomplete information among two agents, this generalizes naturally to the requirement that whenever some *aspects* of both agents beliefs are commonly known, these must coincide (Bonanno-Nehring (1999)). If the notion of "aspect of belief" is understood sufficiently broadly to include the sign of an agent's expectation about arbitrary random variables, this suffices to obtain a common prior.

Two agents i and j *agree about the random-variable* f (at state α) if the signs of the agents' expectations $\text{sgn}(E_i f)$ and $\text{sgn}(E_j f)$ coincide when they are commonly known;[19] we say "agree" as short for "fail to agree to disagree" in the sense of Aumann (1976). For example, if it is common knowledge that the agents "agree to disagree" about the probability of some event A in that it is common knowledge that i's probability of A is q_i while j's is q_j, with $q_i \neq q_j$, then in our language they disagree about the random-variable $1_A - q_i$. The agents i and j are in **Comprehensive Agreement** at α if they agree at α about any random-variable f on Ω. It follows from Theorem 1 below that the agents i and j are in Comprehensive Agreement at α if and only if there exists a common prior at α for the two.

Comprehensive Agreement is a condition on *pairs* of agents. It is clear that Comprehensive Agreement among all pairs already entails its natural multi-agent generalization (equality of the signs of all agents' expectations whenever these are commonly known among all of them), which is therefore equivalent to the existence of pairwise common priors. Belief hierarchies with pairwise common priors, but no joint common prior, arise easily, as illustrated by the following example involving three agents and three states.

Example 1.

The following is common knowledge:

- There is a skeleton in the closet of exactly one of the agents Ann, Bob, and Yang.
- Every agent knows the content of his/her closet.

[19] The sign operator is assumed to take on the values +,-, and 0; the added complexity in the current definition simplifies things later on.

- If her closet is empty, Ann believes the skeleton to be in Bob's with $\frac{2}{3}$ probability, similarly that if Bob's (resp. Yang's) closet is empty, (s)he believes the skeleton to be in Yang's (resp. Ann's) with $\frac{2}{3}$ probability.

This is represented by the following state space $\Omega = \{\alpha, \beta, \gamma\}$; at α, the skeleton is in Ann's closet, at β in Bob's, at γ in Yang's. In the following matrix, each entry specifies an agent's beliefs at a state.

	α	β	γ
Ann	$(1,0,0)$	$(0,\frac{2}{3},\frac{1}{3})$	$(0,\frac{2}{3},\frac{1}{3})$
Bob	$(\frac{1}{3},0,\frac{2}{3})$	$(0,1,0)$	$(\frac{1}{3},0,\frac{2}{3})$
Yang	$(\frac{2}{3},\frac{1}{3},0)$	$(\frac{2}{3},\frac{1}{3},0)$	$(0,0,1)$

Note that it follows from the structure of the type partitions alone that there is a common prior for every pair of agents; the common prior for Ann and Bob, for example, is given by $(\frac{1}{7},\frac{4}{7},\frac{2}{7})$. Since this not a prior for Yang, no jointly common prior for the three agents exists. Intuitively, one feels that the three agents do disagree; apparently, the criterion of "Comprehensive Agreement" fails to elicit all true disagreements. An extension of the criterion is needed that brings to light those latent disagreements.

4. MUTUAL CALIBRATION

Comprehensive Agreement rules out in particular disagreement about some agent's ("Ann's') surprise $f - E_{Ann}f$ about any variable f. Since Ann does not expect to be surprised herself (in the net) due to Introspection, i.e. since $E_{Ann}(f - E_{Ann}f) = 0$ at every state, Comprehensive Agreement in effect rules out common knowledge of Bob's expectation of her being surprised, i.e. of the event that $E_{Bob}(f - E_{Ann}f) > 0$. This makes perfect intuitive sense: in itself, the event is unexceptional, being naturally attributed to the possession of "implicit private information" by Bob. However, *common knowledge* of Bob's expectation of Ann's surprise $E_{Bob}(f - E_{Ann}f) > 0$ cannot similarly be explained in terms of implicit private information among Ann and Bob. It is naturally read as reflecting Bob's view of Ann as holding a "too pessimistic"[20], negatively biased attitude about f.

[20] "Pessimistic" and "optimistic" expectations simply mean low respectively high ones; no psychological connotations intended !

Enter a third agent, Yang[21]. Again, while it is quite natural that Yang *expects* Bob to expect Ann to be positively surprised about f, i.e. that $E_{Yang}E_{Bob}(f - E_{Ann}f) > 0$ due to implicit private information, this explanation ceases to work (at a preformal, heuristic level!) if Yang's expectation is common knowledge among all three agents. To rule out also such higher-order expectations of surprise, we propose the following condition of "Mutual Calibration".

Definition 3 (Mutual Calibration) *The agents $i \in I$ are* mutually calibrated *with respect to f at α if, for no finite sequence $(i_1, ..., i_K)$ in I it is common knowledge at α among I that $E_{i_K}E_{i_{K-1}}...E_{i_2}(f - E_{i_1}f) > 0$ or it is common knowledge that $E_{i_K}E_{i_{K-1}}...E_{i_2}(f - E_{i_1}f) < 0$. The agents $i \in I$ are* **mutually calibrated** *if they are* mutually calibrated *with respect to any random-variable f on Ω.*

To see how Mutual Calibration excludes interactive beliefs that do not admit a common prior, consider again Example 1. There is mutual miscalibration about the probability of any state; for example, with $f = 1_{\{\beta\}}$, it is common knowledge everywhere that $E_{Yang}E_{Bob}(f - E_{Ann}f) > 0$.

Remark. With only two agents, higher-order instances of Mutual Calibration are equivalent to first-order instances, and thus directly entailed by Comprehensive Agreement.[22] To illustrate, consider Ann's expectation of $E_{Ann}E_{Bob}(f - E_{Ann}f)$. One can write

$$E_{Ann}E_{Bob}(f - E_{Ann}f) = E_{Ann}\left[(E_{Ann}f - f) - E_{Bob}(E_{Ann}f - f)\right],$$

since $E_{Ann}(E_{Ann}f - f)$ is identically zero. Thus common knowledge of the event that $E_{Ann}E_{Bob}(f - E_{Ann}f) > 0$ amounts to straight disagreement in the sense of section 3 on the random variable $(E_{Ann}f - f) - E_{Bob}(E_{Ann}f - f)$.

The following is the main result of the paper.

Theorem 1 *There is a common prior at state α if and only if the agents are mutually calibrated at α.*

Proof. If there is a common prior μ, then by observation 1 the prior expectation of any agent's expectation is simply the prior expectation, i.e. $E_\mu E_i g = E_\mu g$ for any agent i and any

[21] In flesh and blood, Yang emphatically disapproves of the CPA, on normative as well as descriptive grounds.

[22] Note the contrast to Samet (1998a) who motivates his use of higher-order iterated expectations in the context of two agents.

random-variable g. For any finite sequence $(i_1, ..., i_K)$ in I and any random variable f, $K-$fold application of this equation implies that

$$E_\mu E_{i_K} E_{i_{K-1}} ... E_{i_2}(f - E_{i_1} f) = 0 \ , \tag{2}$$

ensuring Mutual Calibration.

The converse relies on a fundamental result due to Samet (1998a) on the limit of infinitely iterated expectations. A sequence $s = (i_1, i_2, ...)$ of elements in I is an $I-$ *sequence* if, for every $i \in I$, $i_k = i$ for infinitely many k's.

Proposition 1 (Samet) *For each random-variable f on Ω and $I-$sequence s, the limit of the iterated expectations* $\lim_{k\to\infty}(E_{i_k} ... E_{i_1} f)$ *exists and its value is common knowledge at each state.*

We will write $E_s f$ for $\lim_{k\to\infty}(E_{i_k} ... E_{i_1} f)$ and refer to it as the *asymptotic iterated expectation of* f with respect to s. Proposition 1 has the following intuitive content. Under incomplete information, finitely iterated expectations typically reflect some private information; in particular, their value will not be common knowledge. However, each time expectations are taken with respect to agent i_k, some of the private information about f that may have remained after the first $k-1$ iterations among the agents $i \neq i_k$ is removed. In the limit of iterating expectations infinitely deeply, all private information is removed, hence the infinitely iterated expectation will be common knowledge.

The converse follows immediately from the following two lemmata. The first can be viewed as a restatement of the definition of a common prior in terms of asymptotic iterated expectations; only the if-part is needed in the proof.[23]

Lemma 1 *Fix any $I-$sequence s. There is a common prior μ at α if and only if, for all $i \in I$ and random-variables f on Ω,*

$$E_s(f - E_i f) = 0. \tag{3}$$

The common prior μ is given by $\mu(A) = E_s(1_A)$ for $A \subseteq \Omega$.

The second lemma derives this asymptotic version of the CPA from Mutual Calibration. Its proof exploits Proposition 1 to conclude than any violation of equation (3) must show up as high-order expectation of surprise.

[23] The second statement is part of Samet's (1998a) Theorem 1.

Lemma 2 *If the agents agree in iterated expectations about f at α then, for any $i \in I$ and any $I-$sequence s, $E_s(f - E_i f) = 0$ at α.*

This concludes the proof of Theorem 1.

5. I CANNOT PREDICT FOREVER THAT YOU WILL BE SURPRISED

Mutual Calibration expresses the Harsanyi Doctrine in a negative way: "differences in belief", specifically: iterated expectations of surprise, cannot be public. Beyond this one would also like to see a rendering of the *positive* intuition behind it: that agents treat other agents' beliefs as "*fully conveying information*".[24] This can be done using a dynamic extension of the notion of Mutual Calibration. We will not be fully explicit formally.

Consider first situations with two agents who are mutually calibrated, Ann and Bob. Suppose that Bob expects Ann to be surprised, e.g. to be too pessimistic, $E_{Bob}(f - E_{Ann}f) > 0$. If Bob reveals his expectation $E_{Bob}(f - E_{Ann}f)$ to Ann, then Ann may and indeed typically would revise her own expectation of f upwards in response to her glimpse of Bob's private information about f. Suppose that Bob keeps revealing his expectation of Ann's remaining surprise, *without Ann ever revealing anything.* Eventually, Ann will stop revising her beliefs, and due to their Mutual Calibration Bob will no longer expect Ann to be surprised. In other words, it eventually will be common knowledge that $E_{Bob}(f - E_{Ann}f) = 0$. This can be proved formally in the manner of the demonstration by Geanakoplos-Polemarchakis (1982) that two agents with a common prior "cannot disagree forever", replacing their dialogue with a monologue.[25] Since Bob does not learn anything in the process, his beliefs remain unchanged; thus, in Bob's estimate, Ann's expectation of f converges fully to his own, i.e. $E_{Bob}E_{Ann}f$ converges to $E_{Bob}f$. In this sense, Ann *fully* incorporates the information content of those aspects of Bob's interactive beliefs about f that have been revealed in the process. Note that for this convergence to occur, it would not be enough if Bob merely revealed his own expectation of f. Rather, Bob needs to reveal his expectation relative to his estimate of her's; intuitively, the latter allows Ann to assess the information content that Bob's beliefs *add* to her own.

[24] The existence of a common prior does not seem adequate for this purpose, since it only clarifies the formal structure of a class of belief-hierarchies via a representation theorem but does not make the notion of a prior complete-information stage meaningful.

[25] The proof idea is straightforward: a common prior, and thus Mutual Calibration, are preserved under Bayesian updating. Thus, at any stage in the revelation process, as long as an expected surprise cannot be ruled out in terms of common knowledge, the common-knowledge cell shrinks due to Mutual Calibration. By finiteness of the state-space, the shrinking has stop in finite time; it will then have become common knowledge that $E_{Bob}(f - E_{Ann}f) = 0$.

Re-enter a third agent, Yang, in Mutual Calibration with Ann and Bob! Now let Yang announce repeatedly her estimate of Bob's expectation of Ann's surprise $E_{Yang}E_{Bob}(f - E_{Ann}f)$ to both Ann and Bob. Then Ann and/or Bob keep revising their beliefs upon hearing Yang's monologue until it is common knowledge that Yang's iterated expectation of Ann's surprise is eliminated, i.e. $E_{Yang}E_{Bob}(f - E_{Ann}f) = 0$. To illustrate, modify Example 1 by assigning to Yang the beliefs $\mu_{\{Ann,Bob\}}/\Pi_{Yang}$ required by consistency with the common prior between Ann and Bob, i.e. $(\frac{1}{5}, \frac{4}{5}, 0)$ at states α or β and $(0, 0, 1)$ at γ. Consider the effect of Yang announcing (at α or β) that $E_{Yang}E_{Bob}(f - E_{Ann}f) = \frac{8}{45}$. At state β, Ann can infer from the announcement the true state, β; Bob knew this already. Similarly at state α, Ann already knows α to be the true state, and Bob can infer from the announcement that the true state is α. Thus, depending on the state, Ann or Bob revise their beliefs; after the announcement, the true state is common knowledge among Ann and Bob. Yang knows *this* (but not the true state); in particular, after having made the announcement, she knows that then $E_{Bob}(f - E_{Ann}f) = 0$.

6. COMPARISON TO SAMET CONSISTENCY

Samet (1998a) characterizes the existence of a common prior in terms of the following condition.

Definition 4 *The agents I are* **Samet consistent** *if, for any random-variable f on Ω and any two $I-$sequences s and s', $E_s f = E_{s'} f$.*

To compare Samet Consistency to Mutual Calibration, it is helpful to introduce the following intermediate condition.

Definition 5 *The agents I are* **asymptotically calibrated** *if there is an $I-$sequence s such that for any random-variable f on Ω and any $i \in I$, $E_s(f - E_i f) = 0$.*

Recall from Lemma 1 that Asymptotic Calibration yields a straightforward characterization of the existence of a common prior. While it is a special case of Samet Consistency, the two notions are rather distinct conceptually. Asymptotic Calibration makes a statement about agents' expectations from an asymptotic perspective expressed by the asymptotic expectation operator E_s; by contrast, Samet Consistency is naturally read as postulating *uniqueness* of the asymptotic perspective by requiring equality of all asymptotic expectation operators.

Both conditions are meaningful in admitting a non-vacuous paraphrase, but they are also marked by a certain epistemic opaqueness. 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 object of expectation" is given, the random-variable f. As a result, for beliefs inconsistent with the CPA, the asymptotic expectations $E_s f$ lack interpretable content; in particular, they are not "priors" of any kind. By consequence, Samet Consistency cannot be read as expressing commonness of priors. While appealing as a "definition" or paraphrase *of* the CPA in terms of belief hierarchies as intended by Samet, it seems ill-suited for our purpose of providing an epistemic rationale *for* the CPA.

Some of the opaqueness simply results from the non-trivial amount of mathematics built into the very statement of these conditions; this is reflected in the need to restrict the conditions to sequences of expectations of a special kind. By contrast, when one replaces Asymptotic Calibration by Mutual Calibration as its finite counterpart[26], no such restriction is necessary. Moreover, as described in section 4, the conceptual content of the latter can be articulated with a minimal depth of iterated expectations (Yang's of Bob's of Ann's); iterated expectations of high or even infinite order play no special conceptual role.

Samet consistency has the following finite analogue: "For no two finite sequences s and s' and no random-variable f on Ω it is commonly known that $E_s f > E_{s'} f$." While a striking consequence of the CPA, this condition is substantially stronger than Mutual Calibration and clearly less compelling as a primitive assumption underlying the CPA.

7. MUTUAL CALIBRATION AND NO BETTING

In trying to generalize Comprehensive Agreement to more than two agents, the prevalent move in the literature has been to equate Comprehensive Agreement with absence of commonly known mutually profitable betting opportunities, which has a canonical definition for any number of agents. Let $(t_i)_{i \in I}$ be an $I-$tuple of random variables; $(t_i)_{i \in I}$ is *feasible* (as a trade among the agents) if the t_i sum up to zero in each state, $\sum_{i \in I} t_i(\omega) = 0$ for all $\omega \in \Omega$; t_i can be viewed as the betting position taken by agent i.

Definition 6 *The agents* $i \in I$ fail to bet on $(t_i)_{i \in I}$ *at* α *if it is not common knowledge at* α *that each agents' expectation* $E_i t_i$ *of his own bet is non-negative, and at least one strictly positive. There is* **No Betting** *at* α *among* I *if the agents in* I *fail to bet on any feasible* $(t_i)_{i \in I}$.[27]

[26] Recall that Asymptotic Calibration can be derived from Mutual Calibration with the help of Proposition 1 (Lemma 2).

[27] In reading this condition as one of "No Betting", agents are assumed to be risk-neutral for simplicity.

Versions of the following result have been obtained by a variety of authors (Morris (1994), Feinberg (2000), Samet (1998b), Halpern (1998)):

Theorem 2 *There is a common prior for I at α if and only if there is No Betting among I at α.*

This result characterizes the CPA in terms of a condition on agents' beliefs at a particular state, thereby establishing its *meaningfulness* convincingly; moreover, it has clear behavioral content. On the other hand, on careful reflection, the *epistemic rationale* for No Betting is not really clear when there are more than two agents[28]. Specifically, it is not transparent how No Betting expresses the Harsanyi Doctrine as a *primitive* assumption on belief hierarchies,[29] since it elicits agents' beliefs about *different* random variables t_i, *one for each agent*; as a result, is unclear how the existence of a profitable bet $(t_i)_{i \in I}$ reflects a genuine "epistemic inconsistency" between the agents. Consider Example 1, where it is common knowledge that Ann believes β to be more likely than γ, that Bob believes γ to be more likely than α, and that Yang believes α to be more likely than β. Is there an obvious inconsistency? While it would clearly be inconsistent for *one* agent to believe these three things simultaneously, why should it be inconsistent for different agents to have these beliefs?[30]

At a primitive level, epistemic inconsistency should be expressed as conflict of belief among agents about the same random-variable f, with "conflict of belief" being formalized in an appropriate way. This is achieved naturally by conditions with the logical form "It is not common knowledge that Γ concerning f", with Γ formalizing epistemic conflict. It is this logical form which qualifies a condition as expressing the Harsanyi doctrine, since it says that "differences in belief" in the manner specified by Γ "cannot be public". Comprehensive Agreement and Mutual Calibration have the desired logical form. So does the finite version of Samet Consistency, if in a less compelling manner.[31]

Note also that in view of the finiteness of Ω, one could equivalently require all expectations to be strictly positive.

[28] With only two agents, this is issue is moot since No Betting translates directly into Comprehensive Agreement, the existence of a profitable bet $(t_1, -t_1)$ being equivalent to disagreement about the random variable t_1.

[29] If the Harsanyi Doctrine were indeed coherent under incomplete information conceptually, as feared by the Dekel-Gul (1997) skepticism, then the existence or non-existence of profitable trading opportunities may just be a brute fact of life, signifying nothing, and so would be the entailed existence of a common prior. Then no well-defined demarcation between differences in beliefs due to raw attitude / due to differences in information could be made.

[30] In response, one might appeal to the possibility of an Agreement of the agents with an outside observer whose beliefs are commonly known, but then the entire issue of characterizing the CPA becomes fairly trivial.

[31] Note that there is no claim that Mutual Calibration is "more intuitive" than No Betting, as is no claim that it is more intuitive than the CPA. As the CPA has "no *right*" to being intuitive on its own (under incomplete information), No Betting has no *right* to being epistemically intuitive on its own.

In line with this analysis, the non-profitability of a feasible bet $(t_i)_{i \in I}$ can be derived from the *conjunction* of Mutual Calibration conditions about the random-variables t_i. Each of these expresses a requirement of epistemic consistency of agents' interactive beliefs concerning a particular random variable, supplying the missing epistemic interpretation of No Betting.

Proposition 2 *If the agents $i \in I$ are mutually calibrated with respect to each t_i at α, they fail to bet on $(t_i)_{i \in I}$*

In complementary fashion, Mutual *Mis*calibration about some random variable gives rise to a mutually beneficial bet based on agents' higher-order expectations about that random variable.

Proposition 3 *If the agents are mutually miscalibrated with respect to f, there exists a profitable trade $(t_i)_{i \in I}$ such that $t_i \in [f]$ for every $i \in I$.*

Here, $[f]$ is defined as the smallest linear subspace L of $\mathbb{R}^\Omega$ containing f with the property that $L \ni E_i g$ whenever $L \ni g$, for any agent i and any random variable g. By Proposition 3, Mutual Calibration amounts to a No Betting condition specialized to trades with a particular structure, and thus, even from a purely formal, interpretation-free point of view, to a substantially weaker condition. The Proposition is a straightforward consequence of the following lemma which states that iterated expectations of surprise are equivalent to sums of first-order expectations of surprise.[32]

Lemma 3 *Given any finite sequence $(i_1, ..., i_K)$ in I (with $K \geq 2$) and any f on Ω, there exist random variables $\{g_i\}_{i \in I}$ in $[f]$ such that*

$$E_{i_K} E_{i_{K-1}} ... E_{i_2} (f - E_{i_1} f) = \sum_{i \in I} E_{i_K} (g_i - E_i g_i).$$

Combining Proposition 3 with our main result, Theorem 1 above, one obtains a new proof for the existence of a common prior due to No Betting, Theorem 2. The existing proofs are all based on versions of separation arguments for non-intersecting convex sets; while mathematically elegant, these proofs lack epistemic content. In particular, they fail to explain in epistemic terms *why* non-existence of a common prior necessarily leads to the existence of a mutually profitable bet. This seems in fact rather remarkable, as it implies that sufficiently many aspects of agents' beliefs must always be common knowledge so that some disagreement can be elicited. The alternative proof via Mutual Calibration rectifies this by exploiting the rich array of commonly known facts established

[32] The Proposition follows immediately from the lemma, setting $t_i = g_i - E_i g_i$ for $i \neq i_K$ and $t_{i_K} = \sum_{i \in I} (g_i - E_i g_i)$, multiplying all of these jointly by -1 if necessary.

by Proposition 1. These cause the absence of a common prior to show up as mutual at sufficiently deep iterations of taking expectations. Proposition 3 then constructs a mutually beneficial trade on its basis of this miscalibration.

8. APPLICATION TO SINGLE-AGENT BELIEF DYNAMICS: BAYESIAN UPDATING WITHOUT A PRIOR

Consider the evolution of a single individual's probabilistic beliefs over time who may not only be uncertain about the future but also about the past; indeed, no real-world individual remembers more than a small fraction of her past, especially of her past beliefs. Type spaces are the natural framework to describe such beliefs, with different "agents" representing the individual at different dates, $I = \{0, 1, ...T\}$. A "state" describes now a complete evolution of the individual's beliefs over time, including beliefs about beliefs such as remembering what she expected and foreseeing what she might forget.

Example 1 can be given an intertemporal interpretation in which Dan[33] rotates in a monthly rhythm between three residences Alpha, Beta and Gamma, one of which has a skeleton in the closet. Dan knows the skeleton's location if it is under her nose, but forgets it otherwise; we leave it to the reader to fill in the details of the translation. Suppose that the true state is α. Then, in January, when residing in Alpha, Dan knows of the skeleton in her closet, and foresees that she will have forgotten its whereabouts in February, etc.. Moreover, the rhythm of her rediscovering and forgetting the location of the skeleton is intertemporal common knowledge among her dated "agents". As forgetful people know, intertemporal belief hierarchies have non-trivial complexity; for example, if Dan asks herself where to look for the skeleton she tried to hide, she would try to remember where she might have predicted that she would look in the future, and presumably anticipated *that*.

A central feature of intertemporal belief hierarchies / type spaces is the absence of independently meaningful information partitions among the primitives. This is essential to the modelling of belief dynamics characterized by *genuinely incomplete* and not merely "asymmetric" information. Even in the context of perfect recall, it has frequently (and validly) been argued that belief revision cannot exclusively be understood as updating based on the receipt of certain information. For example, sensory information may be "impressionistic" and uncertain, and beliefs may change through intro-

[33] "Dan" is short for my sister Daniela who in real life faces similar problems even when she is not looking for skeletons.

spection and thinking.[34] In the philosophical literature, this claim is associated particularly with R. Jeffrey under the headings of "belief kinematics", "radical probabilism", and "Bayesianism with a human face". With characteristic verve, he puts it thus: "Remember: The senses are not telegraph lines on which the external world sends observation sentences for us to condition upon" (Jeffrey 1992, p. 78).[35] The case for type space modelling seems even more compelling with imperfect memory: for while restrictive, it at least makes intuitive sense to postulate a belief change that is fully driven by newly acquired information (i.e. by updating); by comparison, what would it mean to postulate a belief change that is *fully driven* by newly lost information?[36]

The problem, then, is how to define dynamic consistency of probabilistic beliefs in a type-space framework. At a mathematical level, the goal seems clear: consistency should be identifiable with the CPA applied intertemporally. *Once the CPA is justified,* intertemporal beliefs are mathematically isomorphic to those obtained from updating on an initial prior remembered throughout time. However, just as in the interpersonal case, the justification itself has to be given in the general incomplete information setting. Indeed, the above discussion shows that the same foundational issues concerning its meaningfulness and epistemic rationale arise. In particular, the common "prior" cannot generally be credited with historical reality, for this would simply amount to postulating perfect recall of the individual's original beliefs at "epistemic birth", date 0. Not only is this terribly ad hoc, it will also often be wholly implausible, especially since the original beliefs to be remembered must be defined on the entire type space of belief histories and not merely on the space of external uncertainty.

In spite of these difficulties of interpreting the CPA directly, the conceptual content of "dynamic belief consistency" is still quite clear: it is simply the (extended) Harsanyi Doctrine, read intertemporally rather than interpersonally. Here the CPA enjoys an advantage in a temporal setting, since it seems natural to rule out purely attitudinal changes in beliefs over time in view of the identity of an individual over time as a person. And indeed, in the restricted context of perfect recall, the

[34] Binmore and Brandenburger's (1990, pp. 142-147) charge in the context of a critique of the CPA that genuine learning can rarely be modelled as Bayesian updating can also be read as supporting this point.

[35] At a formal level, this insight boils down to the recognition that the relevant intertemporal state space is a type space on which standard Bayesian conditioning can be performed; unfortunately, the discussion is not normally conducted in a type-space setting. A type-space modelling is advocated in Samet (1998c); see also Gaifman (1988).

[36] Notice the following simple disanalogy between aquiring and loosing information. In the case of acquiring information, hence of restricting the support, the current belief contains enough information to determine posterior beliefs via conditioning. In the case of loosing information, i.e. of enlarging the support, current beliefs clearly do not contain enough information to canonically determe a unique posterior.

CPA has enjoyed preeminent status in the form of Bayesian updating. Furthermore, the "reflection principle" (proposed in the philosophical literature as an axiomatic underpinning of Bayesian updating) can be read as a formalizing the Harsanyi Doctrine in an intertemporal context. The reflection principle has not gone unchallenged, however, not even normatively; see in particular Maher (1993) for a survey of the recent discussion.

Based on the above interpersonal analysis of the CPA, we propose intertemporal Mutual Calibration as defining dynamic coherence of probabilistic beliefs. From Theorem 1, one obtains a representation of an individual's beliefs as posteriors derived from Bayesian updating of the (common=) *intertemporal prior* on the dated agent's type. Just as in the interpersonal case, the intertemporal "prior" is a mathematical construct that is not embodied as the belief of a dated agent.

With perfect recall, the characterization of an intertemporal prior simplifies greatly. The interesting point here is the natural role of iterated expectations. Formally, there is *perfect recall* at state α if $p_{t'}^{\alpha}(\Pi_t(\alpha)) = 1$ for all $t, t' \in I$ such that $t' \geq t$ (the individual remembers her past types at any date). Clearly, there is common knowledge of perfect recall if and only if the sequence $\{\Pi_t\}_{t=0,\ldots,T}$ is a filtration, i.e. iff the type partitions of later agents refine those of earlier agents; note that then in particular $\Pi_0 = \Pi_I = \{\Omega\}$.

Definition 7 *The individual's expectations are* **sequentially calibrated** *at α, if, for any $t, t' \in I$ such that $t < t'$ and any random-variable f on Ω, $E_t E_{t'} f(\alpha) = E_t f(\alpha)$.*[37]

Sequential Calibration says that future surprises $f - E_{t'} f$ as well as belief changes $E_{t'} f - E_t f$ always have zero expectation. Note that it imposes a consistency condition on the earlier beliefs at t, rather than on the later beliefs at t'. This reflects the nature of a type space in which beliefs *define* types, which precludes an understanding of later beliefs as derived from earlier beliefs by updating on the type. The following result is straightforward; the only slight surprise is the necessity of perfect recall for Sequential Calibration.

Proposition 4 *There is a common prior and common knowledge of perfect recall at state α if and only if there is common knowledge of Sequential Calibration at α.*

Proposition 4 shows that Mutual Calibration can be viewed as a natural generalization of Sequential Calibration to situations with imperfect memory.

[37] Sequential Calibration is due to Goldstein (1983) under the name of "Iteration"; see also Hild (1996).

APPENDIX: REMAINING PROOFS

Proof of Observation 1.

Only if: If μ is a common prior for I then, for any random-variable f, $\omega \in \Omega$, and $i \in I$: $E_i f = E_{\mu(./\Pi_i)} f$ in natural notation. Hence $E_\mu E_i f = E_\mu E_{\mu(./\Pi_i)} f = E_\mu f$; the latter is simply the law of iterated expectations for conditional expectations.

If: Take any $i \in I$, $\omega \in \Omega$ and $A \subseteq \Omega$. Since $p_i^\omega(A) = p_i^\omega(A \cap \Pi_i(\omega))$ and $\mu_i(A/\Pi_i(\omega)) = \mu_i(A \cap \Pi_i(\omega)/\Pi_i(\omega))$, it is w.l.o.g. to assume that $A \subseteq \Pi_i(\omega)$. But then by assumption, setting $f = 1_A$, one obtains $\mu_i(A) = \mu_i(\Pi_i(\omega)) p_i^\omega(A)$, hence $p_i^\omega(A) = \mu_i(A/\Pi_i(\omega))$. □

Proof of Lemma 1.

(If) For any state $\omega \in \Omega$, define a common-prior probability measure μ on Ω by $\mu(\{\omega\}) = E_\mu(1_{\{\omega\}}) := E_s(1_{\{\omega\}})$ (at any $\beta \in \Omega$), which is well-defined in view of Proposition 1. By the linearity of the expectation operators E_μ and E_s, they must coincide, in that $E_\mu(f) = E_s(f)$ for all f on Ω and any $\beta \in \Omega$. It follows that, for any $i \in I$ and any f on Ω, $E_\mu(f - E_i f) = 0$, which shows that μ must be a common prior in view of observation 1.

(Only if) Straightforward from equation (2) and Proposition 1. □

Proof of Lemma 2.

By contradiction. Consider some $i \in I$ and $I-$sequence $s = (i_1, ..., i_k, ...)$ such that $E_s(f - E_i f)(\alpha) \neq 0$. This limit exists and is constant on Ω by Proposition 1; w.l.o.g. $\varepsilon := E_s(f - E_i f)(\alpha) > 0$. By the finiteness of the state space Ω, this implies for sufficiently high K that $E_{i_K} ... E_{i_1}(f - E_i f)(\omega) \geq \frac{\varepsilon}{2} > 0$ for all $\omega \in \Omega$, contradicting Mutual Calibration. □

Proof of Lemma 3.

The proof of the lemma is by induction. The validity of the claim for $K \leq 3$ follows from the identity $E_{i_3} E_{i_2}(f - E_{i_1} f) = E_{i_3}(f - E_{i_1} f) - E_{i_3}\left[(f - E_{i_1} f) - E_{i_2}(f - E_{i_1} f)\right].$

Thus, suppose the claim to hold for all $K \leq K^* - 1$. Then $E_{i_{K^*}} E_{i_{K^*-1}} ... E_{i_2}(f - E_{i_1} f) = E_{i_{K^*}}\left(\sum_{i \in I} E_{i_{K^*-1}}(t_i - E_i t_i)\right) = \sum_{i \in I} E_{i_{K^*}} E_{i_{K^*-1}}(t_i - E_i t_i)$ for appropriate $\{t_i\}_{i \in I}$. Validity of the claim for $K = K^*$ thus follows from its validity for $K = 3$, applied to each term $E_{i_{K^*}} E_{i_{K^*-1}}(t_i - E_i t_i)$ separately. □

Proof of Proposition 2.

Suppose that it is common knowledge that $E_i t_i \geq 0$ for all $i \in I$, and that $E_j t_j \geq \varepsilon > 0$. Take any $I-$sequence s; by Mutual Calibration with respect to each of the t_i, we obtain from lemma 2 $E_s(t_i - E_i t_i) = 0$, hence that $E_s t_i \geq 0$ for all $i \in I$ and $E_s t_j \geq \varepsilon > 0$. By the linearity of the operator E_s, this implies $E_s(\sum_{i \in I} t_i) > 0$, but $\sum_{i \in I} t_i$ is the null-vector. □

Proof of Proposition 4

If there is a common prior and common knowledge of perfect recall at state α, then common knowledge of Sequential Calibration is a straightforward consequence of the law of iterated expectations for conditional expectations.

For the converse, take any $t \leq t'$, $\alpha \in \Omega$. From the common knowledge of Sequential Calibration, one infers $E_t E_{t'} 1_{\Pi_t(\alpha)} = E_t 1_{\Pi_t(\alpha)} = 1_{\Pi_t(\alpha)}$, whence also $E_{t'} 1_{\Pi_t(\alpha)} = 1_{\Pi_t(\alpha)}$, and thus in particular $p^{\alpha}_{t'}(\Pi_t(\alpha)) = 1$; since α was arbitrary, this shows common knowledge of perfect recall.

Finally, setting $t = 0$, Sequential Calibration at α implies, for any random-variable f and any t', $E_0 E_{t'} f(\alpha) = E_0 f(\alpha)$. In view of observation 1, this establishes p^{α}_{o} as the intertemporal (common) prior. □

REFERENCES

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

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

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

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

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

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

[7] Brandenburger, A. and E. Dekel (1986), Hierarchies of Beliefs and Common Knowledge, *Journal of Economic Theory,* 59, 189-198.

[8] Dawid, A. (1982), The Well-Calibrated Bayesian, *Journal of the American Statistical Association,* 77, 605-613.

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

[10] Feinberg, Y. (2000), Characterizing the Existence of a Common Prior in the Form of Posteriors, *Journal of Economic Theory,* forthcoming.

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

[12] Gaifman, H. (1988), A Theory of Higher Order Probabilities. In: B. Skyrms and W. Harper (eds.): *Causation, Chance and Credence.* Kluwer.

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

[14] Goldstein, M. (1983), The Prevision of a Prevision, *Journal of the American Statistical Association* 78, 817-819.

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

[16] Halpern, J. (1998), Characterizing the Common Prior Assumption, in I. Gilboa (ed.): *Theoretical Aspects of Rationality and Knowledge, Proceedings of the Seventh Conference.* Morgan Kaufmann

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

[18] Hild, M. (1996), Auto-Epistemology and Updating, mimeo, Oxford.

[19] Jeffrey, R. (1992), *Probability and the Art of Judgment,* Cambridge University Press.

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

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

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

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

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

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

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

[27] Samet, D. (1998c), Bayesianism without Learning, forthcoming in *Research in Economics.*