

Article

A Foundational Approach to Generalising the Maximum Entropy Inference Process to the Multi-Agent Context

George Wilmers

School of Mathematics, University of Leeds, Leeds LS2 9JT, UK;

E-Mail: george.wilmers@gmail.com; Tel: +44-1614455124

Academic Editors: Juergen Landes and Jon Williamson

Received: 1 December 2014 / Accepted: 13 January 2015 / Published: 2 February 2015

Abstract: The present paper seeks to establish a logical foundation for studying axiomatically multi-agent probabilistic reasoning over a discrete space of outcomes. We study the notion of a *social inference process* which generalises the concept of an inference process for a single agent which was used by Paris and Vencovská to characterise axiomatically the method of maximum entropy inference. Axioms for a social inference process are introduced and discussed, and a particular social inference process called the *Social Entropy Process*, or **SEP**, is defined which satisfies these axioms. **SEP** is justified heuristically by an information theoretic argument, and incorporates both the maximum entropy inference process for a single agent and the multi-agent normalised geometric mean pooling operator.

Keywords: inference process; maximum entropy; social entropy; Kullback-Leibler; probabilistic reasoning; pooling operator; discrete probability function; probabilistic merging; multi-agent reasoning

1. Introduction

In this introduction we describe briefly the context to the conceptual framework first sketched in [1], and which is developed further in the present work. In section 1.1 we explain how the present paper is structured, while in the remaining sections of the chapter we introduce some necessary background ideas

and technical prerequisites. We also indicate in some places in this chapter details which can be omitted by readers who are only interested in some aspects of the present work.

1.1. Overall Structure

Intuitively a *social inference process* is just a general method for aggregating the partially defined probabilistic beliefs of a finite number of agents into a *single* probabilistic belief function. While the probabilistic beliefs of each individual agent are assumed to be consistent, it is not assumed that the union of the beliefs of any two or more agents is consistent.

The notion of a social inference process includes as special cases two much older, but quite distinct, concepts from probabilistic reasoning: the notion of a single agent *inference process* of [2] and [3], and the notion of a multi-agent discrete probabilistic *pooling* operator familiar from decision theory (see [4] or [5]). Both of these older notions have been studied intensively from an axiomatic standpoint, with some considerable success, particularly in the case of inference processes.

One aim of this paper is to illustrate how the use of the axiomatic method applied to social inference processes can illuminate the study of particular examples of such processes. In particular it can illustrate how an initially attractive, but fundamentally *ad hoc*, definition of a social inference process may fail some quite basic desideratum. On the other hand the formalisation inherent in the axiomatic study of social inference processes may perhaps dissuade researchers from naively criticising a social inference process for failing to satisfy a combination of desiderata which cannot in fact be satisfied by *any* social inference process. There is an interesting historic parallel here with the case of (single agent) inference processes: the centre of mass inference process **CM** was well-known and popular 25 years ago for presumably pragmatic reasons, yet in [3] it was shown that it fails to satisfy some quite elementary desiderata such as Language Invariance, which had not previously been formulated. On the other hand at the time of the first rigorous axiomatic treatment of the notion of inference process in [2] and [3] the maximum entropy inference process **ME** was often criticised for its failure to satisfy a desideratum known as *representation independence*, a superficially attractive principle which Paris [3] showed with a simple proof to be incoherent, since it cannot be satisfied by *any* inference process. The historical point being made here is that had the axiomatic approach to inference processes been formulated earlier this would have spared extensive, but pointless, criticisms of **ME** on the grounds that it was “representation dependent”¹.

The necessary background material and notation covering inference processes and pooling operators respectively is covered briefly in sections 1.2 and 1.3 below, while in section 1.4 the notion of a social inference process is formally introduced.

¹ See Paris [3] where representation independence is called *Atomicity*, and [15,16] where representation independence is discussed at greater length. It should be mentioned that even prior to Paris’s result there were very good reasons to distrust the notion of representation independence as E.T.Jaynes, the longtime champion of **ME**, frequently pointed out to critics of **ME**. However mere good reasons are never as effective in silencing critics as a proof that their arguments are incoherent, in this case because the proposed desideratum is vacuous.

Chapter 2 is devoted to developing an axiomatic framework in order to capture the intended intuitive notion corresponding to the formal representation of a social inference process. This requires some considerable care in first formulating *informally* exactly what notion it is that we are trying to capture. We may then test the consequences of our subsequent axiomatic formalisation against our intuitions and experience in a process which may later be refined and iterated. Such an approach to the foundations of a mathematically tractable domain of thought is sometimes referred to by logicians and philosophers of mathematics as *informal rigour*². Accordingly the first two sections of Chapter 2 are devoted to a detailed analysis of the heuristics and assumptions lying behind our approach. Although these sections are important in terms of justifying and explaining our methodology, they are not required for the formal development, and may therefore be omitted by readers who are only interested in the latter. In section 2.3 we develop a set of principles which we believe that any social inference process should satisfy on the basis of the assumptions explained in sections 2.1 and 2.2.

Chapter 3 is devoted to the particular social inference process SEP, the Social Entropy Process, first defined in [1]. In section 3.1, we formally define SEP, making clear the information theoretic intuitions behind the definition, and establishing certain structural properties, including the relationship of SEP to ME, minimum cross-entropy, and the normalised geometric mean pooling operator. A number of technical results are necessary to this development to ensure that it makes sense mathematically. The reader who wishes to skip these details on a first reading can glean the bare definition of SEP from definitions 8, 10 and 12.

In section 3.2 we prove that SEP satisfies all the principles formulated in Section 2.3. In section 3.3 we consider briefly certain other principles for an inference process resulting from possible generalisations of principles satisfied by ME.

Our definition of SEP in 3.1 proceeds in two stages. At the first stage the probabilistic information $\vec{\mathbf{K}}$ from all the agents is *merged* by a natural and informationally conservative process Δ to form a non-empty closed convex set of probability functions $\Delta(\vec{\mathbf{K}})$, which can be considered as the preferred set of *possible* probabilistic belief functions of the collective, or “collective knowledge base”. At the second stage the unique probability function from the set $\Delta(\vec{\mathbf{K}})$ which has maximum entropy is chosen to be the definitive belief function of the collective, and is denoted by $\text{ME}(\Delta(\vec{\mathbf{K}}))$. At first sight the necessity to impose such a second stage in order to extract a *unique* probabilistic belief function might seem like an ad hoc artifice. However in Chapter 4 we show that the second stage of the definition can be eliminated by imagining that the agents collectively appoint a new, unbiased, and self-effacing agent as a chairman, whose own personal belief function assigns equal probability to each possible outcome. The chairman then seeks to minimise her own influence by imagining that each of the *other* agents has been replaced by n clones where n large. If the chairman then calculates the first stage procedure for the entire virtual set of agents including herself, and lets n tend to infinity, the result converges to the same single probability function as that defined by SEP, thus eliminating any direct use of ME. The technical theorem corresponding to this result is stated and proved in 4.2.

² The terminology is due to the logician Georg Kreisel, see e.g. [36], although I do not claim that I use the terminology in the sense that Kreisel intends. Similar ideas can also be found in the work of Lakatos [37].

In Chapter 5 we give a brief critical evaluation of our work, suggest directions for future research, and list a number of open problems.

1.2. Basic Concepts and Notation for Single Agent Probabilistic Inference

The framework and terminology which we introduce in this section is in essence that of Paris and Vencovská [2,3], which we will extend in section 1.4 to the multi-agent context.

In order to fix notation let $\mathbf{At} = \{\alpha_1, \alpha_2, \dots, \alpha_J\}$ denote some fixed finite set of mutually exclusive and exhaustive atomic events, or, as we prefer to think of them in a logical framework, atoms of some finite Boolean algebra of propositions. We shall refer to $\mathbf{At} = \{\alpha_1, \alpha_2, \dots, \alpha_J\}$ as *atoms*. A probability function w on \mathbf{At} is a function $w : \mathbf{At} \rightarrow [0, 1]$ such that $\sum_{j=1}^J w(\alpha_j) = 1$. Slightly abusing notation we will identify w with the vector of values $\langle w_1 \dots w_J \rangle$ where w_j denotes $w(\alpha_j)$ for $j = 1 \dots J$. The set of all such vectors is denoted by \mathbb{D}_J . All other more complex events considered are equivalent to disjunctions of the α_j and are represented by the Greek letters θ, ϕ, ψ etc. A probability function w is assumed to extend so as to take values on complex events in the standard way, i.e., for any θ

$$w(\theta) = \sum_{\alpha_j \models \theta} w(\alpha_j)$$

where \models denotes the classical notion of logical implication, and whenever a sentence $\theta \in SL$ is not satisfiable we set $w(\theta) = 0$. Conditional probabilities are defined in the usual manner

$$w(\theta | \phi) = \frac{w(\theta \wedge \phi)}{w(\phi)}$$

when $w(\phi) \neq 0$ and are left undefined otherwise.

If some $w \in \mathbb{D}_J$ represents the subjective belief of an individual \mathbf{A} in the outcomes of \mathbf{At} we refer to w as \mathbf{A} 's *belief function*. We note that in this paper the use the term “belief function” will always denote a probability function in the above sense.

Remark 1. We should note that in the framework of Paris and Vencovská the atoms $\alpha_1, \alpha_2, \dots, \alpha_J$ of \mathbf{At} are usually taken to be the atoms of the Boolean (Lindenbaum) algebra generated by a finite language of the propositional calculus $L = \{p_1 \dots p_k\}$ where the p_i are the propositional variables. Thus up to logical equivalence the atoms are just the 2^k sentences of the form

$$\bigwedge_{i=1}^k \pm p_i$$

where $\pm p_i$ denotes either p_i or $\neg p_i$. In such a presentation J is 2^k and so is necessarily a power of 2. More complex “events” are just sentences of the language, which by the disjunctive normal form theorem are logically equivalent to disjunctions of atoms. This addition of an extra semantic layer in the form of an underlying language L which generates the atoms has important conceptual advantages in the formulation and justification of certain natural principles such as the Language Invariance and Irrelevant Information principles of [3]. However since we shall only consider principles of this latter type in sections 3.3 and 5, we may otherwise assume that the mutually exclusive and exhaustive atoms

$\alpha_1, \alpha_2, \dots, \alpha_J$ are given a priori, rather than being generated as atoms of a propositional language L , and we may then allow J to take any positive integral value³.

The problematic of Paris and Vencovská is that of a single individual \mathbf{A} whose belief function \mathbf{w} is in general not completely specified, but whose set of beliefs is instead regarded as a set of constraints \mathbf{K} on the possible values which the vector \mathbf{w} may take. The constraint set \mathbf{K} therefore defines a certain subregion of \mathbb{D}_J , denoted by $\mathbf{V}_{\mathbf{K}}$, consisting of all vectors $\mathbf{w} \in \mathbb{D}_J$ which satisfy the constraints in \mathbf{K} . In the special case when \mathbf{K} is the empty set of constraints, the corresponding region $\mathbf{V}_{\mathbf{K}}$ is just \mathbb{D}_J itself. We say that \mathbf{K} is *consistent* if $\mathbf{V}_{\mathbf{K}} \neq \emptyset$, and that \mathbf{w} is *consistent with \mathbf{K}* if $\mathbf{w} \in \mathbf{V}_{\mathbf{K}}$.

It is assumed that the constraint sets \mathbf{K} which we consider are consistent, and are such that $\mathbf{V}_{\mathbf{K}}$ has pleasant geometrical properties. More precisely, the exact requirement on a set of constraints \mathbf{K} is that the set $\mathbf{V}_{\mathbf{K}}$ forms a non-empty closed convex region of Euclidean space. Throughout the rest of this paper **all constraint sets to which we refer will be assumed to satisfy this requirement**, and we shall refer to such constraint sets as *nice* constraint sets. This formulation ensures that linear equality constraint conditions such as $\mathbf{w}(\theta) = a$, $\mathbf{w}(\phi) = b\mathbf{w}(\psi)$, and $\mathbf{w}(\psi | \theta) = c$, where $a, b, c \in [0, 1]$ and θ, ϕ , and ψ are Boolean combinations of the α_j 's, are all permissible in \mathbf{K} provided that the resulting constraint set \mathbf{K} is consistent. Here a conditional constraint such as $\mathbf{w}(\psi | \theta) = c$ is interpreted as $\mathbf{w}(\psi \wedge \theta) = c\mathbf{w}(\theta)$ which is always a well-defined linear constraint, albeit vacuous when $\mathbf{w}(\theta) = 0$. See e.g. [3] for further details.

We should perhaps remark here that while we have allowed the notion of a nice set of constraints to include more general constraints which do not have the form of linear equalities of the type above, the philosophical justification for the approach which we develop in the present paper is most clearly applicable when the constraints have this form. This observation does not however affect in any way the validity of the formal mathematical results.

A nice set of constraints \mathbf{K} as above is called a *knowledge base*. Where these constraints correspond to an individual \mathbf{A} 's probabilistic beliefs, we say that \mathbf{A} has knowledge base \mathbf{K} . Note that if \mathbf{K}_1 and \mathbf{K}_2 are knowledge bases, then $\mathbf{V}_{\mathbf{K}_1 \cup \mathbf{K}_2} = \mathbf{V}_{\mathbf{K}_1} \cap \mathbf{V}_{\mathbf{K}_2}$, and that $\mathbf{K}_1 \cup \mathbf{K}_2$ is also a knowledge base provided that it is consistent.

Paris and Vencovská ask the question: given that an individual \mathbf{A} 's belief function is subject to the constraint set \mathbf{K} , by what rational principles should \mathbf{A} choose her belief function \mathbf{w} consistent with \mathbf{K} , *in the absence of any other information?*

A rule \mathcal{I} which for every such \mathbf{K} chooses such a $\mathbf{w} \in \mathbf{V}_{\mathbf{K}}$ is called an *inference process*. Given \mathbf{K} we denote the belief function \mathbf{w} chosen by \mathcal{I} by $\mathcal{I}(\mathbf{K})$. The question above can then be reformulated as: what self-evident general principles should an inference process \mathcal{I} satisfy? This question has been intensively studied over the last twenty–five years, and much is known. In particular in [2], Paris and Vencovská found an elegant set of principles which uniquely characterise the maximum entropy inference process⁴.

³ We should note that allowing J to take any positive integral value does not in any real sense generalise the Paris–Vencovská framework, but is sometimes a notational convenience when constructing simple examples.

⁴ This characterisation considerably strengthens earlier work of [38].

ME, which is defined as follows: given \mathbf{K} as above, $\text{ME}(\mathbf{K})$ chooses that unique belief function w which maximises the Shannon entropy of w , defined as

$$-\sum_{j=1}^J w_j \log w_j$$

subject to the condition that $w \in \mathbf{V}_{\mathbf{K}}$. Although some of the principles used to characterise ME may individually be open to philosophical challenge, they are sufficiently convincing overall to give ME the appearance of a gold standard, in the sense that no other known inference process satisfies an equally convincing set of principles. Other popular inference processes which satisfy many, but not all, of these principles are the minimum distance inference process, MD, the limit centre of mass process, CM^∞ , all **Renyi** inference processes, and the **Maximin** process of [6] ⁵. The Paris-Vencovská axiomatic characterisation of ME is particularly striking because it is quite independent of historically much earlier justifications of ME which stem either from ideas in statistical mechanics (see [7–9], or from axiomatic treatments of the concept of information itself (as in [10–12]). While both of the latter kinds of treatment are conceptually attractive it might be argued that they carry more philosophical baggage than does a purely axiomatic treatment of the desiderata to be satisfied by an abstract notion of inference process.

1.3. Pooling Operators

An apparently very different framework of probabilistic inference, this time in the multi-agent context, has been much studied in the decision theoretic literature. Given the set of possible atoms At as before, let $\{\mathbf{A}_i \mid i = 1 \dots m\}$ be a finite set of agents each of whom possesses her own particular probabilistic belief function $w^{(i)}$ on At , and let us suppose that these $w^{(i)}$ have already been determined. How then should these individual belief functions be aggregated so as to yield a single probabilistic belief function v which most accurately represents the *collective* beliefs of the agents? We call such an aggregated belief function a *social belief function*, and a general method of aggregation a *pooling operator*. Again we can ask: what principles should a pooling operator satisfy? In this framework various plausible principles have been investigated extensively in the literature, and have in particular been used to characterise two popular, but very different pooling operators **LinOp** and **LogOp**. **LinOp** takes v to be the arithmetic mean of the $w^{(i)}$

$$v_j = \frac{1}{m} \sum_{i=1}^m w_j^{(i)} \quad \text{for each } j = 1 \dots J$$

whereas **LogOp** chooses v to be the normalised geometric mean given by

$$v_j = \frac{(\prod_{i=1}^m w_j^{(i)})^{\frac{1}{m}}}{\sum_{k=1}^J (\prod_{i=1}^m w_k^{(i)})^{\frac{1}{m}}} \quad \text{for each } j = 1 \dots J$$

⁵ See Paris [3] for a general introduction to inference processes, and also Hawes [6], especially the comparative table in Chapter 9, for an excellent résumé of the current state of knowledge concerning this topic. Renyi inference processes are those which maximise one of the family of generalised notions of entropy due to Alfred Renyi (see [6,11,33,39]).

Various continua of other pooling operators related to **LinOp** and **LogOp** have also been investigated. However the existing axiomatic analysis of pooling operators, while technically simpler than the analysis of inference processes, is also more ambiguous and perhaps less intellectually satisfying in its conclusions than the analysis of inference processes developed within the Paris-Vencovská framework; in the former case one arrives at rival, apparently plausible, axiomatic characterisations of various pooling operators, including in particular **LinOp** and **LogOp**, without any very convincing foundational criteria for deciding, within the limited context of the framework, which operator is justified, if any ⁶. Strictly from a logician's point of view **LogOp** has by far the most attractive invariance properties of pooling operators which have been studied, but it has one major drawback from the perspective of decision theory or AI: it allows a single agent to have a completely disproportionate influence over the social belief function in the case when the agent's belief in some event is zero or close to zero. For this reason **LogOp** and its variants tend to be eschewed by decision theorists in favour of "softer" operators such as **LinOp**. We will argue in this paper that from a foundational point of view such pragmatism is misconceived. The solution to the conundrum lies rather in a deeper analysis of the semantics underlying the notion of a pooling operator. By embedding the concept of a pooling operator in the broader framework of social inference processes, we are able to see where the problem lies, and the outlines of possible solutions, a matter to which we return to in our concluding chapter.

1.4. The Multi-agent Problematic

In the present paper we seek to extend the Paris-Vencovská notion of inference process to the multi-agent case, thereby encompassing both the Paris-Vencovská framework of inference processes and the framework of pooling operators as special, or *marginal*, cases. To this end we consider, for any $m \geq 1$, a set \mathbf{M} consisting of m individuals $\mathbf{A}_1 \dots \mathbf{A}_m$, each of whom possesses her own nice set of constraints, respectively $\mathbf{K}_1 \dots \mathbf{K}_m$, on her possible belief function on the set of outcomes $\{\alpha_1, \alpha_2, \dots, \alpha_J\}$. (Note that we are only assuming here that the beliefs of each individual are consistent, not that the beliefs of different individuals are jointly consistent). We shall refer to such a set \mathbf{M} of individuals as a *college*. The intuitive problem now is how the college \mathbf{M} should choose a single belief function which best represents the totality of information conveyed by $\mathbf{K}_1 \dots \mathbf{K}_m$.

Definition 1. Let \mathcal{C} denote a given fixed class of constraints sets. A **social inference process for \mathcal{C}** is a function, \mathfrak{F} , which chooses, for any $m \geq 1$ and constraint sets $\mathbf{K}_1 \dots \mathbf{K}_m \in \mathcal{C}$, a probability function on \mathbf{At} , denoted by $\mathfrak{F}(\mathbf{K}_1 \dots \mathbf{K}_m)$, which we refer to as the **social belief function** defined by \mathfrak{F} acting on $\mathbf{K}_1 \dots \mathbf{K}_m$.

When considering general properties of unspecified social inference processes, we may not specify exactly what the class \mathcal{C} is, but in general we shall always assume that \mathcal{C} is a class of nice constraint sets.

Note that, trivially, provided that when $m = 1$ $\mathfrak{F}(\mathbf{K}) \in \mathbf{V}_{\mathbf{K}}$ for all $\mathbf{K} \in \mathcal{C}$, \mathfrak{F} marginalises to an inference process. On the other hand, in the special case where $\mathbf{K}_1 \dots \mathbf{K}_m$ are such that $\mathbf{V}_{\mathbf{K}_i}$ is a

⁶ See [4,5,20,35,40–45] for further discussion of the axiomatics of pooling operators.

singleton for all $i = 1 \dots m$, then \mathfrak{F} marginalises to a pooling operator. The new framework therefore encompasses naturally as special cases the two classical frameworks described in sections 1.2 and 1.3 above.

Again we can ask: what principles would we wish such a social inference process \mathfrak{F} to satisfy *in the absence of any further information*? Is there any social inference process \mathfrak{F} which satisfies them? If so, to which inference process and to which pooling operator does such an \mathfrak{F} marginalise? It turns out that merely by posing these questions in the right framework, and by making certain simple mathematical observations, we can gain considerable insight.

2. An Axiomatic Framework for a Social Inference Process

2.1. Background Heuristics: Rational Norms for Collective Probabilistic Reasoning

Our approach to multi-agent probabilistic reasoning is both rational and normative: we are concerned with how an independent external chairman of a college of agents *should by some objective process* aggregate the probabilistic information declared to her by members of the college into an optimal single belief function, on the assumption that the chairman herself has no other information about the agents than that which they declare. However, in order to place ourselves in a position to formulate rational criteria for such a process to satisfy, we are compelled to make certain idealising assumptions analogous to those made in the classical treatment of inference processes in [2,3,13], but with a somewhat more complex analysis owing to the multi-agent context. We present three such assumptions in subsections 2.1.1, 2.1.2 and 2.1.3 below. The first two assumptions are close to those made in the classical framework of inference processes, but the third assumption is specific to the multi-agent context.

We stress that our approach in this paper is *strictly foundational*. We insist on the importance of the qualification above that the chairman, with whose viewpoint we identify, is given no further information than that stated in the problem. In particular the chairman knows nothing about the expertise or reliability of the agents, or about the independence of their opinions. Nor will we be concerned with limitations on computability. However the very fact that we make the qualification above forces us to clarify more precisely the idealising assumptions which the chairman must make about her relationship to the information provided by the agents.

In spite of the fact that the idealising assumptions are unrealistic in practice, in line with Chomsky's criticism of the dominant methodology in artificial intelligence [14], we believe that this is the correct initial foundational approach if a *general* theory of multi-agent probabilistic inference is to have any chance of success. One should start from the simplest theoretical problematic; only when one has understood such a simple case does it make sense to try to deepen our understanding by progressively introducing other factors which make the problematic more realistic. In the present context, examples of such second level factors may be limitations on computational complexity, or the level of trust which we assign to the information conveyed by particular agents. The implications of taking into account the question of trust are discussed briefly in Chapter 5.1.

2.1.1. The Total Evidence Principle

Of crucial importance in our general problematic is the assumption above that *all* the relevant communicable probabilistic knowledge of an individual agent is incorporated in the given formal representation \mathbf{K} of her probabilistic knowledge base. This or a similar assumption is sometimes referred to as the Principle of Total Evidence ⁷. As was pointed out forcefully by Jaynes in his work justifying the use of maximum entropy inference, in order to avoid hopeless confusion, it is essential that an assumption of this kind be studiously respected in any formal study of the general axiomatic or logical characteristics of a mode of probabilistic inference: otherwise the intrinsic meaning of a formalised problem can be surreptitiously changed by sleight of hand, resulting in the generation of an inexhaustible supply of phony paradoxes or inconsistencies (cf. [7,8,15,16]). However as pointed out by Adamčík and the author in [17] the practical exigencies demanded in the study of particular probabilistic problems arising from the real world have tended to result in a lack of attention being paid to more foundational studies which would require the total evidence principle to be taken seriously:

“...when applied to the formalisation of any real life problem considered by a human agent, the Principle of Total Evidence is never observed in practice. This banal fact of life has historically bedevilled theoretical discussion of probabilistic inference, because it is often extremely hard to give any real world example to illustrate an abstract principle of probabilistic inference without an opponent being tempted to challenge one’s reasoning using implicit or intuitive background information concerning the example, which has not been included in its formal representation. In the context of multi-agent probabilistic inference this situation has resulted in a heavy concentration of research on computationally pragmatic approaches to specialised problems of probabilistic inference, and a notable neglect of the study of more abstract axiomatic or foundational frameworks. This neglect appears to the authors to be unfortunate, not least because the foundations of artificial intelligence would seem to demand that the Principle of Total Evidence be taken seriously.”

Note that the Total Evidence Principle is also assumed in the justification of the classical inference process framework of [3].

2.1.2. Assumption of Logical and Computational Closure

We assume that there are no restrictions on the ability of individual agents to calculate the probabilistic consequences of any given constraint set \mathbf{K} , and that consequently there is no essential semantic difference between the status of the probabilistic knowledge represented by \mathbf{K} and that represented by its representation $\mathbf{V}_{\mathbf{K}}$ in Euclidean space. Consequently if \mathbf{K} and \mathbf{K}' are constraint sets such that $\mathbf{V}_{\mathbf{K}} = \mathbf{V}_{\mathbf{K}'}$ we shall regard them as equivalent knowledge bases from the point of view of any agent. Under this assumption we may therefore informally identify an agent’s knowledge base \mathbf{K} with its representation $\mathbf{V}_{\mathbf{K}}$. Notice that under this assumption an agent will be aware whether or not a set of

⁷ This terminology is due to Carnap [46]. The principle is also known as Bernoulli’s Maxim, while in [3] it is called the Watts Assumption.

constraints is consistent, and from this point of view our previously stated requirement that a knowledge base be consistent seems reasonable. Of course, as is well known, unaided individual *human* agents' assessments are notoriously inconsistent in practice [18], and furthermore if we assume that $P \neq NP$ then the calculations which are required for the present assumption are in general infeasible (cf. Chapter 10 of [3]). Nevertheless, as in the case of inference processes, this does not diminish the value of our assumption as a normative tool.

2.1.3. The Intersubjectivity Assumption

For the rest of this paper we will assume for ease of exposition that the college \mathbf{M} appoints an independent chairman \mathbf{A}_0 , whom we may suppose to be a mathematically trained philosopher, and whose only task is to aggregate the knowledge bases of the agents in the college into a social belief function v according to strictly rational criteria, but ignoring any personal beliefs which \mathbf{A}_0 herself may hold.

The Intersubjectivity Assumption states that in performing the function above the chairman treats the knowledge base provided by each agent *as if* it represented intersubjective probabilistic information⁸. By this we mean that whatever the unknown background observations or introspections might be from which any particular agent \mathbf{A}_i 's knowledge base \mathbf{K}_i arises, the *process* by which \mathbf{K}_i arose is assumed to be in conformity with the laws of probability, and intersubjective in the sense that any other agent with exactly the same background information and experience as, say, \mathbf{A}_i would arrive at a set of constraints equivalent to \mathbf{K}_i . The fact that the union of the knowledge bases \mathbf{K}_1 and \mathbf{K}_2 of respective agents \mathbf{A}_1 and \mathbf{A}_2 may be inconsistent does not in any way contradict this assumption, since the limited background observations or introspections of each agent may be different, which may result in the agents' probabilistic assessments being incompatible. While the intersubjectivity assumption might seem grossly unrealistic, particularly in the case of human agents, it is nevertheless a valuable idealisation, not least because it helps to identify which features of an agent's possible relation to her reported information we are *not* taking into account.

The information which the agents report is thus taken at face value and treated as if it were totally trustworthy by chairman \mathbf{A}_0 , even though the chairman may recognise that such trust is not merited. While this fact of itself indicates one of the principal limitations of the initial framework of rational collective reasoning which we are attempting to formulate, as we will outline in Chapter 5 it also suggests natural ways in which a notion of *degree of trust* could later be incorporated into the framework, thus mitigating the effects of this limitation. Incorporating such a notion would allow a social inference process to accommodate the less than complete trust which a chairman might actually hold in the information provided by individual agents.

2.2. Towards a Framework for Rational Collective Probabilistic Reasoning

⁸ For an interesting account of intersubjective probability see Gillies [47].

While particular examples of special social inference processes can be found in many places in the literature (see e.g. [19–25]), the abstract idea of a social inference process was first formulated in [1] and has been recently studied further in [17,26,27] where the properties of a number of different social inference processes are considered. However most of the earlier work published on particular social inference processes has, with few exceptions, been pragmatically motivated, and has not considered broader foundational questions or logical justifications. This is due in some cases to a concern to find a computationally practical solution to a more specialised problem, and in other cases to a tempting reductionism, which would see the problem of finding a social inference process as a two stage process in which a favored classical inference process \mathcal{I} is first chosen and applied to the constraints \mathbf{K}_i of each agent i to yield a belief function $w^{(i)}$ appropriate to that agent, and a preferred pooling operator is then applied to the set of $w^{(i)}$ to yield a social belief function. Following the terminology of Adamčík [26] we shall call a social inference process which has this special reductionist form *obdurate*. Of course from a reductionist point of view the concept of a social inference process is not particularly interesting foundationally, since we could hardly expect an analysis of such social inference processes to tell us anything fundamentally new about collective probabilistic reasoning⁹. A notable exception to such approaches is found in the work of Williamson [28], which offers a detailed philosophical analysis of the principles underlying the merging of probabilistic evidence from an objective Bayesian perspective, which is not reductionist in the sense above, but which is somewhat different from our own¹⁰.

Our approach here is radically non-reductionist. We reject the two stage approach above on the grounds that the classical notion of an inference process applies to an isolated single individual, and is valid only on the assumption that that individual has absolutely no knowledge or beliefs other than those specified by her personal constraint set. Indeed the preliminary point should be made that in the case of an *isolated* individual A , whereas A 's constraint set \mathbf{K} is subjective and personal to that individual, the *actual passage* from \mathbf{K} to A 's assumed belief function w via an inference process should be made using rational or normative principles, and should therefore be considered to have an *objective* character. Nor should we confuse the epistemological status of w with that of \mathbf{K} . By hypothesis \mathbf{K} represents the sum total of A 's beliefs; *ipso facto* \mathbf{K} also represents, in general, a description of the extent of A 's *ignorance*. While w may be regarded as the *belief function which best represents* A 's subjective beliefs, it must not be confused with those beliefs themselves, since in the passage from \mathbf{K} to w it is clear that

⁹ An interesting example of a social inference process which is *de facto* obdurate, but whose initial definition does not appear intentionally reductionist, is that given by Kern-Isberner and Rödder in [23]. In effect the social inference process which they define applies the ME inference process to each \mathbf{K}_i and then applies a weighted arithmetic mean pooling operator to the resulting points, where the weights are proportional to the exponential entropies of the respective points. See also Adamčík [26] for an account of which principles are satisfied by this social inference process.

¹⁰ Technically Williamson is concerned in [28] with the question as to how a single agent should rationally arrive at a unique belief function on the basis of probabilistic evidence which is derived from different “sources” which, taken together, may be inconsistent. Williamson’s proposed solution is to consider the convex hull of the regions of \mathbb{D}_J defined by considering maximal consistent unions of the sets of constraints corresponding to the individual sources, and then to choose the maximum entropy solution from the resulting convex region. Of course if we treat the different sources of the agent’s evidence as themselves being agents, then this procedure in effect defines a social inference process. See also [48] and [26].

certain “information” has been discarded¹¹; thus, while w is determined by \mathbf{K} once an inference process is given and applied, neither \mathbf{K} nor $\mathbf{V}_{\mathbf{K}}$ can be recaptured from w . As a trivial example we may note that specifying that \mathbf{A} 's constraint set \mathbf{K} is empty, *i.e.*, that \mathbf{A} claims total ignorance, is informationally very different from specifying that \mathbf{K} is such that $\mathbf{V}_{\mathbf{K}} = \{ \langle \frac{1}{J}, \frac{1}{J} \dots \frac{1}{J} \rangle \}$, although the application of ME, or of any other reasonable inference process, yields $w = \langle \frac{1}{J}, \frac{1}{J} \dots \frac{1}{J} \rangle$ in both cases. This example of an agent who is totally ignorant has an illustrative force which we return to later.

From this point of view the situation of an individual who is a member of a college whose members seek to elicit an optimal “social” belief function to best represent the belief of the collective seems quite different from that of an isolated individual. Indeed in the collective context it appears more natural to assume as a normative principle that, if the social belief function is to be optimal, then each individual member \mathbf{A}_i *should be deemed* to choose her personal belief function $w^{(i)}$ so as to take account of the information provided by the other individuals, in such a way that $w^{(i)}$ is *consistent* with her own knowledge base \mathbf{K}_i , while being informationally as close as possible to the social belief function $\mathfrak{F}(\mathbf{K}_1 \dots \mathbf{K}_m)$ which is to be defined. We will show in chapter 3 that this suggestive, but imprecise, idea can be made mathematically coherent, and can be used to define a particular social inference process with pleasing properties. Notice however that it is not necessary to assume that a given \mathbf{A}_i subjectively or consciously holds the particular personal belief function $w^{(i)}$ which is attributed to her by the procedure above: such an $w^{(i)}$ is viewed as nothing more than the belief function which \mathbf{A}_i *ought rationally to hold*, given the personal knowledge base \mathbf{K}_i which represents her own beliefs, together with the *extra information* which would be available to her if she were to be made aware of the knowledge bases of the remaining members of the college. Just as in the case of an isolated individual, the passage from \mathbf{A}_i 's actual subjective belief set \mathbf{K}_i to her notional subjective belief function $w^{(i)}$ has an intersubjective or normative character: however the calculation of $w^{(i)}$ now depends not only on \mathbf{K}_i but on the knowledge bases of all the other members of the college.

Considerations similar to the above also give rise to an important general principle which we believe a social inference process should satisfy, which we will call Collegiality. In the next section we shall introduce this principle together with some other desiderata for a social inference process to satisfy. The latter are either natural symmetry principles or fairly straightforward generalisations of familiar desiderata from the Paris-Vencovská framework of inference processes.

2.3. Desiderata for a Social Inference Process

The Equivalence Principle

If for all $i = 1 \dots m$ $\mathbf{V}_{\mathbf{K}_i} = \mathbf{V}_{\mathbf{K}'_i}$ then

$$\mathfrak{F}(\mathbf{K}_1 \dots \mathbf{K}_m) = \mathfrak{F}(\mathbf{K}'_1 \dots \mathbf{K}'_m)$$

¹¹ The word “information” is clearly used here in a different sense from that of Shannon information; it would perhaps be more accurate to say that what is being discarded is in this case is information about the extent of \mathbf{A} 's ignorance.

□

Otherwise expressed the Equivalence Principle states that substituting constraint sets which are equivalent in the sense that the set of belief functions which satisfy them is unchanged will leave the values of \mathfrak{F} invariant. This principle is a familiar one adopted from the theory of inference processes (cf. [3]), and is in line with our assumption in section 2.1.2. In this paper we shall always consider only social inference processes (or inference processes) which satisfy the Equivalence Principle. For this reason we may occasionally allow a certain sloppiness of notation in the sequel by identifying a constraint set \mathbf{K} with its set of solutions $\mathbf{V}_{\mathbf{K}}$ where the meaning is clear and this avoids an awkward notation. In particular if Δ is a non-empty closed convex set of belief functions then we may write $\text{ME}(\Delta)$ to denote the unique $w \in \Delta$ which maximises the Shannon entropy function.

The Anonymity Principle

For any permutation σ of $1, \dots, m$

$$\mathfrak{F}(\mathbf{K}_{\sigma(1)} \dots \mathbf{K}_{\sigma(m)}) = \mathfrak{F}(\mathbf{K}_1 \dots \mathbf{K}_m)$$

□

A consequence of the above principle is that $\mathfrak{F}(\mathbf{K}_1 \dots \mathbf{K}_m)$ depends only on the multiset of knowledge bases $\{\mathbf{K}_1 \dots \mathbf{K}_m\}$ and not on the order in which the \mathbf{K}_i 's are listed.

The following natural principle ensures that \mathfrak{F} does not choose a belief function which violates the beliefs of some member of the college unless there is no alternative. The principle also ensures that \mathfrak{F} behaves like a classical inference process in the special case when $m = 1$.

The Consistency Principle

If $\mathbf{K}_1 \dots \mathbf{K}_m$ are such that

$$\bigcap_{i=1}^m \mathbf{V}_{\mathbf{K}_i} \neq \emptyset$$

then

$$\mathfrak{F}(\mathbf{K}_1 \dots \mathbf{K}_m) \in \bigcap_{i=1}^m \mathbf{V}_{\mathbf{K}_i}$$

□

Let σ denote a permutation of the atoms of At . Such a σ induces a corresponding permutation on the coordinates of probability distributions $\langle w_1 \dots w_J \rangle$, and on the corresponding coordinates of variables occurring in the constraints of constraint sets \mathbf{K}_i , which we denote below with an obvious notation. The following principle is again a familiar one satisfied by classical inference processes (see [3]):

The Atomic Renaming Principle

For any permutation σ of the atoms of \mathbf{At} , and for all $\mathbf{K}_1 \dots \mathbf{K}_m$

$$\mathfrak{F}(\sigma(\mathbf{K}_1) \dots \sigma(\mathbf{K}_m)) = \sigma(\mathfrak{F}(\mathbf{K}_1 \dots \mathbf{K}_m))$$

□

The following principle is characteristic of the non-reductionist approach which we described in section 2.2:

The Collegiality Principle

A social inference process \mathfrak{F} satisfies the Collegiality Principle (abbreviated to *Collegiality*) if for any $m \geq 1$ and $\mathbf{A}_1 \dots \mathbf{A}_m$ with respective knowledge bases $\mathbf{K}_1 \dots \mathbf{K}_m$, if for some $k < m$ $\mathfrak{F}(\mathbf{K}_1 \dots \mathbf{K}_k)$ is consistent with $\mathbf{K}_{k+1} \cup \mathbf{K}_{k+2} \cup \dots \cup \mathbf{K}_m$, then

$$\mathfrak{F}(\mathbf{K}_1 \dots \mathbf{K}_m) = \mathfrak{F}(\mathbf{K}_1 \dots \mathbf{K}_k)$$

□

Collegiality may be interpreted as stating the following: if the social belief function v generated by some subset of the college is consistent with the individual beliefs of the remaining members, then v is also the social belief function of the whole college. The following immediate consequence of collegiality is worth a special mention:

Corollary 2 (The Ignorance Principle). For any $m \geq 1$ and all knowledge bases $\mathbf{K}_1 \dots \mathbf{K}_m$

$$\mathfrak{F}(\mathbf{K}_1 \dots \mathbf{K}_m) = \mathfrak{F}(\mathbf{K}_1 \dots \mathbf{K}_m, \emptyset)$$

where \emptyset denotes the knowledge base with empty set of constraints.

Proof. This follows at once from the collegiality principle. □

The ignorance principle just states that adding to the college a new agent who declares that she has no probabilistic knowledge concerning \mathbf{At} will leave the social belief function unchanged. The ignorance principle is of interest firstly because it seems particularly hard to challenge, and secondly because it seems to encapsulate the essence of the difference in information between an agent asserting that she has empty knowledge base, and the same agent asserting that her knowledge base is $\alpha_1 = \alpha_2 = \dots = \alpha_J = \frac{1}{J}$. Indeed this observation leads to the conclusion that obdurate social inference processes have serious credibility problems, since any obdurate \mathfrak{F} which satisfies atomic renaming must either fail to satisfy

the ignorance principle or else must marginalise to a pooling operator with pathological behaviour. In particular the social inference process of Kern-Isberner and Rödder defined¹² in [23] is easily shown not to satisfy the ignorance principle. Furthermore in [29] Adamčík shows that a very large class of obdurate social inference processes, including that of [23] cannot satisfy the consistency principle either¹³.

The consistency and collegiality principles together immediately imply that \mathfrak{F} satisfies the following unanimity property:

Lemma 3 (Unanimity Principle). *If \mathfrak{F} satisfies Consistency and Collegiality then for any \mathbf{K}*

$$\mathfrak{F}(\mathbf{K} \dots \mathbf{K}) = \mathfrak{F}(\mathbf{K}).$$

Proof. Immediate from definitions. \square

Our next axiom goes to the heart of certain basic intuitions concerning probability. For expository reasons we will consider first the case when $m = 1$, in which case we are essentially discussing a principle to be satisfied by a classical inference process. First we introduce some fairly obvious terminology.

Let \mathbf{w} denote \mathbf{A}_1 's belief function. (Since we are considering the case when $m = 1$ we will drop the superscript from $\mathbf{w}^{(1)}$ for ease of notation). For some non-empty set of atoms $\{\alpha_{j_1} \dots \alpha_{j_t}\}$ let ϕ denote the event $\bigvee_{r=1}^t \alpha_{j_r}$. Suppose that \mathbf{K} denotes a set of constraints on the variables $w_{j_1} \dots w_{j_t}$ which defines a non-empty closed convex region of t -dimensional Euclidean space with $\sum_{r=1}^t w_{j_r} \leq 1$ and all $w_{j_r} \geq 0$. We shall refer to such a \mathbf{K} as a *nice set of constraints about ϕ* . Such a set of constraints \mathbf{K} may also be thought of as a constraint set on the \mathbf{w} which determines a closed convex region $\mathbf{V}_{\mathbf{K}}$ of \mathbb{D}_J defined by

$$\mathbf{V}_{\mathbf{K}} = \{\mathbf{w} \in \mathbb{D}_J \mid \langle w_{j_1} \dots w_{j_t} \rangle \text{ satisfies } \mathbf{K}\}.$$

Now let \hat{w}_r denote $\mathbf{w}(\alpha_{j_r} \mid \phi)$ for $r = 1 \dots t$, with the \hat{w}_r undefined if $\mathbf{w}(\phi) = 0$. Then $\hat{\mathbf{w}} = \langle \hat{w}_1 \dots \hat{w}_t \rangle$ is a probability distribution provided that $\mathbf{w}(\phi) \neq 0$. Let \mathbf{K} be a nice set of constraints on the probability distribution $\hat{\mathbf{w}}$: we shall refer to such a \mathbf{K} as a *nice set of constraints conditioned on ϕ* . In line with our previous conventions we shall consider such \mathbf{K} to be trivially satisfied in the case when $\mathbf{w}(\phi) = 0$.

Again an important point here is that while a nice set of constraints \mathbf{K} conditioned on ϕ as above is given as a set of constraints on $\hat{\mathbf{w}}$ it can equally well be interpreted as defining a certain equivalent set of constraints on \mathbf{w} instead, and it is easy to see that, with a slight abuse of notation, the corresponding region $\mathbf{V}_{\mathbf{K}}$ of \mathbb{D}_J defined by

$$\mathbf{V}_{\mathbf{K}} = \{\mathbf{w} \mid \hat{\mathbf{w}} \text{ satisfies } \mathbf{K}\}$$

is both convex and closed.

¹² See also footnote 9.

¹³ We should note that while from our point of view, collegiality appears a very natural principle, this fact depends heavily on our underlying assumptions; from the very different viewpoint of Williamson [48] collegiality is too strong a principle, and it is not even clear if Williamson would accept the ignorance principle.

In what follows we may regard both a nice set of constraints conditioned on some event ϕ , and a nice set of constraints about some event ϕ , as if they defined constraints on the probability function w , as explained above.

Notice that while a nice set of constraints *conditioned* on ϕ can say nothing about the value of belief in ϕ itself, a nice set of constraints *about* ϕ may do so, and may even fix belief in ϕ at a particular value.

The following principle captures a basic intuition about probabilistic reasoning which is valid for all standard inference processes:

The Locality Principle (for an Inference Process)

An inference process \mathcal{I} satisfies the locality principle if for all sentences ϕ and θ , every nice set of constraints \mathbf{K} *conditioned on* ϕ , and every nice set of constraints \mathbf{K}^* *about* $\neg\phi$,

$$\mathcal{I}(\mathbf{K} \cup \mathbf{K}^*)(\theta | \phi) = \mathcal{I}(\mathbf{K})(\theta | \phi)$$

provided that $\mathcal{I}(\mathbf{K} \cup \mathbf{K}^*)(\phi) \neq 0$ and $\mathcal{I}(\mathbf{K})(\phi) \neq 0$

□

Let us refer to the set of all events which logically imply the event ϕ as *the world of* ϕ . Then the Locality Principle may be roughly paraphrased as saying that if \mathbf{K} contains only information about the *relative size* of probabilistic beliefs about events in the world of ϕ , while \mathbf{K}^* contains only information about beliefs concerning events in the world of $\neg\phi$, then the values which the inference process \mathcal{I} calculates for probabilities of events conditioned on ϕ should be unaffected by the information in \mathbf{K}^* , except in the trivial case when belief in ϕ is forced to take the value 0. Put rather more more succinctly: beliefs about the world of $\neg\phi$ should not affect beliefs conditioned on ϕ . Note that we cannot expect to satisfy a strengthened version of this principle which would have belief in the events in the world of ϕ unaffected by \mathbf{K}^* since the constraints in \mathbf{K}^* may well affect belief in ϕ itself. Thus the Locality Principle asserts that, *ceteris paribus*, rationally derived *relative probabilities between events inside a “world”* are unaffected by information about what happens strictly outside that world.

The Locality Principle is in essence a combination of both the Relativisation Principle¹⁴ of Paris [3] and the Homogeneity Axiom of Hawes [6]. The following theorem, which demonstrates that the most commonly accepted inference processes all satisfy Locality, is very similar to results proved previously, especially to results in [6]. It follows from the theorem below that if we reject the Locality Principle for an inference process, then we are in effect forced to reject not just ME, but also all currently known plausible inference processes, including all inference processes derived by maximising a generalised

¹⁴ We note here that Csiszár in [49], [50], introduces a property which he calls locality, but which corresponds to the relativisation principle of Paris [3] and is much weaker than the notion of locality in the present paper.

notion of entropy. This is an important point heuristically when we come to extend the Locality Principle to the multi-agent case ¹⁵.

Theorem 4. *The inferences processes ME, CM[∞], MD (minimum distance), together with all Renyi inference processes ¹⁶, and the Maximin inference process of [6], all satisfy the Locality Principle.*

Proof. Let F be a real valued function defined on the domain

$$\bigcup_{J \in \mathbb{N}^+} [0, 1]^J$$

by

$$F(\mathbf{w}) = \sum_{j=1}^J f(w_j) \tag{1}$$

for some function $f : [0, 1] \rightarrow \mathbb{R}$.

We will say that F is *deflation proof* if for every $J \in \mathbb{N}^+$, all $\mathbf{w}, \mathbf{u} \in \mathbb{D}_J$, and every $\lambda \in (0, 1)$

$$F(\lambda \mathbf{w}) < F(\lambda \mathbf{v}) \text{ if and only if } F(\mathbf{w}) < F(\mathbf{v}) \tag{2}$$

Here $\lambda \mathbf{w}$ denotes the scalar multiplication of \mathbf{w} by λ . Note that $\lambda \mathbf{w}$ will not be a vector in \mathbb{D}_J in the above case since its coordinates sum to λ instead of 1.

We will see below that any inference process \mathcal{I} such that $\mathcal{I}(\mathbf{K})$ is defined to be that point $\mathbf{v} \in \mathbf{V}_{\mathbf{K}}$ which maximises a strictly convex deflation proof function F of the above form satisfies the locality principle.

We first note the following lemma:

Lemma 5. *The inference processes listed in the statement of Theorem 4, with the exception of CM[∞] and Maximin, may all be defined by the maximisation of deflation proof strictly convex functions of the form (1) above.*

Proof. The inference process ME is defined by maximising

$$F(\mathbf{w}) = - \sum_{j=1}^J w_j \log w_j$$

subject to the given constraints. Now for $\mathbf{w} \in \mathbb{D}_J$

$$F(\lambda \mathbf{w}) = - \sum_{j=1}^J \lambda w_j \log \lambda w_j = - \lambda \log \lambda + \lambda F(\mathbf{w})$$

¹⁵ The proof of Theorem 4 is not however germane to understanding the remainder of this paper and may safely be skipped if the reader so wishes.

¹⁶ See footnote 5. A definition of Renyi processes is given below in the proof of Theorem 4.

from which (2) follows at once.

The Renyi inference process \mathbf{REN}_r , where r is a fixed positive real parameter not equal to 1, is given by maximising the function

$$F(\mathbf{w}) = - \sum_{j=1}^J (w_j)^r$$

for $\mathbf{w} \in \mathbf{V}_K$ in the case when $r > 1$, and by maximising

$$F(\mathbf{w}) = \sum_{j=1}^J (w_j)^r$$

for $\mathbf{w} \in \mathbf{V}_K$ in the case when $0 < r < 1$.

Since for the above functions $F(\lambda\mathbf{w}) = \lambda^r F(\mathbf{w})$, they also trivially satisfy (2) and so are deflation proof. Note that the minimum distance inference process MD is just \mathbf{REN}_2 . The functions F defined above are all strictly convex (see e.g. [3]) and so the lemma follows. \square

Returning to the main proof, let \mathcal{I} be an inference process such that $\mathcal{I}(\mathbf{K})$ is defined by the maximisation of a deflation proof strictly convex function F of the form as in (1) above. Let ϕ , \mathbf{K} , and \mathbf{K}^* be as in the statement of the locality principle. Without loss of generality we may assume for notational convenience that the atoms are so ordered that for some k with $1 \leq k < J$

$$\phi \equiv \bigvee_{j=1}^k \alpha_j \quad \text{and} \quad \neg\phi \equiv \bigvee_{j=k+1}^J \alpha_j$$

Let $\mathbf{u} = \mathcal{I}(\mathbf{K})$ and let $\mathbf{v} = \mathcal{I}(\mathbf{K} \cup \mathbf{K}^*)$. Let $\mathbf{u}(\phi) = a$ and let $\mathbf{v}(\phi) = b$. By hypothesis we know that a and b are non-zero. It suffices for us to show that

$$\left\langle \frac{v_1}{b} \dots \frac{v_k}{b} \right\rangle = \left\langle \frac{u_1}{a} \dots \frac{u_k}{a} \right\rangle \tag{3}$$

Now notice that since the constraints of \mathbf{K}^* refer only to coordinates $k + 1 \dots J$ while the constraints of \mathbf{K} refer only to coordinates $1 \dots k$, the solution \mathbf{v} which by definition maximises $\sum_{j=1}^J f(w_j)$ subject to the condition that $\mathbf{w} \in \mathbf{V}_{\mathbf{K} \cup \mathbf{K}^*}$, must also satisfy the condition that $\langle v_1 \dots v_k \rangle$ is that vector $\langle w_1 \dots w_k \rangle$ which maximises $\sum_{j=1}^k f(w_j)$ subject to $\langle \frac{w_1}{b} \dots \frac{w_k}{b} \rangle$ satisfying the constraints of \mathbf{K} together with the constraint that $\sum_{j=1}^k w_j = b$. Now changing variables by setting $y_j = \frac{w_j}{b}$, with $\mathbf{y} = \langle y_1 \dots y_k \rangle$ this is equivalent to maximising

$$F(b\mathbf{y}) = \sum_{j=1}^k f(by_j)$$

subject to $\mathbf{y} \in \mathbb{D}_k$ and \mathbf{y} satisfying the constraints of \mathbf{K} . However since F is deflation proof (and strictly convex) the unique $\mathbf{y} \in \mathbb{D}_k$ which achieves this maximisation does not depend on b and by setting $b = 1$ we see that it is just the unique vector $\mathbf{y} \in \mathbb{D}_k$ maximising $F(\mathbf{y})$ and satisfying the constraints in \mathbf{K} . Since this definition is independent of both \mathbf{K}^* and b , it follows by replacing \mathbf{K}^*

by the empty set of constraints and b by a that equation (3) holds, which completes the proof for the case of inference processes defined by the maximisation of a deflation proof strictly convex function of the form (1) above. By lemma 5 the theorem follows for all the inference processes mentioned except for CM^∞ and **Maximin**.

The fact that the limit centre of mass inference process, CM^∞ , satisfies locality may either be proved using the standard definition of CM^∞ in [3], and slightly modifying the idea of the proof above, or simply by observing that by a result of Hawes [6], for any knowledge base \mathbf{K}

$$CM^\infty(\mathbf{K}) = \lim_{r \rightarrow 0^+} REN_r(\mathbf{K})$$

and then applying the results above already proved for REN_r .

The result for **Maximin** also follows easily from results in Hawes [6]. This completes the proof of Theorem 4. \square

While Theorem 4 above merely provides very strong corroborating evidence in favour of accepting the Locality Principle for an inference process, an interesting aspect of the *intuition* underlying the principle is that the justification for it appears no less cogent when we attempt to generalise it to the context of a social inference process. If we accept the intuition in favour of the Locality Principle in the case of a single individual then it is hard to see why we should reject analogous arguments in the case of a social belief function which is derived by considering the beliefs of m individuals each of whom has knowledge bases of the type considered above. The argument is a general informational one: if information about probabilities conditioned on ϕ is unaffected by information about the world of $\neg\phi$, then, *ceteris paribus*, this should be true regardless of whether the information is obtained from one agent or from many agents. Accordingly we may formulate more generally

The General Locality Principle (for a social inference process \mathfrak{F})

For any $m \geq 1$ let \mathbf{M} be a college of m individuals $\mathbf{A}_1 \dots \mathbf{A}_m$. If for each $i = 1 \dots m$ \mathbf{K}_i is a nice set of constraints conditioned on ϕ , and \mathbf{K}_i^* is a nice set of constraints about $\neg\phi$, then for every event θ

$$\mathfrak{F}(\mathbf{K}_1 \cup \mathbf{K}_1^*, \dots, \mathbf{K}_m \cup \mathbf{K}_m^*)(\theta | \phi) = \mathfrak{F}(\mathbf{K}_1, \dots, \mathbf{K}_m)(\theta | \phi)$$

provided that $\mathfrak{F}(\mathbf{K}_1 \cup \mathbf{K}_1^*, \dots, \mathbf{K}_m \cup \mathbf{K}_m^*)(\phi) \neq 0$

and $\mathfrak{F}(\mathbf{K}_1, \dots, \mathbf{K}_m)(\phi) \neq 0$. \square

At this point we make a simple observation. In the very special marginal case when for each i the knowledge bases $\mathbf{K}_i \cup \mathbf{K}_i^*$ are such as to completely determine \mathbf{A}_i 's belief function, so that the task of \mathfrak{F} reduces to that of a pooling operator, the locality principle above reduces to a condition closely related to the well-known condition on a pooling operator that it be *externally Bayesian*¹⁷. We will not discuss

¹⁷ This condition was first formulated by Madansky [51] in 1964 and further analyzed in [52] and [43], where it is shown that the only externally Bayesian pooling operators are closely related to **LogOp**. See also [4] for related properties of pooling operators.

this further here except to note the important point that if \mathfrak{F} is taken to satisfy General Locality, then this fact alone seriously restricts those pooling operators to which it is possible for \mathfrak{F} to marginalise. Thus while **LogOp** satisfies the relevant cases of General Locality, as follows from Theorem 14 below, the popular pooling operator **LinOp** does not do so. The following provides a simple counterexample:

Example 1 (Counterexample to General Locality for **LinOp**).

Proof. Let $J = 3$, let $\theta = \alpha_1 \vee \alpha_2$ and let

$$\mathbf{K}_1 = \{ w(\alpha_1 | \theta) = \frac{2}{3} \} \text{ and } \mathbf{K}_1^* = \{ w(-\theta) = \frac{1}{4} \}$$

$$\mathbf{K}_2 = \{ w(\alpha_1 | \theta) = \frac{1}{3} \} \text{ and } \mathbf{K}_2^* = \{ w(-\theta) = \frac{5}{6} \}$$

Then the unique belief function satisfying $\mathbf{K}_1 \cup \mathbf{K}_1^*$ is $\mathbf{w}^{(1)} = \langle \frac{1}{2}, \frac{1}{4}, \frac{1}{4} \rangle$ while the unique belief function satisfying $\mathbf{K}_2 \cup \mathbf{K}_2^*$ is $\mathbf{w}^{(2)} = \langle \frac{1}{18}, \frac{1}{9}, \frac{5}{6} \rangle$.

Applying **LinOp** we obtain

$$\mathbf{LinOp}(\mathbf{K}_1 \cup \mathbf{K}_1^*, \mathbf{K}_2 \cup \mathbf{K}_2^*)(\alpha_1 | \theta) = \frac{20}{33}$$

If we now set

$$\mathbf{K}_1^{**} = \{ w(-\theta) = \frac{3}{4} \} \text{ and } \mathbf{K}_2^{**} = \{ w(-\theta) = \frac{1}{2} \}$$

then the unique belief function satisfying $\mathbf{K}_1 \cup \mathbf{K}_1^{**}$ is $\mathbf{w}^{(1)} = \langle \frac{1}{6}, \frac{1}{12}, \frac{3}{4} \rangle$ while the unique belief function satisfying $\mathbf{K}_2 \cup \mathbf{K}_2^{**}$ is $\mathbf{w}^{(2)} = \langle \frac{1}{6}, \frac{1}{3}, \frac{1}{2} \rangle$.

Applying **LinOp** gives

$$\mathbf{LinOp}(\mathbf{K}_1 \cup \mathbf{K}_1^{**}, \mathbf{K}_2 \cup \mathbf{K}_2^{**})(\alpha_1 | \theta) = \frac{4}{9} \neq \mathbf{LinOp}(\mathbf{K}_1 \cup \mathbf{K}_1^*, \mathbf{K}_2 \cup \mathbf{K}_2^*)(\alpha_1 | \theta)$$

showing that General Locality fails for any \mathfrak{F} which marginalises to the pooling operator **LinOp**. By contrast it is easily verified that

$$\mathbf{LogOp}(\mathbf{K}_1 \cup \mathbf{K}_1^{**}, \mathbf{K}_2 \cup \mathbf{K}_2^{**})(\alpha_1 | \theta) = \frac{1}{2} = \mathbf{LogOp}(\mathbf{K}_1 \cup \mathbf{K}_1^*, \mathbf{K}_2 \cup \mathbf{K}_2^*)(\alpha_1 | \theta)$$

as expected. \square

Related facts concerning **LinOp** and **LogOp** have been widely noted in the literature on pooling operators; what we are noting that is new here is that arguments in favour of the General Locality Principle in the far broader context of a social inference process give a quite new perspective on the relative acceptability of classical pooling operators such as **LogOp** and **LinOp**.

Our final axiom relates to a hypothetical situation where several exact copies of a college are amalgamated into a single college.

A *clone* of a member \mathbf{A}_i of \mathbf{M} is a member $\mathbf{A}_{i'}$ whose set of belief constraints on her belief function is identical to that of \mathbf{A}_i : i.e., $\mathbf{K}_i = \mathbf{K}_{i'}$. Suppose now that each member \mathbf{A}_i of \mathbf{M} is replaced by n clones

of \mathbf{A}_i , so that we obtain a new college \mathbf{M}^* with nm members. \mathbf{M}^* may equally be regarded as k copies of \mathbf{M} amalgamated into a single college; so since the social belief function associated with each of these copies of \mathbf{M} would be the same, we may argue that surely the result of amalgamating the copies into a single college \mathbf{M}^* should again yield the same social belief function.

For any knowledge base \mathbf{K} let $n\mathbf{K}$ stand for a sequence of n copies of \mathbf{K} . Then the heuristic argument above generates the following:

The Proportionality Principle

For any integer $n \geq 1$

$$\mathfrak{F}(n\mathbf{K}_1, n\mathbf{K}_2, \dots, n\mathbf{K}_m) = \mathfrak{F}(\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_m)$$

□

Notice that for the single agent case $m = 1$ this principle reduces to the Unanimity Principle 3. The Proportionality Principle looks rather innocent. Nevertheless we shall see in Theorem 17 of chapter 4 that a slight variant of the same idea, formulated as a limiting version, has some unexpected consequences.

3. The Social Entropy Process SEP

3.1. Definition of SEP

In this section we introduce a natural social inference process, **SEP**, which extends both the inference process **ME** and the pooling operator **LogOp**. Our heuristic derivation of **SEP** will be purely information theoretic. We prove certain important structural properties necessary to show that **SEP** is well-defined, and we show in Theorem 14 that **SEP** satisfies the seven principles introduced in the previous section.

In order to avoid problems with our definition of **SEP** however, we are forced to add a slight further restriction to the set of m knowledge bases $\mathbf{K}_1 \dots \mathbf{K}_m$ which respectively represent the beliefs sets of the individuals $\mathbf{A}_1 \dots \mathbf{A}_m$. We assume in this section that the constraints are such that there exists at least one atom α_{j_0} such that no knowledge base \mathbf{K}_i forces α_{j_0} to take belief 0. In the special case when each \mathbf{K}_i specifies a unique probability distribution, the condition corresponds to that necessary to ensure that **LogOp** is well-defined.

In order to motivate the definition of **SEP** heuristically, let us consider again the task of the college chairman \mathbf{A}_0 . Following the reasoning elaborated in sections sections 2.1.3 and 2.2, \mathbf{A}_0 decides that as an initial criterion she will choose a social belief function $\mathbf{v} = \langle v_1 \dots v_J \rangle$ in such a manner as to minimize the average informational distance between $\langle v_1 \dots v_J \rangle$ and the m belief functions $\mathbf{w}^{(i)} = \langle w_1^{(i)} \dots w_J^{(i)} \rangle$ of the members of \mathbf{M} , where the $\mathbf{w}^{(i)}$ are all simultaneously chosen in such a manner as to minimize this quantity subject to the relevant sets of belief constraints \mathbf{K}_i of the members of the college.

The standard measure of informational distance between probability distributions \mathbf{v} and \mathbf{u} is the well-studied notion of *Kullback-Leibler divergence* [30], sometimes known as cross-entropy, given by

$$\text{KL}(\mathbf{v}, \mathbf{u}) = \sum_{j=1}^J v_j \log \frac{v_j}{u_j}$$

where the convention is observed that $v_j \log \frac{v_j}{u_j}$ takes the value 0 if $v_j = 0$ and the value $+\infty$ if $v_j \neq 0$ and $u_j = 0$.

We recall that Kullback-Leibler divergence is *not* a symmetric function; intuitively in the context of updating for a single agent $\text{KL}(\mathbf{v}, \mathbf{u})$ represents the informational distance from old belief function \mathbf{u} to new belief function \mathbf{v} . Using this notion of informational distance \mathbf{A}_0 's idea is therefore to choose \mathbf{v} and $\mathbf{w}^{(1)} \dots \mathbf{w}^{(m)}$ with each $\mathbf{w}^{(i)}$ satisfying \mathbf{K}_i , so as to minimize

$$\frac{1}{m} \sum_{i=1}^m \text{KL}(\mathbf{v}, \mathbf{w}^{(i)})$$

We will see below that, while such a procedure will not by itself always produce unique belief functions for \mathbf{v} and the associated $\mathbf{w}^{(1)} \dots \mathbf{w}^{(m)}$, the set of possible belief functions satisfying these criteria has both a pleasant characterisation and a tight mathematical structure.

A fundamental property of Kullback-Leibler divergence which we shall need is

Lemma 6 (Gibbs Inequality). *For all belief functions \mathbf{v} and \mathbf{u}*

$$\text{KL}(\mathbf{v}, \mathbf{u}) \geq 0$$

with equality holding if and only if $\mathbf{v} = \mathbf{u}$.

Proof. See [30] or [3]. \square

The next lemma allows us to express \mathbf{A}_0 's criterion above in a much more convenient mathematical form.

Lemma 7. *Let $\mathbf{K}_1 \dots \mathbf{K}_m$ be constraint sets on belief functions $\mathbf{w}^{(1)} \dots \mathbf{w}^{(m)}$ respectively. Then the following are equivalent:*

(i) *The belief functions $\mathbf{v}, \mathbf{w}^{(1)}, \dots, \mathbf{w}^{(m)}$ minimize the quantity*

$$\frac{1}{m} \sum_{i=1}^m \text{KL}(\mathbf{v}, \mathbf{w}^{(i)}) \tag{4}$$

subject to the given constraints.

(ii) *The belief functions $\mathbf{w}^{(1)} \dots \mathbf{w}^{(m)}$ maximize the quantity*

$$\sum_{j=1}^J \left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}} \tag{5}$$

subject to the given constraints, and

$$v_j = \frac{\left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}}}{\sum_{j=1}^J \left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}}} \tag{6}$$

for all $j = 1 \dots J$.

Proof. We note first that by our assumptions concerning the constraint sets, the minimum value of (4) must be finite. For by assumption there exists some j_0 and some $\mathbf{u}^{(i)} \in \mathbf{V}_{\mathbf{K}_i}$ such that $u_{j_0}^{(i)} \neq 0$ for all $i = 1 \dots m$; then by replacing each $\mathbf{w}^{(i)}$ by $\mathbf{u}^{(i)}$ and setting $v_{j_0} = 1$ and all other v_j equal to zero gives (4) a finite value. From this it follows that for any j if v_j is non-zero then $w_j^{(i)}$ is non-zero for all $i = 1 \dots m$. Thus we can rewrite (4) as

$$\sum_{j=1}^J v_j \log \frac{v_j}{\left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}}}$$

or, equivalently, as

$$\sum_{j=1}^J v_j \log \left(\frac{v_j}{\frac{\left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}}}{\sum_{j'=1}^J \left[\prod_{i=1}^m w_{j'}^{(i)} \right]^{\frac{1}{m}}}} \right) - \log \sum_{j=1}^J \left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}} \tag{7}$$

which, by the Gibbs inequality, will for any given $\mathbf{w}^{(1)} \dots \mathbf{w}^{(m)}$ take its minimum value when the first term vanishes and \mathbf{v} is given by the expression at (6). On the other hand the second term is minimized when $\sum_{j'=1}^J \left[\prod_{i=1}^m w_{j'}^{(i)} \right]^{\frac{1}{m}}$ is maximized. It follows that the minimum possible value of (4) is obtained by *first* maximising $\sum_{j'=1}^J \left[\prod_{i=1}^m w_{j'}^{(i)} \right]^{\frac{1}{m}}$ subject to the constraints, and then letting \mathbf{v} be determined by the equation (6).

□

The above lemma shows that Chairman \mathbf{A}_0 's initial criterion for selecting appropriate \mathbf{v} for consideration as the social belief function can be reduced to the problem of finding those sequences of belief functions $\mathbf{w}^{(1)} \dots \mathbf{w}^{(m)}$ which maximize $\sum_{j=1}^J \left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}}$, subject to each $\mathbf{w}^{(i)}$ satisfying the relevant set of constraints \mathbf{K}_i . Notice that the function being maximized above is just a sum of geometric means. Since this function is bounded and continuous and the space over which it is being maximized is by assumption closed, a maximum value is certainly attained.

In order to make our presentation more readable we shall in future abbreviate $\mathbf{K}_1 \dots \mathbf{K}_m$ by $\vec{\mathbf{K}}$.

Definition 8. For a sequence of knowledge bases $\vec{\mathbf{K}}$ we define

$$M_{\vec{K}} = \text{Max} \left\{ \sum_{j=1}^J \left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}} \mid \mathbf{w}^{(i)} \in \mathbf{V}_{\mathbf{K}_i} \text{ for all } i = 1 \dots m \right\}$$

□

It is now easy to see that

Lemma 9. Given knowledge bases $\mathbf{K}_1 \dots \mathbf{K}_m$ and $M_{\vec{K}}$ defined as above then $0 < M_{\vec{K}} \leq 1$.

Furthermore the value $M_{\vec{K}} = 1$ occurs if and only if for every $j = 1 \dots J$ and for all $i, i' \in \{1 \dots m\}$ $w_j^{(i)} = w_j^{(i')}$. Hence given $\mathbf{K}_1 \dots \mathbf{K}_m$ the following are equivalent:

1. $M_{\vec{K}} = 1$
2. Every $\mathbf{w}^{(1)} \dots \mathbf{w}^{(m)}$ which generates the value $M_{\vec{K}}$ satisfies $\mathbf{w}^{(1)} = \dots = \mathbf{w}^{(m)} = \mathbf{v}$.
3. The knowledge bases $\mathbf{K}_1 \dots \mathbf{K}_m$ are jointly consistent: i.e., there exists some belief function which satisfies all of them.

Proof. Let $\mathbf{w}^{(1)} \dots \mathbf{w}^{(m)}$ be belief functions satisfying $\mathbf{K}_1 \dots \mathbf{K}_m$ respectively, and which generate the value $M_{\vec{K}}$. First note that by assumption for some j_0 no \mathbf{K}_i forces the probability given to atom α_{j_0} to be zero, and hence $M_{\vec{K}} > 0$, since it is possible to choose belief functions $\mathbf{u}^{(1)} \dots \mathbf{u}^{(m)}$ respectively consistent with $\mathbf{K}_1 \dots \mathbf{K}_m$ such that $\left[\prod_{i=1}^m u_{j_0}^{(i)} \right]^{\frac{1}{m}} > 0$.

Now by applying the arithmetic-geometric mean inequality m times we get

$$M_{\vec{K}} = \sum_{j=1}^J \left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}} \leq \sum_{j=1}^J \frac{1}{m} \sum_{i=1}^m w_j^{(i)} = 1 \quad \text{since} \quad \sum_{i=1}^m \sum_{j=1}^J w_j^{(i)} = m.$$

Moreover since equality for any of the arithmetic-geometric mean inequalities occurs just when all the terms are equal, the case $M_{\vec{K}} = 1$ occurs if and only if $\mathbf{w}^{(1)} = \mathbf{w}^{(1)} = \dots = \mathbf{w}^{(m)} = \mathbf{v}$. This suffices to prove the lemma. □

Now it is obvious from the above that Chairman \mathbf{A}_0 's proposed method of choosing \mathbf{v} will not in general result in a uniquely defined social belief function. Indeed if $\bigcap_{i=1}^m \mathbf{V}_{\mathbf{K}_i} \neq \emptyset$ then any point \mathbf{w} in this intersection, if adopted as the belief function of each member, will generate the maximum possible value for $M_{\vec{K}}$ of 1 and so will be a possible candidate for a social belief function \mathbf{v} . Moreover even if $\bigcap_{i=1}^m \mathbf{V}_{\mathbf{K}_i} = \emptyset$ the process above may not result in a unique choice of either the $\mathbf{w}^{(i)}$ or of \mathbf{v} .

Chairman \mathbf{A}_0 now reasons as follows: if the result of the above operation of minimizing the average Kullback-Leibler divergence does not result in a unique solution for \mathbf{v} , then the best rational recourse which she has left is to choose that \mathbf{v} which has maximum entropy from the set of possible \mathbf{v} previously obtained, assuming of course that such a choice is well-defined. Chairman \mathbf{A}_0 reasons that by adopting this procedure she is treating the set of \mathbf{v} defined by minimizing the average Kullback-Leibler divergence

of v with possible belief functions of college members *as if* that were the set of her own possible belief functions, and then choosing a belief function from that set by applying the ME inference process, as she would if that were indeed the case.

However in order to show that this procedure is well-defined, Chairman A_0 needs to prove certain technical results.

Definition 10. For knowledge bases \vec{K} we define

$$\Gamma(\vec{K}) = \{ \langle w^{(1)} \dots w^{(m)} \rangle \in \bigotimes_{i=1}^m V_{K_i} \mid \sum_{j=1}^J \left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}} = M_{\vec{K}} \}$$

By Lemma 7, each point $\langle w^{(1)} \dots w^{(m)} \rangle$ in $\Gamma(\vec{K})$ gives rise to a uniquely determined corresponding social belief function v whose j 'th coordinate is given by

$$v_j = \frac{1}{M_{\vec{K}}} \left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}}$$

We will refer to the v thus obtained from $\langle w^{(1)} \dots w^{(m)} \rangle$ as

$$\text{LogOp}(w^{(1)} \dots w^{(m)})$$

and we let

$$\Delta(\vec{K}) = \{ \text{LogOp}(w^{(1)} \dots w^{(m)}) \mid \langle w^{(1)} \dots w^{(m)} \rangle \in \Gamma(\vec{K}) \}$$

□

$\Delta(\vec{K})$ is thus the candidate set of possible social belief functions from which Chairman A_0 wishes to make her final choice by selecting the point in this set which has maximum entropy.

From now on we shall abbreviate a typical point $\langle w^{(1)} \dots w^{(m)} \rangle$ in $\bigotimes_{i=1}^m V_{K_i}$ by \vec{w} . For any such \vec{w} we denote the vector $\langle w_j^{(1)} \dots w_j^{(m)} \rangle$ by w_j . Thus we may think of \vec{w} as an $m \times J$ matrix with rows $w^{(i)}$, columns w_j , and individual entries $w_j^{(i)}$.

Our problem is to analyze the linked structures of $\Gamma(\vec{K})$ and $\Delta(\vec{K})$, and in particular to show that $\Delta(\vec{K})$ is convex. A slight complicating factor in this analysis turns out to be the possibility that some entries in a matrix $\vec{w} \in \Gamma(\vec{K})$ may turn out to be zero. Notice that the corresponding social belief function v will have j 'th coordinate v_j equal to zero if and only if some entry in the column vector w_j is equal to zero. Such zero entries v_j may be classified as of two possible kinds: either $v_j = 0$ because for some i the knowledge base K_i forces $w_j^{(i)} = 0$, or, when this is not the case, because for some i it just so happens that $w_j^{(i)} = 0$. The first case is in a certain sense trivial since for an arbitrary $\vec{w} \in \bigotimes_{i=1}^m V_{K_i}$ the columns w_j corresponding to such j will make zero contribution to the function to be maximised. For this reason it is convenient to introduce a notation which allows us to

eliminate such j from consideration. Accordingly, for given $\vec{\mathbf{K}}$, we define the set of significant j , $\text{Sig}_{\vec{\mathbf{K}}}$ by:

$$\text{Sig}_{\vec{\mathbf{K}}} = \{ j \mid \text{for no } i \text{ is it the case that } w_j^{(i)} = 0 \text{ for all } \mathbf{w}^{(i)} \in \mathbf{V}_{\mathbf{K}_i} \}$$

Notice that by our initial assumption about $\vec{\mathbf{K}}$ at the beginning of this section $\text{Sig}_{\vec{\mathbf{K}}}$ is non-empty.

For any $\vec{\mathbf{w}} \in \otimes_{i=1}^m \mathbf{V}_{\mathbf{K}_i}$ we now define $\vec{\mathbf{w}}_{\text{Sig}_{\vec{\mathbf{K}}}}$ to be the projection of $\vec{\mathbf{w}}$ on to those coordinates (i, j) such that $j \in \text{Sig}_{\vec{\mathbf{K}}}$; i.e., $\vec{\mathbf{w}}_{\text{Sig}_{\vec{\mathbf{K}}}}$ may be viewed as the matrix obtained from the matrix $\vec{\mathbf{w}}$ by deleting those columns j for which $j \notin \text{Sig}_{\vec{\mathbf{K}}}$. Similarly for any probability function \mathbf{w} we define $\mathbf{w}_{\text{Sig}_{\vec{\mathbf{K}}}}$ to be the projection of \mathbf{w} to a vector obtained by deleting those coordinates which are not in $\text{Sig}_{\vec{\mathbf{K}}}$. (Notice however that the effect of this is that the sum of the components of such a $\mathbf{w}_{\text{Sig}_{\vec{\mathbf{K}}}}$ may be less than unity). Similarly we define

$$\Gamma^{\text{Sig}}(\vec{\mathbf{K}}) = \{ \vec{\mathbf{w}}_{\text{Sig}_{\vec{\mathbf{K}}}} \mid \vec{\mathbf{w}} \in \Gamma(\vec{\mathbf{K}}) \}$$

and

$$\Delta^{\text{Sig}}(\vec{\mathbf{K}}) = \{ \mathbf{v}_{\text{Sig}_{\vec{\mathbf{K}}}} \mid \mathbf{v} \in \Delta(\vec{\mathbf{K}}) \}$$

Note that in contrast to the situation for the row vectors of a matrix in $\Gamma^{\text{Sig}}(\vec{\mathbf{K}})$, the components of any vector in $\Delta^{\text{Sig}}(\vec{\mathbf{K}})$ do sum to unity, and that there is therefore a trivial homeomorphism between $\Delta^{\text{Sig}}(\vec{\mathbf{K}})$ and $\Delta(\vec{\mathbf{K}})$.

The next theorem¹⁸, which guarantees that Chairman \mathbf{A}_0 's plan is realisable, provides a crucial structure theorem for $\Gamma(\vec{\mathbf{K}})$ and $\Delta(\vec{\mathbf{K}})$, which depends strongly on the concavity properties of the geometric mean function and of sums of such functions.

Theorem 11 (Structure of $\Gamma(\vec{\mathbf{K}})$ and $\Delta(\vec{\mathbf{K}})$).

Let $\vec{\mathbf{K}}$ be a fixed vector of knowledge bases such that $\Delta(\vec{\mathbf{K}})$ is not a singleton.

(i) Let $\vec{\mathbf{w}} \in \Gamma^{\text{Sig}}(\vec{\mathbf{K}})$, and let \mathbf{v} be the corresponding point in $\Delta^{\text{Sig}}(\vec{\mathbf{K}})$.

Then for each $j \in \text{Sig}_{\vec{\mathbf{K}}}$ then **either** $w_j^{(i)} = 0$ for all $i = 1 \dots m$ **or** $w_j^{(i)}$ is nonzero for all $i = 1 \dots m$.

Furthermore in the case when $w_j^{(i)}$ is nonzero for all $i = 1 \dots m$, if $\vec{\mathbf{w}}'$ is any other point in $\Gamma^{\text{Sig}}(\vec{\mathbf{K}})$ with corresponding point \mathbf{v}' in $\Delta^{\text{Sig}}(\vec{\mathbf{K}})$, then

$$\mathbf{w}'_j = (1 + \mu_j)\mathbf{w}_j \quad \text{for some } \mu_j \in \mathbb{R} \text{ with } \mu_j \geq -1.$$

¹⁸ An earlier version of this theorem which was stated without proof in [1] contains an error because the statement of the result is incorrect for cases in which 0's appear in the coordinates.

and hence also

$$v'_j = (1 + \mu_j)v_j$$

(ii) There is a point $\vec{w} \in \Gamma^{Sig}(\vec{K})$ with corresponding $v \in \Delta^{Sig}(\vec{K})$ such that for every other point $\vec{w}' \in \Gamma^{Sig}(\vec{K})$ with corresponding $v' \in \Delta^{Sig}(\vec{K})$, for each $j \in Sig_{\vec{K}}$ there exists $\mu_j \geq -1$ such that

$$w'_j = (1 + \mu_j)w_j$$

and

$$v'_j = (1 + \mu_j)v_j$$

(iii) The regions $\Gamma^{Sig}(\vec{K})$, $\Delta^{Sig}(\vec{K})$, $\Gamma(\vec{K})$, and $\Delta(\vec{K})$ are all compact and convex.

(iv) If LogOp^{Sig} denotes the function defined on $\Gamma^{Sig}(\vec{K})$ by restricting the definition of the LogOp function defined on $\Gamma(\vec{K})$ in 3.5 above to those j which are in $Sig_{\vec{K}}$, then

$$\text{LogOp}^{Sig} : \Gamma^{Sig}(\vec{K}) \rightarrow \Delta^{Sig}(\vec{K})$$

is a continuous bijection.

Proof. Define the function $F : \otimes_{i=1}^m \mathbb{D}_J \rightarrow \mathbb{R}$: by

$$F(\vec{w}) = \sum_{j=1}^J \left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}}$$

This is this function which is to be maximised for $\vec{w} \in \otimes_{i=1}^m V_{K_i}$ in order to define the points in the region $\Gamma(\vec{K})$. We note first of all that for non-negative arguments the geometric mean function is always concave (see e.g. [31]) and hence a sum of such functions is also concave. Since the region $\otimes_{i=1}^m V_{K_i}$ is convex and compact by its definition, it follows that F attains a maximum value and hence that $\Gamma(\vec{K})$ is non-empty. Moreover it is an easy consequence of the definition of a concave function that the set of points which give maximal value to such a function over a compact convex region itself forms a compact convex set. Thus $\Gamma(\vec{K})$ is compact and convex. Since both compactness and convexity are preserved by projections in Euclidean space it follows that $\Gamma^{Sig}(\vec{K})$ is also compact and convex.

Let $[\otimes_{i=1}^m V_{K_i}]^{Sig_{\vec{K}}}$ denote the projection of $\otimes_{i=1}^m V_{K_i}$ onto those coordinates with $j \in Sig_{\vec{K}}$. This region is also compact and convex. Then if we define F^{Sig} for any $\vec{w} \in \otimes_{i=1}^m V_{K_i}$ by

$$F^{Sig}(\vec{w}_{Sig_{\vec{K}}}) = \sum_{j \in Sig_{\vec{K}}} \left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}}$$

then it is clear that

$$F^{Sig}(\vec{w}_{Sig_{\vec{K}}}) = F(\vec{w})$$

so that it suffices for us to confine our analysis to F^{Sig} acting on the points in $[\otimes_{i=1}^m V_{K_i}]^{Sig_{\vec{K}}}$.

Now let us consider a general point $\vec{a} \in \Gamma^{Sig}(\vec{K})$. We will show that for every $j \in Sig_{\vec{K}}$ we cannot have that $a_j^{(i)} = 0$ while $a_j^{(i')} \neq 0$ for some $i, i' \in \{1 \dots m\}$. Suppose for contradiction that

such j, i and i' exist. We first note that there exists some $\vec{b} \in [\otimes_{i=1}^m \mathbf{V}_{\mathbf{K}_i}]^{\text{Sig}_{\vec{K}}}$ such that $b_j^{(i)} \neq 0$ for all $i = 1 \dots m$ and all $j \in \text{Sig}_{\vec{K}}$. This follows from the convexity of $[\otimes_{i=1}^m \mathbf{V}_{\mathbf{K}_i}]^{\text{Sig}_{\vec{K}}}$ since for each particular i and j we can by our assumptions choose some $\vec{x} \in [\otimes_{i=1}^m \mathbf{V}_{\mathbf{K}_i}]^{\text{Sig}_{\vec{K}}}$ such that $x_j^{(i)} \neq 0$ and by convexity we can then form a suitable \vec{b} by taking the arithmetic mean of all these. So let us fix some such \vec{b} .

Let $\vec{u} = \vec{b} - \vec{a}$. Then by convexity, for any $\lambda \in [0, 1]$, the point $\vec{a} + \lambda \vec{u}$ is in $[\otimes_{i=1}^m \mathbf{V}_{\mathbf{K}_i}]^{\text{Sig}_{\vec{K}}}$. Note that by the definition of \vec{b} , for all i and j if $a_j^{(i)} = 0$ then $u_j^{(i)} > 0$.

Consider the behaviour of $\mathbf{F}^{\text{Sig}}(\vec{a} + \lambda \vec{u})$ as $\lambda \rightarrow 0$. Now differentiating with respect to λ we get

$$\frac{d\mathbf{F}^{\text{Sig}}}{d\lambda}(\vec{a} + \lambda \vec{u}) = \frac{1}{m} \sum_{j \in \text{Sig}_{\vec{K}}} \left[\left[\prod_{i=1}^m (a_j^{(i)} + \lambda u_j^{(i)}) \right]^{\frac{1}{m}} \sum_{i=1}^m \frac{u_j^{(i)}}{a_j^{(i)} + \lambda u_j^{(i)}} \right]$$

As $\lambda \rightarrow 0^+$ we see that all terms on the right hand side are bounded except in the case of those i, j where $a_j^{(i)} = 0$ and at least one $a_j^{(i')}$ is non-zero for some $i' \neq i$, in which case that term tends to $+\infty$. Since we are supposing that such j, i and i' do exist, it follows that \mathbf{F}^{Sig} is increasing as $\vec{a} + \lambda \vec{u}$ moves away from \vec{a} , and hence since \mathbf{F}^{Sig} is continuous at \vec{a} , \vec{a} cannot be a maximum point of $[\otimes_{i=1}^m \mathbf{V}_{\mathbf{K}_i}]^{\text{Sig}_{\vec{K}}}$, contradicting hypothesis. Thus we have shown that for any point \vec{w} in $\Gamma^{\text{Sig}}(\vec{K})$ if some column vector of \vec{w} has a zero entry then that column vector is identically zero, which establishes the first part of (i).

The second part of (i) follows directly from (ii), so we will prove (ii) instead.

By (i) and the convexity of $\Gamma^{\text{Sig}}(\vec{K})$ there exists an \vec{a} such that if there exists any \vec{b} in $\Gamma^{\text{Sig}}(\vec{K})$ for which for some j in $\text{Sig}_{\vec{K}}$ b_j is not a zero vector then all the entries of a_j are non-zero. Let us fix such an \vec{a} and let \vec{b} be any other point in $\Gamma^{\text{Sig}}(\vec{K})$. Again we consider $\vec{u} = \vec{b} - \vec{a}$ for $\lambda \in [0, 1]$, noting that in this case by the convexity of $\Gamma^{\text{Sig}}(\vec{K})$, $\vec{a} + \lambda \vec{u}$ is a point of $\Gamma^{\text{Sig}}(\vec{K})$, and hence $\mathbf{F}^{\text{Sig}}(\vec{a} + \lambda \vec{u}) = M_{\vec{K}}$ and so has constant value.

Let $\text{Sig}_{\vec{K}}^*$ denote $\{j \mid j \in \text{Sig}_{\vec{K}} \text{ and } a_j \neq \mathbf{0}\}$. Then by the definition of \vec{a} and of $\text{Sig}_{\vec{K}}^*$

$$\mathbf{F}^{\text{Sig}}(\vec{a} + \lambda \vec{u}) = \sum_{j \in \text{Sig}_{\vec{K}}^*} \left(\prod_{i=1}^m (a_j^{(i)} + \lambda u_j^{(i)}) \right)^{\frac{1}{m}}$$

Noting that all the $a_j^{(i)}$ occurring on the right are by definition non-zero, differentiating twice with respect to λ we have

$$\frac{d^2\mathbf{F}^{\text{Sig}}}{d\lambda^2}(\vec{a} + \lambda \vec{u}) = \frac{1}{m^2} \sum_{j \in \text{Sig}_{\vec{K}}^*} \left[\prod_{i=1}^m (a_j^{(i)} + \lambda u_j^{(i)}) \right]^{\frac{1}{m}} \left[\left[\sum_{i=1}^m \frac{u_j^{(i)}}{a_j^{(i)} + \lambda u_j^{(i)}} \right]^2 - m \sum_{i=1}^m \frac{(u_j^{(i)})^2}{(a_j^{(i)} + \lambda u_j^{(i)})^2} \right]$$

Since \mathbf{F}^{Sig} is constant for $\lambda \in [0, 1]$, setting the above expression equal to 0 for $\lambda = 0$ we get

$$\frac{1}{m^2} \sum_{j \in \text{Sig}_{\vec{K}}^*} \left[\prod_{i=1}^m a_j^{(i)} \right]^{\frac{1}{m}} \left[\left[\sum_{i=1}^m \frac{u_j^{(i)}}{a_j^{(i)}} \right]^2 - m \sum_{i=1}^m \left[\frac{u_j^{(i)}}{a_j^{(i)}} \right]^2 \right] = 0$$

from which we obtain

$$-\sum_{j \in \text{Sig}_{\vec{K}}^*} \left[\prod_{i=1}^m a_j^{(i)} \right]^{\frac{1}{m}} \sum_{i, i'=1 \dots m} \left[\frac{u_j^{(i)}}{a_j^{(i)}} - \frac{u_j^{(i')}}{a_j^{(i')}} \right]^2 = 0$$

From the negative definite form of the above expression we deduce that for all $j \in \text{Sig}_{\vec{K}}^*$ and all $i, i' = 1 \dots m$

$$\frac{u_j^{(i)}}{a_j^{(i)}} = \frac{u_j^{(i')}}{a_j^{(i')}}$$

whence for all $j \in \text{Sig}_{\vec{K}}^*$ and all $i, i' = 1 \dots m$

$$\frac{b_j^{(i)}}{a_j^{(i)}} = \frac{b_j^{(i')}}{a_j^{(i')}}$$

which suffices to establish part(ii) of the theorem.

To show (iv) note that the function $\text{LogOp}^{\text{Sig}} : \Gamma^{\text{Sig}}(\vec{K}) \rightarrow \Delta^{\text{Sig}}(\vec{K})$ is by definition continuous and surjective. However by (ii) it is also clearly injective. Finally to show part (iii) we have already noted that $\Gamma(\vec{K})$ and $\Gamma^{\text{Sig}}(\vec{K})$ are compact and convex. Since $\Delta(\vec{K})$ and $\Delta^{\text{Sig}}(\vec{K})$ are the continuous images of these compact sets under LogOp and $\text{LogOp}^{\text{Sig}}$ respectively, it follows that $\Delta(\vec{K})$ and $\Delta^{\text{Sig}}(\vec{K})$ are also compact. From the convexity of $\Gamma^{\text{Sig}}(\vec{K})$ the convexity of $\Delta^{\text{Sig}}(\vec{K})$ follows by (ii), while the convexity of $\Delta(\vec{K})$ follows immediately from that of $\Delta^{\text{Sig}}(\vec{K})$. This completes the proof of Theorem 11. \square

Now since $\Delta(\vec{K})$ is a compact convex set by 11(iii) and since the entropy function

$$-\sum_{j=1}^J v_j \log v_j$$

is strictly concave and bounded over this set, the set contains a unique point v^{ME} at which the entropy function achieves its maximum value. It follows at once that the following formal definition of the social inference process SEP defines, for every \vec{K} satisfying the conditions of this section, a unique social belief function.

Definition 12. *The Social Entropy Process, SEP, is the social inference process defined by*

$$\text{SEP}(\vec{K}) = \text{ME}(\Delta(\vec{K}))$$

where $\text{ME}(\Delta(\vec{K}))$ denotes the unique maximum entropy point in $\Delta(\vec{K})$.

We remark that it follows immediately from the definition above that the social inference process SEP marginalises to the inference process ME and to the pooling operator LogOp.

It is worth noting that Theorem 11(i) at once provides a simple sufficient condition for $\Delta(\vec{K})$ to be a singleton and thus for the application of ME in the definition of $\text{SEP}(\vec{K})$ to be redundant:

Corollary 13. *If $K_1 \dots K_m$ are such that for each $j = 1 \dots J$ except possibly at most one there exists some i with $1 \leq i \leq m$ such that the condition $w^{(i)} \in V_{K_i}$ forces $w_j^{(i)}$ to take a unique value, then $\Delta(K_1 \dots K_m)$ is a singleton. In particular this occurs if for some i V_{K_i} is a singleton.*

3.2. Principles Satisfied by SEP

Theorem 14. *SEP satisfies the seven principles of section 2.3: Equivalence, Anonymity, Atomic Renaming, Consistency, Collegiality, General Locality, and Proportionality.*

Proof. The fact that principles of Equivalence, Anonymity, and Atomic Renaming hold for SEP follows easily from the basic symmetry properties of the definition of SEP.

To prove that SEP satisfies Consistency, suppose that $\vec{K} = K_1 \dots K_m$ are such that

$$\bigcap_{i=1}^m V_{K_i} \neq \emptyset$$

Then for any $u \in \bigcap_{i=1}^m V_{K_i}$, if we set

$$v = w^{(1)} = \dots = w^{(m)} = u$$

Then

$$\sum_{j=1}^J \left(\prod_{i=1}^m w_j^{(i)} \right)^{\frac{1}{m}} = 1$$

and since by Lemma 9 $M_{\vec{K}} \leq 1$, it follows that $M_{\vec{K}} = 1$, and hence that $u \in \Delta(\vec{K})$. Conversely by Lemma 9, since $M_{\vec{K}} = 1$, then for any $v \in \Delta(\vec{K})$ if some $\vec{w} \in \Gamma(\vec{K})$ generates v , then $v = w^{(1)} = \dots = w^{(m)}$, and so $v \in \bigcap_{i=1}^m V_{K_i}$. It follows that

$$\text{SEP}(K_1 \dots K_m) \in \bigcap_{i=1}^m V_{K_i}$$

as required.

To prove Collegiality suppose that $K_1 \dots K_m$ are such that for some k with $1 < k < m$

$$\text{SEP}(K_1 \dots K_k) \in \bigcap_{i=k+1}^m V_{K_i}$$

Let $\tilde{v} = \text{SEP}(K_1 \dots K_k)$ and let $\hat{v} = \text{SEP}(K_1 \dots K_m)$.

Let $\langle \tilde{w}^{(1)} \dots \tilde{w}^{(k)} \rangle \in \Gamma(K_1 \dots K_k)$ be such that $\tilde{v} = \text{LogOp}(\tilde{w}^{(1)} \dots \tilde{w}^{(k)})$.

Similarly let $\langle \hat{w}^{(1)} \dots \hat{w}^{(m)} \rangle \in \Gamma(K_1 \dots K_m)$ be such that $\hat{v} = \text{LogOp}(\hat{w}^{(1)} \dots \hat{w}^{(m)})$ Then by definition

$$\sum_{i=1}^k \sum_{j=1}^J v_j \log \frac{v_j}{w_j^{(i)}}$$

takes its minimum possible value for $\mathbf{w}^{(1)} \dots \mathbf{w}^{(k)}$ subject to the constraints $\mathbf{K}_1 \dots \mathbf{K}_k$ when $\langle \mathbf{w}^{(1)} \dots \mathbf{w}^{(k)} \rangle = \langle \check{\mathbf{w}}^{(1)} \dots \check{\mathbf{w}}^{(k)} \rangle$ and $\mathbf{v} = \text{LogOp}(\check{\mathbf{w}}^{(1)} \dots \check{\mathbf{w}}^{(k)})$. We denote this value by Min_1 . Similarly

$$\sum_{i=1}^m \sum_{j=1}^J v_j \log \frac{v_j}{w_j^{(i)}}$$

takes its minimum possible value for $\mathbf{w}^{(1)} \dots \mathbf{w}^{(m)}$ subject to the constraints $\mathbf{K}_1 \dots \mathbf{K}_m$ when $\langle \mathbf{w}^{(1)} \dots \mathbf{w}^{(m)} \rangle = \langle \hat{\mathbf{w}}^{(1)} \dots \hat{\mathbf{w}}^{(m)} \rangle$ and $\mathbf{v} = \text{LogOp}(\hat{\mathbf{w}}^{(1)} \dots \hat{\mathbf{w}}^{(m)})$. We denote this value by Min_2 .

We now define $\check{\mathbf{w}}^{(i)}$ to be equal to $\check{\mathbf{v}}$ for $k + 1 \leq i \leq m$. Notice that by hypothesis $\check{\mathbf{w}}^{(1)} \dots \check{\mathbf{w}}^{(m)}$ now satisfy respectively the constraints $\mathbf{K}_1 \dots \mathbf{K}_m$. Hence we have by the definitions above

$$Min_2 \leq \sum_{i=1}^m \sum_{j=1}^J \check{v}_j \log \frac{\check{v}_j}{\check{w}_j^{(i)}} = \sum_{i=1}^k \sum_{j=1}^J \check{v}_j \log \frac{\check{v}_j}{\check{w}_j^{(i)}} = Min_1. \tag{8}$$

Similarly we also have

$$Min_2 = \sum_{i=1}^m \sum_{j=1}^J \hat{v}_j \log \frac{\hat{v}_j}{\hat{w}_j^{(i)}} \geq \sum_{i=1}^k \sum_{j=1}^J \hat{v}_j \log \frac{\hat{v}_j}{\hat{w}_j^{(i)}} \geq Min_1. \tag{9}$$

It follows that the six quantities appearing in (8) and (9) above are all equal, and hence that

$$\check{\mathbf{v}} \text{ and } \hat{\mathbf{v}} \text{ are both in } \Delta(\mathbf{K}_1 \dots \mathbf{K}_k) \cap \Delta(\mathbf{K}_1 \dots \mathbf{K}_m).$$

However by definition $\check{\mathbf{v}}$ is the unique belief function with the highest entropy in $\Delta(\mathbf{K}_1 \dots \mathbf{K}_k)$, while $\hat{\mathbf{v}}$ is the unique belief function with the highest entropy in $\Delta(\mathbf{K}_1 \dots \mathbf{K}_m)$. Hence $\check{\mathbf{v}} = \hat{\mathbf{v}}$ as required.

To prove General Locality, consider a college with members $\mathbf{A}_1 \dots \mathbf{A}_m$ initially having respective knowledge bases $\mathbf{K}_1 \dots \mathbf{K}_m$, where each \mathbf{K}_i is a nice set of constraints conditioned on some fixed non-contradictory sentence ϕ . Now for each $i = 1 \dots m$ let \mathbf{K}_i^* be a nice set of constraints about $\neg\phi$. We are given that

$\text{SEP}(\mathbf{K}_1 \cup \mathbf{K}_1^*, \dots, \mathbf{K}_m \cup \mathbf{K}_m^*)(\phi) \neq 0$ and that $\text{SEP}(\mathbf{K}_1, \dots, \mathbf{K}_m)(\phi) \neq 0$. We must show that for any sentence θ

$$\text{SEP}(\mathbf{K}_1 \cup \mathbf{K}_1^*, \dots, \mathbf{K}_m \cup \mathbf{K}_m^*)(\theta \mid \phi) = \text{SEP}(\mathbf{K}_1, \dots, \mathbf{K}_m)(\theta \mid \phi).$$

Clearly for this purpose it suffice to show that for any atom α such that $\alpha \models \phi$

$$\text{SEP}(\mathbf{K}_1 \cup \mathbf{K}_1^*, \dots, \mathbf{K}_m \cup \mathbf{K}_m^*)(\alpha \mid \phi) = \text{SEP}(\mathbf{K}_1, \dots, \mathbf{K}_m)(\alpha \mid \phi).$$

Notice that while we assume about each \mathbf{K}_i that it determines a closed convex set of probability functions *conditioned* on ϕ , such a \mathbf{K}_i when interpreted as a set of constraints about beliefs in the original atoms $\alpha_1, \alpha_2, \dots, \alpha_J$ also determines a closed convex region of \mathbb{D}_J which as usual we denote by $\mathbf{V}_{\mathbf{K}_i}$. Hence $\mathbf{V}_{\mathbf{K}_i \cup \mathbf{K}_i^*}$ is also a closed convex region of \mathbb{D}_J . Furthermore the conditions imply that for each $i = 1 \dots m$ $\mathbf{K}_i \cup \mathbf{K}_i^*$, is consistent, and hence the above applications of **SEP** are legitimately made.

Without loss of generality we may assume as in the proof of Theorem 4 that the atoms are so ordered that for some k with $1 \leq k < J$

$$\phi \equiv \bigvee_{j=1}^k \alpha_j \quad \text{and} \quad \neg\phi \equiv \bigvee_{j=k+1}^J \alpha_j$$

Let $u = \text{SEP}(\mathbf{K}_1, \dots, \mathbf{K}_m)$ be generated by $\vec{x} \in \Gamma(\mathbf{K}_1, \dots, \mathbf{K}_m)$, and let $v = \text{SEP}(\mathbf{K}_1 \cup \mathbf{K}_1^*, \dots, \mathbf{K}_m \cup \mathbf{K}_m^*)$ be generated by

$\vec{y} \in \Gamma(\mathbf{K}_1 \cup \mathbf{K}_1^*, \dots, \mathbf{K}_m \cup \mathbf{K}_m^*)$. For each $i = 1 \dots m$, let $\sum_{j=1}^k x_j^{(i)} = a^{(i)}$, and let $\sum_{j=1}^k y_j^{(i)} = b^{(i)}$. Note that $a^{(i)}$ and $b^{(i)}$ are non-zero for all i since otherwise ϕ would get social belief zero contradicting hypotheses.

Now consider the point $\vec{z} \in \bigotimes_{i=1}^m \mathbf{V}_{\mathbf{K}_i}$ given for each $i = 1 \dots m$ by

$$z_j^{(i)} = \begin{cases} y_j^{(i)} \frac{a^{(i)}}{b^{(i)}} & \text{for } j = 1, \dots, k \\ x_j^{(i)} & \text{for } j = k + 1, \dots, J \end{cases}$$

By the definition of the point \vec{x} we know that

$$\sum_{j=1}^J \left[\prod_{i=1}^m z_j^{(i)} \right]^{\frac{1}{m}} \leq \sum_{j=1}^J \left[\prod_{i=1}^m x_j^{(i)} \right]^{\frac{1}{m}}$$

from which it follows that

$$\sum_{j=1}^k \left[\prod_{i=1}^m y_j^{(i)} \frac{a^{(i)}}{b^{(i)}} \right]^{\frac{1}{m}} \leq \sum_{j=1}^k \left[\prod_{i=1}^m x_j^{(i)} \right]^{\frac{1}{m}}$$

Dividing both sides by $\left[\prod_{i=1}^m a^{(i)} \right]^{\frac{1}{m}}$ we obtain that

$$\sum_{j=1}^k \left[\prod_{i=1}^m \frac{y_j^{(i)}}{b^{(i)}} \right]^{\frac{1}{m}} \leq \sum_{j=1}^k \left[\prod_{i=1}^m \frac{x_j^{(i)}}{a^{(i)}} \right]^{\frac{1}{m}} .$$

However by repeating a similar argument, but this time with \vec{x} and \vec{y} interchanged we obtain the reverse inequality, from which it follows that

$$\sum_{j=1}^k \left[\prod_{i=1}^m \frac{y_j^{(i)}}{b^{(i)}} \right]^{\frac{1}{m}} = \sum_{j=1}^k \left[\prod_{i=1}^m \frac{x_j^{(i)}}{a^{(i)}} \right]^{\frac{1}{m}} = M_1 \text{ say.} \tag{10}$$

Note that the above equality implies that the value M_1 does not depend on the \mathbf{K}_i^* in any way.

Let $\sum_{j=k+1}^J \left[\prod_{i=1}^m y_j^{(i)} \right]^{\frac{1}{m}} = M_2$ and let $\left[\prod_{i=1}^m b^{(i)} \right]^{\frac{1}{m}} = B$.

Then from (3) we know that $\sum_{j=1}^k \left[\prod_{i=1}^m y_j^{(i)} \right]^{\frac{1}{m}} = M_1 B$.

Let us denote by C the quantity

$$\sum_{j=1}^J \left[\prod_{i=1}^m y_j^{(i)} \right]^{\frac{1}{m}} = M_1 B + M_2 \tag{11}$$

and we note that by definition C is the maximal value which can be taken by $\sum_{j=1}^J \left[\prod_{i=1}^m t_j^{(i)} \right]^{\frac{1}{m}}$ for any $\vec{t} \in \otimes_{i=1}^m \mathbf{V}_{\mathbf{K}_i \cup \mathbf{K}_i^*}$. We now consider those \vec{t} of this form for which $t_j^{(i)} = y_j^{(i)}$ for all $j = k + 1, \dots, J$ and all $i = 1, \dots, m$. Then, since for each $j = 1 \dots k$ $v_j = C^{-1} \left[\prod_{i=1}^m y_j^{(i)} \right]^{\frac{1}{m}}$ the definition of \vec{y} ensures that the column vectors $\mathbf{y}_1 \dots \mathbf{y}_k$ are of the form $\mathbf{t}_1 \dots \mathbf{t}_k$ where

$$- \sum_{j=1}^k C^{-1} \left[\prod_{i=1}^m t_j^{(i)} \right]^{\frac{1}{m}} \log \left(C^{-1} \left[\prod_{i=1}^m t_j^{(i)} \right]^{\frac{1}{m}} \right) \tag{12}$$

is maximised subject to the conditions that for each i the probability distribution $\langle \frac{t_1^{(i)}}{b^{(i)}} \dots \frac{t_k^{(i)}}{b^{(i)}} \rangle$ satisfies the knowledge base \mathbf{K}_i , that

$$\sum_{j=1}^k \left[\prod_{i=1}^m t_j^{(i)} \right]^{\frac{1}{m}} = M_1 B \tag{13}$$

and that for each i

$$\sum_{j=1}^k t_j^{(i)} = b^{(i)}. \tag{14}$$

Using some elementary algebra and (13) above we can rewrite the quantity in (12) which is to be maximised as

$$- \frac{M_1 B}{C} \log \frac{B}{C} - \frac{B}{C} \sum_{j=1}^k \left[\prod_{i=1}^m \frac{t_j^{(i)}}{b^{(i)}} \right]^{\frac{1}{m}} \log \left[\prod_{i=1}^m \frac{t_j^{(i)}}{b^{(i)}} \right]^{\frac{1}{m}} \tag{15}$$

Now since B , C , and M_1 , are positive constants for the \vec{t} under consideration, it follows that maximising (15), or equivalently (12), under the given constraints, is equivalent to maximising

$$- \sum_{j=1}^k \left[\prod_{i=1}^m \frac{t_j^{(i)}}{b^{(i)}} \right]^{\frac{1}{m}} \log \left[\prod_{i=1}^m \frac{t_j^{(i)}}{b^{(i)}} \right]^{\frac{1}{m}} \tag{16}$$

Hence, writing $w_j^{(i)}$ for $\frac{t_j^{(i)}}{b^{(i)}}$, it follows that this is in turn equivalent to maximising

$$- \sum_{j=1}^k \left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}} \log \left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}} \tag{17}$$

subject to the constraints that each k -dimensional row vector $\mathbf{w}^{(i)} = \langle w_1^{(i)} \dots w_k^{(i)} \rangle$ sums to 1 and satisfies \mathbf{K}_i when interpreted as a probability function conditioned on ϕ , and that

$$\sum_{j=1}^k \left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}} = M_1 \tag{18}$$

Now by the remark following (10), the value M_1 must be the largest possible which can be attained by $\sum_{j=1}^k \left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}}$ for the $\mathbf{w}^{(i)}$ probability functions satisfying the \mathbf{K}_i . Hence since the \mathbf{K}_i are

nice constraint sets, it follows by the fact that SEP is well-defined that any solution for \vec{w} to the above maximisation problem generates the unique $\text{SEP}(\mathbf{K}_1, \dots, \mathbf{K}_m)$ solution given by

$$\text{SEP}(\mathbf{K}_1, \dots, \mathbf{K}_m) (\alpha_j | \phi) = \frac{\left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}}}{\sum_{r=1}^k \left[\prod_{i=1}^m w_r^{(i)} \right]^{\frac{1}{m}}}$$

for $j = 1 \dots k$.

However by the definition of the above $w_j^{(i)}$ and the uniqueness of the SEP values, it follows that for such a solution \vec{w} , for each $j = 1 \dots k$

$$\left[\prod_{i=1}^m w_j^{(i)} \right]^{\frac{1}{m}} = \left[\prod_{i=1}^m \frac{y_j^{(i)}}{b^{(i)}} \right]^{\frac{1}{m}} \tag{19}$$

whence for each $j = 1 \dots k$

$$\begin{aligned} \text{SEP}(\mathbf{K}_1, \dots, \mathbf{K}_m) (\alpha_j | \phi) &= \frac{\left[\prod_{i=1}^m \frac{y_j^{(i)}}{b^{(i)}} \right]^{\frac{1}{m}}}{\sum_{r=1}^k \left[\prod_{i=1}^m \frac{y_r^{(i)}}{b^{(i)}} \right]^{\frac{1}{m}}} \\ &= \frac{C^{-1} \left[\prod_{i=1}^m y_j^{(i)} \right]^{\frac{1}{m}}}{C^{-1} \sum_{r=1}^k \left[\prod_{i=1}^m y_r^{(i)} \right]^{\frac{1}{m}}} \\ &= \text{SEP}(\mathbf{K}_1 \cup \mathbf{K}_1^*, \dots, \mathbf{K}_m \cup \mathbf{K}_m^*) (\alpha_j | \phi) \end{aligned}$$

as required. This concludes the proof of General Locality.

It remains for us to prove that SEP satisfies Proportionality.

Let $\mathbf{K}_1, \dots, \mathbf{K}_m$ be knowledge bases and for each $r = 1 \dots n$ let \mathbf{K}_{ir} denote a copy of the knowledge base \mathbf{K}_i , so that $\mathbf{V}_{\mathbf{K}_{ir}} = \mathbf{V}_{\mathbf{K}_i}$. As a shorthand we denote the sequence $\mathbf{K}_{i1} \dots \mathbf{K}_{in}$ by $n\mathbf{K}_i$. Clearly it suffices for us to prove that

$$\Delta(n\mathbf{K}_1, n\mathbf{K}_2, \dots, n\mathbf{K}_m) = \Delta(\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_m) \tag{20}$$

Let $\mathbf{v} \in \Delta(n\mathbf{K}_1, n\mathbf{K}_2, \dots, n\mathbf{K}_m)$ be generated by some $\vec{w} \in \Gamma(n\mathbf{K}_1, n\mathbf{K}_2, \dots, n\mathbf{K}_m)$. Then letting

$$\sum_{r=1}^n \sum_{i=1}^m \sum_{j=1}^J v_j \log \frac{v_j}{w_j^{(ir)}} = D \tag{21}$$

by definition D is minimal subject only to the constraints that $\mathbf{w}^{(ir)} \in \mathbf{V}_{\mathbf{K}_i}$ for all $r = 1 \dots n$ and $i = 1 \dots m$, (but no constraints on \mathbf{v}). Then for each $r = 1 \dots n$

$$\sum_{i=1}^m \sum_{j=1}^J v_j \log \frac{v_j}{w_j^{(ir)}} = \frac{D}{n} \tag{22}$$

and

$\frac{D}{n}$ is the minimum value which can be taken by

$$\sum_{i=1}^m \sum_{j=1}^J v_j \log \frac{v_j}{z_j^{(i)}} \quad \text{for } \mathbf{z}^{(i)} \in \mathbf{V}_{\mathbf{K}_i} \tag{23}$$

(22) holds because otherwise we would have that for some r_0 with $1 \leq r_0 \leq n$

$$\sum_{i=1}^m \sum_{j=1}^J v_j \log \frac{v_j}{w_j^{(ir_0)}} < \frac{D}{n}$$

and if we then define \vec{y} by

$$y_j^{(ir)} = w_j^{(ir_0)}$$

for all i, j and r , we would have that $\sum_{r=1}^n \sum_{i=1}^m \sum_{j=1}^J v_j \log \frac{v_j}{y_j^{(ir)}} < D$ contradicting the definition of D in (21). The same argument shows also that (23) holds. From (22) and (23) it follows that $\mathbf{v} \in \Delta(\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_m)$.

Conversely if some $\mathbf{v} \in \Delta(\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_m)$ is generated by a $\vec{z} \in \Gamma(\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_m)$ then it is easy to see that

$$\sum_{i=1}^m \sum_{j=1}^J v_j \log \frac{v_j}{z_j^{(i)}} = \frac{D}{n} \tag{24}$$

where D is the minimal value defined at (14) since the value of $\sum_{i=1}^m \sum_{j=1}^J v_j \log \frac{v_j}{z_j^{(i)}}$ cannot be smaller than $\frac{D}{n}$ by the same argument used to show (22) and (23). However if we now define \vec{w} by $w_j^{(ir)} = z_j^{(i)}$ then by (24) the equation (21) holds and so $\mathbf{v} \in \Delta(n\mathbf{K}_1, n\mathbf{K}_2, \dots, n\mathbf{K}_m)$.

Thus $\Delta(n\mathbf{K}_1, n\mathbf{K}_2, \dots, n\mathbf{K}_m) = \Delta(\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_m)$ as required.

This concludes the proof of Theorem 14. \square

Remark 2. We note that Savage [32] has shown that a certain form of converse of Collegiality holds for SEP. Namely

if $\mathbf{K}_1 \dots \mathbf{K}_m$ are such that for each $j = 1 \dots J$ $\text{SEP}(\mathbf{K}_1 \dots \mathbf{K}_m)(\alpha_j) \neq 0$, then if $\text{SEP}(\mathbf{K}_1 \dots \mathbf{K}_{m-1}) \notin \mathbf{V}_{\mathbf{K}_m}$ then $\text{SEP}(\mathbf{K}_1 \dots \mathbf{K}_{m-1}) \neq \text{SEP}(\mathbf{K}_1 \dots \mathbf{K}_m)$.

Further related properties of a social inference process may be found in [17] and [26].

3.3. Some Other Possible Principles

We end this chapter with some brief remarks concerning possible generalisations to the context of a social inference process, and in particular to SEP, of some remaining key principles which were identified by Paris and Vencovská ([3], [2]) as characterising the ME inference process.

One such key principle satisfied by ME, is that of **Open Mindedness**. An inference process \mathcal{I} satisfies Open Mindedness if for every knowledge base \mathbf{K} , for all $j = 1 \dots J$ $\mathcal{I}(\mathbf{K})(\alpha_j) \neq 0$ unless $w_j = 0$ for all $\mathbf{w} \in \mathbf{V}_{\mathbf{K}}$. The most obvious way of extending this principle to the case of a social inference process \mathfrak{F} would seem to be to propose that for all $j = 1 \dots J$ and for all $\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_m$, $\mathfrak{F}(\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_m)(\alpha_j) \neq 0$ unless for some i $w_j^{(i)} = 0$ for all $\mathbf{w}^{(i)} \in \mathbf{V}_{\mathbf{K}^{(i)}}$. It is easy to see however that such a principle cannot hold for any \mathfrak{F} which satisfies the Consistency Principle (cf. [26]). For if we take the example where there are three atoms $\alpha_1, \alpha_2, \alpha_3$, and $\mathbf{K}_1 = \{\mathbf{w}(\alpha_1) = \frac{1}{3}\}$,

while $\mathbf{K}_2 = \{w(\alpha_2) = \frac{2}{3}\}$, then by the consistency principle the only possible social belief function is given by $v = \langle \frac{1}{3}, \frac{2}{3}, 0 \rangle$, despite the fact that neither \mathbf{K}_1 nor \mathbf{K}_2 on their own force belief in α_3 to be zero. Furthermore, at least in the case of the inference process SEP, it is easy to show that similar counterexamples \mathbf{K}'_1 and \mathbf{K}'_2 to such a principle can be found where the union of the constraint sets \mathbf{K}'_1 and \mathbf{K}'_2 is not consistent. It seems reasonable to conclude therefore that Open Mindedness, at least in this formulation, is not a reasonable principle for a social inference process. Nevertheless SEP does satisfy the following weak form of Open Mindedness:

Theorem 15 (Weak Open Mindedness for SEP). *For any atom α and vector of knowledge bases $\vec{\mathbf{K}}$, if $\text{SEP}(\vec{\mathbf{K}})(\alpha) = 0$ then at least one of the following holds:*

- (1) For some $i = 1 \dots m$ $w(\alpha) = 0$ for all $w \in V_{\mathbf{K}_i}$
- (2) For every $i = 1 \dots m$ $w(\alpha) = 0$ for some $w \in V_{\mathbf{K}_i}$

Proof. Since by its definition $\text{SEP}(\vec{\mathbf{K}})$ is obtained by applying LogOp to a certain point in $\Gamma(\vec{\mathbf{K}})$, the result follows easily from part (i) of theorem 11, the structure theorem for $\Gamma(\vec{\mathbf{K}})$ and $\Delta(\vec{\mathbf{K}})$. \square

The above result can be rephrased by saying that $\text{SEP}(\vec{\mathbf{K}})(\alpha)$ will be non-zero unless either some \mathbf{K}_i forces $w(\alpha)$ to be zero, or for each i it is consistent with \mathbf{K}_i that $w(\alpha) = 0$. In the converse direction it is clear that the first condition suffices to ensure that $\text{SEP}(\vec{\mathbf{K}})(\alpha) = 0$, but that the second condition does not. Note that we have formulated Weak Open Mindedness for SEP in terms that would make sense (but would not necessarily hold) for an arbitrary social inference process, since we do not explicitly refer to $\Gamma(\vec{\mathbf{K}})$ or $\Delta(\vec{\mathbf{K}})$. However it is worth noting that in the case of SEP, theorem 15 still holds if we replace the second condition by the much stronger condition:

$$2'. \text{ For all } i = 1 \dots m \text{ and all } \vec{w} \in \Gamma(\vec{\mathbf{K}}) \quad w^{(i)}(\alpha) = 0 ,$$

and moreover, in the converse direction, this condition obviously implies the statement that $\text{SEP}(\vec{\mathbf{K}})(\alpha) = 0$.

Another pleasing property of the inference process ME, identified in [3] is that of continuity with respect to the Blaschke topology. At present we do not know whether an analogous formulation of this continuity principle holds for SEP, although it seems likely that this is the case.

The **Obstinacy Principle** for an inference process \mathcal{I} states that if \mathbf{K} and \mathbf{K}' are knowledge bases such that $\mathcal{I}(\mathbf{K}) \in V_{\mathbf{K}'}$ then $\mathcal{I}(\mathbf{K} \cup \mathbf{K}') = \mathcal{I}(\mathbf{K})$. While this principle is satisfied by ME, and indeed by nearly all standard inference processes (see [3]), an appropriate straightforward generalisation of the principle to the multi-agent context is not apparent. On the other hand it may be noted that the Collegiality Principle bears a certain formal resemblance to Obstinacy.

The important remaining principles characteristic of ME, as formulated in [2], [3] are those of **Language Invariance**, **Irrelevant Information**, and **Independence**. These important properties have in common the fact that they are most naturally stated in a context where the atoms of **At** arise as the Boolean atoms of a finite propositional language L as in Remark 1. In the formulations which follow we shall therefore assume this context.

The Language Invariance Principle

For an inference process \mathcal{I} this principle states that, for any \mathbf{K} , $\mathcal{I}(\mathbf{K})$ does not depend on the underlying language L in which \mathbf{K} is formulated. A fine point here is that strictly speaking an inference process or social inference process is formulated for a fixed language (or set of atoms). So more formally we should refer to an inference process \mathcal{I}_L rather than just \mathcal{I} in order to make explicit the underlying language L . However it is clear that any general formulation of an inference process, such as **ME**, will in reality represent a family of inference processes \mathcal{I}_L , one for each possible L . Then the following problem naturally arises. Suppose the knowledge base \mathbf{K} is formulated in the language L , and the language L is now extended to a larger propositional language L' by the addition of some new propositional variables. Then \mathbf{K} is also a knowledge base for the language L' . Intuitively if θ is a sentence (*i.e.*, disjunction of atoms) of L , we would expect that the mere expansion of the language from L to L' without the addition of any new constraints should not change the belief which is accorded to θ . So following [2] we say that \mathcal{I} satisfies Language Invariance if for all L, L', \mathbf{K} and θ as above

$$\mathcal{I}_{L'}[\mathbf{K}](\theta) = \mathcal{I}_L[\mathbf{K}](\theta)$$

A large number of inference processes, including **ME**, satisfy Language Invariance (cf. [6] p.213). Moreover Language Invariance seems as natural a principle for social inference processes as it does for inference processes: we say that \mathfrak{F} satisfies Language Invariance if

$$\mathfrak{F}_{L'}[\vec{\mathbf{K}}](\theta) = \mathfrak{F}_L[\vec{\mathbf{K}}](\theta)$$

for all $L \subset L'$, sequences of knowledge bases $\vec{\mathbf{K}}$ formulated in L , and sentences θ of L .

Pleasingly **SEP** does satisfy Language Invariance, as was shown in [27].

The Irrelevant Information Principle

An inference process \mathcal{I} satisfies the Irrelevant Information Principle if whenever a finite propositional language L is the union of two disjoint languages L_1 and L_2 , and \mathbf{K}_1 and \mathbf{K}_2 are knowledge bases for the languages L_1 and L_2 respectively, then

$$\mathcal{I}_L[\mathbf{K}_1 \cup \mathbf{K}_2](\theta) = \mathcal{I}_L[\mathbf{K}_1](\theta)$$

for all sentences θ of the language L_1 .

Intuitively this is also a very natural principle which roughly says that adding additional knowledge formulated in a language L_2 should not affect our beliefs in sentences formulated in a disjoint language L_1 formed on the basis of knowledge about L_1 . It is however a principle which is hard to satisfy, and excluding some artificial constructions, only two inferences processes are known to satisfy it: **ME** and the **Maximin** process of [6].

The principle has a natural generalisation to social inference processes. We say that a social inference process \mathfrak{F} satisfies Irrelevant Information if for any L, L_1 and L_2 as above, and for any sequences of knowledge bases $\mathbf{K}_1, \dots, \mathbf{K}_m$ and $\mathbf{K}'_1, \dots, \mathbf{K}'_m$ for L_1 and L_2 respectively

$$\mathfrak{F}_L[\mathbf{K}_1 \cup \mathbf{K}'_1 \dots \mathbf{K}_m \cup \mathbf{K}'_m](\theta) = \mathfrak{F}_L[\mathbf{K}_1 \dots \mathbf{K}_m](\theta)$$

for all sentences θ of the language L_1 .

However SEP does not satisfy the Irrelevant Information Principle, although it is known to satisfy a weak form; specifically, with the above notation, the extra conditions are required that (i) $\mathbf{K}'_1 \cup \dots \cup \mathbf{K}'_m$ is consistent and that (ii) $\Delta_{L_1}[\mathbf{K}_1 \dots \mathbf{K}_m]$ is a singleton (see [27]). It is not known if the second condition can be dropped.

Independence

The inference process ME satisfies very natural independence properties which can be expressed in a number of different ways. When considering a formulation which it might be appropriate to generalise to the case of a social inference process, the following property of an inference process \mathcal{I} , which is satisfied by ME, seems particularly natural:

Let $L = L_1 \cup L_2$ where L_1 and L_2 are disjoint propositional languages. Let \mathbf{K}_1 and \mathbf{K}_2 be knowledge bases in L_1 and L_2 respectively. Then

$$\mathcal{I}_L[\mathbf{K}_1 \cup \mathbf{K}_2] = \mathcal{I}_{L_1}[\mathbf{K}_1] \cdot \mathcal{I}_{L_2}[\mathbf{K}_2]$$

where the multiplication of the two belief functions on the right is performed in the obvious way to yield a belief function on L .

If an inference process \mathcal{I} has the above property for all $L, L_1, L_2, \mathbf{K}_1$ and \mathbf{K}_2 as above then we say that \mathcal{I} satisfies **Strong Independence**.

The fact that ME satisfies strong independence is proved in [3]. Notice that provided that \mathcal{I} satisfies Language Invariance it follows easily that if \mathcal{I} satisfies Strong Independence then it satisfies Irrelevant Information. However even in the presence of Language Invariance the converse implication does not hold as the Maximin inference process of [6] attests.

We may generalise the property to a social inference \mathfrak{F} by defining \mathfrak{F} to satisfy **Strong Independence** if for any L, L_1 and L_2 as above, and for any sequences of knowledge bases $\mathbf{K}_1, \dots, \mathbf{K}_m$ and $\mathbf{K}'_1, \dots, \mathbf{K}'_m$ for L_1 and L_2 respectively

$$\mathfrak{F}_L[\mathbf{K}_1 \cup \mathbf{K}'_1 \dots \mathbf{K}_m \cup \mathbf{K}'_m] = \mathfrak{F}_{L_1}[\mathbf{K}_1 \dots \mathbf{K}_m] \cdot \mathfrak{F}_{L_2}[\mathbf{K}'_1 \dots \mathbf{K}'_m]$$

Again, provided that \mathfrak{F} satisfies Language Invariance, it follows easily that if \mathfrak{F} satisfies Strong Independence, then it satisfies Irrelevant Information. Since SEP satisfies Language Invariance but not Irrelevant Information, SEP does not satisfy Strong Independence. However, as is noted in [26],

there do exist \mathfrak{F} which satisfy both Language Invariance and Strong Independence, but the only ones known so far are obdurate.

4. An Alternative Definition of SEP

4.1. The Self-Effacing Chairman

A remarkable characteristic of SEP is that the use of maximum entropy at the second second stage of the defining process, which is included in order to force the choice of a social belief function to be unique in cases when this would not otherwise hold, can actually be eliminated by insisting that the social inference process satisfy a limiting variant of the axiom of proportionality. Such an argument counters a possible objection that the invocation of maximum entropy at the second stage of the definition is somewhat artificial. To be precise it is possible to substitute the following procedure to define SEP. We will explain and justify the procedure heuristically before formally stating and proving the corresponding theorem.

In order to calculate a unique social belief function v for a college \mathbf{M} with vector of knowledge bases $\vec{\mathbf{K}} = \mathbf{K}_1 \dots \mathbf{K}_m$, Chairman \mathbf{A}_0 recognises that she may have to use a casting knowledge base of her own in order to eliminate ambiguities caused by the failure of the agreed process of minimising the sum of Kullback-Leibler divergences to provide a *unique* social belief function. However, as a good chairman, she wishes to intervene in a manner which (a) demonstrates that she is completely unbiased, and (b) reduces to an absolute minimum the effect which her own opinion may have on the outcome. In order to fulfil (a) it seems clear to her that she should choose her casting knowledge base \mathbf{K}_0 to be a constraint set \mathbf{I} with

$$\mathbf{V}_{\mathbf{I}} = \{ \langle \frac{1}{J}, \frac{1}{J} \dots \frac{1}{J} \rangle \}$$

Her only other possible choice would seem to be to take \mathbf{K}_0 to be the empty set of constraints, but by Collegiality this would clearly not resolve any ambiguity. On the other hand Chairman \mathbf{A}_0 worries that if she simply adds in her knowledge base \mathbf{I} as a single extra member of the opinion forming body, she may be exerting more influence than is necessary or appropriate, if other opinions are finely balanced. She therefore resolves to dilute her influence in the following manner. Inspired by the Proportionality Principle, she imagines that, for some large finite number n , each member of the college except herself is replaced by exactly n clones, each clone having exactly the same set of constraints as the member replaced; and to this virtual new college of nm members \mathbf{A}_0 adds herself as a single additional member with knowledge base \mathbf{I} as above.

The vector of sets of constraints of the members of the new college of $nm + 1$ members now looks as follows:

$$\mathbf{K}_1, \dots, \mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_2, \dots, \mathbf{K}_m, \dots, \mathbf{K}_m, \mathbf{I}$$

Chairman \mathbf{A}_0 notices that since $\mathbf{V}_{\mathbf{I}}$ is a singleton, by Corollary 13 the result of minimising the sum of Kullback-Leibler divergences subject to these constraint sets will, for any given n always yield a

unique social belief function. She reasons that if as $n \rightarrow \infty$ the resulting sequence of social belief functions converges to a belief function v then this should be an optimal choice of social belief function since her own influence on the process will surely have become as diluted as possible, thus satisfying the condition (b) above. We will prove in Theorem 17 below that not only does this sequence of belief functions converge, but that the resulting limiting belief function v will in fact *always be* $\text{SEP}(\vec{K})$. This is true whether or not $\Delta(\vec{K})$ is a singleton. Consequently Chairman A_0 can reason that her use of ME in the definition of SEP is fully justified by the above heuristic.

4.2. Weak SEP and The Chairman’s Theorem

In order to state formally and prove the result stated above we introduce the following definition:

Definition 16. *The Weak Social Entropy Process, WSEP, is defined by*

$$\text{WSEP}(\vec{K}) = \begin{cases} v & \text{if } \Delta(\vec{K}) \text{ is the singleton } \{v\}, \\ \text{undefined} & \text{otherwise.} \end{cases}$$

□

WSEP is of course not a true social inference process since it is only partially defined. Obviously however $\text{WSEP}(\vec{K}) = \text{SEP}(\vec{K})$ whenever the former is defined.

We will denote the knowledge bases of the college of $nm + 1$ members

$$K_1, \dots, K_1, K_2, \dots, K_2, \dots, K_m, \dots, K_m, I$$

in abbreviated form by

$$n\vec{K}, I.$$

Theorem 17 (The Chairman’s Theorem). *For any \vec{K} and any $n \in \mathbb{N}^+$*

$$\text{WSEP}(n\vec{K}, I) = \text{SEP}(n\vec{K}, I)$$

and furthermore

$$\lim_{n \rightarrow \infty} \text{WSEP}(n\vec{K}, I) = \text{SEP}(\vec{K})$$

Proof. Since V_I is a singleton, by Corollary 13 $\Delta(n\vec{K}, I)$ is always a singleton, and so $\text{WSEP}(n\vec{K}, I)$ is a well-defined point for any $n \in \mathbb{N}^+$, from which the first part of the theorem follows trivially. It does not follow from this that $\Gamma(n\vec{K}, I)$ is a singleton, but we will show below that nevertheless “significant” coordinates are uniquely determined.

For now let us fix n . Then if $\text{WSEP}(n\vec{\mathbf{K}}, \mathbf{I}) = \mathbf{v}$ say, and noting that

$$\text{Sig}_{n\vec{\mathbf{K}}, \mathbf{I}} = \text{Sig}_{\vec{\mathbf{K}}} ,$$

then for every $j \in J$

$$v_j = 0 \text{ if and only if } j \notin \text{Sig}_{\vec{\mathbf{K}}}$$

This is true because if $\vec{\mathbf{w}}$ is a point in $\Gamma(n\vec{\mathbf{K}}, \mathbf{I})$ which generates \mathbf{v} , then if $j \in \text{Sig}_{\vec{\mathbf{K}}}$ then since $w_j^{(mn+1)} = \frac{1}{j}$ it follows from Theorem 11(i) that every entry in the column vector \mathbf{w}_j is non-zero, so v_j is non-zero.

Furthermore for any such $\vec{\mathbf{w}}$ in $\Gamma(n\vec{\mathbf{K}}, \mathbf{I})$ it is clear that the first n rows, *i.e.*, with $i = 1 \dots n$, which correspond to the members with knowledge base \mathbf{K}_1 , must all be identical for those entries $w_j^{(i)}$ with $j \in \text{Sig}_{\vec{\mathbf{K}}}$. For if two of these rows were not so identical then, if they differed in the j 'th entry for some $j \in \text{Sig}_{\vec{\mathbf{K}}}$, we could interchange them to obtain a different point $\vec{\mathbf{w}}'$ in $\Gamma(n\vec{\mathbf{K}}, \mathbf{I})$: however the j 'th column \mathbf{w}'_j could not then be a multiple of \mathbf{w}_j , contradicting Theorem 11. Moreover exactly the same argument works for the second and subsequent blocks of n rows, up to the m 'th block of n rows.

From the above observations it follows that finding an $\vec{\mathbf{w}}$ in $\Gamma(n\vec{\mathbf{K}}, \mathbf{I})$ is essentially the same problem as finding an $\vec{\mathbf{x}} \in \otimes_{i=1}^m \mathbf{V}_{\mathbf{K}_i}$ for which

$$\sum_{j=1}^J \left[\frac{1}{J} \prod_{i=1}^m (x_j^{(i)})^n \right]^{\frac{1}{nm+1}} \text{ is maximal,}$$

or equivalently, for which the function defined by

$$\mathbf{H}_{\epsilon(n)}(\vec{\mathbf{x}}) = \sum_{j=1}^J \left[\left[\prod_{i=1}^m x_j^{(i)} \right]^{\frac{1}{m}} \right]^{(1-\epsilon(n))} \text{ is maximal,}$$

where $\epsilon(n) = \frac{1}{mn+1} \rightarrow 0$ as $n \rightarrow \infty$.

Note that for any $\epsilon(n)$ as above the values of $\left[\prod_{i=1}^m x_j^{(i)} \right]^{\frac{1}{m}}$ for which $\mathbf{H}_{\epsilon(n)}(\vec{\mathbf{x}})$ is maximal are uniquely determined for each $j = 1 \dots J$ and are non-zero if and only if $j \in \text{Sig}_{\vec{\mathbf{K}}}$.

In order to make what follows more readable, we shall temporarily write ϵ instead of $\epsilon(n)$ and suppress the dependence of ϵ on n .

For any such ϵ as above we denote the vector of unique values of $\left[\prod_{i=1}^m x_j^{(i)} \right]^{\frac{1}{m}}$ as defined above by

$$\mathbf{y}_\epsilon = \langle y_{1,\epsilon} \dots y_{J,\epsilon} \rangle \tag{25}$$

and we denote the maximal value of $\mathbf{H}_\epsilon(\vec{\mathbf{x}})$ by m_ϵ , so that

$$m_\epsilon = \sum_{j=1}^J (y_{j,\epsilon})^{1-\epsilon} \tag{26}$$

and let

$$M_\epsilon = \sum_{j=1}^J y_{j,\epsilon} \tag{27}$$

We need to examine the behaviour of y_ϵ as $\epsilon \rightarrow 0$, i.e., as $n \rightarrow \infty$.

Define M_0 to be $M_{\vec{K}}$, i.e the maximum possible value of $\sum_{j=1}^J \left[\prod_{i=1}^m x_j^{(i)} \right]^{\frac{1}{m}}$. By our initial assumptions $M_0 > 0$.

A straight forward consequence of the above definitions is the following:

Lemma 18.

$$M_\epsilon \leq M_0 \leq m_\epsilon \text{ for all } \epsilon \in (0, 1)$$

□

Lemma 19.

$$M_\epsilon \rightarrow M_0 \text{ as } \epsilon \rightarrow 0^+.$$

Proof. We show first that the function $y^{1-\epsilon}$ converges uniformly to y as $\epsilon \rightarrow 0^+$ in the sense that there exists some positive real valued function $T(\epsilon)$ such that for all $y \in [0, 1]$ and all ϵ with $0 < \epsilon < \frac{1}{2}$

$$y^{1-\epsilon} < y + T(\epsilon) \text{ and } \lim_{\epsilon \rightarrow 0^+} T(\epsilon) = 0.$$

Now

$$y^{1-\epsilon} = y e^{-\epsilon \log y}$$

whence, expanding the exponential function as a power series, multiplying by y , and taking out a factor of ϵ , we get

$$y^{1-\epsilon} - y = \epsilon \sum_{k=1}^{\infty} \frac{\epsilon^{k-1} y (-\log y)^k}{k!}$$

The absolute value of $y (\log y)^k$ is at a maximum when $y = e^{-k}$ and hence the absolute value of the k 'th term of the above series is bounded by $\frac{\epsilon^k}{k!} \left[\frac{k}{e} \right]^k$. Since this bound decreases for decreasing ϵ , we have that for all ϵ with $0 < \epsilon < \frac{1}{2}$ and $y \in [0, 1]$

$$y^{1-\epsilon} - y < \epsilon \sum_{k=1}^{\infty} \left[\frac{1}{2} \right]^{k-1} \left[\frac{k}{e} \right]^k \frac{1}{k!}$$

and since the sum converges by d'Alembert's ratio test, the right hand side provides the required function $T(\epsilon)$.

To complete the proof of 4.4 we note that, using 4.3 and the above,

$$M_\epsilon \leq M_0 \leq m_\epsilon = \sum_{j=1}^{\infty} (y_{j,\epsilon})^{1-\epsilon} \leq \sum_{j=1}^{\infty} y_{j,\epsilon} + T(\epsilon) = M_\epsilon + T(\epsilon).$$

Hence, letting ϵ tend to zero, we obtain the required result. □

We now note that for fixed ϵ an equivalent definition of \mathbf{y}_ϵ is as that vector of values which maximises the function

$$\mathbf{G}_\epsilon(\mathbf{y}) = \frac{1}{\epsilon} \log \left[\frac{\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} (y_j)^{(1-\epsilon)}}{M_\epsilon} \right] \tag{28}$$

subject to the conditions that

$$y_j = \left[\prod_{i=1}^m (x_j^{(i)}) \right]^{\frac{1}{m}} \text{ for } j \in \text{Sig}_{\vec{\mathbf{K}}} \text{ , and } \vec{\mathbf{x}} \in \bigotimes_{i=1}^m \mathbf{V}_{\mathbf{K}_i} . \tag{29}$$

For fixed ϵ we will now consider the behaviour of $\mathbf{G}_\epsilon(\mathbf{y})$ for general \mathbf{y} satisfying conditions (29) above. Actually we are only interested in those \mathbf{y} which are either of the form $\mathbf{y}_{\epsilon(n)}$ for some n or which are such that $\sum_{j=1}^J y_j = M_0$, and from now on we shall assume that \mathbf{y} is of this kind. We note that for $j \in \text{Sig}_{\vec{\mathbf{K}}}$ $0 < y_j \leq 1$, and that for such y_j for any $k \in \mathbb{N}^+$ $|y_j(\log y_j)^k|$ is uniformly bounded above by $(\frac{k}{e})^k$, (as in the proof of 4.4).

By 4.4 it follows that

$$J \geq \sum_{j=1}^J y_j > c > 0 \tag{30}$$

for some fixed bound c for all such \mathbf{y} .

Now

$$(y_j)^{1-\epsilon} = (y_j) e^{-\epsilon \log y_j} \tag{31}$$

$$= y_j - \epsilon y_j \log y_j + \sum_{k=2}^{\infty} y_j (-\epsilon \log y_j)^k \tag{32}$$

whence

$$\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} (y_j)^{1-\epsilon} = \sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_j \left[1 - \epsilon \frac{\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_j \log y_j}{\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_j} + O(\epsilon^2) \right] \tag{33}$$

where the term $O(\epsilon^2)$ is such that its modulus is, by the argument in the proof of 4.4, uniformly bounded by $\epsilon^2 D$ for some positive constant D .

Rewriting the equation (4) we now have

$$\mathbf{G}_\epsilon(\mathbf{y}) = \frac{1}{\epsilon} \left[\log \left[1 - \epsilon \frac{\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_j \log y_j}{\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_j} + O(\epsilon^2) \right] + \log \frac{\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_j}{M_\epsilon} \right] \tag{34}$$

Expanding the logarithm as a power series and using (6) we obtain

$$\mathbf{G}_\epsilon(\mathbf{y}) = - \frac{\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_j \log y_j}{\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_j} + \epsilon \mathbf{R}(\epsilon, \mathbf{y}) + \frac{1}{\epsilon} \log \frac{\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_j}{M_\epsilon} \tag{35}$$

where $|\mathbf{R}(\epsilon, \mathbf{y})|$ has a uniform bound independent of \mathbf{y} and of ϵ .

Now notice the following facts about equation (35):

1. For given $\epsilon = \epsilon(n)$ corresponding to a specific value of n , the vector \mathbf{y}_ϵ satisfies

$$\mathbf{G}_\epsilon(\mathbf{y}_\epsilon) = -\frac{\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_{j, \epsilon} \log y_{j, \epsilon}}{M_\epsilon} + \epsilon \mathbf{R}(\epsilon, \mathbf{y}_\epsilon) \tag{36}$$

since the final term vanishes.

2. For any \mathbf{y} for which $\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_j = M_0$ the final term of (35) is positive since $M_\epsilon \leq M_0$ by 4.3.

Let us denote by \mathbf{z} that unique \mathbf{y} for which

$$\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_j = M_0 \quad \text{and} \quad \sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} -y_j \log y_j \quad \text{is maximal.}$$

Then

$$\text{SEP}(\vec{\mathbf{K}}) = \left\langle \frac{z_1}{M_0} \dots \frac{z_J}{M_0} \right\rangle \tag{37}$$

since

$$\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} -y_j \log y_j \quad \text{is maximal if and only if} \quad \sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} -\frac{y_j}{M_0} \log \frac{y_j}{M_0} \quad \text{is maximal.}$$

To complete the proof of the theorem we need to show that

$$\lim_{n \rightarrow \infty} \left\langle \frac{y_{1, \epsilon(n)}}{M_{\epsilon(n)}} \dots \frac{y_{J, \epsilon(n)}}{M_{\epsilon(n)}} \right\rangle = \left\langle \frac{z_1}{M_0} \dots \frac{z_J}{M_0} \right\rangle \tag{38}$$

Since by 4.4 $M_\epsilon \rightarrow M_0$ as $\epsilon \rightarrow 0$, it suffices to show that $\mathbf{y}_\epsilon \rightarrow \mathbf{z}$ as $\epsilon \rightarrow 0$.

Now since all the \mathbf{y} are in $[0, 1]^J$, by compactness the sequence of $\mathbf{y}_{\epsilon(n)}$ for $n \in \mathbb{N}$ has a convergent subsequence, say $\mathbf{y}_{\epsilon(\rho(n))}$, where $\epsilon(\rho(n)) \rightarrow 0$ as $n \rightarrow \infty$.

Let

$$\lim_{n \rightarrow \infty} \mathbf{y}_{\epsilon(\rho(n))} = \mathbf{y}^*. \tag{39}$$

Then from (36) above and the fact that $M_\epsilon \rightarrow M_0$, it follows that

$$\lim_{n \rightarrow \infty} \mathbf{G}_{\epsilon(\rho(n))}(\mathbf{y}_{\epsilon(\rho(n))}) = -\frac{\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_j^* \log y_j^*}{M_0} \tag{40}$$

and

$$\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_j^* = M_0 \tag{41}$$

We now show that $\mathbf{y}^* = \mathbf{z}$.

For suppose for contradiction that this were not so. Let

$$\frac{1}{M_0} \left[-\sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} z_j \log z_j + \sum_{j \in \text{Sig}_{\vec{\mathbf{K}}}} y_j^* \log y_j^* \right] = d \tag{42}$$

Then $d > 0$ since \mathbf{y}^* and \mathbf{z} both have sum M_0 and \mathbf{z} is the unique maximum entropy point. Now by (35)

$$\begin{aligned} \mathbf{G}_{\epsilon(\rho(n))}(\mathbf{z}) &= d - \frac{\sum_{j \in \text{Sig} \vec{\mathbf{K}}} y_j^* \log y_j^*}{M_0} + \epsilon(\rho(n)) \mathbf{R}(\epsilon(\rho(n)), \mathbf{z}) + \frac{1}{\epsilon(\rho(n))} \log \frac{M_0}{M_{\epsilon(\rho(n))}} \\ &\geq d - \frac{\sum_{j \in \text{Sig} \vec{\mathbf{K}}} y_j^* \log y_j^*}{M_0} + \epsilon(\rho(n)) \mathbf{R}(\epsilon(\rho(n)), \mathbf{z}) \end{aligned} \tag{43}$$

However for large enough the right hand-side is strictly greater than $\mathbf{G}_{\epsilon(\rho(n))}(\mathbf{y}_{\epsilon(\rho(n))})$ by (36),(40), and the boundedness of \mathbf{R} .

This is impossible since then $\mathbf{G}_{\epsilon(\rho(n))}(\mathbf{z}) > \mathbf{G}_{\epsilon(\rho(n))}(\mathbf{y}_{\epsilon(\rho(n))})$ which contradicts the definition of $\mathbf{y}_{\epsilon(\rho(n))}$. Thus we have shown that $\mathbf{y}^* = \mathbf{z}$.

It remains to show that the whole sequence of the $\mathbf{y}_{\epsilon(n)}$ converges to \mathbf{z} as $n \rightarrow \infty$. If this were not the case then there would be some $\delta > 0$ such that there exists an infinite subsequence $\mathbf{y}_{\epsilon(\tau(n))}$ of the $\mathbf{y}_{\epsilon(n)}$ such that the $\mathbf{y}_{\epsilon(\tau(n))}$ are bounded away from \mathbf{z} by Euclidean distance $|\mathbf{y}_{\epsilon(\tau(n))} - \mathbf{z}| > \delta$ for all $n \in \mathbb{N}$. However now by compactness again this subsequence $\mathbf{y}_{\epsilon(\tau(n))}$ itself has an infinite convergent subsequence which converges to a point say \mathbf{y}^{**} . By the same argument as for \mathbf{y}^* we must have that $\mathbf{y}^{**} = \mathbf{z}$; on the other hand by its definition \mathbf{y}^{**} is bounded away from \mathbf{z} by distance at least δ , which gives a contradiction. Thus we have established (38) and the proof of Theorem 17 is complete. \square

Remark 3. It is worth noting that in the very special case when there is only a single member \mathbf{A}_1 of the college apart from the Chairman \mathbf{A}_0 , the explanation of Theorem 17 given at the beginning of this section provides a new interpretation of an old technical result. For in this special case, for any $n \in \mathbb{N}$, $\text{WSEP}(n\mathbf{K}_1, \mathbf{I})$ returns that probability function \mathbf{v} which satisfies the constraints \mathbf{K}_1 and which maximises the function

$$\sum_{j=1}^J v_j^{\binom{n}{n+1}}$$

In other words for a given $n \in \mathbb{N}$ this gives the same result as applying the Renyi inference process REN_r with parameter $r = \binom{n}{n+1}$. Now it is an old result (see e.g. [33] or [6]) that as $r \rightarrow 1$ the result of applying the Renyi process REN_r to a given set of constraints \mathbf{K}_1 tends to the maximum entropy solution for \mathbf{K}_1 . So the heuristic explanation underlying Theorem 17 may be regarded as a generalised interpretation of this classical result, albeit from a quite new perspective.

Remark 4. The reader might wonder whether perhaps a rather more general “limit proportionality” theorem than theorem 17 might hold for SEP , which would assert that for any vectors of knowledge bases $\vec{\mathbf{K}}$ and $\vec{\mathbf{K}}'$

$$\lim_{n \rightarrow \infty} \text{SEP}(n\vec{\mathbf{K}}, \vec{\mathbf{K}}') = \text{SEP}(\vec{\mathbf{K}})$$

Such an assertion is however easily seen to be false even in the simplest of cases, simply by choosing $\vec{\mathbf{K}}$ to comprise a single knowledge base which is empty, and $\vec{\mathbf{K}}'$ to comprise a single knowledge base which specifies a single belief function distinct from $\langle \frac{1}{j}, \frac{1}{j} \dots \frac{1}{j} \rangle$.

5. Conclusion

5.1. A Critical Evaluation and Future Directions

In the present paper we first introduced the general notion of a Social Inference Process, which provides an axiomatic framework for the study of how to elicit a single optimally representative probability function derived from partial information about the probabilistic beliefs of several different agents. We have examined in some detail the properties of the particular social inference process **SEP**. In particular we have noted that **SEP** satisfies eight important desiderata: Equivalence, Anonymity, Atomic Renaming, Consistency, Collegiality, General Locality, Proportionality, and Language Invariance¹⁹. Some of these desiderata are relatively easy for a Social Inference Process to satisfy, but Collegiality, and in particular General Locality, are harder to satisfy.

SEP was initially defined in two stages: a merging stage consisting of a merging operator Δ which, given a vector of knowledge bases $\vec{\mathbf{K}}$, yields a merged knowledge base $\Delta(\vec{\mathbf{K}})$. From this merged knowledge base $\Delta(\vec{\mathbf{K}})$ the social belief function is then extracted by an application of **ME**. In Chapter 4 we proved a technical result which shows that the application of **ME** at the second stage fits in a very natural way with the first stage operator $\Delta(\vec{\mathbf{K}})$. The merging operator $\Delta(\vec{\mathbf{K}})$ itself has interesting properties which we have not considered here, but which have been analyzed in [17,26,29]. In particular in [17] it is shown that $\Delta(\vec{\mathbf{K}})$ satisfies a set of conditions for a merging operator close to those defined by Konieczny and Pino Pérez in [34].

Our work on **SEP** raises many unsolved problems. The definition of **SEP** illustrates the information theoretic connections between **ME**, minimum cross entropy, and **LogOp** in the context of multi-agent reasoning. However, whether or not one accepts **SEP** as being the ideal generalisation of **ME** to a social inference process, there are independent reasons to believe that, if such a correct generalisation \mathfrak{F} exists, \mathfrak{F} should marginalise to the **LogOp** pooling operator. One such reason is the remarkable manner in which **ME** and **LogOp** “fit together”, which can be seen by considering results concerning the obdurate social inference process, **OSEP**, defined by

$$\text{OSEP}[\mathbf{K}_1, \dots, \mathbf{K}_m] = \text{LogOp}[\text{ME}(\mathbf{K}_1), \dots, \text{ME}(\mathbf{K}_m)]$$

Adamčík showed in [26] that this simple amalgam of **ME** and **LogOp** satisfies the Strong Independence Principle and Language Invariance, and hence, a fortiori, the Irrelevant Information Principle. The reason why Strong Independence holds is interesting. The Strong Independence property satisfied by **ME** states that if \mathbf{K} and \mathbf{K}' are constraint sets formulated in disjoint propositional languages L and L' respectively and if $\mathbf{K} \cup \mathbf{K}'$ denotes the combined constraint sets in the language $L \cup L'$, then $\text{ME}(\mathbf{K} \cup \mathbf{K}')$ is just the product of the two probability functions $\text{ME}(\mathbf{K})$ and $\text{ME}(\mathbf{K}')$. Now Adamčík observed that the product condition

$$\mathbf{w}^{(i)}(\alpha_j \wedge \alpha'_{j'}) = \mathbf{w}^{(i)}(\alpha_j) \cdot \mathbf{w}^{(i)}(\alpha'_{j'})$$

¹⁹ The first seven of these were announced, but not proved, in the author’s earlier work [1], while Language Invariance was proved in [27].

where α_j and $\alpha'_{j'}$ range over all the respective atoms of the languages L and L' , is actually preserved by the **LogOp** pooling operation, which suffices to prove the Irrelevant Information result. This independence preservation property may surprise some, owing to a general belief that **LogOp** does not “preserve independence” because of an old result of Genest and Wagner [35]. However while the result of [35] is technically perfectly correct, that paper misses the interesting *logical* property of **LogOp** because the authors did not formulate the independence condition appropriately in terms of belief functions on *distinct propositional languages*. In fact, from a logician’s viewpoint, it can be argued that there is no good intuitive reason why the more general “formula-wise” notion of independence preservation used in [35] should actually hold. Indeed the fact noted in [35] that the “formula-wise” notion *is* preserved in the special case of a language with four atoms, which appears to be just an anomalous result in that paper, can easily be seen from a logician’s perspective to be just a special case of Adamčík’s observation above. The situation here is comparable to the controversy over representation independence cited in section 1.1, and again illustrates the value of a foundational approach.

The observations above leave us however with a puzzling situation. We know that there does exist a social inference extending **ME** which satisfies the Strong Independence Principle: namely **OSEP**. This fact is significant because the Strong Independence Principle is hard to satisfy even for an inference process. (Indeed **ME** is the only reasonable inference process known to satisfy it).

Moreover in addition to Strong Independence, **OSEP** also satisfies six of the eight principles satisfied by **SEP** as listed at the beginning of section 5.1: Equivalence, Anonymity, Atomic Renaming, General Locality, Proportionality, and Language Invariance²⁰. However **OSEP** suffers from the usual, and in the opinion of the author, fatal, drawbacks of obdurate inference processes: it satisfies neither Consistency nor Collegiality, nor even the (weaker) Ignorance Principle.

On the other hand, while **SEP** satisfies most of our desiderata for a social inference principle, it fails to satisfy Strong Independence or the Irrelevant Information Principle, although it does satisfy a rather weak version of the latter (cf. [27] and section 3.3 above). This raises the obvious question as to whether there exists any social inference process \mathfrak{F} extending **ME** which satisfies the same desiderata as those which **SEP** satisfies and which also satisfies the Strong Independence, or even the weaker Irrelevant Information Principle. It seems very likely that if such an \mathfrak{F} exists it would have to marginalise to **LogOp**. However a non-existence proof would also be very interesting and would certainly strengthen any claim of **SEP** to foundational optimality.

At this point we should address another obvious criticism of **SEP** which derives from the previously lauded fact that it marginalises to the pooling operator **LogOp**. As a consequence however **SEP** inevitably inherits any criticisms attached to **LogOp**. The most obvious criticism of **LogOp** is its extreme behaviour in the case when any of the agents has belief in some particular atom α close to zero. Indeed an agent A_i with belief $w^{(i)}(\alpha) = 0$ has dictatorial powers, forcing the social belief function v to give the value zero to α . This is clearly useless from any practical point of view. We now examine briefly how this phenomenon can be explained and perhaps remedied.

²⁰ These properties are all noted by Adamčík in [26], with the exception of General Locality and Proportionality. Proportionality holds trivially since it is satisfied by **LogOp**, while General Locality holds for **OSEP** because **ME** satisfies Locality (cf. Theorem 4) and **LogOp** satisfies General Locality (an immediate corollary of Theorem 14 above).

Let us recall the Intersubjectivity Assumption that the chairman treats the knowledge base provided by each agent *as if* it represented intersubjective probabilistic information. This means in effect that the chairman is treating the reported information as if it were intersubjectively trustworthy: in particular the agent is assumed not to be cheating nor to have miscalculated, and any priors which might have been used by her in her calculations are assumed to be hypothetically common to all the agents if they were privy to the same background information. This gives rise to several observations.

The chairman might reason that if pathological priors are ruled out as irrational, and if the number of background observational or mental states of agent is assumed to be finite, then on the basis of the chairman's Intersubjectivity Assumption, no agent should definitively assign probability zero to an event unless the agent considers that event logically impossible. On the other hand the same assumption implies that if one agent considers an event to be logically impossible, then all the agents should do likewise. Thus if the chairman is going to be able to abide by her Intersubjectivity Assumption it is necessary that for any atom α , either *each* of the knowledge bases in $\vec{\mathbf{K}}$ separately forces belief in α to be zero, or none of them do so. The chairman might therefore reasonably insist that for any atom α which is not ruled by prior universal agreement to be logically impossible, no agent shall specify for α a definitive value zero in her knowledge base \mathbf{K}_i . However while the extreme problem caused by zeros might be evaded in this way, the general problem caused by an agent's specified belief value close to zero still remains.

From the chairman's highly idealised viewpoint it now seems that the extreme influence over the social belief which SEP or LogOp gives to an agent who has belief close to zero in a particular atom α is, *given her normative assumptions*, not quite so unreasonable as it at first appeared. The phenomenon arises *precisely because* SEP treats all knowledge bases at face value: since some knowledge bases may be providing much more information than others to the social belief function, the chairman may obviously be in trouble if she does not actually have full trust in each agent's input. It is intuitively clear that an agent who ascribes belief 2^{-40} to some particular propositional variable p is providing more information to the college than an agent who ascribes belief $\frac{1}{2}$ to p , while an agent whose constraint set is empty supplies no information at all. While for normative reasons the chairman has decided to treat each agent's knowledge base as if it represented intersubjective knowledge in the sense described, nevertheless because she actually recognises that the information of the agent may be untrustworthy, she may still wish to *limit the influence* of the agent on the social belief function. Her *revised* attitude could then be summed up as: "I will take the information you supply at face value, but I will attempt to limit your influence on the social belief function in some proportionate manner".

The reasons for the chairman's desire to limit the influence of a particular agent may be of two kinds: *intrinsic* and *extrinsic*. By an *extrinsic reason* we mean some extra information about the agent or the nature of her knowledge, or about the nature of any intended application of the social belief function. By an *intrinsic reason* we mean a natural caution on the part of the chairman to limit the effects on the social belief function of the *knowledge bases* of particular agents, based solely on the combined properties of the knowledge bases, and independently of any information of an extrinsic nature. Here we shall consider only the problem of an intrinsic analysis of influence limitation.

There are a some obvious *ad hoc* methods by which the chairman could attempt to limit the influence of individual agents by modifying SEP. For example, given some $\epsilon > 0$ with $\epsilon \ll \frac{1}{j}$ let \mathbb{D}_ϵ denote the

convex subset of \mathbb{D}_J consisting of all $w \in \mathbb{D}_J$ for which $w_j > \epsilon$ for all $j = 1 \dots J$. Then the chairman, having chosen an ϵ , could replace each V_{K_i} for which $V_{K_i} \cap \mathbb{D}_\epsilon = \emptyset$ by some convex $V_{K_i}^*$ consisting of those points in \mathbb{D}_ϵ which are “informationally closest” to V_{K_i} . (We deliberately leave this imprecise because there are several different plausible interpretations). However, while such *ad hoc* methods may prevent extreme influence by a single agent in simple situations such as when each agent specifies a single belief function, it is not clear that they will have the desired effect in other situations such as that of the counterexample to Open Mindedness given in section 3.3 above.

The “information” contained in a vector of knowledge bases \vec{K} is in general extremely complex, indeed far more so than in the case addressed by pooling operators, because the knowledge bases are likely to contain much information about each agent’s ignorance, a situation which one would expect information theoretic methods to be good at dealing with. It appears to the author that what is needed here in order to deal with the intrinsic problem of trust and influence limitation outlined above, is a foundational study from an information theoretic point of view of the intuitive notion of *degree of influence* of an agent’s knowledge base K_i on the social belief function, relative to the vector of knowledge bases \vec{K} . If such an analysis can be carried out it may suggest naturally justifiable ways in which SEP can be modified in order to limit the influence of an agent to a prescribed degree, which could reflect the degree of trust which we are prepared to place on the information provided by the agent.

5.2. Open Problems

- (1) Is there a social inference process extending ME which satisfies the principles known to be satisfied by SEP (*i.e.*, Language Invariance and those of Theorem 14), and which also satisfies the Strong Independence Principle, or at least the Irrelevant Information Principle?
- (2) Is there any mathematical “number of possible states” argument similar in spirit to that of statistical mechanics which can be used to derive SEP or some other social inference process, in a manner analogous to the classical derivation of ME as in [9]?
- (3) Is there some set of principles which can be used to characterise SEP uniquely in a manner similar to that in which ME is characterised in [2] ?
- (4) Is it possible to develop an information theoretic theory of *influence* and of *trust* along the lines suggested in the previous section, which could be applied to adapt SEP for practical use?
- (5) Are there algorithms for the calculation of SEP which are of comparable efficiency to those available for ME ?
- (6) The quantity $M_{\vec{K}}$ of Definition 8 appears to be a natural measure of the joint consistency of knowledge bases \vec{K} . What are its properties when it is viewed as such a measure?

Acknowledgements

My thanks are due to Alena Vencovská and Martin Adamčík for many stimulating discussions over the last four years, and for pointing out some errors in earlier versions of the text. Thanks are also due

to Jon Williamson for some insightful criticism, and to two anonymous referees for helpful suggestions. I am very grateful to Dugald Macpherson and the School of Mathematics of Leeds University for their collegiality in granting me an academic refuge after my retirement from Manchester University. Finally my gratitude is also due to Jeff Paris for his inspiration and steadfast support at Manchester over an academic lifetime.

Conflicts of Interest

The author declares no conflict of interest.

References

1. Wilmers, G.M. The Social Entropy Process: Axiomatising the Aggregation of Probabilistic Beliefs. *Probability, Uncertainty and Rationality*, Hosni, H. and Montagna, F., Eds.; CRM series, Edizioni Della Normale, Scuola Normale Superiore, Pisa, Italy, **2010**, *10*, 87–104.
2. Paris, J.B.; Vencovská, A. A Note on the Inevitability of Maximum Entropy. *Int. J. Approximate Reasoning* **1990**, *4*, 183–224.
3. Paris, J.B. *The Uncertain Reasoner's Companion - A Mathematical Perspective*; Cambridge University Press: Cambridge, UK, 1994.
4. Genest, C.; Zidek, J.V. Combining probability distributions: A critique and an annotated bibliography. *Stat. Sci.* **1986**, *1*, 114–148.
5. French, S. Group Consensus Probability Distributions: A Critical Survey. *Bayesian Statistics*; Bernardo, J. M., De Groot, M.H., Lindley, D.V., Smith, A.F.M., Eds.; North Holland: Amsterdam, The Netherlands, 1985; pp. 183–201.
6. Hawes, P. An Investigation of Properties of Some Inference Processes. PhD Thesis, Manchester University, Manchester, UK, 2007; MIMS eprints, available from <http://eprints.ma.man.ac.uk/1304/>, accessed on 13 January 2015.
7. Jaynes, E.T. Where do we Stand on Maximum Entropy. *The Maximum Entropy Formalism*; Levine, R.D., Tribus, M., Eds.; 1979; MIT Press: Cambridge, MA, USA.
8. Jaynes, E.T. The Well-Posed Problem. *Found. Phys.* **1973**, *3*, 477–493
9. Paris, J.B.; Vencovská, A. On the Applicability of Maximum Entropy to Inexact Reasoning. *Int. J. Approximate Reasoning* **1989**, *3*, 1–34.
10. Shannon, C.E.; Weaver, W. *The Mathematical Theory of Communication*; University of Illinois Press, Urbana, IL, USA, 1949.
11. Renyi, A. On Measures of Entropy and Information. In *Proceedings of the 4th Berkeley Symposium in Mathematical Statistics*; University of California Press: Oakland, CA, USA, 1961; Volume 1, pp. 547 - 561.
12. Fadeev, D.K. Zum Begriff der Entropie einer endlichen Wahrscheinlichkeitsschemas. *Arbeiten zur Informationstheorie* **1957**, *1*, 85–90. Deutscher Verlag der Wissenschaften, Berlin.
13. Paris, J.B. Common Sense and Maximum Entropy. *Synthese* **1999**, *16*, 75–93.
14. Chomsky, N. Interviewed by Katz, Y. Noam Chomsky on where Artificial Intelligence Went Wrong. *The Atlantic* 2012.

15. Paris, J.B. What You See Is What You Get. *Entropy* **2014**, *16*, 6186–6194.
16. Paris, J.B.; Vencovská, A. In Defense of the Maximum Entropy Inference Process. *Int. J. Approximate Reasoning* **1997**, *17*, 77–103.
17. Adamčík, M.; Wilmers, G.M. Probabilistic Merging Operators. *Logique et Analyse* **2014**, in press.
18. Tversky, A.; Kahnemann, D. Judgement under Uncertainty: Heuristics and Biases. *Science* **1974**, *185*, 1124–1131.
19. Levy, W.B.; Delic, H. Maximum entropy aggregation of individual opinions. *IEEE Trans. Syst. Man. Cybern.* **1994**, *24*, 606–613.
20. Osherson, D.; Vardi, M. Aggregating Disparate Estimates of Chance. *Game Econ. Behav.* **2006**, 148–173.
21. Kracík, J. On composition of probability density functions. *Multiple Participant Decision Making*; In Workshop on Computer-Intensive Methods in Control and Data Processing, Prague, Czech, 12–14 May 2004; pp.113–121.
22. Kracík, J. Cooperation Methods in Bayesian Decision Making with Multiple Participants, Ph.D. Thesis, Czech Technical University, Prague, Czech, 2009.
23. Kern-Isberner, G.; Rödder, W. Belief Revision and Information Fusion on Optimum Entropy. *Int. J. Intell. Syst.* **2004**, *19*, 837–857.
24. Yue, A.; Liu, W. A Syntax-based Framework for Merging Imprecise Probabilistic Logic Programs. *Int. Joint Conf. Artif. Intell.* **2009**, 1990–1995.
25. Myung, J.; Ramamoorti, S.; Bailey, A.D., Jr.; Maximum Entropy Aggregation of Expert Predictions. *Manag. Sci.* **1996**, *42*, 1420–1436.
26. Adamčík, M. Collective Reasoning under Uncertainty and Inconsistency. PhD. Thesis, The University of Manchester, Manchester, UK, 2014.
27. Adamčík, M.; Wilmers, G.M. The Irrelevant Information Principle for Collective Probabilistic Reasoning. *Kybernetika* **2014**, *50*, 175–188.
28. Williamson, J. *Defence of Objective Bayesianism*; Oxford University Press: Oxford, UK, 2010.
29. Adamčík, M. The Information Geometry of Bregman Divergences and Some Applications in Multi-Expert Reasoning. *Entropy* **2014**, *16*, 6338–6381.
30. Kullback, S. *Information Theory and Statistics*; Wiley: New York, NY, USA, 1959.
31. Boyd, S.; Vandenberghe, L. *Convex Optimization*, Cambridge University Press: Cambridge, UK, 2004.
32. Savage, S.D. The Logical and Philosophical Foundations of Social Inference Processes. MSc Dissertation, University of Manchester, Manchester, UK, 2010.
33. Mohamed, I.A.M. Some Properties of the Class of Renyi Generalized Entropies in the Discrete Case. MPhil. Thesis, School of Mathematics, Manchester University, Manchester, UK, 1998.
34. Konieczny, S.; Pino Pérez, R. On the Logic of Merging. In Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning, San Francisco, USA, 1998; pp. 488–498.
35. Genest, C.; Wagner, C.G. Further evidence against independence preservation in expert judgement synthesis. *Aequationes Mathematicae* **1987**, *32*, 74–86.

36. Kreisel, G. Church's Thesis and the Ideal of Informal Rigour. *Notre Dame J. Formal Logic* **1987**, *28*, 499–518.
37. Lakatos, I. *Proofs and Refutations*; Cambridge University Press: Cambridge, UK, 1976.
38. Shore, J.E.; Johnson, R.W. Axiomatic Derivation of the Principle of Maximum Entropy and the Principle of Minimum Cross-Entropy. *IEEE Trans. Inform. Theor.* **1980**, *IT-26*, 26–37.
39. Renyi, A. *Wahrscheinlichkeitsrechnung*. Deutscher Verlag der Wissenschaften, Berlin, Germany, 1962.
40. Cooke, R.M. *Experts in Uncertainty: Opinion and Subjective Probability*; Science, Environmental Ethics and Science Policy Series. Oxford University Press: New York, NY, USA, 1991.
41. Garg, A.; Jayram, T.S.; Vaithyanathan, S.; Zhu, H. Generalized Opinion Pooling. In *Proceedings of the 8th Intl. Symp. on Artificial Intelligence and Mathematics*, Fort Lauderdale, Florida, USA, 4–6 January 2004.
42. Genest, C. A conflict between two axioms for combining subjective distributions. *J. Roy. Stat. Soc.* **1984**, *46*, 403–405.
43. Genest, C.; McConway, K.J.; Schervish, M.J. Characterization of externally Bayesian pooling operators. *Ann. Math. Stat.* **1986**, *14*, 487–501.
44. Wagner, C. Aggregating Subjective Probabilities: Some Limitative Theorems. *Notre Dame J. Formal Logic* **1984**, *25*, 233–240.
45. Wallsten, T.S.; Budescu, D.V.; Erev, I.; Diederich, A. Evaluating and Combining Subjective Probability Estimates. *J. Behav. Decis. Making* **1997**, *10*, 243–268.
46. Carnap, R. On the application of inductive logic. *Philosophy and Phenomenological Research* **1947**, *8*, 133–148.
47. Gillies, D. *Philosophical Theories of Probability*; Routledge: London, UK, 2000.
48. Williamson, J. Deliberation, Judgement and the Nature of Evidence. *Economics and Philosophy* **2014**, in press.
49. Csiszár, I. Why Least Squares and Maximum Entropy? An Axiomatic Approach to Inference for Linear Inverse Problems. *Ann. Stat.* **1991**, *19*, 2032–2066.
50. Csiszár, I. Axiomatic Characterisations of Information Measures. *Entropy* **2008**, *10*, 261–273.
51. Madansky, A. *Externally Bayesian Groups*. Technical Report RM-4141-PR, RAND Corporation, 1964.
52. Genest, C. A characterization theorem for externally Bayesian groups. *Ann. Stat.* **1984**, *12*, 1100–1105.