



The Rational Group

Franz Dietrich

► To cite this version:

| Franz Dietrich. The Rational Group. 2020. halshs-02905409v2

HAL Id: halshs-02905409

<https://shs.hal.science/halshs-02905409v2>

Submitted on 23 Jul 2020 (v2), last revised 26 Jan 2021 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

CES

Centre d'Économie de la Sorbonne
UMR 8174

The Rational Group

Franz DIETRICH

2020.14



The Rational Group

December 2019, revised June 2020

Franz Dietrich

Paris School of Economics & CES & CNRS fd@franzdietrich.net

www.FranzDietrich.net

Abstract

Can a group be a standard rational agent? This would require the group to hold aggregate preferences which maximise expected utility and change only by Bayesian updating. Group rationality is possible, but the only preference aggregation rules which support it (and are minimally Paretian and continuous) are the *linear-geometric rules*, which combine individual tastes linearly and individual beliefs geometrically.

Keywords: Bayesian aggregation, preference aggregation under uncertainty, expected-utility hypothesis for groups, Bayesian revision, rational group agents, linear versus geometric opinion pooling

JES codes: D71, D81

Acknowledgments: Personal acknowledgments to be added. An earlier version was circulated and presented under the title ‘Fully Bayesian Aggregation’. This research was supported by the French National Research Agency through the grants “Coping With Heterogeneous Opinions” (ANR-17-CE26-0003) and "Collective Attitude Formation" (ANR-16-FRAL-0010) and an EUR.

1 Introduction

Economics and other social sciences work with a well-established paradigm of a rational agent. They routinely apply this paradigm to groups such as households, firms, governments, or entire societies. These group agents make decisions, form and revise plans, engage in interactions, compete on markets, or entertain international relations. They hold and revise preferences and beliefs in just the same rational way as individuals. But is a rational group agent actually possible and meaningful, given heterogeneous (rational) group members? That is, could any rational group agent emerge from combining conflicting attitudes of group members?

This problem is of obvious interest, but for two very different reasons. Firstly, we aim for aggregational micro-foundations for the hypothesis of rational groups, to legitimize our modelling practice. Working with rational groups without modelling individuals is useful and convenient, but it would be comforting, to say the least, if group agents could be interpreted as aggregations of (unmodelled) group members. Secondly, we are sometimes explicitly interested in individuals, and need to combine their attitudes into rational group attitudes, for instance in order to determine ‘fair’ group choices, which might be carried out by a group representative or ‘planner’. This goal is constructive: we wish to build a rational group agent that respects some given group members.

Existing aggregation theories provide powerful results that take us some way towards a rational group. But they have never aimed for a full-fledged rational group agent, to the best of my knowledge. For instance, Arrovian preference aggregation ignores uncertainty, while Bayesian preference aggregation captures uncertainty, but ignores the group’s response to information. Our question is therefore alive: can a group be a standard rational agent?

The theory of Bayesian preference aggregation offers the right conceptual and formal tools for addressing our question. This theory seeks to combine individual expected-utility preferences under uncertainty. Perhaps surprisingly, this apparently static framework has enough resources for studying the rational group in the full sense of orthodox rationality, including the (neglected) dynamic dimension of group rationality, Bayesian updating. But first, what does this theory already teach us? If group members have identical beliefs, combining their expected-utility preferences is perfectly possible, but the Pareto principle implies that group utility must be a linear combination of individual utilities, by Harsanyi’s Theorem (Harsanyi 1955). The picture changes under heterogeneous beliefs: we can then no longer construct any group expected-utility preferences which meet the Pareto principle (Mongin 1995). Is this already the end of group rationality under uncertainty? No. Gilboa, Samet and Schmeidler (2004) (‘GSS’) proved the literature’s central possibility result: a suitably restricted Pareto principle makes it again possible to combine expected-utility preferences, but firstly group utility must be a linear combination of individual utilities and secondly group probability must be a linear combination of individual probabilities. I call such preference aggregation *linear-linear*. Further developments are discussed in Section 5.

GSS’s theory has two shortcomings: it does not make the group fully rational, and

relies on a Pareto axiom that is arguably still too strong, as we shall see. Fixing both problems will lead us to *linear-geometric* aggregation: tastes are combined linearly, but beliefs geometrically.

How are linear-linear groups irrational? Despite its name, Bayesian preference aggregation theory has pursued only a semi-Bayesian (‘semi-rational’) agenda. It has aimed for the group agent to follow the static Bayesian norm, setting aside the dynamic Bayesian norm. That is, the group should be an expected-utility (EU) maximiser, but can violate Bayes’ rule when information arrives. The theory is simply silent on the group’s revision behaviour; it does not discipline revision. Bayesian updating is however a cornerstone of classic rationality. A household or other group which updates its preferences irrationally conflicts with our models, and with our paradigm of ‘rational households’ and, more generally, ‘rational group agents’. Such a group agent displays dynamically incoherent behaviour, and runs into the very same well-known problems and paradoxes as dynamically incoherent individuals. It suffers preference reversals during dynamic decision problems and games. It can no longer form and execute stable plans, jeopardizing intertemporal budget planning. It becomes vulnerable to Dutch books, i.e., engages in sequential betting behaviour that leads to sure loss. So it can be exploited – by third parties or even group members.

Linear-linear aggregation creates a statically, but not dynamically rational group agent, as has been complained (e.g., Mongin and Pivato 2020). How else must preferences be combined to make the group rational? Two reasons might explain why this natural problem is open. One is that full group rationality might seem to be an unreachable goal, since already static group rationality is so hard to reach. The other reason lies in the literature’s predominant single-profile approach of working with fixed individual preferences, not a domain of possible preference profiles (a notable exception is Blackorby, Donaldson and Weymark, 2008). This approach makes preference *change* unaddressable.

This paper contributes the following theorem: only linear-geometric preference aggregation creates a rational group agent that (for non-triviality) ‘reflects’ group members, i.e., depends on them in a minimally Paretian and continuous way. A first lesson is that full group rationality is non-trivially possible, supporting the standard hypothesis of rational group agents. Another lesson is that group rationality requires combining beliefs non-classically (i.e., geometrically), but combining tastes classically (i.e., linearly).

Section 2 gives an example, Section 3 states the theorem, Sections 4 and 5 discuss the theorem and the literature, and the Appendix proves the theorem in a generalized version.

2 Illustration of group-preference change

We reconsider GSS’s classic story, but in a dynamic variant. The story is pure fiction, but its structure is typical for real group agents. The group consists of two gentlemen 1 and 2. They have a dispute and must decide whether to fight a duel. The outcome of a duel is either that 1 wins (and 2 loses) or that 2 wins (and 1 loses), depending on

a state of nature. There are three states:

- in state s_1 , 1 is stronger than 2, so would win a duel.
- in state s_2 , 2 is stronger than 1, so would win a duel.
- in state s_3 , 2 has a superior weapon (and is equally strong), so would win a duel.

In all states, having no duel has the outcome that nobody wins. Both gentlemen are fully rational: they hold EU preferences and update them via Bayes' rule. We consider two time points: before and after learning the event $E = \{s_1, s_2\}$ that 2 does not have a superior weapon. Table 1 displays the gentlemen's utilities of both outcomes of a duel,

	utility of		old prob. of			old EU	new prob. of			new EU
	1 wins	2 wins	s_1	s_2	s_3	of duel	s_1	s_2	s_3	of dual
gentleman 1	1	-5	.85	.05	.1	.1	.94	.06	0	.67
gentleman 2	0	1	.15	.15	.7	.85	.5	.5	0	.5
linear-linear group	.5	-2	.5	.1	.4	-.75	.72	.28	0	-.19
linear-geometric group	.5	-2	.51	.12	.37	-.74	.80	.20	0	.01

Table 1: Tastes, beliefs, and expected utilities before and after learning the event $E = \{s_1, s_2\}$ (numbers are rounded to two decimal digits)

the probabilities of states, and the expected utilities of a duel, before and after learning E . Note different things. Each gentleman most prefers winning himself (utility 1). While gentleman 1 fears dying (utility -5), the reckless and honour-obsessed gentleman 2 does not mind dying (utility 0). Each gentleman initially believes strongly that he would win a duel, and updates his probabilities rationally via Bayes' rule. At each moment, the gentlemen have conflicting utilities and conflicting beliefs, yet unanimously prefer duelling, as duelling gives positive expected utility while not duelling gives zero expected utility.

Table 1 also displays *group* utilities, probabilities, and expected utilities, under two alternative aggregation rules for forming group EU preferences:

- The *linear-linear rule* defines group utility as the (unweighted) linear average of individual utilities; and similarly for group probability.
- The *linear-geometric rule* defines group utility as the (unweighted) linear average of individual utilities; but it defines group probability as the (unweighted) geometric average of individual probabilities, normalized to a probability function. For instance, the old group probability of s_3 is $k(.1)^.5(.7)^.5 \approx 0.37$, where k is the normalisation factor $1/[(.85)^.5(.15)^.5 + (.05)^.5(.15)^.5 + (.1)^.5(.7)^.5]$.

Under both rules, not duelling is initially collectively better than duelling – against the gentlemen's unanimous preference. Such Pareto violations have been at the heart of Bayesian aggregation theory, but this paper instead asks whether the group updates its preferences rationally. This is not the case under the linear-linear rule: according to Bayes' rule, the new group probabilities of s_1 and s_2 should have been $\frac{0.5}{0.5+0.1} \approx .83$ and $\frac{0.1}{0.5+0.1} \approx .17$ rather than .72 and .28, and the new group expected utility of duelling should have been $\frac{0.5}{0.5+0.1}0.5 + \frac{0.1}{0.5+0.1}(-2) \approx .08$ rather than -0.19. So the group should have come to prefer duelling. By contrast, under the linear-geometric rule the new group probabilities and expected utility in Table 1 arise from the old ones via Bayes'

rule, as one can check and as our theorem will imply generally. As the new expected utility of duelling is $.01 > 0$, the information makes duelling collectively superior (in our fictional setting which ignores the unacceptability of duels).

The dynamic rationality of linear-linear aggregation cannot be restored by using *weighted* linear averages and allowing weights to depend on the preference profile, hence on information states.¹

3 The theorem

We consider a group of individuals $N = \{1, \dots, n\}$ ($n \geq 2$). I use the Anscombe-Aumann framework. Let X be a finite set X of *outcomes* ($|X| \geq 2$), and \mathcal{X} the set of *lotteries*, i.e., probability functions over X , capturing objective risk (and formally defined on the power set 2^X). The probability of an outcome x under a lottery a is $a(x) = a(\{x\})$. A *utility function* is a function $u : X \rightarrow \mathbb{R}$, representing tastes or values. It is *normalized* if minimal utility is $\min_{x \in X} u(x) = 0$ and maximal utility is $\max_{x \in X} u(x) = 1$. As usual, it is extended to lotteries by taking expectations: $u(a) := \mathbb{E}_a(u)$ for lotteries $a \in \mathcal{X}$.

To allow for subjective uncertainty, let S be a non-empty finite set of *states*. Sets of states are *events*. We allow the single-state case $|S| = 1$, but exclude the two-state case $|S| = 2$, in which our theorem curiously does not hold.²

Choice options are functions $a : S \rightarrow \mathcal{X}$ (*acts*), representing the prospect of facing lottery $a(s)$ in state s . I use similar symbols for acts and lotteries (*'a'*, *'b'*, ...), because lotteries are identifiable with special acts: constant acts, containing no subjective uncertainty. A *preference relation* is a binary relation \succeq over acts, formally $\succeq \subseteq \mathcal{X}^S \times \mathcal{X}^S$; we write \succ for its asymmetric component (representing strict preference) and \sim for its symmetric component (representing indifference). A state $s \in S$ is *null* under \succeq if the outcome at s is irrelevant, i.e., acts that agree on $S \setminus \{s\}$ are indifferent.

A preference relation \succeq has *expected-utility* type – ‘is EU’ – if it maximizes some expected-utility function, i.e., there are a non-constant utility function $u : X \rightarrow \mathbb{R}$ and a probability function p on 2^S such that

$$a \succeq b \Leftrightarrow \mathbb{E}_p(u(a)) \geq \mathbb{E}_p(u(b)) \text{ for all acts } a, b \in \mathcal{X}^S.$$

Here, p is unique, and u is unique if we impose normalization. The unique p and normalized u are denoted p_\succeq and u_\succeq , respectively. Let \mathcal{P} be the set of EU preference relations.

Distributions of tastes, beliefs, or preferences across individuals are captured by ‘profiles’. A *taste profile* is a vector $\mathbf{u} = (u_i)$ of normalized utility functions of individuals $i \in N$. A *belief profile* is a vector $\mathbf{p} = (p_i)$ of probability functions on the set 2^S of events. An *EU preference profile* is a vector $(\succeq_i) \in \mathcal{P}^N$ of EU preference relations; it induces a taste profile and a belief profile, given by (u_{\succeq_i}) and (p_{\succeq_i}) , respectively.

¹Such ‘linear-linear rules with variable weights’ still violate Bayes’ rule, except if beliefs are combined dictatorially by concentrating all weight on some individual.

²In the two-state case, our axioms are necessary, but no longer sufficient for linear-geometric aggregation, as they for instance permit linear-linear aggregation.

Bayesian aggregation theory usually works with a fixed preference profile. To study preference *change*, we take a multi-profile approach: we consider an entire domain \mathcal{D} of possible profiles. For generality, we assume little about \mathcal{D} . Formally, a *domain* is any set of profiles $\mathcal{D} \subseteq \mathcal{P}^N$ satisfying two natural conditions. First, each profile $(\succeq_i) \in \mathcal{D}$ is *coherent*: individuals have mutually consistent beliefs, i.e., at least one state $s \in S$ is non-null under each \succeq_i ($i \in N$). Coherence is plausible because presumably some state in S is ‘true’ and hence not excluded by any rational individual.³ Second, \mathcal{D} is *closed under belief change*: whenever \mathcal{D} contains (\succeq_i) then \mathcal{D} also contains all other coherent profiles in \mathcal{P}^N which differ from (\succeq_i) in beliefs, not in tastes. This assumption allows us to study belief change.

An *EU preference aggregation rule*, or simply a *rule*, is a function $F : \mathcal{D} \rightarrow \mathcal{P}$, transforming EU preference profiles from some domain \mathcal{D} into group EU preference relations.

Definition 1 A rule $F : \mathcal{D} \rightarrow \mathcal{P}$ is *linear-geometric* if there exist individual weights $\alpha_i \in \mathbb{R}$ and $\beta_i \in \mathbb{R}_+$ ($i \in N$) where $\sum_{i \in N} \beta_i = 1$ such that for each preference profile $(\succeq_i) \in \mathcal{D}$ the group preference relation $\succeq = F((\succeq_i))$ has

- u_\succeq given by $\sum_{i \in N} \alpha_i u_{\succeq_i}$ up to an additive constant,
- p_\succeq given on S by $\prod_{i \in N} [p_{\succeq_i}]^{\beta_i}$ up to a multiplicative constant.⁴

Linear-geometric rules differ from classic ‘linear-linear’ rules, which are definable analogously, by replacing the second bullet point by ‘ $p_\succeq = \sum_{i \in N} \beta_i p_{\succeq_i}$ ’.

Our key axiom on a rule $F : \mathcal{D} \rightarrow \mathcal{P}$ requires rational revision of group preferences. A rational agent conditionalizes preferences after learning an event. The *conditionalization* of a preference relation \succeq on an event $E \subseteq S$ is the preference relation \succeq_E obtained by applying Bayesian revision to \succeq . What does this mean? If \succeq is EU (and E is non-null), \succeq_E arises from conditionalizing beliefs and leaving utilities unchanged: $p_{\succeq_E} = p_\succeq(\cdot|E)$ and $u_{\succeq_E} = u_\succeq$. The general definition of \succeq_E , due to Savage, needs no reference to utilities and probabilities: for any acts a, b , we define $a \succeq_E b$ to mean that $a' \succeq_E b'$ for some acts a', b' such that (i) within the event E , a' agrees with a and b' agrees with b , and (ii) outside E , a' agrees with b' . So, $a \succeq_E b$ means that a becomes weakly preferred to b after equalizing (‘ignoring’) outcomes outside E . This is the axiom:

Bayesian Updating: For all preference profiles $(\succeq_i), (\succeq'_i) \in \mathcal{D}$ and events $E \subseteq S$, if $\succeq'_i = \succeq_{i,E}$ for all $i \in N$, then $\succeq' = \succeq_E$ (where $\succeq = F((\succeq_i))$ and $\succeq' = F((\succeq'_i))$).

Informally, if information E arrives (i.e., if the profile changes from (\succeq_i) to $(\succeq_{i,E})$), then group preferences are updated rationally. Violation of this axiom makes the group dynamically incoherent and unable to execute stable plans. This may also create opportunities to manipulate group preferences and decisions through delaying information.

³As nobody can possess conclusive evidence against the truth.

⁴In the representation of a linear-geometric rule, the belief weights β_i and the multiplicative constant are unique (except in the single-state case $|S| = 1$, in which beliefs are trivial). The taste weights α_i and the additive constant are unique under the diversity condition defined in the appendix.

Recall the standard Pareto indifference axiom:

Pareto: For all profiles $(\succeq_i) \in \mathcal{D}$ and acts $a, b \in \mathcal{X}^S$, if $a \sim_i b$ for each $i \in N$, then $a \sim b$ (where $\succeq = F((\succeq_i))$).

This axiom is notoriously criticized because, under uncertainty, unanimities can be *spurious*: they can rest on conflicting beliefs (Mongin 1997/2016). In the introductory example, the unanimous preference for duelling is spurious, as the gentlemen hold (very) different subjective probabilities. GSS have therefore restricted Pareto to acts that depend only on events of uncontroversial probability. Formally, given a profile (\succeq_i) , an act a is *common-belief-determined* if all individual probability functions p_{\succeq_i} ($i \in N$) agree on each event relevant to a , i.e., on the event that a results in l ($\{s \in S : a(s) = l\}$) for each $l \in \mathcal{X}$. Here is GSS's axiom, translated into our Anscombe-Aumann-type framework:

Restricted Pareto: For all profiles $(\succeq_i) \in \mathcal{D}$ and common-belief-determined acts $a, b \in \mathcal{X}^S$, if $a \sim_i b$ for each $i \in N$, then $a \sim b$ (where $\succeq = F((\succeq_i))$).

In fact, this axiom is not fully safe against spurious unanimities. It excludes ‘directly’ spurious unanimities, but not ‘indirectly’ spurious unanimities. I call a unanimous preference or indifference *directly spurious* if it is based on conflicting beliefs, and *indirectly spurious* if it is based on unanimous beliefs which are themselves based on conflicting beliefs. As an example of a unanimity that is only indirectly spurious, let some individuals unanimously prefer the UK to remain in rather than leave the EU, based on a unanimous strong belief in the event E that ‘Brexit’ harms the economy (and a unanimous concern for the economy). But let them believe E for conflicting reasons, i.e., conflicting beliefs about epistemically prior events: some strongly believe E because they strongly believe the government’s advisers said so (event E') and advisers tell the truth (event E''), while others strongly believe E because they strongly believe the advisers denied E (event $\overline{E'}$) and advisers lie (event $\overline{E''}$). So, some reason from E' and E'' to E (using $E' \cap E'' \subseteq E$), while others reason from $\overline{E'}$ and $\overline{E''}$ to E (using $\overline{E'} \cap \overline{E''} \subseteq E$).

To also avoid indirectly spurious unanimities, I restrict the axiom further, namely to *common-belief profiles* $(\succeq_i) \in \mathcal{D}$, in which every individual i has same probability function p_{\succeq_i} . In such profiles, not only beliefs underlying the given acts are unanimous, but also beliefs underlying those beliefs, beliefs underlying beliefs underlying those beliefs, etc. This excludes spurious unanimities *of any order of indirectness*. Here is the new axiom:

Minimal Pareto: For all common-belief profiles $(\succeq_i) \in \mathcal{D}$ and acts $a, b \in \mathcal{X}^S$, if $a \sim_i b$ for each $i \in N$, then $a \sim b$ (where $\succeq = F((\succeq_i))$).

Remark 1 *Minimal Pareto weakens Restricted Pareto, which weakens Pareto.*

Our last axiom requires group preferences to depend continuously on individual preferences, where ‘convergence’ is defined in the natural way.⁵

Continuity: If $(\succeq_i^1), (\succeq_i^2), \dots \rightarrow (\succeq_i)$ in \mathcal{D} , then $F((\succeq_i^1)), F((\succeq_i^2)), \dots \rightarrow F((\succeq_i))$.

I now state the theorem. In the main text I do this only for *fixed-taste domains*, i.e., domains \mathcal{D} whose profiles have same taste profile. This allows variation in beliefs, but not in tastes. Fixing tastes is perhaps not a great loss, as Bayesian learning changes beliefs, not tastes; and it is a ‘half-way’ concession to the fixed-profile approach of standard Bayesian aggregation theory, which fixes tastes *and beliefs*. The appendix re-states the theorem without fixing tastes, and gives the proof.

Theorem 1 *An EU preference aggregation rule $F : \mathcal{D} \rightarrow \mathcal{P}$ on a fixed-taste domain \mathcal{D} satisfies Bayesian Updating, Minimal Pareto, and Continuity if and only if it is linear-geometric.*

Two extensions and the uncertainty-free special case are of interest.

Extension 1: non-public information. The axiom of Bayesian Updating covers ‘public’ information (observed by everyone). It can be re-stated in two ways, to cover private information (observed by just one individual) or to cover information of arbitrary spread (observed by at least one individual). It suffices to replace the clause ‘if $\succeq'_i = \succeq_{i,E}$ for all $i \in N$ ’ by ‘if $\succeq'_i = \succeq_{i,E}$ for one (respectively at least one) $i \in N$ and $\succeq'_i = \succeq_i$ for all other $i \in N$ ’. The two modified axioms are logically stronger: for instance, the private-information axiom implies the original axiom since public learning of E can be decomposed into n steps of private learning of E by each individual in turn. The stronger axioms can still be met, yet by slightly fewer rules. We must exclude rules that ignore someone’s beliefs, to avoid that private learning by someone is collectively ignored. Formally, we can strengthen Bayesian Updating in the theorem in one of the ways – no matter which – if we require strict positivity of the belief weights β_i .

Extension 2: respecting individual tastes. Theorem 1 imposes no constraints on the sign of the taste weights α_i . Individual tastes can be ignored ($\alpha_i = 0$) or even counted negatively ($\alpha_i < 0$). We can easily enforce existence of non-negative or even positive taste weights, by strengthening the Minimal Pareto axiom in Theorem 1 in plausible ways. For non-negativity, replace the axiom’s indifferences by weak preferences. For positivity, further add that the group’s weak preference becomes strict whenever some individual’s weak preference becomes strict.

Harsanyi’s Theorem as the uncertainty-free special case. Theorem 1 reduces to Harsanyi’s Theorem in the uncertainty-free case, i.e., the single-state case $|S| = 1$.

⁵A sequence of relations $\succeq^1, \succeq^2, \dots \in \mathcal{P}$ converges to $\succeq \in \mathcal{P}$ – written $\succeq^1, \succeq^2, \dots \rightarrow \succeq$ – if expected utilities of acts converge: $\mathbb{E}_{p_{\succeq^1}}(u_{\succeq^1}(a)), \mathbb{E}_{p_{\succeq^2}}(u_{\succeq^2}(a)), \dots \rightarrow \mathbb{E}_{p_{\succeq}}(u_{\succeq}(a))$ for all acts $a \in \mathcal{X}^S$. This is equivalent to convergence of tastes and beliefs: $u_{\succeq^1}, u_{\succeq^2}, \dots \rightarrow u_{\succeq}$ and $p_{\succeq^1}, p_{\succeq^2}, \dots \rightarrow p_{\succeq}$ (see Lemma 5). A sequence of profiles $(\succeq_i^1), (\succeq_i^2), \dots \in \mathcal{D}$ converges to $(\succeq_i) \in \mathcal{D}$ – written $(\succeq_i^1), (\succeq_i^2), \dots \rightarrow (\succeq_i)$ – if $\succeq_i^1, \succeq_i^2, \dots \rightarrow \succeq_i$ for each individual i .

This is because in that case acts in \mathcal{X}^S reduce to lotteries in \mathcal{X} (which contain no subjective uncertainty), the domain \mathcal{D} becomes singleton (which amounts to fixing the profile), our Minimal Pareto axiom reduces to Harsanyi's Pareto indifference axiom, and our axioms of Bayesian Updating and Continuity drop out (as they hold trivially for a single-profile domain). To state Harsanyi's Theorem, let an *EU preference relation* on \mathcal{X} (rather than \mathcal{X}^S) be a binary relation \succeq on \mathcal{X} which maximizes the expectation of some non-constant utility function u on X ; the unique normalized version of u is denoted u_\succeq . A group preference relation \succeq on \mathcal{X} satisfies the *Pareto indifference principle* w.r.t. individual preference relations \succeq_i on \mathcal{X} ($i \in N$) if $a \sim a'$ whenever $a \sim_i a'$ for all $i \in N$.

Corollary (Harsanyi's Theorem) *A group EU preference relation \succeq on \mathcal{X} satisfies the Pareto indifference principle w.r.t. individual EU preference relations \succeq_i on \mathcal{X} ($i \in N$) if and only if $u_\succeq = \sum_{i \in N} \alpha_i u_{\succeq_i} + \gamma$ for some $\alpha_i \in \mathbb{R}$ ($i \in N$) and $\gamma \in \mathbb{R}$.*

4 Linear or geometric pooling of beliefs?

Linear-geometric aggregation differs from classical linear-linear aggregation only in how *beliefs* are pooled. Should beliefs be pooled linearly or geometrically after all? This is the notorious debate in statistical opinion pooling theory about merging beliefs (cf. Section 5), a debate that should finally reach Bayesian preference aggregation. Opinion pooling theory recognizes geometric pooling as having superior dynamic behaviour, a conclusion similar to ours, though reached through a different revision axiom (cf. Section 5). Meanwhile, it recognizes linear pooling as enjoying the following (conditional) robustness property, which also holds in our preference context: if states are refined or coarsened, then group beliefs are not perturbed, i.e., merely 'extended' or 'restricted', assuming individual beliefs are not perturbed (a questionable assumption unfortunately, as explained shortly).

How important is such robustness in our preference context? Much depends on which of the two goals mentioned in the introduction we pursue. Robustness matters less for the methodological goal of knowing whether the social-scientific hypothesis of 'rational group agents' is coherent. A paradigmatic rational agent holds beliefs defined on fixed, maximally informative states. This requires a maximal state space S , and forbids coarsening states as done in robustness arguments. Robustness however matters if we aim to construct group preferences in practice. In practice, one cannot work with maximally fine states, and so questions of state design and robustness arise.

Unfortunately, not even linear pooling is realistically robust, because *individual* beliefs are systematically non-robust to start with. Why? The states of a model are normally linked at least partly to the (atomic) scenarios which individuals subjectively *consider*, or are *aware* of, or *reason* with – either because the design of states (the 'framing of contingencies', in Ahn and Ergin's 2010 words) affects people's reasoning and conceptualisation of options, or conversely because people's reasoning guides the modeller, who tailors states to the (atomic) scenarios considered by individuals. Given this correlation between states and awareness – between model ontology and individual

ontology – a state refinement reflects growing awareness. But growingly aware individuals do not only extend their old beliefs to new events: they also *revise* their beliefs about *old* events, a process that is rationally governed by a well-defined revision rule, axiomatized by Ahn and Ergin (2010) as ‘partition-dependence’, by Karni and Viero (2013, 2020) as ‘reverse Bayesianism’, and by Dietrich (2018) as ‘proportional rescaling’. This makes individual beliefs non-robust. Hence also group beliefs are non-robust, even if formed linearly.

5 Discussion in relation to the literature

I have proposed to model groups as full rational agents, by applying to group agents what we normally require from individuals. Group rationality is uniquely achieved by linear-geometric aggregation, if group preferences are minimally Paretian and continuous in individual preferences. Linear-geometric aggregation makes the group agent more rational, but less Paretian, than GSS’s classic linear-linear aggregation – a new instance of the classic trade-off between group rationality and Paretianism (see below). Which form of Paretianism is justified under uncertainty is an open debate (see below). Minimal Paretianism avoids objections by being safe against directly and indirectly spurious unanimities.

From here, important questions open up. How about weakening *static* group Bayesianism into non-EU directions while preserving Bayesian Updating? Such a dynamically (not statically) Bayesian approach would be the dual of the statically (not dynamically) Bayesian programme of Harsanyi, Mongin, and Gilboa-Samet-Schmeidler. And how about *geometric-geometric* rules, which pool even utilities geometrically? Such rules remain fully Bayesian, but become radically non-Paretian.

Bayesian preference aggregation theory was born with Harsanyi’s (1955) spectacular theorem: in groups with heterogeneous preferences under risk, group utility must be linear in individual utilities if the group is Paretian and EU rational. Harsanyi regarded his result as an ‘economic derivation’ of philosophical utilitarianism, a controversial thesis ever since (Weymark 1991, Fleurbaey and Mongin 2016). Harsanyi’s Theorem enjoys some robustness within the limited world of objective uncertainty; see generalizations by Fleurbaey (2009, 2014) and Danan, Gajdos and Tallon (2015). The picture reverses for heterogeneous beliefs: EU rationality then becomes incompatible with Paretianism (e.g., Mongin 1995), and meanwhile Paretianism becomes less compelling because unanimities can be spurious (Mongin 1997/2016). Gilboa et al. (2004) restore possibility and rehabilitate Paretianism by restricting the Pareto principle to unanimities that are not (directly) spurious, obtaining linear-linear group preferences. A lively literature follows, exploring the trade-off between group rationality – in the semi-Bayesian sense setting aside Bayes’ rule – and Paretianism. The general direction has been to combine more or less strong group rationality with more or less strong Paretianism, usually working within some classical model of choice under uncertainty (see Chateauneuf, Cohen and Jaffray 2008 for a review).

Some works emphasize impossibility, often by working with non-EU preferences, i.e., abandoning even static Bayesianism. In particular, Chambers and Hayashi (2006)

show that full Paretianism already conflicts with minimal group rationality, i.e., with transitive and complete group preferences satisfying Savage’s P3 *or* his P4. Another threat to preference aggregation comes from individual irrationality: Gajdos, Tallon and Vergnaud (2008) and Zuber (2016) show that, unless individuals have EU preferences, group preferences cannot even be mildly rational and Paretian – also if individuals have identical beliefs.

Other works stress possibility. For instance, Chambers and Hayashi (2006) show the possibility of state-dependent fully Paretian group preferences (see already Mongin 1998). Another positive result is due to Danan et al. (2016): incomplete preferences based on imprecise beliefs can be aggregated in a Paretian way.

Over the years, new Pareto principles have been proposed and defended, such as principles restricted to unanimities that are shared-belief rationalizable (Gilboa, Samuelson and Schmeidler 2014), principles sensitive to whether acts depend on objective or subjective uncertainties (Mongin and Pivato 2020), and principles restricted to unanimities that are common knowledge (Nehring 2004, Chambers and Hayashi 2014). Attempts to make Paretianism immune to spurious unanimities face a general difficulty: beliefs become empirically underdetermined once we remove the (unfalsifiable) hypothesis of state-independent utility (Karni 1993, Wakker and Zank 1999, Baccelli 2019). There are different possible reactions to this intriguing diagnosis, such as: becoming more cautious about Pareto axioms out of fearing hidden spurious unanimities, or on the contrary reverting to full-blown Paretianism out of rejecting the very notion of belief and spurious unanimity. Sprumont (2018) offers such a sceptical view.

The Bayesian Updating axiom is a counterpart for preference aggregation of the classic *External Bayesianity* axiom for probability aggregation or ‘opinion pooling’ (Madansky 1964, Bordley 1982, Genest and Zidek 1986). External Bayesianity differs not only in the objects that are revised and aggregated (i.e., probabilities, not preferences), but also in the notion of information, given by a likelihood function rather than an event. Likelihood functions capture information outside the domain (algebra) of beliefs – the information concept relevant in Bayesian statistics.⁶ As is well-known, geometric probability aggregation satisfies External Bayesianity, but there are other well-behaved rules doing so (e.g., Baccelli and Stewart 2019). By contrast, as shown here, the only well-behaved *preference* aggregation rules satisfying Bayesian Updating are geometric in beliefs. As recently shown, the only opinion pooling rules satisfying a version of External Bayesianity (based on events rather than likelihood functions) and satisfying two regularity conditions are generalized geometric rules in which the weights need not sum to one (Russell et al. 2015, Dietrich 2019).

⁶ A likelihood function maps states to numbers, representing probabilities of information given states. If states are statistical hypotheses, then statistical information (‘data’) is captured by a likelihood function, not an event. If states are weather conditions, then the information ‘it rains’ is captured by an event, but ‘the radio forecasts rain’ could be captured by a likelihood function (which also reflects the reliability of the weather forecast).

A Generalisation and proof

The theorem can be extended beyond fixed-taste domains. The extended theorem will make a standard assumption on profiles (because of which it is, strictly speaking, not logically stronger). An EU preference profile (\succeq_i) satisfies *diversity* if for each individual i there are lotteries $a, a' \in \mathcal{X}$ between which only individual i is non-indifferent (i.e., $a \not\sim_i a'$ while $a \sim_j a'$ for $j \neq i$). In the extended theorem, individual weights can vary with tastes:

Definition 2 A rule $F : \mathcal{D} \rightarrow \mathcal{P}$ is linear-geometric with taste-dependent weights if there exist individual weights $\alpha_{i,\mathbf{u}} \in \mathbb{R}$ and $\beta_{i,\mathbf{u}} \in \mathbb{R}_+$ ($i \in N$) which depend continuously on the taste profile $\mathbf{u} \in \{(u_{\succeq_i}) : (\succeq_i) \in \mathcal{D}\}$, where $\sum_{i \in N} \beta_{i,\mathbf{u}} = 1$ for each \mathbf{u} , such that at each profile $(\succeq_i) \in \mathcal{D}$ the group preference relation $\succeq = F((\succeq_i))$ has

- u_{\succeq} given by $\sum_{i \in N} \alpha_{i,\mathbf{u}} u_{\succeq_i}$ up to an additive constant,
- p_{\succeq} given on S by $\prod_{i \in N} [p_{\succeq_i}]^{\beta_{i,\mathbf{u}}}$ up to a multiplicative constant,

where \mathbf{u} denotes the current taste profile (u_{\succeq_i}) .⁷

Theorem 1⁺ An EU preference aggregation rule $F : \mathcal{D} \rightarrow \mathcal{P}$ on any domain \mathcal{D} of diverse profiles satisfies Bayesian Updating, Minimal Pareto and Continuity if and only if it is linear-geometric with taste-dependent weights.

Extensions 1 and 2 apply analogously to Theorem 1⁺. The domain in Theorem 1⁺ is flexible. Maximally, it contains *all* diverse coherent profiles $(\succeq_i) \in \mathcal{P}^N$. Minimally, it is a fixed-taste domain. Theorem 1 is the fixed-taste special case of Theorem 1⁺ (diversity aside), because the taste index \mathbf{u} in ‘ $\alpha_{i,\mathbf{u}}$ ’ and ‘ $\beta_{i,\mathbf{u}}$ ’ can be dropped if it is fixed. As weights depend on tastes, but not on beliefs, we still have a limited form of profile-dependence. By contrast, the weights in GSS’s linear-linear result can implicitly vary with both tastes and beliefs. This full form of profile-dependence would become visible if GSS’s result were re-stated in a multi-profile framework.

I now prove Theorem 1⁺. Theorem 1 follows immediately, noting that the proof needs no diversity assumption for fixed-taste domains.

Notation: I write ‘ \sum_i ’ for ‘ $\sum_{i \in N}$ ’, ‘ \bigcup_i ’ for ‘ $\bigcup_{i \in N}$ ’, etc. The group relation obtained by aggregating individual relations is denoted using the same symbol as for individuals, but without individual index: so I denote $F((\succeq_i))$ by \succeq , $F((\succeq'_i))$ by \succeq' , etc. Let $\mathcal{U} := \{(u_{\succeq_i}) : (\succeq_i) \in \mathcal{D}\}$ be the set of occurring taste profiles, and let $\mathcal{D}_{Com\ Bel} \subseteq \mathcal{D}$ be the subdomain of common-belief profiles in \mathcal{D} .

Consider two separate conditions on an aggregation rule $F : \mathcal{D} \rightarrow \mathcal{P}$:

LIN: There exist real weights $(\alpha_{i,\mathbf{u}})_{i \in N, \mathbf{u} \in \mathcal{U}}$ such that, at each profile $(\succeq_i) \in \mathcal{D}$, u_{\succeq} is given by $\sum_i \alpha_{i,(u_{\succeq_j})} u_{\succeq_i}$ up to an additive constant.

GEO: There exist non-negative weights $(\beta_{i,\mathbf{u}})_{i \in N, \mathbf{u} \in \mathcal{U}}$ with $\sum_i \beta_{i,\mathbf{u}} = 1$ for all $\mathbf{u} \in \mathcal{U}$ such that, at each profile $(\succeq_i) \in \mathcal{D}$, p_{\succeq} is given on S by $\prod_i [p_i]^{\beta_{i,(u_j)}}$ up to a multiplicative constant.

⁷Footnote 4 about uniqueness of weights generalizes.

We prove Theorem 1 by showing four facts, of which the first three establish sufficiency of the axioms and the fourth establishes necessity of the axioms:

Fact 1: *Bayesian Updating and Minimal Pareto imply LIN.*

Fact 2: *Bayesian Updating, Minimal Pareto, and Continuity imply GEO.*

Fact 3: *LIN, GEO and Continuity imply that the rule is linear-geometric with taste-dependent weights.*

Fact 4: *All linear-geometric rules with taste-dependent weights satisfy Bayesian Updating, Minimal Pareto, and Continuity.*

A.1 Proof of Fact 1

We start with a technical lemma:

Lemma 1 *For all $(\succeq_i) \in \mathcal{D}$, $\bigcap_i \text{supp}(p_{\succeq_i}) \neq \emptyset$, and if Bayesian Updating holds then $\bigcap_i \text{supp}(p_{\succeq_i}) \subseteq \text{supp}(p_{\succeq}) \subseteq \bigcup_i \text{supp}(p_{\succeq_i})$.*

Proof. Let $(\succeq_i) \in \mathcal{D}$. By coherence, $\bigcap_i \text{supp}(p_{\succeq_i}) \neq \emptyset$. Now assume Bayesian Updating. We first show $\text{supp}(p_{\succeq}) \subseteq \bigcup_i \text{supp}(p_{\succeq_i})$. If $E := \bigcup_i \text{supp}(p_{\succeq_i})$, then, $\succeq_{i,E} = \succeq_i$ for all i ; so by Bayesian Updating $\succeq = \succeq_E$, whence $p_{\succeq} = p_{\succeq_E}$. Thus $\text{supp}(p_{\succeq}) \subseteq E$.

Finally, let $s \in \bigcap_i \text{supp}(p_{\succeq_i})$; we show $s \in \text{supp}(p_{\succeq})$. Define $E = \{s\}$. The profile $(\succeq')_{i \in N} := (\succeq_{i,E})_{i \in N}$ is coherent, hence in \mathcal{D} ; so by Bayesian Updating $\succeq' = \succeq_E$. Hence $\succeq_E \in \mathcal{P}$. Thus $s \in \text{supp}(p_{\succeq})$, as otherwise E would be \succeq -null, and \succeq_E would be the all-indifferent relation, which is not in \mathcal{P} . ■

We next prove that Bayesian Updating implies this familiar axiom:

Independence of Group Tastes on Individual Beliefs (IGTIB): For all $(\succeq_i), (\succeq'_i) \in \mathcal{D}$, if $u_{\succeq_i} = u_{\succeq'_i}$ for all $i \in N$, then $u_{\succeq} = u_{\succeq'}$.

Lemma 2 *Bayesian Updating implies IGTIB.*

Proof. Assume Bayesian Updating. Note that if profiles $(\succeq_i), (\succeq'_i) \in \mathcal{D}$ are ‘Bayes neighbours’ in the sense that for some $E \subseteq S$ we have $(\succeq_i) = (\succeq'_{i,E})$ or $(\succeq_{i,E}) = (\succeq'_i)$, then by Bayesian Updating $\succeq = \succeq'_E$ or $\succeq_E = \succeq'$, and so $u_{\succeq} = u_{\succeq'}$ as conditionalization preserves tastes.

Now consider any $(\succeq_i), (\succeq'_i) \in \mathcal{D}$ such that $u_{\succeq_i} = u_{\succeq'_i}$. By the previous observation, it suffices to construct a finite sequence of profiles in \mathcal{D} starting with (\succeq_i) and ending with (\succeq'_i) such that any two adjacent profiles are Bayes neighbours. To do so, pick $s \in \bigcap_i \text{supp}(p_{\succeq_i})$ and $s' \in \text{supp}(p_{\succeq'_i})$ (via Lemma 1). Let $(\succeq_i^1), \dots, (\succeq_i^5)$ be the five profiles such that $(\succeq_i^1) = (\succeq_i)$ and $(\succeq_i^5) = (\succeq'_i)$, and such that $(\succeq_i^2), (\succeq_i^3), (\succeq_i^4)$ have taste profiles given by $u_{\succeq_i^2} = u_{\succeq_i^3} = u_{\succeq_i^4} = u_{\succeq_i}$ ($i \in N$) and belief profiles given by $p_{\succeq_i^2}(s) = 1$, $p_{\succeq_i^3}(s) = p_{\succeq_i^3}(s') = \frac{1}{2}$ and $p_{\succeq_i^4}(s') = 1$ ($i \in N$). These profiles belong to \mathcal{D} as \mathcal{D} is closed under belief change. To check for Bayes neighbourhood, note that, for all $i \in N$, $\succeq_{i,\{s\}}^1 = \succeq_i^2$, $\succeq_i^2 = \succeq_{i,\{s\}}^3$, $\succeq_{i,\{s\}}^3 = \succeq_{i,\{s'\}}^4$, and $\succeq_i^4 = \succeq_{i,\{s'\}}^5$. ■

The proof of Fact 1 is completed by two lemmas. The first uses Harsanyi's Theorem to show that Minimal Pareto alone implies a much weaker linearity property than LIN, which is restricted to common-belief profiles and allows weights to depend arbitrarily on the profile. The second strengthens the linearity conclusion to LIN by adding IGTIB.

Lemma 3 *Under Minimal Pareto, there are weights $\alpha_{i,(\succeq_j)} \in \mathbb{R}$ across $i \in N$ and $(\succeq_j) \in \mathcal{D}_{Com\ Bel}$ such that, at each $(\succeq_j) \in \mathcal{D}_{Com\ Bel}$, $u_{\succeq} = \sum_i \alpha_{i,(\succeq_j)} u_{\succeq_i} + c$ for some $c \in \mathbb{R}$.*

Proof. Let $(\succeq_j) \in \mathcal{D}_{Com\ Bel}$. Under Minimal Pareto, the restriction of the group relation to lotteries, $\succeq|_{\mathcal{X}}$, satisfies Harsanyi's (1955) Pareto indifference condition w.r.t. the restricted individual relations $\succeq_i|_{\mathcal{X}}$ ($i \in N$). So our linearity conclusion holds by Harsanyi's Theorem (Harsanyi 1955). ■

Lemma 4 *IGTIB and Minimal Pareto jointly imply LIN.*

Proof. This result follows from Lemma 3, since under IGTIB the linearity conclusion in Lemma 3 extends to arbitrary profiles $(\succeq_i) \in \mathcal{D}$ (since each $(\succeq_i) \in \mathcal{D}$ has the same taste profile as some $(\succeq'_i) \in \mathcal{D}_{Com\ Bel}$), where we can take the weights $\alpha_{i,(\succeq_j)}$ to depend on (\succeq_j) only through the taste profile (u_{\succeq_j}) . ■

A.2 Proof of Fact 2

We begin by a simple characterization of preference convergence:

Lemma 5 *A sequence \succeq^k converges to \succeq in \mathcal{P} (i.e., $\mathbb{E}_{p_{\succeq^k}}(u_{\succeq^k}(a)) \rightarrow \mathbb{E}_{p_{\succeq}}(u_{\succeq}(a))$ for all $a \in \mathcal{X}^S$) if and only if $u_{\succeq^k} \rightarrow u_{\succeq}$ and $p_{\succeq^k} \rightarrow p_{\succeq}$, where ' \rightarrow ' denotes pointwise convergence or (equivalently) uniform convergence.*

Proof. Consider \succeq^k ($k = 1, 2, \dots$) and \succeq in \mathcal{P} . First, if $u_{\succeq^k} \rightarrow u_{\succeq}$ and $p_{\succeq^k} \rightarrow p_{\succeq}$, then $\succeq^k \rightarrow \succeq$ because for all acts $a \in \mathcal{X}^S$

$$\begin{aligned} \mathbb{E}_{p_{\succeq^k}}(u_{\succeq^k}(a)) &= \sum_{s \in S} p_{\succeq^k}(s) \sum_{x \in X} a(s)(x) u_{\succeq^k}(x) \\ &\rightarrow \sum_{s \in S} p_{\succeq}(s) \sum_{x \in X} a(s)(x) u_{\succeq}(x) = \mathbb{E}_{p_{\succeq}}(u_{\succeq}(a)). \end{aligned}$$

Conversely, assume $\succeq^k \rightarrow \succeq$. Then $u_{\succeq^k} \rightarrow u_{\succeq}$ since for each $x \in X$ we can use the constant act $a \equiv x$ to infer $u_{\succeq^k}(x) = \mathbb{E}_{p_{\succeq^k}}(u_{\succeq^k}(a)) \rightarrow \mathbb{E}_{p_{\succeq}}(u_{\succeq}(a)) = u_{\succeq}(x)$. Now we fix $s \in S$ and show $p_{\succeq^k}(s) \rightarrow p_{\succeq}(s)$. Pick $x, x' \in X$ such that $u_{\succeq}(x) = 1$ and $u_{\succeq}(x') = 0$, and consider the act $a \in \mathcal{X}^S$ mapping s to x and all other states to x' . Since $\succeq^k \rightarrow \succeq$, we have $\mathbb{E}_{p_{\succeq^k}}(u_{\succeq^k}(a)) \rightarrow \mathbb{E}_{p_{\succeq}}(u_{\succeq}(a))$, i.e.,

$$p_{\succeq^k}(s) u_{\succeq^k}(x) + (1 - p_{\succeq^k}(s)) u_{\succeq^k}(x') \rightarrow p_{\succeq}(s) 1 + (1 - p_{\succeq}(s)) 0 = p_{\succeq}(s).$$

Since $u_{\succeq^k}(x) \rightarrow u_{\succeq}(x) = 1$ and $u_{\succeq^k}(x') \rightarrow u_{\succeq}(x') = 0$, we can infer $p_{\succeq^k}(s) \rightarrow p_{\succeq}(s)$. ■

Let $\mathbf{P} = \{\mathbf{p} \in \Delta(S)^N : \bigcap_i \text{supp}(p_i) \neq \emptyset\}$ be the set 'coherent' belief profiles. A belief aggregation rule is a function $\pi : \mathbf{P} \rightarrow \Delta(S)$. Consider three conditions on such a rule:

- (C1) *Unanimity Preservation*: $\pi(p, \dots, p) = p$ for all $p \in \Delta(S)$.
(C2) *Continuity*: if $\mathbf{p}^k \rightarrow \mathbf{p}$ in $\Delta(S)^N$ (i.e., $p_i^k \rightarrow p_i$ for all $i \in N$), then $\pi(\mathbf{p}^k) \rightarrow \pi(\mathbf{p})$.
(C3) *Bayesian Belief Updating*: if $\mathbf{p} \in \mathbf{P}$ and $E \subseteq S$ with $\mathbf{p}(E) \gg 0$, then $\pi(\mathbf{p})(E) > 0$ and $\pi(\mathbf{p}(\cdot|E)) = \pi(\mathbf{p})(\cdot|E)$. (Here and elsewhere, $\mathbf{p}(\cdot|E)$ stands for the updated profile $(p_i(\cdot|E))$.)

Our preference aggregation rule F induces a family of belief aggregation rules $\pi_{\mathbf{u}} : \mathbf{P} \rightarrow \Delta(S)$, where $\mathbf{u} \in \mathcal{U}$, defined as follows. Let $\mathbf{u} \in \mathcal{U}$. For each $\mathbf{p} \in \mathbf{P}$, form the preference profile $(\succeq_i) \in \mathcal{D}$ with taste profile \mathbf{u} and belief profile \mathbf{p} , then form the group relation $\succeq = F((\succeq_i))$, and let $\pi_{\mathbf{u}}(\mathbf{p}) := p_{\succeq}$. The rules $\pi_{\mathbf{u}}$ ($\mathbf{u} \in \mathcal{U}$) inherit several properties from F :

- Lemma 6** (a) *If F satisfies Minimal Pareto, then each $\pi_{\mathbf{u}}$ ($\mathbf{u} \in \mathcal{U}$) satisfies C1.*
(b) *If F is continuous, then each $\pi_{\mathbf{u}}$ ($\mathbf{u} \in \mathcal{U}$) satisfies C2.*
(c) *If F satisfies Bayesian Updating, then each $\pi_{\mathbf{u}}$ ($\mathbf{u} \in \mathcal{U}$) satisfies C3.*

Proof. The result is obvious if $|S| = 1$, as then there is only one belief aggregation rule, which trivially satisfies C1-C3. Now suppose $|S| \neq 1$, and $\mathbf{u} \in \mathcal{U}$.

(a) Assume Minimal Pareto. Let $\mathbf{p} = (p, \dots, p) \in \Delta(S)^N$ be a unanimous belief profile. Let $(\succeq_i) \in \mathcal{D}$ be the preference profile with taste profile \mathbf{u} and (unanimous) belief profile \mathbf{p} . Form $\succeq = F((\succeq_i))$. We fix a state $s \in S$ and must show that $\pi_{\mathbf{u}}(\mathbf{p})(s) = p(s)$, i.e., that $p_{\succeq}(s) = p(s)$. Pick outcomes $x, y \in X$ such that $x \succ y$, i.e., $u_{\succeq}(x) > u_{\succeq}(y)$. Let a be the act which yields x at s and y on $S \setminus \{s\}$. Let b be the constant act which at all states yields the lottery $p(s)\delta_x + (1 - p(s))\delta_y$, where δ_x (δ_y) denotes the lottery with sure outcome x (y). Then $a \sim_i b$ for all $i \in N$. So, as (\succeq_i) is a common-belief profile, $a \sim b$ by Minimal Pareto. Hence, $p_{\succeq}(s)u_{\succeq}(x) + (1 - p_{\succeq}(s))u_{\succeq}(y) = p(s)u_{\succeq}(x) + (1 - p(s))u_{\succeq}(y)$. As $u_{\succeq}(x) > u_{\succeq}(y)$, this implies $p_{\succeq}(s) = p(s)$. Q.e.d.

(b) Assume F is continuous, and $\mathbf{p}^k \rightarrow \mathbf{p}$ in $\Delta(S)^N$. Let (\succeq_i^k) ($k = 1, 2, \dots$) and (\succeq_i) be the preference profiles in \mathcal{D} with the belief profiles \mathbf{p}^k and \mathbf{p} , respectively, and with same taste profile \mathbf{u} . As $\mathbf{p}^k \rightarrow \mathbf{p}$ and $\mathbf{u} \rightarrow \mathbf{u}$, we have $(\succeq_i^k) \rightarrow (\succeq_i)$ by Lemma 5. So $\succeq^k \rightarrow \succeq$ by Continuity of F , and thus by Lemma 5 $p_{\succeq^k} \rightarrow p_{\succeq}$, i.e., $\pi_{\mathbf{u}}(\mathbf{p}^k) \rightarrow \pi_{\mathbf{u}}(\mathbf{p})$. Q.e.d.

(c) Let F satisfy Bayesian Updating. Assume $\mathbf{p} \in \mathbf{P}$ and $E \subseteq S$ with $\mathbf{p}(E) \gg 0$. Let (\succeq_i) have taste profile \mathbf{u} and belief profile \mathbf{p} . As $\mathbf{p}(E) \gg 0$, each of (\succeq_i) and $(\succeq_{i,E})$ is coherent, so in \mathcal{D} . By Bayesian Updating, $\succeq' = \succeq_E$. Hence, $p_{\succeq'} = p_{\succeq}(\cdot|E)$, i.e., $\pi_{\mathbf{u}}(\mathbf{p}(\cdot|E)) = \pi_{\mathbf{u}}(\mathbf{p})(\cdot|E)$. ■

Given Lemma 6, Fact 2 follows from the following result:

Lemma 7 *Each belief aggregation rule $\pi : \mathbf{P} \rightarrow \Delta(S)$ satisfying C1-C3 is geometric, i.e., there exist weights $\beta_i \geq 0$ ($i \in N$) of sum one such that, for each belief profile $\mathbf{p} \in \mathbf{P}$, $\pi(\mathbf{p})$ is on S given by $\prod_i [p_i]^{\beta_i}$ up to a multiplicative constant.*

This result and its proof are variants of results in Russell et al. (2015) and Dietrich (2019), who use only slightly different conditions than C1-C3 and derive a generalized

geometric formula.⁸ For completeness, we give a self-contained proof here.

Proof. The result is trivial if $|S| = 1$. So, as $|S| \neq 2$, we can assume without loss of generality that $|S| \geq 3$. Let $\pi : \mathbf{P} \rightarrow \Delta(S)$ satisfy C1-C3.

Claim 1: for all $\mathbf{p} \in \mathbf{P}$ and all $s, t \in S$, $\mathbf{p}(s) = \mathbf{p}(t) \neq 0 \Rightarrow \pi(\mathbf{p})(s) = \pi(\mathbf{p})(t) \neq 0$.

Assume $\mathbf{p}(s) = \mathbf{p}(t) \neq 0$. For non-triviality, $s \neq t$. Let $E = \{s, t\}$, and let $\mathbf{p}' = (p', \dots, p')$ be the unanimous profile with $p'(s) = p'(t) = \frac{1}{2}$. Applying C3 and then C1, we have $\pi(\mathbf{p})(E) \neq 0$ and $\pi(\mathbf{p})(\cdot|E) = \pi(\mathbf{p}'(\cdot|E)) = \pi(\mathbf{p}') = p'$. So $\pi(\mathbf{p})(s) = \pi(\mathbf{p})(t) \neq 0$. Q.e.d.

Claim 2: For all $s \neq t$ in S there is a unique $f_{s,t} : (0, \infty)^N \rightarrow (0, \infty)$ such that $\frac{\pi(\mathbf{p})(s)}{\pi(\mathbf{p})(t)} = f_{s,t} \left(\left(\frac{p_i(s)}{p_i(t)} \right) \right)$ for all $\mathbf{p} \in \mathbf{P}$ with $\mathbf{p}(s), \mathbf{p}(t) \gg 0$.

Let $s \neq t$ in S . Uniqueness holds as each $\mathbf{x} \in (0, \infty)^N$ equals $\left(\left(\frac{p_i(s)}{p_i(t)} \right) \right)$ for a $\mathbf{p} \in \mathbf{P}$. As for existence, let $\mathbf{p}, \mathbf{p}' \in \mathbf{P}$ with $\mathbf{p}(s), \mathbf{p}'(t) \gg 0$ and $\left(\frac{p_i(s)}{p_i(t)} \right) = \left(\frac{p'_i(s)}{p'_i(t)} \right)$. We show $\frac{\pi(\mathbf{p})(s)}{\pi(\mathbf{p})(t)} = \frac{\pi(\mathbf{p}')(s)}{\pi(\mathbf{p}')(t)}$. Put $E = \{s, t\}$. Note $\mathbf{p}(\cdot|E) = \mathbf{p}'(\cdot|E)$. So $\pi(\mathbf{p}(\cdot|E)) = \pi(\mathbf{p}'(\cdot|E))$, whence by C3 $\pi(\mathbf{p})(\cdot|E) = \pi(\mathbf{p}')(\cdot|E)$. So $\frac{\pi(\mathbf{p})(s)}{\pi(\mathbf{p})(t)} = \frac{\pi(\mathbf{p}')(s)}{\pi(\mathbf{p}')(t)}$, where both ratios are well-defined and non-zero as $\pi(\mathbf{p})(s), \pi(\mathbf{p})(t), \pi(\mathbf{p}')(s), \pi(\mathbf{p}')(t) \neq 0$ by C3. Q.e.d.

Claim 3: $f_{s,r}(\mathbf{xy}) = f_{s,t}(\mathbf{x})f_{t,r}(\mathbf{y})$ for all $\mathbf{x}, \mathbf{y} \in (0, \infty)^N$ and pairwise distinct $s, t, r \in S$.

Use that for all $\mathbf{x}, \mathbf{y} \in (0, \infty)^N$ and pairwise distinct $s, t, r \in S$ one can construct a $\mathbf{p} \in \mathbf{P}$ such that $\mathbf{x} = \left(\frac{p_i(s)}{p_i(t)} \right)$, $\mathbf{y} = \left(\frac{p_i(t)}{p_i(r)} \right)$, and so $\mathbf{xy} = \left(\frac{p_i(s)}{p_i(r)} \right)$. Q.e.d.

Claim 4: All $f_{s,t}$ for $s \neq t$ are the same function, to be denoted f .

Let $s, s', t, t' \in S$ with $s \neq t$ and $s' \neq t'$, and $\mathbf{x} \in (0, \infty)^N$. I must show $f_{s,t}(\mathbf{x}) = f_{s',t'}(\mathbf{x})$.

Case 1: $s = s'$. Pick $\mathbf{p} \in \mathbf{P}$ such that $\mathbf{p}(t), \mathbf{p}(t') \gg 0$ and $\mathbf{x} = \left(\frac{p_i(s)}{p_i(t)} \right) = \left(\frac{p_i(s)}{p_i(t')} \right)$. So $\mathbf{p}(t) = \mathbf{p}(t') \neq 0$. Hence $\pi(\mathbf{p})(t) = \pi(\mathbf{p})(t') \neq 0$ by Claim 1. So $\frac{\pi(\mathbf{p})(s)}{\pi(\mathbf{p})(t)} = \frac{\pi(\mathbf{p})(s)}{\pi(\mathbf{p})(t')}$, whence $f_{s,t}(\mathbf{x}) = f_{s',t'}(\mathbf{x})$.

Case 2: $t = t'$. By an argument analogous to that in Case 1, $f_{s,t}(\mathbf{x}) = f_{s',t'}(\mathbf{x})$.

Case 3: $s \neq s'$ and $t \neq t'$. I show $f_{s,t}(\mathbf{x}) = f_{s',t'}(\mathbf{x})$ by drawing on Cases 1 and 2. If $s \neq t'$, then $f_{s,t}(\mathbf{x}) = f_{s,t'}(\mathbf{x}) = f_{s',t'}(\mathbf{x})$. If $s' \neq t$, then $f_{s,t}(\mathbf{x}) = f_{s',t}(\mathbf{x}) = f_{s',t'}(\mathbf{x})$. If $s = t'$ and $s' = t$, then, choosing any $r \in S \setminus \{s, t\}$, $f_{s,t}(\mathbf{x}) = f_{s,r}(\mathbf{x}) = f_{t,r}(\mathbf{x}) = f_{t,s}(\mathbf{x})$. Q.e.d.

Claim 5: $f(\mathbf{xy}) = f(\mathbf{x})f(\mathbf{y})$ for all $\mathbf{x}, \mathbf{y} \in (0, \infty)^N$, and $f(\mathbf{1}) = 1$.

The functional equation holds by Claims 3 and 4. The identity $f(\mathbf{1}) = 1$ follows because $f(\mathbf{1}) = f(\mathbf{1})f(\mathbf{1})$. Q.e.d.

Claim 6: $\frac{\pi(\mathbf{p})(s)}{\pi(\mathbf{p})(t)} = f \left(\left(\frac{p_i(s)}{p_i(t)} \right) \right)$ for all $s, t \in S$ and $\mathbf{p} \in \mathbf{P}$ with $\mathbf{p}(s), \mathbf{p}(t) \gg 0$.

For $s \neq t$ this holds by Claims 2 and 4, while for $s = t$ it holds as $\frac{\pi(\mathbf{p})(s)}{\pi(\mathbf{p})(t)} = 1$ and $f \left(\left(\frac{p_i(s)}{p_i(t)} \right) \right) = f(\mathbf{1}) = 1$. Q.e.d.

Claim 7: There exist $\beta_1, \dots, \beta_n \in \mathbb{R}$ such that $f(\mathbf{x}) = x_1^{\beta_1} \dots x_n^{\beta_n}$ for all $\mathbf{x} \in (0, \infty)^N$.

⁸These authors assume C2, C3 and a third condition, and they derive that group beliefs are given by a generalized geometric formula in which the weights need not sum to one.

The function $g : \mathbf{x} \mapsto \ln(f((\exp x_i)))$ on \mathbb{R}^N obeys Cauchy's functional equation ' $g(\mathbf{x} + \mathbf{y}) = g(\mathbf{x}) + g(\mathbf{y})$ ' by Claim 5 and is continuous by C2. So there are $\beta_i \in \mathbb{R}$ ($i \in N$) such that $g(\mathbf{x}) = \sum_i \beta_i x_i$ for all $\mathbf{x} \in \mathbb{R}^n$ (Aczél 1966). Hence,

$$f(\mathbf{x}) = \exp g((\ln x_i)) = \exp \sum_i \beta_i \ln x_i = x_1^{\beta_1} \cdots x_n^{\beta_n} \text{ for all } \mathbf{x} \in (0, \infty)^N.$$

Claim 8: Consider the subdomain of full-support profiles $\mathbf{P}^* := \{\mathbf{p} \in \mathbf{P} : \text{supp}(p_i) = S \text{ for all } i \in N\}$. For each $\mathbf{p} \in \mathbf{P}^*$, $\pi(\mathbf{p})$ is on S given by $\prod_i [p_i]^{\beta_i}$ up to a multiplicative constant.

Let $\mathbf{p} \in \mathbf{P}^*$. Fix any $t \in S$, and define $k' := \pi(\mathbf{p})(t)$ and $k'' := \prod_i [p_i(t)]^{\beta_i}$. We have $k', k'' > 0$, because $\text{supp}(\pi(\mathbf{p})) = S$ as by Lemma 1 $\text{supp}(\pi(\mathbf{p})) \supseteq \bigcap_i \text{supp}(p_i) = S$. For all $s \in S$,

$$\pi(\mathbf{p})(s) = k' \frac{\pi(\mathbf{p})(s)}{\pi(\mathbf{p})(t)} = k' f\left(\left(\frac{p_i(s)}{p_i(t)}\right)\right) = k' \prod_i \left(\frac{p_i(s)}{p_i(t)}\right)^{\beta_i} = \frac{k'}{k''} \prod_i [p_i(s)]^{\beta_i}. \text{ Q.e.d.}$$

Claim 9: $\beta_i \geq 0$ for all $i \in N$ and $\sum_i \beta_i = 1$.

We proceed by contradiction. First, assume $\sum_i \beta_i \neq 1$. Pick any $\mathbf{p} \in \mathbf{P}^*$ in which all p_i are the same p such that $p(s)$ is neither identical for all $s \in S$ nor 1 at any $s \in S$. By C1, $\pi(\mathbf{p}) = p$. So by Claim 8, p is proportional to $\prod_i p^{\beta_i} = p^{\sum_i \beta_i}$, a contradiction as $\sum_i \beta_i \neq 1$.

Second, assume $i \in N$ such that $\beta_i < 0$. Pick $s \in S$ and a sequence $\mathbf{p}^1, \mathbf{p}^2, \dots \in \mathbf{P}^*$ converging to a $\mathbf{p} \in \mathbf{P} \setminus \mathbf{P}^*$ such that $\text{supp}(p_i) = S \setminus \{s\}$ while $\text{supp}(p_j) = S$ for all $j \in N \setminus \{i\}$. By $\beta_i < 0$ and Claim 8, the sequence $\pi(\mathbf{p}^1), \pi(\mathbf{p}^2), \dots$ converges to the measure assigning probability 1 to s . Meanwhile by C2 the limit must be $\pi(\mathbf{p})$. So $\pi(\mathbf{p})(s) = 1$, whence $\text{supp}(\pi(\mathbf{p})) = \{s\}$. Yet by Lemma 1 $\text{supp}(\pi(\mathbf{p})) \supseteq \bigcap_i \text{supp}(p_i) = S \setminus \{s\}$, a contradiction. Q.e.d.

Claim 10: π coincides with the geometric rule with weights β_i , $i \in N$.

Note that the geometric rule in question is well-defined by Claim 9. As π and this geometric rule are two continuous functions on \mathbf{P} which by Claim 8 coincide on the topologically dense subdomain \mathbf{P}^* , the two functions coincide globally. ■

A.3 Proof of Fact 3

Assume $F : \mathcal{D} \rightarrow \mathcal{P}$ is continuous and satisfies LIN and GEO, say w.r.t. weights $(\alpha_{i,\mathbf{u}})_{i \in N, \mathbf{u} \in \mathcal{U}}$ and $(\beta_{i,\mathbf{u}})_{i \in N, \mathbf{u} \in \mathcal{U}}$, respectively. Without loss of generality, we assume that in the single-state case $|S| = 1$ (in which the geometric weights $\beta_{i,\mathbf{u}}$ are arbitrary) each $\beta_{i,\mathbf{u}}$ ($i \in N$) is constant in \mathbf{u} . By LIN, group utility only depends on the taste profile. For each taste profile $\mathbf{u} \in \mathcal{U}$, denote the corresponding normalized group utility function by $u_{\mathbf{u}}$; it equals $\sum_i \alpha_{i,\mathbf{u}} u_i$ up to an additive constant.

Claim 1: The mapping $\mathbf{u} \mapsto u_{\mathbf{u}}$ on \mathcal{U} is continuous.

We assume $\mathbf{u}^k \rightarrow \mathbf{u}$ in \mathcal{U} and show $u_{\mathbf{u}^k} \rightarrow u_{\mathbf{u}}$. Pick profiles $(\succeq_i^k) \in \mathcal{D}$ ($k = 1, 2, \dots$) and $(\succeq_i) \in \mathcal{D}$ with identical belief profiles and taste profiles \mathbf{u}^k and \mathbf{u} , respectively. By Lemma 5, $\succeq_i^k \rightarrow \succeq_i$ for each i . Hence, by continuity of F , $\succeq^k \rightarrow \succeq$. So, again by Lemma 5, $u_{\succeq^k} \rightarrow u_{\succeq}$, i.e., $u_{\mathbf{u}^k} \rightarrow u_{\mathbf{u}}$. Q.e.d.

Claim 2: The mapping $\mathbf{u} \mapsto (\alpha_{i,\mathbf{u}})$ from \mathcal{U} to \mathbb{R}^N is continuous.

Let $m := |X|$, and label the outcomes in X by x_1, \dots, x_m . For each $\mathbf{u} = (u_i) \in \mathcal{U}$, we identify $u_{\mathbf{u}}$ with the column vector $(u_{\mathbf{u}}(x_1), \dots, u_{\mathbf{u}}(x_m))^T \in \mathbb{R}^{m \times 1}$, which can be written as $U_{\mathbf{u}} a_{\mathbf{u}}$ where:

$$U_{\mathbf{u}} := \begin{pmatrix} u_1(x_1) & \cdots & u_n(x_1) & 1 \\ \vdots & & \vdots & \vdots \\ u_1(x_m) & \cdots & u_n(x_m) & 1 \end{pmatrix} \in \mathbb{R}^{m \times (n+1)}, \quad a_{\mathbf{u}} := \begin{pmatrix} \alpha_{1,\mathbf{u}} \\ \vdots \\ \alpha_{n,\mathbf{u}} \\ c_{\mathbf{u}} \end{pmatrix} \in \mathbb{R}^{(n+1) \times 1},$$

with $c_{\mathbf{u}}$ defined as the normalization constant such that the minimal entry of $U_{\mathbf{u}} a_{\mathbf{u}}$ is zero (so $c_{\mathbf{u}} = -\min_{x \in X} \sum_i \alpha_{i,\mathbf{u}} u_i(x)$). Since $u_{\mathbf{u}} = U_{\mathbf{u}} a_{\mathbf{u}}$, we have $U_{\mathbf{u}}^T u_{\mathbf{u}} = U_{\mathbf{u}}^T U_{\mathbf{u}} a_{\mathbf{u}}$, where $U_{\mathbf{u}}^T$ is the transpose of $U_{\mathbf{u}}$. By diversity, the functions $u_{i,\mathbf{u}}$ on X ($i \in N$) are affinely independent, and so the columns of $U_{\mathbf{u}}$ are linearly independent. Hence the square matrix $U_{\mathbf{u}}^T U_{\mathbf{u}} \in \mathbb{R}^{(n+1) \times (n+1)}$ is invertible, whence $(U_{\mathbf{u}}^T U_{\mathbf{u}})^{-1} U_{\mathbf{u}}^T u_{\mathbf{u}} = a_{\mathbf{u}}$. To see why the mapping $\mathbf{u} \mapsto a_{\mathbf{u}} = (U_{\mathbf{u}}^T U_{\mathbf{u}})^{-1} U_{\mathbf{u}}^T u_{\mathbf{u}}$ on \mathcal{U} is continuous, note that it is the composition of various continuous functions and operations: $\mathbf{u} \mapsto u_{\mathbf{u}}$ is continuous by Claim 1, $\mathbf{u} \mapsto U_{\mathbf{u}}$ is continuous, and the operations of matrix transposition, matrix inversion and matrix multiplication are continuous. As $\mathbf{u} \mapsto (\alpha_{i,\mathbf{u}})$ is a subfunction of the continuous function $\mathbf{u} \mapsto a_{\mathbf{u}}$, it is itself continuous. Q.e.d.

Claim 3: The mapping $\mathbf{u} \mapsto (\beta_{i,\mathbf{u}})$ from \mathcal{U} to \mathbb{R}^N is continuous.

If $|S| = 1$, then $\mathbf{u} \mapsto (\beta_{i,\mathbf{u}})$ is constant, hence continuous. Now assume $|S| \neq 1$, and let $\mathbf{u}^k \rightarrow \mathbf{u}$ in \mathcal{U} . Fix $j \in N$. We show $\beta_{j,\mathbf{u}^k} \rightarrow \beta_{j,\mathbf{u}}$. Pick distinct $s, s' \in S$. Let (\succeq_i^k) ($k = 1, 2, \dots$) and (\succeq_i) be the profiles in \mathcal{D} with taste profiles \mathbf{u}^k and \mathbf{u} , respectively, and with a same belief profile $\mathbf{p} = (p_i)_{i \in N}$ such that $p_j(s) = \frac{2}{3}$ and $p_j(s') = \frac{1}{3}$ while for $i \neq j$ $p_i(s) = \frac{1}{3}$ and $p_i(s') = \frac{2}{3}$. By Lemma 5 and the fact that $\mathbf{u}^k \rightarrow \mathbf{u}$, we have $(\succeq_i^k) \rightarrow (\succeq_i)$. So, by continuity of F , $\succeq^k \rightarrow \succeq$, whence by Lemma 5 $p_{\succeq^k} \rightarrow p_{\succeq}$. So, as $s, s' \in \text{supp}(p_{\succeq^k}), \text{supp}(p_{\succeq})$ by GEO, $\frac{p_{\succeq^k}(s)}{p_{\succeq^k}(s')} \rightarrow \frac{p_{\succeq}(s)}{p_{\succeq}(s')}$. Since

$$\frac{p_{\succeq^k}(s)}{p_{\succeq^k}(s')} = \frac{(2/3)^{\beta_{j,\mathbf{u}^k}} (1/3)^{1-\beta_{j,\mathbf{u}^k}}}{(1/3)^{\beta_{j,\mathbf{u}^k}} (2/3)^{1-\beta_{j,\mathbf{u}^k}}} = 2^{2\beta_{j,\mathbf{u}^k}-1} \quad \text{and} \quad \frac{p_{\succeq}(s)}{p_{\succeq}(s')} = \frac{(2/3)^{\beta_{j,\mathbf{u}}} (1/3)^{1-\beta_{j,\mathbf{u}}}}{(1/3)^{\beta_{j,\mathbf{u}}} (2/3)^{1-\beta_{j,\mathbf{u}}}} = 2^{2\beta_{j,\mathbf{u}}-1},$$

it follows that $2^{2\beta_{j,\mathbf{u}^k}-1} \rightarrow 2^{2\beta_{j,\mathbf{u}}-1}$. So, $\beta_{j,\mathbf{u}^k} \rightarrow \beta_{j,\mathbf{u}}$. ■

A.4 Proof of Fact 4

Suppose $F : \mathcal{D} \rightarrow \mathcal{P}$ is linear-geometric with taste-dependent weights, say w.r.t. weights $(\alpha_{i,\mathbf{u}}, \beta_{i,\mathbf{u}})_{i \in N, \mathbf{u} \in \mathcal{U}}$.

F satisfies Minimal Pareto: Assume $(\succeq_i) \in \mathcal{D}$ has common belief $p_{\succeq_i} \equiv p$. Then $p_{\succeq} = p$, as p_{\succeq} is proportional on S to $\prod_i p^{\beta_{i,\mathbf{u}}} = p^{\sum_i \beta_{i,\mathbf{u}}} = p$ where $\mathbf{u} := (u_{\succeq_i})$. So, whenever acts a, b satisfy $a \sim_i b$ for all $i \in N$, then $a \sim b$ because

$$\mathbb{E}_{p_{\succeq}}(u_{\succeq}(a)) - \mathbb{E}_{p_{\succeq}}(u_{\succeq}(b)) = \mathbb{E}_p \left(\sum_i \alpha_{i,\mathbf{u}} \underbrace{(u_{\succeq_i}(a) - u_{\succeq_i}(b))}_{=0} \right) = 0. \quad \text{Q.e.d.}$$

F satisfies Continuity: We assume $(\succeq_i^k) \rightarrow (\succeq_i)$ in \mathcal{D} and show $\succeq^k \rightarrow \succeq$. For all $i \in N$, $u_{\succeq_i^k} \rightarrow u_{\succeq_i}$ and $p_{\succeq_i^k} \rightarrow p_{\succeq_i}$ by Lemma 5. Hence $(u_{\succeq_j^k}) \rightarrow (u_{\succeq_j})$, and thus by continuity of the weights $\alpha_{i,(u_{\succeq_j^k})} \rightarrow \alpha_{i,(u_{\succeq_j})}$ and $\beta_{i,(u_{\succeq_j^k})} \rightarrow \beta_{i,(u_{\succeq_j})}$ for all $i \in N$. So,

$$\sum_i \alpha_{i,(u_{\succeq_j^k})} u_{\succeq_i^k} \rightarrow \sum_i \alpha_{i,(u_{\succeq_j})} u_{\succeq_i} \text{ and } \prod_i [p_{\succeq_i^k}]^{\beta_{i,(u_{\succeq_j^k})}} \rightarrow \prod_i [p_{\succeq_i}]^{\beta_{i,(u_{\succeq_j})}}.$$

Hence $p_{\succeq^k} \rightarrow p_{\succeq}$ and $p_{\succeq^k} \rightarrow p_{\succeq}$ by definition of F . So $\succeq^k \rightarrow \succeq$ by Lemma 5. Q.e.d.

F satisfies Bayesian Updating: Consider $(\succeq_i), (\succeq'_i) \in \mathcal{D}$ and $E \subseteq S$ such that $\succeq'_i = \succeq_{i,E}$ for all i . So (*) $u_{\succeq'_i} = u_{\succeq_i}$ for all i , (**) E is non-null under each \succeq_i (as otherwise some \succeq'_i would be the full-indifference relation, so that $(\succeq'_i) \notin \mathcal{D}$), and (***) $p_{\succeq'_i} = p_{\succeq_i}(\cdot|E)$ for all i . To show that, $\succeq' = \succeq_E$, we prove that $u_{\succeq'} = u_{\succeq}$, E is non-null under \succeq , and $p_{\succeq'} = p_{\succeq}(\cdot|E)$. Write $\mathbf{u} := (u_{\succeq_i}) = (u_{\succeq'_i})$. First, by (*) $\sum_i \alpha_{i,\mathbf{u}} u_{\succeq_i} = \sum_i \alpha_{i,\mathbf{u}} u_{\succeq'_i}$, i.e., $u_{\succeq} = u_{\succeq'}$. Second, E is non-null under \succeq , i.e., $\text{supp}(p_{\succeq}) \cap E \neq \emptyset$, because

$$\text{supp}(p_{\succeq}) \cap E \supseteq \left[\bigcap_i \text{supp}(p_{\succeq_i}) \right] \cap E = \bigcap_i [\text{supp}(p_{\succeq_i}) \cap E] = \bigcap_i \text{supp}(p_{\succeq'_i}) \neq \emptyset,$$

where the ' \supseteq ' holds by Lemma 1 and the ' \neq ' holds by coherence of (\succeq'_i) . Third, note that in (****) each $p_{\succeq'_i}$ is zero on $S \setminus E$ and proportional to p_{\succeq_i} on S , by (**) and (***). To show that $p_{\succeq'} = p_{\succeq}(\cdot|E)$, it suffices to prove that $p_{\succeq'}$ is zero on $S \setminus E$ and proportional to p_{\succeq} on E . This holds because $p_{\succeq'}$ is proportional on S to $\prod_i [p_{\succeq'_i}]^{\beta_{i,\mathbf{u}}}$, which by (****) is zero on $S \setminus E$ and proportional on E to $\prod_i [p_{\succeq_i}]^{\beta_{i,\mathbf{u}}}$, hence to p_{\succeq} . Q.e.d.

References

- Aczél, J. (1966) *Lectures on Functional Equations and their Applications*, New York and London: Academic Press
- Ahn, D., Ergin, H. (2010) Framing contingencies, *Econometrica* 78: 655–695
- Baccelli, J. (2019) The problem of state-dependent utility: a reappraisal, *British Journal for the Philosophy of Science*, forth.
- Baccelli, J., Stewart, R. (2019) Support for geometric pooling, working paper
- Blackorby, C., Donaldson, D., Weymark, J. Social aggregation and the Expected Utility Hypothesis. In: Marc Fleurbaey, Maurice Salles, John A. Weymark (eds.) *Justice, Political Liberalism, and Utilitarianism: Themes from Harsanyi and Rawls*, Cambridge University Press: Cambridge and New York, 2008: 136–183
- Bordley, R. (1982) A multiplicative formula for aggregating probability assessments, *Management Science* 28(10): 1137–48
- Chambers, C. P., Hayashi, T. (2006) Preference aggregation under uncertainty: Savage vs. Pareto, *Games and Economic Behavior* 54: 430–440
- Chambers, C. P., Hayashi, T. (2014) Preference aggregation with incomplete information, *Econometrica* 82: 589–599
- Chateauneuf, A., Cohen, M., Jaffray, J.-Y. (2008) Decision under uncertainty: the classical models, halshs-00348818

- Danan, E., Gajdos, T., Tallon, J. M. (2015) Harsanyi's aggregation theorem with incomplete preferences, *American Economic Journal: Microeconomics* 7: 61–69
- Danan, E., Gajdos, T., Hill, B., Tallon, J.-M. (2016) Robust social decisions, *American Economic Review* 106: 2407–25
- Dietrich, F. (2018) Savage's theorem under changing awareness, *Journal of Economic Theory* 176: 1–54
- Dietrich, F. (2019) A theory of Bayesian groups, *Noûs* 53: 708–736
- Fleurbaey, M. (2009) Two variants of Harsanyi's Aggregation Theorem, *Economics Letters* 105: 300–302
- Fleurbaey, M. (2010) Assessing risky social situations, *Journal of Political Economy* 118: 649–680
- Fleurbaey, M., Mongin, P. (2016) The utilitarian relevance of the Social Aggregation Theorem, *American Economic Journal: Microeconomics* 8: 289–306
- Gajdos, T., Tallon, J. M., Vergnaud, J. C. (2008) Representation and aggregation of preferences under uncertainty, *Journal of Economic Theory* 141, p. 68–99
- Genest, C., Zidek, J. V. (1986) Combining probability distributions: a critique and annotated bibliography, *Statistical Science* 1(1): 113–135
- Gilboa, I., Samet, D., Schmeidler, D. (2004) Utilitarian aggregation of beliefs and tastes, *Journal of Political Economy* 112: 932–938
- Gilboa, I., Samuelson, L., Schmeidler, D. (2014) No-betting pareto dominance, *Econometrica* 82: 1405–1442
- Harsanyi, J. (1955) Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility, *J Political Economy* 63: 309–21
- Karni, E. (1993) Subjective expected utility theory with state-dependent preferences, *Journal of Economic Theory*: 60, 428–438
- Karni, E., Viero, M.-L. (2013) “Reverse Bayesianism”: A Choice-Based Theory of Growing Awareness, *American Economic Review* 103: 2790–2810
- Karni, E., Viero, M.-L. (2020) Reverse Bayesianism: a Generalization, working paper
- Madansky, A. (1964) Externally Bayesian groups, *Technical Report RM-4141-PR*, RAND Corporation
- Mongin, P. (1995) Consistent Bayesian aggregation, *Journal of Economic Theory* 66: 313–51
- Mongin, P. (1997/2016) Spurious unanimity and the Pareto Principle, *THEMA Working Paper*, University of Cergy-Pontoise, 1997. Published in *Economics and Philosophy* 32(3): 511–532, 2016
- Mongin, P., Pivato, M. (2015) Ranking multidimensional alternatives and uncertain prospects, *Journal of Economic Theory* 157, 146–171
- Mongin, P., Pivato, M. (2020) Social preference under twofold uncertainty, *Economic Theory*, forth.
- Nehring, K. (2004) The veil of public ignorance, *Journal of Economic Theory* 119: 247–270
- Russell, J. S., Hawthorne, J., Buchak, L. (2015) Groupthink, *Philosophical Studies* 172: 1287–1309
- Savage, L. J. (1954) *The Foundations of Statistics*, New York: Wiley
- Sprumont, Y. (2018) Belief-weighted Nash aggregation of Savage preferences, *Journal*

- of Economic Theory* 178(C): 222–245
- Wakker, P.P., Zank, H. (1999) State-dependent expected utility for Savage’s state space, *Math. Operations Res.* 24: 8–34.
- Weymark, J. (1993) Harsanyi’s Social Aggregation Theorem and the Weak Pareto Principle, *Social Choice and Welfare* 10: 209–21
- Weymark, J. (1991) A reconsideration of the Harsanyi-Sen debate on utilitarianism. In: Elster, J, Roemer J. (eds.) *Interpersonal comparisons of well-being*, Cambridge University Press, pp. 255–320
- Zuber, S. (2016) Harsanyi’s theorem without the sure-thing principle: on the consistent aggregation of Monotonic Bernoullian and Archimedean preferences, *Journal of Mathematical Economics* 63: 78–83