Theory Dec. (2018) 85:21–39
https://doi.org/10.1007/s11238-017-9639-3

Weighted averaging, Jeffrey conditioning and invariance

Denis Bonnay[1] · **Mikaël Cozic**[2]

Published online: 17 January 2018
© Springer Science+Business Media, LLC 2018

Abstract Jeffrey conditioning tells an agent how to update her priors so as to grant a given probability to a particular event. Weighted averaging tells an agent how to update her priors on the basis of testimonial evidence, by changing to a weighted arithmetic mean of her priors and another agent's priors. We show that, in their respective settings, these two seemingly so different updating rules are axiomatized by essentially the same invariance condition. As a by-product, this sheds new light on the question how weighted averaging should be extended to deal with cases when the other agent reveals only parts of her probability distribution. The combination of weighted averaging (for the events whose probability the other agent reveals) and Jeffrey conditioning (for the events whose probability the other agent does not reveal) is a comprehensive updating rule to deal with such cases, which is again axiomatized by invariance under embedding. We conclude that, even though one may dislike Jeffrey conditioning or weighted averaging, the two make a natural pair when a policy for partial testimonial evidence is needed.

Keywords Jeffrey conditioning · Aggregations of opinion · Linear pooling · Weighted averaging · Geometric averaging · Testimony

✉ Mikaël Cozic
mikael.cozic@u-pec.fr

Denis Bonnay
denis.bonnay@u-paris10.fr

1 Département de Philosophie, Université Paris-Ouest Nanterre and IHPST, 200, avenue de la Republique, 92001 Nanterre Cedex, France

2 Département de philosophie, Université Paris-Est Créteil (EA LIS), IUF and IHPST, 61 avenue du général de Gaulle, 94010 Créteil Cedex, France

1 Introduction

A theory of rational belief typically contains two components: a synchronic component which describes the constraints which an agent's doxastic state should obey, and a diachronic component saying how an agent's doxastic state should be updated upon receiving new information. In this paper, we will be concerned with the dynamics of beliefs, under the assumption that an agent's doxastic state can be represented by a probability distribution.

When one is interested in the issue of knowing how an agent should update her prior probabilities upon receiving new information, one may prima facie distinguish two kinds of situation. In the case of an individual update, the agent learns that some event has a given probability and updates her priors so as to grant that probability to the event. Conditioning (when the new probability of the event is 1) and Jeffrey conditioning (when the new probability of the event is an arbitrary new value) are the two most popular updating rules for this kind of situation. In the case of a social update, the agent learns another agent's opinions and updates her own priors so as to take into account the other's opinions. In this case, the most popular rules consist in averaging the two priors, using weighted arithmetic or geometric means.[1]

Let us have a closer look at social updates. An agent A_1 updates her priors on the basis of the testimony of some other agent A_2, who reveals her degree of belief in an event E. This process may be broken down into two different stages. First, some kind of trade-off occurs, between A_1's own prior towards E and A_2's prior. A_1 has to decide how, and how much, she is going to change her belief in E given A_2's degree of belief in E. The output of this process is A_1's posterior probability for E. The second stage consists in A_1's adjusting her probabilities towards the other events. A_1 has to decide what her degrees of beliefs in all other events become, now that her belief in E has changed.

Theories of individual update (like those mentioned above) typically deal with adjustment. What the posterior probability for the target event should be is taken as given, and the problem is precisely how to do the adjustment for the other events in the algebra. By contrast, theories of social update, as usually stated, only deal with the trade-off task. They answer the question what the posterior probability for E should be,[2] but they remain silent as to how the agent should complete her posterior distribution following the trade-off stage. The question how this should be done has recently been raised, by Jehle and Fitelson (2009) and Steele (2012), but a principled supplementation of trade-off rules, such as weighted averaging, is still lacking. Jehle and Fitelson (2009) are concerned with situations of peer-disagreement. As an updating rule to handle such situations, they consider unweighted arithmetic means supplemented with a constraint of distance minimization (using Euclidean distance) with respect to the agent's priors, but they do not claim to provide a vindication of this particular way to supplement the averaging. Steele (2012) considers supplementing

[1] For surveys, see Genest and Zidek (1986) and Dietrich and List (2016).

[2] We are simplifying a bit: A_2 may well disclose her subjective probabilities for other events, and maybe for all events in the algebra under consideration. But the point is that A_2 may not do that and may only reveal parts of her subjective probabilities.

weighted averaging with Jeffrey conditioning, but ends up rejecting both, on account of failure of commutativity. Neither is there to be found a joint approach to individual and social updates: characterizations of individual updating rules are usually spelled out in terms which are alien to the axiomatic approach widespread in the literature on combining probability distributions. In this paper, we wish to propose such a unified view of individual and social updates on the one hand, and of trade-off and adjustment on the other hand. Our approach will be axiomatic: we aim to find principled axioms from which one can derive adjustment rules for individual updates, trade-off rules for restricted social updates (which assume that A_2's priors are revealed for all events in the algebra), and, finally, trade-off and adjustment rules for unrestricted social updates (which generalize over restricted social updates by allowing for cases when agent A_2 reveals her priors about some but not all events in the algebra).

Surprisingly enough, a single invariance axiom does the job for the three kinds of situations we wish to consider. In Sect. 2, we introduce invariance under embedding and show that it axiomatizes Jeffrey conditioning (J), as an individual updating rule for an agent setting her prior for an event to some new value. This result is a discrete version of a previous characterization of Jeffrey conditioning by Teller (1973) and van Fraassen (1990) which allows itself with the resources of real analysis. Section 3 sets the stage for social updates, in their restricted and unrestricted forms. We then show in Sect. 4 that weighted averaging (WA), as a restricted social updating rule for an agent who wishes to mitigate her priors with another agent's fully disclosed priors, is also axiomatized by invariance under embedding (this is merely a variation on known results). Fruits are ripe to show at the end of Sect. 4 that invariance under embedding axiomatizes weighted averaging extended with Jeffrey conditioning (EWA), as an unrestricted social updating rule for an agent who updates with respect to the partially revealed priors of an another agent. Finally, in Sect. 5, we lay the basis for further work on invariance and updating rules. Building on earlier results by Gilardoni (2002), we suggest how the same strategy may be applied to trade-offs based on geometric rather than linear averaging, leaving open the question what the corresponding adjustment rule would be.

2 Jeffrey conditioning and invariance

We first consider individual update functions, as ways to update one's priors, given that one wishes to set one's subjective probability for a given event to a particular value.[3] Without loss of generality, we first consider updates with respect to only one event at a time. For simplicity, we shall consider a fixed, infinite set Ω of possible worlds. The agents' priors and posteriors shall be represented as probabilities p defined on a finite

[3] We leave aside the motivations for such a wish and the question whether perceptual evidence provides us with that kind of information.

subalgebra S of 2^Ω.[4] The "belief state" of an agent is thus a pair $\mathcal{S} = \langle S, p\rangle$,[5] and we note $\mathbb{S}$ the set of possible belief states.

An update instruction $U = \langle A, r\rangle$ for the belief state $\mathcal{S}$ consists in an event $A \in S$ ($A \neq \Omega, \emptyset$), and a new probability $r \in [0, 1]$ to be attached to that event. We note $\mathbb{U}$ the set of update instructions for some belief state.

Definition 1 An individual updating rule is a continuous[6] function $F : \mathcal{D} \to \mathbb{S}$ where

1. $\mathcal{D} = \{(\mathcal{S} = \langle S, p\rangle, \mathbb{U} = \langle A, r\rangle) : \mathcal{S}$ is a belief state, $\mathbb{U}$ is an instruction for $\mathcal{S}$ and $0 < p(A) < 1\}$.[7]
2. $F(U, \mathcal{S}) = \langle S, p'\rangle$ for some probability distribution p' such that $p'(A) = r$.[8]

For a given instruction U, $F_U(.)$ will refer to $F(U, .)$. F_U maps an old belief states $\langle S, p\rangle$ to a new one $\langle S, p'\rangle$. Since the algebra S is implicitly encoded in p and p' (by being their domain), F_U can equivalently be regarded as a function mapping a probability measure p (on some finite algebra S) to a new one. We shall indeed often regard F_U as a mapping between probability measures p (on finite algebras) rather than belief states $\langle S, p\rangle$. So, for any belief state $\mathcal{S} = \langle S, p\rangle$ with $r \in [0, 1]$, $F_U(\langle S, p\rangle)$ stands for the new belief state and $F_U(p)$ stands for the new probability measure on S, i.e. $F_U(\langle S, p\rangle) = (S, F_U(p))$.
In this framework, Jeffrey conditioning (Jeffrey 1983) is the individual updating rule defined by

$$J_U(p)(B) = p(B|A) \times r + p(B|\neg A) \times (1 - r)$$

with $\mathcal{S} = \langle S, p\rangle$, $U = \langle A, r\rangle$, $A, B \in S$ and $r \in [0, 1]$.

The result of this section will relate Jeffrey conditioning to a property of invariance under embedding of belief states. An embedding of a "smaller‘" belief state into a "bigger" one is to be thought of as an isomorphism between the smaller state and a substate of the bigger one.

Definition 2 (*Embedding*) An embedding of $\mathcal{S} = \langle S, p\rangle$ into $\mathcal{S}' = \langle S', p'\rangle$ is an injective map $f : S \to S'$ such that for all $B, C \in S$

- $f(B \cup C) = f(B) \cup f(C)$ and $f(\overline{B}) = \overline{f(B)}$ (preservation of set-theoretic operators), and
- $p(B) = p'(f(B))$ (preservation of probability)

[4] The results could be extended to infinite σ-algebras, but we find it interesting that they already hold working just with finite ones.

[5] In the following, varying S will prove to be crucial. This is the reason why we make it explicit, instead of merely working with probability distributions on a fixed algebra, as is commonly done in the literature on opinion pooling.

[6] Continuity applies only to one of the arguments of an updating rule, namely the probability function. It means that, given a sequence of belief states of the form $\langle S, p_i\rangle$ and a belief state $\langle S, p\rangle$ such that $\langle S, p\rangle = \lim_{i\to\infty} \langle S, p_i\rangle$, F_U commutes with limits, that is $F_U(\langle S, p\rangle) = \lim_{i\to\infty} F_U(\langle S, p_i\rangle)$. $\langle S, p\rangle = \lim_{i\to\infty} \langle S, p_i\rangle$ is short for $p(A) = \lim_{i\to\infty} p_i(A)$, for all $A \in S$.

[7] We require that $0 < p(A) < 1$ to guarantee that Jeffrey conditioning is well-defined on the whole domain.

[8] This condition is reminiscent of the so-called Success Postulate in AGM-type belief revision.

When f is an embedding and $U = \langle A, r \rangle$ an update instruction, we note $f(U)$ for $\langle f(A), r \rangle$, the update instruction corresponding to U in the bigger belief state. Invariance under embedding says that the result of an update should be the same for a given belief state and for an isomorphic copy of that belief state within a bigger belief state:

Definition 3 (*Invariance under embedding*) An individual updating rule F is invariant under embedding (IE) if, whenever f is an embedding from $\mathcal{S}$ into $\mathcal{S}'$ and U is an updating instruction on $\mathcal{S}$, f is also an embedding from $F_U(\mathcal{S})$ into $F_{f(U)}(\mathcal{S}')$.

In other words, if an updating rule F is invariant under embedding, then the following diagram commutes:

$$\begin{array}{ccc} \mathcal{S} & \overset{f}{\hookrightarrow} & \mathcal{S}' \\ \downarrow{\scriptstyle F_U} & & \downarrow{\scriptstyle F_{f(U)}} \\ F_U(\mathcal{S}) & \overset{f}{\hookrightarrow} & F_{f(U)}(\mathcal{S}') \end{array}$$

Why should we want F to be so invariant? Intuitively, invariance asks for similar results across similar situations. Slightly more precisely, invariance of updatings under a given class of transformations requires that no superfluous information be used in order to determine what the result of the updating is. Information is deemed superfluous if it distinguishes between two scenarios represented by two similar structures, two structures being similar if one can be transformed into the other by a transformation in the class. Now embeddings relate a belief state and an isomorphic copy of that belief state (possibly within a bigger state). Hence invariance under embedding requires that updatings do not take into account information regarding the nature of the events under scrutiny: the updating rule should be formal. Embeddings also relate a belief state 'on its own' and that belief state as part of a broader algebra (possibly modulo an isomorphism). Hence invariance under embedding requires that updatings do not take into account the context surrounding some given events: the rule should be local. Thus, invariance under embedding captures a twofold requirement of formality and locality: the updating rule should be insensitive to content and context.[9] Is such frugality desirable per se? The question goes beyond the scope of the present paper. However, let us remark that frugality is desirable if the updating is to make minimal assumption regarding what the agent knows. When the agent just knows that she should set her degree of belief for some event to a given value, and does not know how

[9] Invariance under embedding could be unpacked into two distinct requirements of invariance under isomorphism and of invariance under substates (we are thankful to the editor in charge for this remark). This will also show up in Sect. 4, in connection with standard axioms for the aggregation of probability distributions, which similarly distinguish the formality and the locality constraints. In the present paper, we wish to push the intuition, familiar in model theory, of invariance under embedding as a sui generis idea: updatings should be insensitive to moving around the belief state in which they operate.

the nature of particular events or the context in which the updating takes place should impact that updating, then she should go for a frugal rule, as captured by invariance under embedding.[10]

From now on, we shall take for granted that invariance under embedding is at least an interesting requirement, possibly vindicated under minimality assumptions on the information available to the updating agent. Our first result is that invariance under embedding exactly characterizes Jeffrey conditioning.

Theorem 1 *Jeffrey conditioning is the only individual updating rule satisfying (IE).*

Proof ($\Rightarrow$). First, let us prove that Jeffrey conditioning satisfies (IE). Let $\mathcal{S} = \langle S, p\rangle$ and $\mathcal{S}' = \langle S', p'\rangle$ be two belief states, $U = \langle A, r\rangle$ an updating instructions for $\mathcal{S}$, and $f : S \to S'$ an embedding of $\mathcal{S}$ into $\mathcal{S}'$.
We want to show that f is an embedding of $\langle S, J_U(p)\rangle$ into $\langle S', J_{f(U)}(p')\rangle$. Let $C \in S$. We need to show that $J_U(p)(C) = J_{f(U)}(p')(f(C))$. If $C = A$, this is obviously the case since $J_{f(U)}(p')(f(A)) = r = J_U(p)(A)$ by definition of an updating rule. When $C \neq A$, we have, by definition of J:

$$J_{f(U)}(p')(f(C)) = p'(f(C)|f(A)) \times r + p'(f(C)|\neg f(A)) \times (1 - r)$$
$$J_U(p)(C) = p(C|A) \times r + p(C|\neg A) \times (1 - r)$$

One needs to show that

$$p'(f(C)|f(A)) = p(C|A)$$
$$p'(f(C)|\neg f(A)) = p(C|\neg A)$$

By Bayes rule, the first equality amounts to

$$\frac{p'(f(A) \cap f(C))}{p'(f(A))} = \frac{p(A \cap C)}{p(A)}$$

Hence it is sufficient to remark that f is an embedding, since this implies that $p'(f(A)) = p(A)$ and that $p'(f(A) \cap f(C)) = p'(f(A \cap C)) = p(A \cap C)$. The second equality holds for similar reasons.
($\Leftarrow$). We have to show that if an individual updating rule F satisfies (IE), then F is Jeffrey conditioning. So let F be an individual updating rule satisfying (IE), $\mathcal{S} = \langle S, p\rangle$ a belief state and $U = \langle A, r\rangle$ an updating instruction for $\mathcal{S}$. In what follows, we assume that p has values in $[0, 1] \cap \mathbb{Q}$, the general result follows by continuity of F. For $B \in S$ with $B \neq A$, we need to show that

$$F_U(p)(B) = p(B|A) \times r + p(B|\neg A) \times (1 - r)$$

[10] van Fraassen (1990) appeals to invariance under embedding in the wider context of a vindication of basic laws by symmetry principles. He does not provide a detailed vindication in that particular case, but he points out that, when a rule is invariant under embedding, “we are allowed to switch our attention to a more tractable ’equivalent’ probability space” (p. 334).

For simplicity, let B be an atom (that is, a minimal element of the algebra distinct from the empty set) and $B \subset A$. This implies that we will have to show that $F_U(p)(B) = p(B|A) \times r$ since $p(B|\neg A) = 0$. The proof for $B \subset \neg A$ is similar, and the property for other sets in S follows from the property for atoms.
Since p has values in the rational numbers, we can write $\frac{p(B)}{p(A)}$ as $\frac{m}{n}$ for some integers m and n with $m \leq n$. Now let us consider two new belief states $\mathcal{S}' = \langle S', p' \rangle$ and $\mathcal{S}'' = \langle S'', p'' \rangle$. Let (X, Y, Z) be a ternary partition of Ω. S' is defined as the algebra whose atoms are X, Y and Z. In addition,

- $p'(X) = p(B)$,
- $p'(Y) = p(A \backslash B)$, and
- $p'(Z) = p(\neg A)$.

S'' is just like S' except that X is split into m new atoms $X_1, \ldots, X_m$ and Y is split into $n - m$ new atoms $Y_1, \ldots, Y_{n-m}$. p'' is defined by

- $p''(X_i) = p''(Y_j) = \frac{p(A)}{n}$.

Note that $p''(X_1 \cup \cdots \cup X_m) = p'(X) = p(B)$ and $p''(Y_1 \cup \cdots \cup Y_{n-m}) = p'(Y) = p(A \backslash B)$.

There is an obvious embedding from $\mathcal{S}'$ to $\mathcal{S}$ such that the image of X is B, the image of Y is $A \backslash B$ and the updating instruction $U' = \langle X \cup Y, r \rangle$ on $\mathcal{S}'$ is turned into U on $\mathcal{S}$. Similarly, there is an obvious embedding from $\mathcal{S}'$ to $\mathcal{S}''$ such that the image of X is $X_1 \cup \cdots \cup X_m$, the image of Y is $Y_1 \cup \cdots \cup Y_{n-m}$ and U' is turned into $U'' = \langle X_1 \cup \cdots \cup X_m \cup Y_1 \cup \cdots \cup Y_{n-m}, r \rangle$ on $\mathcal{S}''$. Since F satisfies (IE), F treats all three updatings similarly:

$$F_U(\mathcal{S})(B) = F_{U'}(\mathcal{S}')(X) = F_{U''}(\mathcal{S}'')(X_1 \cup \cdots \cup X_m\})$$

Hence, it is sufficient to show that $F_{U''}(\mathcal{S}'')(X_1 \cup \cdots \cup X_m) = p(B|A) \times r$. Consider any $M, N \in \{X_1, \ldots, X_m, Y_1, \ldots, Y_{n-m}\}$ and $f_{M,N} : S'' \rightarrow S''$ be the function which substitutes M for N (and vice-versa) in each set of S''. $f_{M,N}$ is an embedding of $\mathcal{S}''$ into itself. Moreover, $X_1 \cup \cdots \cup X_m \cup Y_1 \cup \cdots \cup Y_{n-m}$ is its own image. Since F satisfies (IE), it follows that

$$F_{U''}(\mathcal{S}'')(M) = F_{U''}(\mathcal{S}'')(N)$$

That is, atoms, which were equiprobable before the updating, still are after the updating. This gives what we wanted: since all the atoms $X_1, \ldots, X_m, Y_1, \ldots, Y_{n-m}$ are equiprobable and $F_{U''}(\mathcal{S}'')(X_1 \cup \cdots \cup X_m \cup Y_1 \cup \cdots \cup Y_{n-m}) = r$,

$$\begin{aligned} F_{U''}(\mathcal{S}'')(X_1 \cup \cdots \cup X_m) &= m \times \frac{F_{U''}(\mathcal{S}'')(X_1 \cup \cdots \cup X_m \cup Y_1 \cup \cdots \cup Y_{n-m})}{n} \\ &= m \times \frac{r}{n} \\ &= p(B|A) \times r \end{aligned}$$

Thus, $F_U(\mathcal{S})(B) = p(B|A) \times r$.

Theorem 1 also generalizes to the case where the updating instruction is partitional, i.e. to the case where the agent's degrees of belief adjust to a change towards the elements of a partition $\mathbf{A} = (A_i)_{i \leq n}$. In this case, an updating instruction has the form $U = \langle \mathbf{A}, (r_i)_{i \leq n} \rangle$ with $r_i \geq 0$ and $\sum_i r_i = 1$ and the prior for each A_i in (0, 1). The exact same proof would show that within each A_i, for an atom $B \subset A_i$, $F_U(\mathcal{S})(B) = p(B|A_i) \times r_i$.

Let us compare our result to other characterizations of Jeffrey conditioning. A simple way to get Jeffrey conditioning is to require that the updating rule be rigid, that is probabilities conditional on the events involved in the updating instruction do not change.[11] Relative to an updating instruction $U = \langle \mathbf{A}, (r_i)_{i \leq n} \rangle$ and for all $A_i \in \mathbf{A}$, rigidity may be spelled out as:

$$p(.|A_i) = F_U(p)(.|A_i) \quad (R)$$

Equivalently, Jeffrey conditioning may be characterized by the preservation of the ratios of probabilities of atoms belonging to the same cell of the relevant partition. In other words, let B_i, B_j be atoms belonging to the cell $A_k \in \mathbf{A}$,

$$p(B_i)/p(B_j) = F_U(p)(B_i)/F_U(p)(B_j) \quad (R')$$

Technically, (R') is the condition closest to invariance under isomorphism, which straightforwardly implies within-cell preservation of equality of equiprobable atoms—the "embedding" condition further giving the full force of (R'). Note that that both (R) and (R') may be viewed as ways of capturing the idea that the belief change achieved by Jeffrey conditioning should be minimal, since they characterize Jeffrey conditioning as an updating rule keeping at least some things unchanged.

Teller (1973) showed that within-cell preservation of equality for atoms characterizes Jeffrey conditioning, under the assumption that the probability measure is 'full'. Fullness is defined for a probability measure p by requiring that for each event A in the algebra, for each real number r lower than $p(A)$, there is an event $B \subset A$ such that $p(B) = r$. In a continuous rather than discrete setting, van Fraassen (1990) defines the same notion of invariance under embedding we have been using in this section and relies on Teller's result to show that it characterizes Jeffrey conditioning. Van Fraassen's proof consists in showing that invariance under embedding (or, in that case, just invariance under isomorphism) implies within-cell preservation of equality of atoms, and that any probability space can be embedded into a full space (this is the part where embeddings, and not just isomorphisms, are needed). We provide a simpler proof which works for the discrete case, where the embedding into an algebra with equiprobable atoms makes it unnecessary to appeal to the theorem of calculus about additive functions used by Teller and van Fraassen.[12]

[11] Dietrich et al. (2016) have recently proposed a characterization of Jeffrey conditioning in terms of a property called "conservativeness", that generalizes rigidity. In a nutshell, conservativeness says that the updating process should leave unchanged the probabilities on which the update instruction is "silent", in that case conditional probabilities.

[12] We are grateful to an anonymous referee for pointing us to Teller and van Fraassen's work, which we were unaware of when we first up with Theorem 1.

There are also several approaches to Jeffrey conditioning in terms of distance minimization. The idea is to think of Jeffrey conditioning as the rule whose output minimizes the distance (formally defined) to the prior distribution (May 1976; Williams 1980; Van Fraassen 1980). These attempts are surveyed by Diaconis and Zabell (1982) (Sect. 5). For instance, Williams (1980) proves that Jeffrey conditioning minimizes the Kullback–Leibler measure of the information in p relative to p^0 defined as $I(p, p^0) = \sum_{\omega \in \Omega} p(\omega) \times \ln(p(\omega)/p^0(\omega))$. Again, albeit with a different approach to minimality, this amounts to requiring that the change in belief is minimal.

The approach in terms of invariance provides a seemingly distinct way of deriving Jeffrey conditioning. As explained above, invariance under embedding simultaneously enforces the ideas of formality and locality. The updating mechanism should be insensitive to content and context. This embodies a view of minimal change which differs from the ones just discussed. Those concern the result of the updating, which would be as close as possible to the agent's priors (by retaining some of their properties, or by being the closest with respect to some notion of distance). By contrast, invariance under embedding puts minimality at play in a different manner: it is the process of belief change itself which is to be minimal, in the sense of being informationally frugal (blindness to content and context). Viewed in this light, our result shows a convergence between two approaches of minimality: requiring that the output of the process is minimally different from the output, or that the mechanism producing that output is informationally frugal, amounts to the same thing.

We have been working in a standard Bayesian setting, where priors are given and the problem is to update those priors. Invariance under embedding has been used by Halpern and Koller (2004) in a different setting, where a "knowledge base" constrains probability distributions, and inferences are drawn by selecting certain probability measures compatible with the knowledge base (e.g. those maximizing entropy). Invariance under embedding is then meant to ensure that those inferences are not sensitive to the format of representation. This suggests a general project of setting down invariance conditions both for the type of inferences considered by Halpern and Koller (2004) and for updating mechanisms. However, in the former case, invariance may prove to be too strong a constraint, since it deals only with the algebra of events and the knowledge base, before any assignment of probabilities. Accordingly, the results in Halpern and Koller (2004) for invariance under embedding as a constraint on the choice of (family of) probability measures are mostly negative, showing that such choices are hard to come by.

3 Weighted averaging and its completion with Jeffrey conditioning

We shall now consider social updating rules, in order to model scenarios where one agent updates her priors when she comes to know another agent's priors. In general, the other agent may not reveal her whole prior distribution, in which case the first agent only learns about a substate of the belief state representing her own priors. As a limit case, she may learn the other's degree of belief only toward some specific event (and, by implication, its complement).

Definition 4 (*Social updating rules*) A social updating rule is a partial function $F : \mathcal{D}' \rightarrow \mathbb{S}$ such that

1. $\mathcal{D}' = \{(\langle S, p\rangle, \langle S', p'\rangle) : \text{(i) } S' \subseteq S \text{ and (ii) for any atom } A_i \text{ of } S', p(A_i) > 0.$
2. if $F(.,.)$ is defined for $(\mathcal{S}, \mathcal{S}')$, the result is a belief state $F(\mathcal{S}, \mathcal{S}') = \langle S, p^*\rangle$ for some probability p^*.

Since F induces a map taking a pair of probability measures as argument, $F(p, p')$ will stand for the new probability measure p^*, i.e. $F(\mathcal{S}, \mathcal{S}') = \langle S, F(p, p')\rangle$.

As discussed in Sect. 1, a social updating rule F is a restricted rule if $F(\mathcal{S}, \mathcal{S}')$ is defined only when $S = S'$, that is when the other agent's priors are totally disclosed. It is unrestricted, accepting partial inputs as well, if it is also defined when $S' \subset S$. Weighted averaging (also called Linear Pooling) is usually defined only as a restricted rule. It is parametrized by a weight λ representing how much impact is granted to the other agent's opinions, the weight for the agent's own opinions being $1 - \lambda$.

Definition 5 The rule of weighted averaging $F_\lambda^W : \mathcal{D}' \rightarrow \mathbb{S}$ is the restricted social rule defined as follows for some $\lambda \in [0, 1]$ and all $A \in S$

$$F_\lambda^W(p, p')(A) = (1 - \lambda) \cdot p(A) + \lambda \cdot p'(A)$$

How can the rule of weighted averaging be extended into an unrestricted rule, dealing with partial updatings as well? The issue has been raised recently by Jehle and Fitelson (2009) and Steele (2012). A natural suggestion is to appeal to Jeffrey conditioning, as follows:

Definition 6 The extended rule of weighted averaging $EF_\lambda^W : \mathcal{D}' \rightarrow \mathbb{S}$ is an unrestricted social rule defined, for $A \in S$, by

$$EF_\lambda^W(p, p')(A) = \sum_i p(A|A_i) \times F_\lambda^W(p, p')(A_i)$$

where $A_1, \ldots, A_n$ are the atoms of S' which partition the domain of $\mathcal{S}$.

Note that when $A \in S \cap S'$, EF_λ^W reduces to F_λ^W. Steele (2012) considers using EF_λ^W but eventually rejects it on account of issues of non-commutativity, and abandons F_λ^W altogether. We shall not aim at a thorough defense of the non-commutativity of Jeffrey conditioning and weighted averaging, but a few words are in order, since our axiomatization also relies on commutativity, albeit of a different kind.

Let us discuss briefly the case for weighted averaging by considering a toy scenario. I have a prior belief in some event A of 0.5. Two witnesses, Joe and John, have different priors for A of, respectively, 0.5 and 0.8. Assume that the doxastic weight I grant to Joe is 0.9 and the doxastic weight I grant to John is 0.5. If I first meet John and then Joe, applying weighted averaging, I end up with a posterior of 0.515 for A. If I first meet Joe and then John, my belief for A will eventually be 0.65. Critics of weighted averaging find this order-dependence unwelcome and contend that the end result should be the same, no matter whether I first meet Joe or John. This relies on the faulty assumption that doxastic weights should themselves be insensitive to the order of encounters. But this is not so. I might be very unconfident about my degree of belief in A to start with, so that I am willing to much defer to John, and average our priors for A by weighting his prior with 0.8 and mine with 0.2. But if I have first met Joe, and taken into account his own prior for A by considering him rather reliable, then my grounds for my belief in A have improved, and it does not seem right to still grant John's a doxastic weight of 0.8: intuitively, the figure should be lower. This leaves open the questions how doxastic weights are to be computed, and whether a different kind of order independence should hold when weights are properly adjusted to the order of encounters. Answering those questions goes well beyond the goal of the present paper, but what has been said suggests that failure of commutativity per se is by no means a knock-down argument against weighted averaging. Interestingly, the non-commutativity of Jeffrey conditioning may be defended along the same lines. As Bradley (2005) puts it, "the same experiences have different effects on my probabilities depending on the order in which they occur". There is no reason to require that conditioning by giving probability 0.9 to A, and then, say, probability 0.3 should yield the same result as giving first giving probability 0.3 to A and then 0.9. This is because, even if this corresponds to the same experiences in a different order, the revision triggered by one and the same experience should not be the same depending on whether this experience occurs first or second. Again, this leaves open whether some order independence should hold for Jeffrey conditioning when order dependent differences in impact are taken into account. But it suffices to show that failure of commutativity per se is not a knock-down objection against Jeffrey conditioning either. In both cases, the objection wrongly takes for granted that the impact of evidence should itself be insensitive to order.

The commutativity property that we have used to axiomatize Jeffrey conditioning in the previous section and that we will use to axiomatize weighted averaging in the next section is of a very different kind. It does not concern the order of experiences, but merely the way experience is represented. What the results to come in the next section will suggest is that, if one is willing to countenance weighted averaging, Jeffrey conditioning appears as a very natural completion, but also, in the other direction, that the same ideas which speak in favor of Jeffrey conditioning also speak in favor of weighted averaging. As a consequence, if one likes Jeffrey conditioning as an individual updating rule, one should also like weighted averaging (at least in similar doxastic conditions), and the other way around.

Another point worth mentioning concerns the interpretation of extended weighted averaging.[13] The interpretation we have favored is in terms of a two-stage process, launched by a social updating stage based on some testimonial probability, say $p'(A)$. However, $p'(A)$ could also be viewed as the probability suggested by "Nature". From this point of view, the "new" probability $(1 - \lambda) \times p(A) + \lambda \times p'(A))$ which serves as input for Jeffrey conditioning (the second stage of the process) is a sort of trade-off between the agent's priors $p(A)$ and Nature's probabilities $p'(A)$. Under this interpretation, extended weighted averaging is not a way of completing weighted

[13] We thank the editor of this issue for having drawn our attention to this interpretive point.

averaging, but rather a way of completing Jeffrey conditioning: Jeffrey conditioning (alone) takes as exogeneous the new probabilities to which it is applied.

4 Axiomatizing weighted averaging and extended weighted averaging

There are well-known axiomatic characterizations of weighted averaging as a restricted social updating rule (see McConway 1981; Wagner 1982). In this section, we will review the properties which are involved in these standard axiomatizations and introduce a new one, which is nothing but invariance under embedding suitably extended to the framework of social updating rules. We will show that this invariance property also characterizes weighted averaging.

Let us begin with two properties which capture the idea that the updating rule operates ‘locally’. The most straightforward expression of this idea is the property of eventwise independence (EI).

Definition 7 (*Eventwise independence*) A social updating rule F satisfies eventwise independence (EI) iff for any belief states $\mathcal{S}_1 = \langle S, p_1\rangle$, $\mathcal{S}'_1 = \langle S, p'_1\rangle$, $\mathcal{S}_2 = \langle S, p_2\rangle$ and $\mathcal{S}'_2 = \langle S, p'_2\rangle$, for any $A \in S$,

$$\text{if } p_1(A) = p_2(A) \text{ and } p'_1(A) = p'_2(A) \text{ then } F(p_1, p'_1)(A) = F(p_2, p'_2)(A)$$

Eventwise independence corresponds to the ‘Weak Setwise Function Property’ (McConway 1981) and to the ‘Irrelevance of Alternatives’ axiom (Wagner 1982). It turns out that (EI) is equivalent to the property called ‘Marginalization’ by McConway (1981).

Definition 8 Given a belief state $\mathcal{S} = \langle S, p\rangle$ and an algebra T with $S \subseteq T$ or $S \supseteq T$, the marginal probability distribution p^T on $T \cap S$ is simply defined as $p^T(A) = p(A)$ for $A \in T \cap S$. Similarly, $\mathcal{S}^T$ denotes the belief state $\langle T \cap S, p^T\rangle$.

Definition 9 (*Marginalization property*) A social updating rule F has the marginalization property (MP) iff, for belief states $\mathcal{S}, \mathcal{S}'$ with S' a sub-algebra of S, for any algebra T on Ω with $S' \subseteq T \subseteq S$ or $T \subseteq S' \subseteq S$,

$$F(p, p')^T(A) = F(p^T, p'^T)(A), \quad \text{for all} \quad A \in T$$

Another important idea is the one according to which the updating rule operates formally, independently of the content of the events at hand. One way to express this idea is as follows:

Definition 10 (*Invariance under isomorphism*) A social updating rule F is invariant under isomorphism (II), iff, for any belief states $\mathcal{S} = \langle S, p\rangle, \mathcal{S}' = \langle S', p'\rangle, \mathcal{T} = \langle T, q\rangle$ and $\mathcal{T}' = \langle T', q'\rangle$, with $S' \subseteq S$ and $T' \subseteq T$, any function f

- which is an isomorphism (i.e., a bijective embedding) from $\mathcal{S}$ to $\mathcal{T}$, and
- whose restriction to S' is an isomorphism from $\mathcal{S}'$ to $\mathcal{T}'$

is also an isomorphism from $\langle S, F(p, p')\rangle$ to $\langle T, F(q, q')\rangle$.

Wagner (1982) considers a weaker version of II, which amounts to invariance under automorphism and which he calls 'Label Neutrality". A different, but maybe more obviously desirable property, is preservation of unanimity, at least when both agents give probability zero.

Definition 11 (*Zero unanimity*) A social updating rule F has the zero unanimity property (ZU) iff, for any belief states $\mathcal{S} = \langle S, p \rangle$ and $\mathcal{S}' = \langle S', p' \rangle$ with $S' \subseteq S$, for any $A \in S \cap S'$,

$$\text{if } p(A) = p'(A) = 0 \quad \text{then} \quad F(p, p')(A) = 0$$

The conjunction of (EI) and either (II) or (ZU) is equivalent to the following property:

Definition 12 (*Neutrality*) A social updating rule F satisfies Neutrality (N) iff for any belief states $\mathcal{S}_1 = \langle S, p_1 \rangle$, $\mathcal{S}_2 = \langle S, p_2 \rangle$ and any belief states $\mathcal{S}'_1 = \langle S', p'_1 \rangle$ and $\mathcal{S}'_2 = \langle S', p'_2 \rangle$, both based on $S' \subseteq S$, for any $A_1, A_2 \in S'$,

$$\text{if } p_1(A_1) = p_2(A_2) \quad \text{and} \quad p'_1(A_1) = p'_2(A_2) \quad \text{then}$$
$$F(p_1, p'_1)(A_1) = F(p_2, p'_2)(A_2)$$

Neutrality corresponds to 'Strong Setwise Function Property' in McConway (1981) and to 'Strong Label Neutrality' in Wagner (1982). To these known conditions, we add a new one, which is nothing but the social version of the invariance condition considered in Sect. 2.

Definition 13 (*Invariance under embedding*) A social updating rule F is invariant under embeddings (IE), iff for any belief states $\mathcal{S} = \langle S, p \rangle$, $\mathcal{S}' = \langle S', p' \rangle$ $\mathcal{T} = \langle T, q \rangle$ and $\mathcal{T}' = \langle T', q' \rangle$, with $S' \subseteq S$ and $T' \subseteq T$, any function f which is an embedding from $\mathcal{S}$ to $\mathcal{T}$ and whose restriction to S' is an embedding from $\mathcal{S}'$ to $\mathcal{T}'$ is also an embedding from $F(\mathcal{S}, \mathcal{S}')$ to $(\mathcal{T}, \mathcal{T}')$.

In other words, if an updating rule F is invariant under embedding, then we obtain the following commuting diagram:

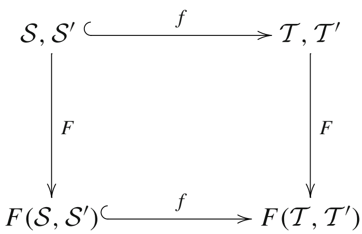

Theorem 2 *If F is a restricted social rule, the following are equivalent for belief states with at least three atoms:*

(i) *$F = F^W_\lambda$ for some λ*
(ii) *F satisfies (N)*

(iii) *F satisfies (MP) and (ZU)*
(iv) *F satisfies (EI) and (ZU)*
(v) *F satisfies (EI) and (II)*
(vi) *F satisfies (MP) and (II)*
(vii) *F satisfies (IE)*[14]

Proof The equivalences between (i), (ii) and (iii) are proven by McConway (1981) (Theorems 3.2 and 3.3). The equivalences between (i), (ii), (iv) and (v) are proven by Wagner (1982). (Theorems 1, 2 and 4[15]) The equivalence between (vi) and (i)–(v) follows from the previous equivalences and the fact proven by McConway (1981) that (MP) is equivalent to (EI). All we need to show is that (IE) implies one of the previous properties and is implied by another.
(vii) implies (vi) This follows from the following two facts: First, (IE) implies (II), since every isomorphism is an embedding. Second, (IE) implies (MP). Let $\mathcal{S}$, $\mathcal{S}'$ with S' a sub-algebra of S, and an algebra T on Ω with $S' \subseteq T \subseteq S$ or $T \subseteq S' \subseteq S$. There is an obvious embedding f from $\mathcal{S}^T = \langle T \cap S, p^T \rangle$ to $\mathcal{S} = \langle S, p \rangle$ and from $\mathcal{S}'^T = \langle T \cap S', p'^T \rangle$ to $\mathcal{S}' = \langle S', p' \rangle$. (IE) implies that $F(p, p')(f(A)) = F(p^T, p'^T)(A)$, for all $A \in T \cap S$ and thus that $F(p, p')^T(A) = F(p^T, p'^T)(A)$, for all $A \in T$.
(i) implies (vii) Let the belief states $\mathcal{S} = \langle S, p \rangle$, $\mathcal{S}' = \langle S', p' \rangle$ $\mathcal{T} = \langle T, q \rangle$ and $\mathcal{T}' = \langle T', q' \rangle$, and a function f which is an embedding from $\mathcal{S}$ to $\mathcal{T}$ and (whose restriction to $S' \subseteq S$ is) an embedding from $\mathcal{S}'$ to $\mathcal{T}'$. We have to show that f is also an embedding from $\langle S, F_\lambda^W(p, p') \rangle$ to $\langle T, F_\lambda^W(q, q') \rangle$. More specifically, we have to show that probabilities are preserved by the updating rule. Let $A \in S$. Clearly, $F_\lambda^W(p, p')(A) = (1 - \lambda) \cdot p(A) + \lambda \cdot p'(A) = (1 - \lambda) \cdot q(f(A) + \lambda \cdot q'(f(A)) = F_\lambda^W(q, q')(f(A))$. □

Incidentally, the proof that (vii) implies (vi) shows how invariance under embedding breaks down into a formality component, captured by invariance under isomorphism, and a locality component, captured by the marginalization property.

Theorem 2 says in particular that weighted averaging is the only restricted social updating rule which satisfies invariance under embedding. Putting this together with Theorem 1, we get a straightforward characterization of extended weighted averaging.

Theorem 3 *Let F be an unrestricted social rule. The following are equivalent:*

(i) *F is* EF_λ^W *for some* λ
(ii) *F satisfies (IE)*

Proof The proof is a simple combination of the proofs for Theorems 1 and 2. □

[14] We have formulated (IE) in a fully general form which applies to unrestricted rules. This is not needed for the present Theorem, which deals with restricted rules (dealing with total inputs, that is cases where $S = S'$ and $T = T'$).

[15] Actually, (ZU) is not implied by (N) in Wagner's framework. The implication holds, however, in McConway's and ours. This is because the empty set is among the considered events. Thus, as soon as there is 0-unanimity towards some arbitrary event A, the updated degree of belief must be zero in virtue of (N).

Thus, the very same requirements of locality and formality which are captured by invariance under embedding characterize Jeffrey conditioning alone (as an individual updating rule), weighted averaging alone (as a restricted social updating rule) and their combination (as an unrestricted social updating rule). However, it should be said that invariance under embedding does not play exactly the same role in both cases. In the case of Jeffrey conditioning, invariance under isomorphism is almost sufficient: if the algebra already has some nice properties, such as fullness, as used by Teller (1973), or equiprobability of atoms, as used in Sect. 2, embeddings are superfluous. In the case of weighted averaging, the appeal to embeddings is more crucial, and cannot be short-circuited by similar assumptions. Embeddings are needed to guarantee that the posterior of the agent for a given even is determined by her prior for that event and the other agent's prior. Also, the global properties are not quite the same. Whereas Jeffrey conditioning is rigid, weighted averaging, as is well known, does not preserve independence.

5 Geometric averaging and invariance

If invariance is to be recognized as a natural constraint on updating rules, the fact that it forces the use of arithmetic means and bans other kind of means is rather worrisome. Geometric averaging, for example, is a rather much praised alternative to (linear) weighted averaging, and one may wonder why it ends up being ruled out by the invariance requirement. After all, is invariance under embedding captures (informational) frugality, why should geometric averaging turn out to be less frudal than arithmetic averaging? The answer is that a lot depends on the framework: when prior beliefs are represented by means of a probability distribution fixing the absolute value of the probability of each event, invariance does end up forcing arithmetic means. But should one choose another way to represent probabilistic beliefs, different things happen. Building on earlier work by Gilardoni (2002), we show that some invariance-type condition can also shed some light on geometric averaging.

In order to ensure the well-definedness of geometric averaging, we will restrict the domain of our updating rules to $\mathcal{D}'' = \{(\langle S, p\rangle, \langle S', p'\rangle) : \text{(i) } S' \subseteq S, \text{ and (ii) } p \text{ is regular}\}$.[16]

Definition 14 The rule of geometric averaging $F^G_\lambda : \mathcal{D}'' \to \mathbb{S}$ is a social rule for total inputs (only) defined for atoms $W \in At(S)$ by

$$F^G_\lambda(p, p')(W) = \frac{p'(W)^\lambda \times p(W)^{1-\lambda}}{\sum\limits_{W' \in At(S)} p'(W')^\lambda \times p(W')^{1-\lambda}}$$

It is well known that geometric averaging does not satisfy eventwise independence (EI) because of the normalization factor $c = 1 \;\; / \sum\limits_{W' \in At(S)} p'(W')^\lambda \times p(W')^{1-\lambda}$. This implies (and it can be easily shown by a numerical example) that it does not satisfy

[16] p is said to be regular if it assigns non-zero probabilities to all non-empty events.

invariance under embedding (IE) either. Interestingly, however, Gilardoni (2002) formulates and studies a condition of independence for ratios of atomic probabilities, which leads to geometric averaging and is defined as follows: Let an algebra S and $At(S)$ its set of atoms. The conditional odds $O_{\mathcal{S}} : At(S) \times At(S) \to (0, \infty)$ associated with a belief state $\mathcal{S} = \langle S, p \rangle$ are defined for all $W, W' \in At(S)$ as

$$O_{\mathcal{S}}(W, W') = \frac{p(W)}{p(W')}$$

We may describe a belief state by taking conditional odds as primitive. Instead of requiring independence of (absolute) probabilities, one may require independence of conditional odds. This means that the conditional odds of the updated probability $O_{F(\mathcal{S},\mathcal{S}')}(W, W')$ depend only on the conditional odds $O_{\mathcal{S}}(W, W')$ and $O_{\mathcal{S}'}(W, W')$.

Definition 15 (*Ratio atomwise independence*) A social updating rule F satisfies ratio atomwise independence (RAI) iff for any belief states $\mathcal{S}_1 = \langle S, p_1 \rangle$, $\mathcal{S}'_1 = \langle S, p'_1 \rangle$, $\mathcal{S}_2 = \langle S, p_2 \rangle$ and $\mathcal{S}'_2 = \langle S, p'_2 \rangle$, for any $W, W' \in At(S)$,

$$\text{if } O_{\mathcal{S}_1}(W, W') = O_{\mathcal{S}_2}(W, W') \text{ and } O_{\mathcal{S}'_1}(W, W') = O_{\mathcal{S}'_2}(W, W'), \text{ then}$$
$$O_{F(\mathcal{S}_1,\mathcal{S}'_1)}(W, W') = O_{F(\mathcal{S}_2,\mathcal{S}'_2)}(W, W')$$

Geometric averaging satisfies ratio atomwise independence. Gilardoni (2002) proves that ratio atomwise independence and invariance under automorphism (also known as "Weak Label Neutrality"), together with a condition of monotonicity and a condition of preservation of unanimity, implies geometric averaging (Proposition 3.4). It turns out that we may define a suitable notion of invariance (that we shall call "Invariance under Ratio-Embedding") which plays with respect to ratio atomwise independence, invariance under automorphism and geometric averaging the same role as the one played by IE with respect to eventwise independence, invariance under automorphism and weighted averaging (see Theorem 2). Specifically, we shall show that (i) invariance under ratio-embedding implies ratio atomwise independence and invariance under automorphism (and, more generally, under isomorphism), and thus that (ii) invariance under ratio-embedding (modulo the conditions of monotonicity and preservation of unanimity) implies geometric averaging.

Definition 16 (*Ratio-embedding*) Let $\mathcal{S} = \langle S, p \rangle$ and $\mathcal{S}' = \langle S', p' \rangle$ be two belief states. $f : S \to S'$ is a ratio-embedding if it is an injective map between S and S' such that

- $f(B \cup C) = f(B) \cup f(C)$ and $f(\overline{B}) = \overline{f(B)}$ (preservation of set-theoretic operators), and
- for all $W, W' \in At(S) \cap f^{-1}(At(S'))$, $O_{\mathcal{S}}(W, W') = O_{\mathcal{S}'}(f(W), f(W'))$ (preservation of ratio of probabilities of atoms)

Note that this definition of ratio-embedding requires the preservation of ratios of atoms only when the images of atoms of $\mathcal{S}$ are themselves atoms of $\mathcal{S}'$.

Definition 17 (*Invariance under ratio-embedding*) A social updating rule F satisfies invariance under ratio-embedding (IRE), iff, for any belief states $\mathcal{S}_1 = \langle S_1, p_1 \rangle$, $\mathcal{S}_1' = \langle S_1', p_1' \rangle$, $\mathcal{S}_2 = \langle S_2, p_2 \rangle$, $\mathcal{S}_2' = \langle S_2', p_2' \rangle$ with $S_1' \subseteq S_1$ and $S_2' \subseteq S_2$, any function f which is a ratio-embedding from $\mathcal{S}_1$ to $\mathcal{S}_2$ and whose restriction to S' is a ratio-embedding from $\mathcal{S}_1'$ to $\mathcal{S}_2'$, is also a ratio-embedding from $F(\mathcal{S}_1, \mathcal{S}_1')$ to $F(\mathcal{S}_2, \mathcal{S}_2')$.

Theorem 4 *Geometric averaging satisfies invariance under ratio-embedding.*

Proof Let the belief states $\mathcal{S}_1 = \langle S_1, p_1 \rangle$, $\mathcal{S}_1' = \langle S_1', p_1' \rangle$, $\mathcal{S}_2 = \langle S_2, p_2 \rangle$, $\mathcal{S}_2' = \langle S_2', p_2' \rangle$. Assume that the function f is a ratio-embedding from $\mathcal{S}_1$ to $\mathcal{S}_2$ and that its restriction to S_1' is a ratio-embedding from $\mathcal{S}_1'$ to $\mathcal{S}_2'$. Let $W, W' \in f^{-1}(At(S_2))$.

$$\begin{aligned} O_{F_\lambda^G(\mathcal{S}_1, \mathcal{S}_1')}(W, W') &= \frac{F_\lambda^G(p_1, p_1')(W)}{F_\lambda^G(p_1, p_1')(W')} \\ &= \frac{p_1'(W)^\lambda \times p_1(W)^{1-\lambda}}{p_1'(W')^\lambda \times p_1(W')^{1-\lambda}} \\ &= O_{\mathcal{S}'_1}(W, W')^\lambda \times O_{\mathcal{S}_1}(W, W')^{1-\lambda} \end{aligned}$$

Similarly,

$$\begin{aligned} O_{F_\lambda^G(\mathcal{S}_2, \mathcal{S}_2')}(f(W), f(W')) &= \frac{F_\lambda^G(\mathcal{S}_2, \mathcal{S}_2')(W)}{F_\lambda^G(\mathcal{S}_2, \mathcal{S}_2')(W')} \\ &= \frac{p_2'(f(W))^\lambda \times p_2(f(W))^{1-\lambda}}{p_2'(f(W'))^\lambda \times p_2(f((W'))^{1-\lambda}} \\ &= O_{\mathcal{S}_2'}(f(W), f(W'))^\lambda \times O_{\mathcal{S}_2}(f(W), f(W'))^{1-\lambda} \end{aligned}$$

By assumption, $O_{\mathcal{S}_2}(f(W), f(W')) = O_{\mathcal{S}_1}(W, W')$ and $O_{\mathcal{S}_2'}(f(W), f(W')) = O_{\mathcal{S}'_1}(W, W')$. Thus,

$$O_{F_\lambda^G(\mathcal{S}_2, \mathcal{S}_2')}(f(W), f(W')) = O_{F_\lambda^G(\mathcal{S}_1, \mathcal{S}_1')}(W, W')$$

Hence f is also a ratio-embedding from $F(\mathcal{S}_1, \mathcal{S}_1')$ to $F(\mathcal{S}_2, \mathcal{S}_2')$. □

Theorem 5 *Let F be a restricted social updating rule which satisfies invariance under ratio-embedding. Then*

1. *F satisfies invariance under isomorphism*
2. *F satisfies ratio atomwise independence*

Proof 1. That IRE implies II follows from the fact that when there is a bijective (and not only an injective) map between S and S' which preserves set-theoretic operators, preservation of absolute probabilities and preservation of conditional odds become equivalent.

2. IRE implies ratio atomwise independence. Let $\mathcal{S}_1 = \langle S, p_1 \rangle$, $\mathcal{S}'_1 = \langle S, p'_1 \rangle$, $\mathcal{S}_2 = \langle S, p_2 \rangle$ and $\mathcal{S}'_2 = \langle S, p'_2 \rangle$ and $W, W' \in At(S)$ such that

$$O_{\mathcal{S}_1}(W, W') = O_{\mathcal{S}_2}(W, W') \text{ and } O_{\mathcal{S}'_1}(W, W') = O_{\mathcal{S}'_2}(W, W')$$

Let T be the algebra generated by W, W' and a new element W^*. Let $\mathcal{T} = \langle T, q \rangle$ where
- $q(W) = p_1(W)$
- $q(W') = p_1(W')$
- $q(W^*) = 1 - p_1(W) - p_1(W')$

Note that we have based $q(.)$ on $p_1(.)$, but we could have chosen any distribution such that the ratios of the probabilities of W and W' is equal to $O_{\mathcal{S}_1}(W, W') = O_{\mathcal{S}_2}(W, W')$. $\mathcal{T}'$ is defined similarly on the basis of $p'_1(.)$. There is an obvious ratio-embedding (but not an embedding) from $\mathcal{T}$ to both $\mathcal{S}_i$ for $i = 1, 2$ and from $\mathcal{T}'$ to $\mathcal{S}'_i$. Since by assumption invariance under ratio-embedding holds,

$$O_{F(\mathcal{T},\mathcal{T}')}(W, W') = O_{F(\mathcal{S}_i,\mathcal{S}'_i)}(W, W')$$

Thus,

$$O_{F(\mathcal{S}_1,\mathcal{S}'_1)}(W, W') = O_{F(\mathcal{S}_2,\mathcal{S}'_2)}(W, W')$$

□

From this Theorem and Gilardoni (2002)'s above mentioned characterization of geometric averaging, it follows that invariance under ratio-embedding (together with monotonicity and preservation of unanimity) also implies geometric averaging.

The next question is how should geometric averaging be extended into an unrestricted rule dealing with partial inputs? It can be shown that Jeffrey conditioning does not satisfy the individual version of invariance under ratio-embedding. Which class of individual updating rules would remain an open question.

6 Conclusion

In this paper, we have studied social updating, a doxastic process which may be decomposed into a "trade-off" stage and an "adjustment" stage. Models of doxastic updating often consider only one of these stages. We have shown that a single invariance axiom, which simultaneously captures intuitions of locality and formality, is powerful enough to axiomatize each stage in isolation, and both in combination as well. Specifically, it implies weighted averaging for the trade-off stage, Jeffrey conditioning for the adjustment stage and a combination of both for the whole process.

The axiomatization of weighted averaging (as a restricted social rule dealing with total inputs) by invariance under embedding should come as no surprise for those familiar with the extant results: the axiom bears close resemblance to neutrality and marginalization. However, the fact that the same axiom suitably applied also yields Jeffrey conditioning and thus the completion of weighted averaging by Jeffrey conditioning is probably much less expected. The main lesson to be drawn from our results is, therefore, that this set of rules form a coherent, unified, package of updating rules.

Finally, invariance conditions are familiar from model theory, and they play a special role in the philosophy of logic, where it has been argued that invariance characterizes what it is to be a logical operation. The fruitfulness of invariance in dealing with updating rules raises intriguing questions. From a conceptual perspective, should we think of updating rules as logical operations? If we should, does their normative import stem from their logical status? From a technical perspective, Sect. 5 paves the way for an extension of the strategy: are there more rational operations which can be characterized as operations which are invariant under natural transformations of the structures to which they apply? Could it be that rational operations in general are nothing but logical operations on doxastic models?

Acknowledgements Earlier versions of this work were presented at seminars and workshops in Paris, Munich and Stockholm, as well as at a symposium at PSA 2014 in Chicago. The present paper greatly benefited from the audiences questions and comments, and special thanks are due to Richard Bradley, Jan-Willem Romeijn and Olivier Roy who were our partners in crime on several of those occasions. We are also particularly thankful to David Etlin, for pointing us to Teller's work, to the reviewers and the editor of the present journal for their careful reading which made for significant improvements in both content and form, and to the members of the *Décision, Rationalité et Interaction* team in Paris. We both acknowledge support from the ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL Grants. Mikael Cozic was also supported by the Institut Universitaire de France (Junior Fellowship).

References

Bradley, R. (2005). Radical probabilism and bayesian conditioning. *Philosophy of Science*, *72*(2), 342–364.

Diaconis, P., & Zabell, S. L. (1982). Updating Subjective Probability. *Journal of the American Statistical Association*, *77*(380), 822–830.

Dietrich, F., List, C. (2016) Probabilistic opinion pooling. In C. Hitchcock and A. Hajek, editors, *Handbook of Probability and Philosophy*. Oxford University Press, forthcoming.

Dietrich, F., List, C., & Bradley, R. (2016). Belief revision generalized: a joint characterization of bayes's and jeffrey's rules. *Journal of Economic Theory*, *162*, 352–371.

Van Fraassen, B. (1980). Rational belief and probability kinematics. *Philosophy of Science*, *47*(2), 165–187.

Genest, Ch., & Zidek, J. (1986). Combining Probability Distributions: A Critique and an Annotated Bibliography. *Statistical Science*, *14*(2), 487–501.

Gilardoni, G. (2002). On Irrelevance of Alternatives and Opinion Pooling. *Brazilian Journal of Probability and Statistics*, *16*, 87–98.

Halpern, J. Y., & Koller, D. (2004). Representation dependence in probabilistic inference. *Journal of Artificial Intelligence Research*, *21*, 319–356.

Jeffrey, R. (1983). *The Logic of Decision* (2nd ed.). USA: University of Chicago Press.

Jehle, D., & Fitelson, B. (2009). What is the Equal Weight View? *Episteme*, *6*(3), 280–293.

May, S. (1976). Probability Kinematics: A Constrained Optimization Problem. *Journal of Philosophical Logic*, *5*, 395–398.

McConway, K. J. (1981). Marginalization and Linear Opinion Pools. *Journal of the American Statistical Association*, *76*(374), 410–414.

Steele, K. (2012). Testimony as Evidence: More Problems for Linear Pooling. *Journal of Philosophical Logic*, *41*, 983–999.

Teller, P. (1973). Conditionalization and observation. *Synthese*, *26*(2), 218–258.

van Fraassen, B. C. (1990). *Laws and Symmetry*. : Clarendon Press.

Wagner, C. (1982). Allocation, Lehrer Models and the Consensus of Probabilities. *Theory and Decision*, *14*, 207–220.

Williams, P. M. (1980). Bayesian Conditionalisation and the Principle of Minimum Information. *British Journal for the Philosophy of Science*, *31*, 131–144.