

Fully Bayesian Aggregation

extended version of January 2021¹

Franz Dietrich²

Abstract

Can a group be an orthodox rational agent? This requires the group's aggregate preferences to follow expected utility (static rationality) and to evolve by Bayesian updating (dynamic rationality). Group rationality is possible, but the only preference aggregation rules which achieve it (and are minimally Paretian and continuous) are the *linear-geometric* rules, which combine individual values linearly and individual beliefs geometrically. Linear-geometric preference aggregation contrasts with classic *linear-linear* preference aggregation, which combines both values and beliefs linearly, and achieves only static rationality. Our characterisation of linear-geometric preference aggregation implies as corollaries a characterisation of linear value aggregation (Harsanyi's Theorem) and a characterisation of geometric belief aggregation.

Keywords: rational group agent, uncertainty, preference aggregation, opinion pooling, static versus dynamic rationality, expected-utility hypothesis, Bayesianism, group rationality versus Paretianism, spurious unanimity, ex-ante versus ex-post Pareto

JEL classification codes: D7, D8

1 Introduction

Economics and other social sciences work with a well-established paradigm of a rational agent. They routinely apply this paradigm to group agents such as households, firms, governments, or entire societies. Such group agents are taken to hold and revise preferences and beliefs, to make decisions, form and revise plans, engage in interactions, compete on markets, or entertain international relations – 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 a rational group agent emerge from combining conflicting attitudes of group members?

¹I dedicate this work to the memory of Philippe Mongin, co-author and friend (1950–2020). The paper has benefited from generous feedback by Jean Baccelli, Marcus Pivato, Martin Rechenauer, and anonymous referees. An earlier version was titled '*The Rational Group*'. The research was supported by the French National Research Agency through three grants (ANR-17-CE26-0003, ANR-16-FRAL-0010 and ANR-17-EURE-0001).

²Paris School of Economics & CNRS, fd@franzdietrich.net, www.franzdietrich.net

This problem matters for two very different reasons. Firstly, it matters to have aggregation-theoretic micro-foundations for the hypothesis of rational groups, so as to legitimize the standard modelling practice of invoking rational groups without modelling group members. This ‘macroscopic’ modelling practice is useful and fruitful, but it would be comforting, to say the least, if group agents could be construed as aggregations of (unmodelled) group members.³ Secondly, one is sometimes explicitly interested in group members, and seeks to combine their attitudes into rational group attitudes, for instance in order to determine ‘fair’ group choices, carried out by a group representative or ‘planner’. This goal is constructive: one seeks to build a rational group agent out of group members, rather than postulating a rational group agent without modelling (or knowing) 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 so far 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. What does it 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: one 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. Following Mongin (1997/2016), the Pareto principle is normatively questionable, because a unanimity can be ‘spurious’: it can rest on conflicting beliefs, and thereby lose its normative force. For example, citizens could unanimously want their state to exit a multi-state union, based on conflicting beliefs: some believe that independence opens their state to the wider world (a consequence they most prefer), while others believe that independence isolates their state (a consequence they most prefer). Taking up this challenge, Gilboa, Samet and Schmeidler (2004) (‘GSS’) propose to restrict the Pareto condition to comparisons between options whose consequences depend only on uncontroversial beliefs, and prove a seminal possibility theorem: the restricted Pareto principle allows the formation of aggregate expected-utility preferences, but only by defining group utility as a linear combination of individual utilities and defining group probability as a linear combination of individual probabilities. I call such preference aggregation *linear-linear*. Further developments are discussed in Section 6.

GSS’s linear-linear approach has two shortcomings. First, it does not make the

³What I call *aggregation-theoretic (micro-)foundations* of group agents differs from so-called *game-theoretic (micro-)foundations* of cooperation. In one case, we assume there is a single actor, the ‘group agent’ (or individual acting on its behalf), whose acts or attitudes we aim to explain or make sense of in terms of aggregation. In the other case, we assume that all group members are agents, whose cooperative actions we aim to explain in terms of individual rationality.

group fully rational, as explained shortly. Second, it relies on a Pareto axiom that is still vulnerable to *indirectly* spurious unanimities, i.e., unanimities that depend on *beliefs that depend on conflicting beliefs*, a problem explained in Section 3 (and identified similarly by Mongin and Pivato 2020). Fixing both problems will lead us to *linear-geometric* aggregation: values are combined linearly, but beliefs geometrically.

How are linear-linear groups irrational? Any complete theory of rationality imposes static and dynamic requirements. Orthodox rational-choice theory is Bayesian; it imposes the static requirement to hold expected-utility preferences and the dynamic requirement to revise preferences by applying Bayesian updating to underlying beliefs.⁴ Despite its name, Bayesian preference aggregation theory has so far pursued only a ‘semi-Bayesian’ or ‘semi-rational’ agenda, by imposing only static rationality on the group agent, not dynamic rationality. That is, the group agent should maximise expected utility, but need not revise à la Bayes. The theory is simply silent on the group’s revision policy; it does nothing to 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. In line with the semi-rational agenda, linear-linear preference aggregation creates a group agent that is statically rational, but dynamically irrational, as has been complained (e.g., Mongin and Pivato 2020).

How else must preferences be aggregated in order to make the group rational? This natural question has remained open. Why? One reason could be that full group rationality seems at first like an unreachable goal, since already static group rationality is so hard to reach. Another reason surely is that the literature predominantly adopts a single-profile approach, i.e., works with a fixed profile of individual preferences and fixed group preferences, instead of adopting a multi-profile approach, i.e., working with a domain of possible profiles and an aggregation function defined on this domain (a notable exception is Blackorby, Donaldson and Weymark 2008). The single-profile

⁴Originally, and for most philosophers, Bayesianism is a theory of (rational) beliefs rather than preferences, imposing the static requirement to hold beliefs in probabilistic form (‘probabilism’) and the dynamic requirement to revise beliefs by Bayesian updating (‘conditionalisation’); see Joyce (2011) and Bradley (2017). Standing in a behaviourist tradition, the economic discipline has recast Bayesianism as a theory of preferences, by recasting the static requirement as the requirement that preferences follow expected utilities, and recasting the dynamic requirement as the requirement that preferences be revised by applying Bayesian updating to underlying beliefs whilst never modifying underlying values (utilities). Also outside the Bayesian world, rationality usually has a static and a dynamic component. In particular, many logical theories of rational beliefs impose the static requirement to hold consistent and deductively closed binary beliefs, and the dynamic requirement to revise those beliefs according to certain belief-revision axioms, e.g., the ‘AGM postulates’ (Alchourron et al. 1985). Also John Broome’s influential philosophical theory of rationality and reasoning takes rationality to impose static and dynamic requirements; see Broome (2013) and Dietrich et al. (2019).

approach makes preference *change* unaddressable, since the profile cannot be transformed. To study dynamic rationality, we therefore adopt a multi-profile approach, which is common elsewhere in social choice theory.⁵

This paper contributes the following theorem: only linear-geometric preference aggregation creates a rational group agent, assuming the group agent depends on group members in a minimally Paretian (and continuous) way. Minimal Paretianism weakens GSS’s Pareto condition further, by restricting the condition to situations (profiles) with common beliefs on *all* events, not just on events on which the two acts under comparison depend. Minimal Paretianism is no longer vulnerable to any spurious unanimities, be they direct or indirect; this addresses the mentioned problem with GSS’s Pareto axiom.

A first lesson is that full group rationality *is* non-trivially possible. This supports the standard modelling hypothesis of rational group agents. Another lesson is that group rationality requires combining beliefs non-classically (i.e., geometrically), but combining values classically (i.e., linearly).

Section 2 gives an example. Sections 3–6 present the theorem, its corollaries and generalisations, and a critical assessment of the findings and their relation to the literature. Proofs are given in the appendix.

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 expected-utility (‘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, 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.

⁵Although a multi-profile setting is not explicitly dynamic (since time is absent), it implicitly opens the door to a dynamic analysis, because one can represent the effect of new information on individual preferences, hence on the profile, and thus on aggregate preferences of the group agent.

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

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, normalised 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.⁶

3 Characterisation of fully rational aggregation

We consider a group of individuals $N = \{1, \dots, n\}$ ($n \geq 2$). I work in the Anscombe-Aumann framework.⁷ So 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 defined on

⁶Such ‘linear-linear rules with variable weights’ still violate Bayes’ rule, except if beliefs are combined dictatorially by concentrating all weight on some individual.

⁷By working in the Anscombe-Aumann framework, we depart from Mongin (1995) and GSS (2004), who work in the Savage framework, and come closer to Harsanyi (1954), who works in the von-Neumann-Neumann framework, a special case of the Anscombe-Aumann framework (i.e., the single-state case).

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 values. It is *normalised* 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 $a \in \mathcal{X}$.

Turning to 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.⁸

Choice options are functions $a : S \rightarrow \mathcal{X}$ ('acts'), representing the prospect of 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 or event is *null* under \succeq if the outcome in it is irrelevant, i.e., all acts that agree outside it are indifferent.

A preference relation \succeq is of *expected-utility* type – 'is EU' – if it maximises 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 one imposes normalisation. The unique p and normalised u are denoted p_\succeq and u_\succeq , respectively. Let \mathcal{P} be the set of EU preference relations.

A *value profile* is a vector $\mathbf{u} = (u_i)$ of normalised utility functions u_i of individuals $i \in N$, summarising the values of individuals. A *belief profile* is a vector $\mathbf{p} = (p_i)$ of probability functions on the set 2^S of events, summarising the beliefs of individuals. An *EU preference profile* is a vector $(\succeq_i) \in \mathcal{P}^n$ of EU preference relations; it summarises the preferences of individuals, and indirectly encodes their values and their beliefs, contained in the value profile (u_{\succeq_i}) and belief profile (p_{\succeq_i}) .

Bayesian aggregation theory usually works with a fixed preference profile. To study preference *change*, we take a multi-profile approach: we consider a set of possible profiles $\mathcal{D} \subseteq \mathcal{P}^n$, the *domain*. For generality, the theorem will assume little about the domain. It will assume that \mathcal{D} is *regular*, by which I mean two things. 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.⁹ Second, to allow belief revision, \mathcal{D} is *closed under belief change*: whenever \mathcal{D} contains a coherent profile $(\succeq_i) \in \mathcal{P}^n$, then \mathcal{D} also contains each coherent profile $(\succeq'_i) \in \mathcal{P}^n$ that differ from (\succeq_i) only in beliefs, not in values.

An *EU preference aggregation rule*, or simply a *rule*, is a function transforming EU preference profiles (from some domain) into group EU preference relations, i.e., a

⁸In the two-state case, our axioms are necessary, but no longer sufficient for linear-geometric aggregation, as they for instance permit linear-linear aggregation.

⁹As nobody can possess conclusive evidence against the truth.

function $F : \mathcal{D} \rightarrow \mathcal{P}$ on some domain $\mathcal{D} \subseteq \mathcal{P}^n$.¹⁰

The following type of rules will emerge from our analysis:

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

- *utility function u_\succeq given by $\sum_i \alpha_i u_{\succeq_i}$ up to an additive constant,*
- *probability function p_\succeq given on states by $\prod_i [p_{\succeq_i}]^{\beta_i}$ up to a multiplicative constant.¹¹*

Such rules contrast with *linear-linear* rules, which are defined analogously except from replacing the second bullet point by:

- *probability function p_\succeq given by $p_\succeq = \sum_i \beta_i p_{\succeq_i}$.*

I now state simple axioms on a rule $F : \mathcal{D} \rightarrow \mathcal{P}$. The core axiom requires rational revision of group preferences. Rational revision is revision by conditionalisation. The *conditionalisation* of a preference relation \succeq on an event $E \subseteq S$ is a new preference relation, denoted \succeq_E , which can be defined in two equivalent ways. Under one definition, which assumes that \succeq is EU (and E is non-null), \succeq_E is the relation obtained from \succeq by conditionalising beliefs and leaving utilities unchanged, i.e., the unique EU relation such that $p_{\succeq_E} = p_\succeq(\cdot|E)$ and $u_{\succeq_E} = u_\succeq$. The general or abstract definition, due to Savage, needs no reference to utilities and probabilities: \succeq_E is the relation such that, for any acts a, b , we have $a \succeq_E b$ if and only if we have $a' \succeq_E b'$ for some acts a', b' such that (i) in the event E , a' agrees with a and b' agrees with b , and (ii) outside E , a' agrees with b' . In short, $a \succeq_E b$ means that a becomes weakly preferred to b after equalizing ('ignoring') outcomes outside E . The axiom can now be stated.

Dynamic Rationality (or Bayesian Updating): When information is learnt by everyone, the new group preferences equal the old ones conditional on the information. Formally, if a profile $(\succeq'_i) \in \mathcal{D}$ arises from another $(\succeq_i) \in \mathcal{D}$ by conditionalising preferences on an event $E \subseteq S$, i.e., $\succeq'_i = \succeq_{i,E}$ for all i , then the new group relation $\succeq' = F((\succeq'_i))$ arises by conditionalising $\succeq = F((\succeq_i))$ on E , i.e., $\succeq' = \succeq_E$.

In short, post-information group preferences must be equal to pre-information group preferences conditional on information. Violation of this axiom makes the group dynamically irrational and unable to execute stable plans. This may also create opportunities to manipulate group preferences and decisions through delaying information.

¹⁰Our notions of 'rule' and 'domain' build in the assumption that preferences (individual or collective) are of EU type, i.e., statically rational. Generalised notions of 'rule' and 'domain' could drop this assumption.

¹¹In 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 value weights α_i and the additive constant are unique under the diversity condition defined in Section 5. Following standard convention, $t^0 = 1$ for all $t \in \mathbb{R}_+$, even for $t = 0$. This convention allows one to interpret $\prod_i [p_{\succeq_i}(s)]^{\beta_i}$ ($s \in S$) even if for some β_i and $p_{\succeq_i}(s)$ are both zero.

Before defining our own Pareto axiom, recall the standard Pareto (indifference) axiom:

Standard 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 criticised because unanimous preferences or indifferences can be *spurious*: they can rest on conflicting beliefs (Mongin 1997/2016). In the example of Section 2, the unanimous preference for duelling is spurious, as the gentlemen hold (very) different subjective probabilities. GSS have therefore restricted Standard Pareto to acts that depend only on events of uncontroversial probability. Formally, given a profile (\succeq_i) , an act a is *common-belief-based* if all individual probability functions p_{\succeq_i} ($i \in N$) agree on the events on which a depend, where the events on which a depend are the events of the form $\{s \in S : a(s) \in L\} = a^{-1}(L)$ for some $L \subseteq \mathcal{L}$, i.e., the events of the form ‘ a has such-and-such result’. An act can be common-belief-based even if beliefs disagree on many events, on which a does not depend.¹²

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-based acts $a, b \in \mathcal{X}^S$, if $a \sim_i b$ for each $i \in N$, then $a \sim b$ (where $\succeq = F((\succeq_i))$).

However, this axiom is only safe against ‘direct’ 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 (possibly unanimous) beliefs that are based on conflicting beliefs (Mongin and Pivato 2020 draw a related distinction to question GSS’s axiom). Here is an example of a unanimous preference that is not directly, but indirectly spurious: a group of individuals unanimously prefers the UK to remain within the EU based on a unanimous strong belief in the event E that ‘Brexit’ harms the British economy (and a unanimous concern for the economy), but this belief is based on conflicting reasons, i.e., conflicting beliefs about epistemically prior events. Some individuals might strongly believe E because they strongly believe that the governmental advisers said so (event E') and that advisers tell the truth (event E''), while others strongly believe E because they strongly believe that the governmental advisers denied E (event $\overline{E'}$) and that they are mistaken (event $\overline{E''}$). The former individuals reason from E' and E'' to E (using $E' \cap E'' \subseteq E$), while the latter individuals reason from $\overline{E'}$ and $\overline{E''}$ to E (using $\overline{E'} \cap \overline{E''} \subseteq E$).

Defining the notion of indirect spuriousness formally would require enriching the model, presumably by a relation of epistemic priority between events or something that generates such a relation, because our current model only captures *logical* relations between events, from which one cannot generally read off the direction of epistemic priority or reasoning.¹³ But for present purposes we need not define indirect spuriousness formally. We only aim to state a Pareto axiom that is safe against all spurious

¹²For instance, even under highly heterogeneous beliefs, all *constant* acts are common-belief-based, because such acts depend only on the trivial events S and \emptyset , on which beliefs must agree.

¹³For instance, in the example the fact that $E' \cap E'' \subseteq E$ does not yet establish that someone who strongly believes E , E' and E'' has reasoned from E' and E'' to E . The person might instead have come

unanimities, direct or indirect. Such a safe Pareto axiom is obtained by restricting the Pareto axiom further, namely to *common-belief profiles*, i.e., profiles (\succeq_i) in which every individual i has same probability function p_{\succeq_i} . In such profiles, not only the beliefs underlying two given acts are unanimous, but also the beliefs underlying those beliefs, the beliefs underlying beliefs underlying those beliefs, etc. This automatically 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 Standard Pareto.*

Our last axiom requires group preferences to depend continuously on individual preferences, where ‘convergence’ is defined in the natural way.¹⁴

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. I do this here only for *fixed-values domains*, i.e., domains \mathcal{D} whose profiles all have the same value profile; Section 5 re-states the theorem without fixing values. Fixed-values domains allow variation in beliefs, but not in values. Fixing values is perhaps not a great loss, as Bayesian learning changes beliefs, not values; and it is a ‘half-way’ concession to the fixed-profile approach of standard Bayesian aggregation theory, which fixes values *and beliefs*.

Theorem 1 *An EU preference aggregation rule $F : \mathcal{D} \rightarrow \mathcal{P}$ on a (regular) fixed-values domain \mathcal{D} is dynamically rational, minimally Paretian, and continuous if and only if it is linear-geometric.*

I now discuss two extensions and the special case of ‘no uncertainty’.

Extension 1: non-public information. The axiom of Dynamic Rationality 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 *semi-private* information (observed by at least one individual). These are the two modified axioms, where I underline the changes from the original:

- *Dynamic Rationality for Private [or: Semi-private] Learning:* When information is learnt by exactly [or: at least] one individual, the new group preferences equal the old ones conditional on the information. Formally, if a profile $(\succeq'_i) \in \mathcal{D}$ arises

up independently with the three beliefs (without any epistemic priorities between these events), or have reasoned from E and some further event G to E' and E'' using that $E \cap G \subseteq E'$ and $E \cap G \subseteq E''$ (in which case the epistemic priority goes the other way round). This illustrates that logical relations underdetermine epistemic priority.

¹⁴A 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 values 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 .

from another $(\succeq_i) \in \mathcal{D}$ by conditionalising the preferences of exactly [or: at least] one individual on an event $E \subseteq S$, then the new group relation $\succeq' = F((\succeq'_i))$ arises by conditionalising $\succeq = F((\succeq_i))$ on E .

The axiom with semi-private information is obviously logically stronger than the original axiom. Even the axiom with private information is logically stronger, because public learning of E can be decomposed into n steps of private learning of E by each individual in turn, assuming the domain \mathcal{D} is regular (to ensure the profile stays in \mathcal{D} in each step). These stronger axioms can still be met, by only slightly fewer rules. We must only exclude rules that ignore someone's beliefs, to avoid that private learning by someone is collectively ignored. More precisely, Theorem 1 still holds if we strengthen Dynamic Rationality in one of the two ways (no matter which) and require strict positivity of the belief weights β_i ($i \in N$).

Extension 2: respecting individual values. Theorem 1 imposes no constraints on the sign of the value weights α_i . Individual values could be ignored ($\alpha_i = 0$) or even counted negatively ($\alpha_i < 0$). We can enforce non-negativity or even positivity of the value weights α_i by strengthening the Minimal Pareto axiom in Theorem 1 in natural ways. To enforce non-negativity, replace the axiom's indifferences by weak preferences. To enforce positivity, do the same and 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 $|S| = 1$. 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 the axioms of Dynamic Rationality 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 maximises the expectation of some non-constant utility function u on X ; the unique normalised 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 $u_\succeq = \sum_i \alpha_i u_{\succeq_i} + \gamma$ 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 *opinion pooling theory* – a debate that should finally reach Bayesian preference aggregation theory. Opinion pooling theory is concerned with merging subjective probabilities, not preferences. It has long recognised that linear and geometric pooling each have one significant argument in its favour.

Pro geometric. Geometric belief aggregation has superior dynamic behaviour (see Section 6 for discussion of the literature). In fact, Theorem 1 implies a new version of this result, namely a ‘belief analogue’ of Theorem 1:

Theorem 2 *A belief aggregation rule (on the domain of coherent belief profiles) is dynamically rational, unanimity-preserving, and continuous if and only if it is geometric.*

The concepts in Theorem 2 (such as ‘dynamically rational’) are belief-theoretic counterparts of the preference-theoretic concepts in Theorem 1. The appendix gives the obvious formal definitions, and proves Theorem 2 by reducing it to a special case of Theorem 1.

Pro linear. Linear opinion pooling enjoys the following (conditional) robustness property: if states are refined or coarsened, then group beliefs are not perturbed, i.e., merely ‘extended’ or ‘restricted’, assuming that individual beliefs are not perturbed, i.e., merely ‘extended’ or ‘restricted’ (a questionable assumption unfortunately, as explained shortly). This robustness property holds similarly in the context of aggregating preferences rather than beliefs.

How important is such robustness in the preference-aggregation 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, i.e., has aggregation-theoretic micro-foundations. An ideal 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. Coarsening states applies to *non*-ideal agents. Robustness however matters if we aim to construct group preferences in practice. In practice, one cannot work with the maximally fine states of ideal rational agents. This opens the door to problems of state design and robustness.

Unfortunately, in practice not even linear pooling is 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 the states of the model 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, axiomatised 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 How stable are the individual weights?

Theorem 1 delivers a type of aggregation in which the impact or ‘say’ of an individual i , given by his pair of weights α_i and β_i , does not depend on the profile, hence is robust to belief change: individual weights remain the same when beliefs change. But what happens to weights when *values* change? This question is bracketed out by Theorem 1 by assuming a fixed-values domain. To answer the question, I now re-state Theorem 1 for domains in which not just beliefs, but also values can change. The finding will be that individual weights can depend on values, but not on beliefs.

Technically, the extended theorem will make a standard assumption on profiles (because of which it is, strictly speaking, not logically stronger than Theorem 1). An EU preference profile (\succeq_i) is *diverse* – one might say, ‘diverse in values’ – 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 values, in the following sense:

Definition 2 *An EU preference aggregation rule $F : \mathcal{D} \rightarrow \mathcal{P}$ ($\mathcal{D} \subseteq \mathcal{P}^n$) is linear-geometric with values-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 value profile $\mathbf{u} \in \{(u_{\succeq_i}) : (\succeq_i) \in \mathcal{D}\}$, where $\sum_i \beta_{i,\mathbf{u}} = 1$ for each \mathbf{u} , such that at each profile $(\succeq_i) \in \mathcal{D}$ the group preference relation $\succeq = F((\succeq_i))$ has*

- *utility function u_{\succeq} given by $\sum_i \alpha_{i,\mathbf{u}} u_{\succeq_i}$ up to an additive constant,*
- *probability function p_{\succeq} given on states by $\prod_i [p_{\succeq_i}]^{\beta_{i,\mathbf{u}}}$ up to a multiplicative constant,*

where \mathbf{u} denotes the current value profile (u_{\succeq_i}) .¹⁵

Theorem 1⁺ *An EU preference aggregation rule $F : \mathcal{D} \rightarrow \mathcal{P}$ on a (regular) domain \mathcal{D} of diverse profiles is dynamically rational, minimally Paretian, and continuous if and only if it is linear-geometric with values-dependent weights.*

Extensions 1 and 2 translate 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-values domain. Aside from the diversity condition, Theorem 1 is a special case of Theorem 1⁺: the case of fixed values. The reason is that by fixing values we fix the index \mathbf{u} in the weights ‘ $\alpha_{i,\mathbf{u}}$ ’ and ‘ $\beta_{i,\mathbf{u}}$ ’; so this index drops out, and we obtain ordinary linear-geometric rules.

As individual weights depend only on values, not on beliefs, weights still enjoy a limited amount of stability or profile-independence. By contrast, the weights in GSS’s linear-linear result can vary with both values and beliefs, even discontinuously. This full profile-dependence would become visible if GSS’s result were re-stated as a multi-profile result, i.e., as a result about an aggregation rule defined on a domain of profiles. The reason why our fully rational approach leads to more weight stability than GSS’s less rational and more Paretian approach is that we adopt inter-profile conditions (Dynamic

¹⁵Footnote 11 about uniqueness of weights applies analogously.

Rationality and Continuity), which connect and structure aggregation across profiles. GSS adopt only *intra*-profile conditions (static rationality and restricted Pareto), which cannot prevent the weights from jumping arbitrarily across profiles.

One may conjecture that combining both approaches leads to an impossibility theorem: there is *no* non-degenerate rule that is fully rational and also Paretian in GSS's stronger sense. Why? By Theorem 1⁺, such a rule must be linear-geometric with values-dependent weights (taking 'non-degenerate' to subsume 'continuous'). On the other hand, the rule must be linear-linear with profile-dependent weights, assuming GSS's theorem translates to our framework, a conjecture whose proof lies outside the scope of this paper.¹⁶ To be both simultaneously, some individual must receive the full belief weight of one, and this individual must (by continuity) be the same at all profiles. So the rule must be dictatorial *in beliefs*, hence degenerate; it can be non-dictatorial (only) *in values*, as (only) the value weights can be distributed freely.

6 Discussion in relation to the literature

I have proposed to model groups as fully 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 seems particularly unobjectionable, as it is safe against directly and indirectly spurious unanimities.

From here, important questions open up. How about weakening *static* group rationality into non-EU directions while preserving Dynamic Rationality, suitably extended? Such a dynamically (not statically) rational approach would be the dual of the statically (not dynamically) rational programme of Harsanyi, Mongin, and Gilboa-Samet-Schmeidler. And how about *geometric-geometric* rules, which pool even values geometrically? Such rules remain fully rational, 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 interpretation ever since (Weymark 1991, Fleurbaey and Mongin 2016). Harsanyi's Theorem enjoys some robustness within the limited world of objective uncertainty; see generalisations by Fleurbaey (2009, 2014) and Danan, Gajdos and Tallon (2015). As mentioned, the picture reverses for heterogeneous beliefs: EU rationality then becomes

¹⁶A version of GSS's theorem within our Anscombe-Aumann and multi-profile framework would assert that (on a domain of such-and-such type) the linear-linear rules with profile-dependent weights are the only restrictedly Paretian rules for generating statically rational group preferences, possibly assuming some non-degeneracy conditions on the rules. Such a result could also be stated in a single-profile format, unless the result involves a non-degeneracy condition of an inter-profile nature.

incompatible with Paretianism (e.g., Mongin 1995), and meanwhile Paretianism becomes less compelling because unanimities can be spurious (Mongin 1997/2016); but Gilboa et al. (2004) restore possibility 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 static EU 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 rationality. 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* Savage’s 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 rationalisable (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). Different possible reactions might be taken in the face of this intriguing underdetermination 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. The former position could be described as mentalist or realist, the latter position as strongly behaviourist or instrumentalist. Sprumont (2018) offers such a behaviourist view.

The Dynamic Rationality axiom is a counterpart for preference aggregation of the classic *External Bayesianity* axiom in the theory of belief aggregation or ‘opinion pooling’ (Madansky 1964, Bordley 1982, Genest and Zidek 1986, Dietrich and List 2016). Like our axiom, External Bayesianity requires aggregation to commute with revision; but it differs firstly in the objects that are aggregated and revised, namely probabilities rather than preferences, and secondly in the information that is learnt and triggers revision, namely information given by a *likelihood function* rather than an event, i.e., a function mapping each state to (what is interpreted as) the probability or probability-density of an information conditional on the state. Informally, likelihood functions

capture information outside the domain (event algebra) relative to which beliefs are held. If states are weather states, then the information ‘it rains’ is representable by an event (the set of states in which it rains), hence lies inside the event algebra, whereas the information ‘the radio forecasts that it rains’ lies outside the event algebra, and is representable by a likelihood function that takes values near 1 in ‘rainy’ states (in which this forecast is likely) and values near 0 in ‘sunny’ states (in which this forecast is unlikely). While events represent the classical information concept in rational choice theory and Bayesianism, likelihood functions represent an information concept that is natural in statistics.¹⁷ This might explain why opinion pooling theory (a field influenced by statistics) uses likelihood-function-type information in its External Bayesianity axiom. As is well-known, geometric opinion pooling rules satisfy External Bayesianity, but are not the only well-behaved rules doing so (e.g., Dietrich and List 2016, Baccelli and Stewart 2019). Geometric opinion pooling rules become the *only* well-behaved solutions if External Bayesianity is re-stated using standard event-type information, thereby making group beliefs classically Bayesian; this is what is shown by Theorem 2 (and, in a preference-aggregation context, by Theorem 1). This move towards classically Bayesian group beliefs is already made in recent work on opinion pooling, in which close variants of Theorem 2 are established (Russell et al. 2015, Dietrich 2019). The upshot of Theorem 2 compared to its precursors is that it addresses geometric opinion pooling in the ordinary sense, while the two earlier classically Bayesian results address a generalised version of geometric pooling, in which the individual weights need not sum to one.

Just as geometric opinion pooling can be generalised by lifting the constraint on the sum of weights, so linear-geometric *preference* aggregation might be generalised by lifting the constraint on the sum of belief weights. Such preference aggregation would still be dynamically rational and continuous, but no longer minimally Paretian, since generalised geometric pooling does not preserve unanimously held probability functions.

A Proof of Theorem 2 on opinion pooling

This appendix provides the formal definitions underlying Theorem 2, and then gives two proofs of Theorem 2, a (short) proof based on Theorem 1, and an independent proof. Having also an independent proof is crucial for us, since Theorem 2 will later serve as a lemma on the way to prove Theorem 1. So, Theorem 2 is both a corollary of, and a proof step towards, our central result. Theorem 2 and its (independent) proof

¹⁷If the states represent the possible values of the unknown parameter in a statistical model, then statistical information (‘data’) is representable by a likelihood function, defined as the data’s parameter-dependent probability (or probability density). *Bayesian* statistics requires prior beliefs over the possible parameter values, and revises them when observing data. This revision is equivalent to conditionalisation on a likelihood function (the one generated by the data), i.e., revision of the type considered in the External Bayesianity axiom. This is the statistical rationale underlying External Bayesianity and the likelihood-function-based concept of information. Seen from this perspective, External Bayesianity is a natural condition in situations where different Bayesian statisticians work together, but disagree on the prior beliefs over the parameters, and hence must aggregate their prior beliefs into compromise prior beliefs (on which the statistical analysis is then based). By contrast, the ordinary event-based concept of information seems more natural from a rational-choice-theoretic perspective.

are close variants of existing ones in Russell et al. (2015) and Dietrich (2019).¹⁸ For completeness, I still give a self-contained proof here.

The definitions underlying Theorem 2. Theorem 2 needs fewer formal primitives than Theorem 1: it only needs our (finite non-empty) set of states S , not our set of outcomes. A *belief aggregation (or opinion pooling) rule* is a function π mapping each belief profile $\mathbf{p} = (p_i)$ from some domain $\mathbf{D} \subseteq \Delta(S)^n$ to a group probability function $\pi(\mathbf{p}) \in \Delta(S)$, hence a function $\pi : \mathbf{D} \rightarrow \Delta(S)$. The rule is *geometric* if there exist weights $\beta_i \geq 0$ ($i \in N$) of sum one such that, at each belief profile $\mathbf{p} \in \mathbf{D}$, $\pi(\mathbf{p})$ is given on states by $\prod_i [p_i]^{\beta_i}$ up to a multiplicative constant. The three axioms used in Theorem 2 are defined as follows:

- *Dynamic Rationality (or Bayesian Updating)*: When information is learnt by everyone, the new group beliefs equal the old ones conditional on the information. Formally, if a belief profile $(p'_i) \in \mathbf{D}$ arises from another $(p_i) \in \mathbf{D}$ by conditionalising beliefs on an event $E \subseteq S$, i.e., for all i , $p'_i = p_i(\cdot|E)$ (in particular $p_i(E) > 0$), then the new group belief $\pi((p'_i))$ arises by conditionalising $\pi((p_i))$ on E , i.e., $\pi((p'_i)) = \pi((p_i))(\cdot|E)$ (in particular $\pi((p_i))(E) > 0$).
- *Unanimity Preservation*: $\pi(p, \dots, p) = p$ for all unanimous belief profiles $(p, \dots, p) \in \mathbf{D}$.
- *Continuity*: If $\mathbf{p}^k \rightarrow \mathbf{p}$ in \mathbf{D} (i.e., $p_i^k \rightarrow p_i$ for all $i \in N$), then $\pi(\mathbf{p}^k) \rightarrow \pi(\mathbf{p})$.

Finally, a belief profile $\mathbf{p} \in \Delta(S)^n$ is *coherent* if some state has positive probability under each individual belief, i.e., $\bigcap_i \text{supp}(p_i) \neq \emptyset$.

Proof of Theorem 2 from Theorem 1. Assume Theorem 1 holds. Fix a (normalised) utility function u . Given u , each probability measure $p \in \Delta(S)$ corresponds uniquely to an EU preference relation $\succeq \in \mathcal{P}$, defined by $p_{\succeq} = p$ and $u_{\succeq} = u$. Hence, any belief aggregation rule $\pi : \mathbf{D} \rightarrow \Delta(S)$ on some domain $\mathbf{D} \subseteq \Delta(S)^n$ corresponds to an EU preference aggregation rule on a corresponding domain $\mathcal{D} \subseteq \mathcal{P}^n$. As in Theorem 2, consider a belief aggregation rule $\pi : \mathbf{D} \rightarrow \Delta(S)$ on the domain \mathbf{D} of coherent belief profiles, and let $F : \mathcal{D} \rightarrow \mathcal{P}$ be the corresponding EU preference aggregation rule. Since \mathbf{D} consists of all coherent belief profiles, \mathcal{D} is a regular fixed-values domain, i.e., a domain of the type assumed in Theorem 1 (more precisely, \mathcal{D} is the fixed-values domain w.r.t. the unanimous value profile (u, \dots, u)). The axioms in Theorem 2 reduce to those in Theorem 1:

- (a) π is dynamically rational if and only if F is dynamically rational, essentially because, as values are fixed, aggregation-revision commutativity for preferences reduces to aggregation-revision commutativity for underlying beliefs.
- (b) π is continuous if and only if F is continuous, essentially because, again as values are fixed, convergence of preferences reduces to convergence of underlying beliefs (drawing on Lemma 5 below).
- (c) π is unanimity-preserving if and only if F is minimally Paretian. Why? First, if π

¹⁸Theorem 2 shares two premises with Russell et al.'s theorem and with Dietrich's previous theorem (i.e., Thm. 2 in Dietrich 2019), namely Dynamic Rationality and Continuity. While Theorem 2 characterises geometric pooling, the two precursors characterise generalised geometric pooling, in which the individual weights need not have sum one.

is unanimity-preserving, then F is minimally Paretian, because at any common-belief profile $(\succeq_i) \in \mathcal{D}$ the group relation $F((\succeq_i))$ has the same probability function as all \succeq_i (by unanimity-preservation) and the same utility function as all \succeq_i (namely u), hence is identical to all \succeq_i , and thus preserves individual indifferences. Conversely, assume F is minimally Paretian. To show that π is unanimity-preserving, consider a unanimous belief profile $(p, \dots, p) \in \mathbf{D}$. We fix a state $s \in S$ and show that $\pi(p, \dots, p)(s) = p(s)$. Let f be a constant act which at each state generates the same lottery having an outcome of utility 1 with probability $p(s)$ and an outcome of utility 0 with probability $1 - p(s)$. Let g be an act which at s has (surely) an outcome of utility 1 and at other states has (surely) an outcome of utility 0. For each individual, f and g both have expected utility $p(s)$, hence are indifferent. So, by minimal Pareto, f and g are collectively indifferent. Since f 's collective expected utility is $p(s)$, so is g 's. Hence the collective probability of s is $p(s)$.

Not only do the axioms of both theorems correspond, but also F is linear-geometric if and only if π is geometric, because linearity in values holds trivially, given that the group has the same utility function as all individuals. So Theorem 2 reduces to (the current instance of) Theorem 1. ■

Direct proof of Theorem 2. Consider a belief aggregation rule π on $\mathbf{D} = \{\mathbf{p} \in \Delta(S)^n : \mathbf{p} \text{ is coherent}\}$. The result is trivial if $|S| = 1$. So, as $|S| \neq 2$, we can assume without loss of generality that $|S| \geq 3$.

Part 1 (necessity of the axioms). First, let F be geometric, say with weights β_i ($i \in N$). Trivially, F is then unanimity-preserving (using that $\sum_i \beta_i = 1$) and continuous. To show dynamic rationality, assume a belief profile $(p'_i) \in \mathcal{D}$ arises from another one $(p_i) \in \mathcal{D}$ by conditionalising each p_i on E , i.e., $p_i(E) \neq 0$ and $p'_i = p_i(\cdot|E)$ for all i . We must show that $\pi((p'_i))$ arises by conditionalising $\pi((p_i))$ on E , i.e., that (*) $\pi((p_i))(E) \neq 0$ and (**) $\pi((p'_i)) = \pi((p_i))(\cdot|E)$. To show (*), note first that $\text{supp}(p_\succeq) \supseteq \bigcap_i \text{supp}(p_i)$ by definition of geometric pooling. So,

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

where the ' \neq ' holds by coherence of (\succeq'_i) . So, $\text{supp}(p_\succeq) \cap E \neq \emptyset$, implying (*). To show (**), note first that $\pi((p'_i))$ is proportional to $\prod_i [p'_i]^{\beta_i}$ on states, where each p'_i ($= p_i(\cdot|E)$) is itself proportional to p_i on states in E and zero on other states. So, $\pi((p'_i))$ is proportional to $\prod_i [p_i]^{\beta_i}$ (hence to $\pi((p_i))$) on states in E and zero on other states. This implies (**). Q.e.d.

Part 2 (sufficiency of the axioms). Now assume π satisfies the three axioms. We prove several claims, the last of which is that π is geometric, as intended.

Claim 1: For all $\mathbf{p} \in \mathbf{D}$ 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 $E = \{s, t\}$. Let $p' \in \Delta(S)$ be given by $p'(s) = p'(t) = \frac{1}{2}$. Form the unanimous profile $\mathbf{p}' = (p', \dots, p')$. Note that $\mathbf{p}(\cdot|E) = \mathbf{p}'$. So, by Dynamic Rationality, $\pi(\mathbf{p})(E) \neq 0$ and $\pi(\mathbf{p})(\cdot|E) = \pi(\mathbf{p}')$. Hence, as $\pi(\mathbf{p}') = p'$ by Unanimity-Preservation, $\pi(\mathbf{p})(s|E) = \pi(\mathbf{p})(t|E) = \frac{1}{2}$. By implication, $\pi(\mathbf{p})(s) = \pi(\mathbf{p})(t) \neq 0$. 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 $\mathbf{p} \in \mathbf{D}$ 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 $\mathbf{p} \in \mathbf{D}$. As for existence, let $\mathbf{p}, \mathbf{p}' \in \mathbf{D}$ 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 Dynamic Rationality $\pi(\mathbf{p})(\cdot|E) = \pi(\mathbf{p}')(\cdot|E)$. 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 because $\pi(\mathbf{p})(s), \pi(\mathbf{p})(t), \pi(\mathbf{p}')(s), \pi(\mathbf{p}')(t) \neq 0$ by Dynamic Rationality. 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 $\mathbf{p} \in \mathbf{D}$ 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 $\mathbf{p} \in \mathbf{D}$ 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 $\mathbf{p} \in \mathbf{D}$ 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 $f(\mathbf{x}) = x_1^{\beta_1} \cdots x_n^{\beta_n}$ for all $\mathbf{x} \in (0, \infty)^n$.

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 Continuity. 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: $\bigcap_i \text{supp}(p_i) \subseteq \text{supp}(\pi(\mathbf{p}))$.

We fix an $s \in \bigcap_i \text{supp}(p_i)$ and show $s \in \text{supp}(\pi(\mathbf{p}))$. Consider the event $E = \{s\}$. The updated profile $\mathbf{p}(\cdot|E)$ is still coherent, i.e., in \mathbf{D} , so that by Dynamic Rationality $\pi(\mathbf{p}(\cdot|E)) = \pi(\mathbf{p})(\cdot|E)$, and in particular $\pi(\mathbf{p})(E) > 0$. So $s \in \text{supp}(\pi(\mathbf{p}))$. Q.e.d.

Claim 9: Consider the subdomain of full-support profiles $\mathbf{D}^* := \{\mathbf{p} \in \mathbf{D} : \text{supp}(p_i) = S \text{ for all } i \in N\}$. For each $\mathbf{p} \in \mathbf{D}^*$, $\pi(\mathbf{p})$ is given on states by $\prod_i [p_i]^{\beta_i}$ up to a multiplicative constant.

Let $\mathbf{p} \in \mathbf{D}^*$. 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 $\text{supp}(\pi(\mathbf{p})) \supseteq \bigcap_i \text{supp}(p_i) = S$ by Claim 8. 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 10: $\beta_i \geq 0$ 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 $\mathbf{p} \in \mathbf{D}^*$ 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 Unanimity-Preservation, $\pi(\mathbf{p}) = p$. So by Claim 9, p is proportional on states 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 $\mathbf{p}^1, \mathbf{p}^2, \dots \in \mathbf{D}^*$ that converges to a $\mathbf{p} \in \mathbf{D} \setminus \mathbf{D}^*$ which satisfies $\text{supp}(p_i) = S \setminus \{s\}$ and $\text{supp}(p_j) = S$ for all $j \in N \setminus \{i\}$. By $\beta_i < 0$ and Claim 9, the sequence $\pi(\mathbf{p}^1), \pi(\mathbf{p}^2), \dots$ converges to the measure assigning probability 1 to s . Meanwhile by Continuity the limit must be $\pi(\mathbf{p})$. So $\pi(\mathbf{p})(s) = 1$, whence $\text{supp}(\pi(\mathbf{p})) = \{s\}$. Yet by Claim 8 $\text{supp}(\pi(\mathbf{p})) \supseteq \bigcap_i \text{supp}(p_i) = S \setminus \{s\}$, a contradiction. Q.e.d.

Claim 11: 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 10. As π and this geometric rule are two continuous functions on \mathbf{D} which by Claim 9 coincide on the topologically dense subdomain \mathbf{D}^* , the two functions coincide globally. ■

B Proof of Theorems 1 and 1⁺ on preference aggregation

I now prove Theorem 1⁺. Although Theorem 1 about fixed-values domains is not quite a corollary (as it does not assume diverse profiles), we need no separate proof of Theorem 1, because in the case of fixed-values domains the proof of Theorem 1⁺ does not require diversity.

Assumptions and notation: Fix a regular domain \mathcal{D} of diverse profiles, and a rule $F : \mathcal{D} \rightarrow \mathcal{P}$. By convention, the group relation obtained by aggregating individual relations is denoted using the same symbol as for individuals, but without individual index: so $F((\succeq_i))$ is denoted by \succeq , $F((\succeq'_i))$ is denoted by \succeq' , etc. Let $\mathcal{U} := \{(u_{\succeq_i}) : (\succeq_i) \in \mathcal{D}\}$ be the set of occurring values profiles. Let $\mathcal{D}_{Com\ Bel} \subseteq \mathcal{D}$ be the subdomain of common-belief profiles in \mathcal{D} .

Consider two separate conditions on the rule F :

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 $\prod_i [p_i]^{\beta_{i,(u_{\succeq_j})}}$ up to a multiplicative constant.

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: *Dynamic Rationality and Minimal Pareto imply LIN.*

Fact 2: *Dynamic Rationality, Minimal Pareto, and Continuity imply GEO.*

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

Fact 4: *All linear-geometric rules with values-dependent weights satisfy Dynamic Rationality, Minimal Pareto, and Continuity.*

B.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 Dynamic Rationality 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 Dynamic Rationality. We first show $\text{supp}(p_{\succeq}) \subseteq \bigcup_i \text{supp}(p_{\succeq_i})$. Let $E = \bigcup_i \text{supp}(p_{\succeq_i})$. Clearly, $\succeq_{i,E} = \succeq_i$ for all i . So by Dynamic Rationality $\succeq = \succeq_E$, whence $p_{\succeq} = p_{\succeq_E}$. Thus $\text{supp}(p_{\succeq}) \subseteq E$. Second, we consider any $s \in \bigcap_i \text{supp}(p_{\succeq_i})$ and show $s \in \text{supp}(p_{\succeq})$. Consider the event $G = \{s\}$. The updated profile $(\succeq'_i) := (\succeq_{i,G})$ is coherent, hence in \mathcal{D} . So Dynamic Rationality applies, and implies that $\succeq' = \succeq_G$. In particular, $\succeq_G \in \mathcal{P}$. Thus $s \in \text{supp}(p_{\succeq})$, as otherwise G would be \succeq -null, and \succeq_G would be the all-indifferent relation, which lies outside \mathcal{P} . ■

We next prove that Dynamic Rationality implies this familiar axiom:

Independence of Group Values on Individual Beliefs (IGVIB): 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 *Dynamic Rationality implies IGVIB.*

Proof. Assume Dynamic Rationality. 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 Dynamic Rationality $\succeq = \succeq'_E$ or $\succeq_E = \succeq'$, and so $u_{\succeq} = u_{\succeq'}$ as conditionalisation preserves values.

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 $(\succ_i^1) = (\succ_i)$ and $(\succ_i^5) = (\succ_i')$, and such that (\succ_i^2) , (\succ_i^3) , (\succ_i^4) have values profiles given by $u_{\succ_i^2} = u_{\succ_i^3} = u_{\succ_i^4} = u_{\succ_i}$ ($i \in N$) and belief profiles given by $p_{\succ_i^2}(s) = 1$, $p_{\succ_i^3}(s) = p_{\succ_i^3}(s') = \frac{1}{2}$ and $p_{\succ_i^4}(s') = 1$ ($i \in N$). These profiles belong to \mathcal{D} as \mathcal{D} is closed under belief change. To check for Bayes neighbourhood, note that, for all $i \in N$, $\succ_{i,\{s\}}^1 = \succ_i^2$, $\succ_i^2 = \succ_{i,\{s\}}^3$, $\succ_{i,\{s'\}}^3 = \succ_i^4$, and $\succ_i^4 = \succ_{i,\{s'\}}^5$. ■

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 IGVIB.

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

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

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

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

B.2 Proof of Fact 2

We begin by a simple characterisation of preference convergence:

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

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

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

Conversely, assume $\succ^k \rightarrow \succ$. Then $u_{\succ^k} \rightarrow u_{\succ}$ since for each $x \in X$ we can use the constant act $a \equiv x$ to infer $u_{\succ^k}(x) = \mathbb{E}_{p_{\succ^k}}(u_{\succ^k}(a)) \rightarrow \mathbb{E}_{p_{\succ}}(u_{\succ}(a)) = u_{\succ}(x)$. Now we fix $s \in S$ and show $p_{\succ^k}(s) \rightarrow p_{\succ}(s)$. Pick $x, x' \in X$ such that $u_{\succ}(x) = 1$ 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 $p_{\succeq^k}(s) \rightarrow p_{\succeq}(s)$. ■

Our preference aggregation rule F induces a family of belief aggregation rules $\pi_{\mathbf{u}} : \mathbf{D} \rightarrow \Delta(S)$, where $\mathbf{u} \in \mathcal{U}$, defined as follows. Let $\mathbf{u} \in \mathcal{U}$. \mathbf{D} consists of all coherent belief profiles. For each $\mathbf{p} \in \mathbf{D}$, form the preference profile $(\succeq_i) \in \mathcal{D}$ with values 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 from F the three properties defined in Appendix A:

Lemma 6 (a) *If F is dynamically rational, then each $\pi_{\mathbf{u}}$ ($\mathbf{u} \in \mathcal{U}$) is dynamically rational.*

(b) *If F is minimally Paretian, then each $\pi_{\mathbf{u}}$ ($\mathbf{u} \in \mathcal{U}$) is unanimity-preserving.*

(c) *If F is continuous, then each $\pi_{\mathbf{u}}$ ($\mathbf{u} \in \mathcal{U}$) is continuous.*

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

(a) Let F be dynamically rational. Assume $\mathbf{p} \in \mathbf{D}$ and $E \subseteq S$ with $\mathbf{p}(E) \gg 0$, i.e., with a well-defined updated profile $\mathbf{p}(\cdot|E) \in \mathbf{D}$. Let (\succeq_i) have values 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 Dynamic Rationality, $\succeq' = \succeq_E$. Hence, $p_{\succeq'} = p_{\succeq}(\cdot|E)$, i.e., $\pi_{\mathbf{u}}(\mathbf{p}(\cdot|E)) = \pi_{\mathbf{u}}(\mathbf{p})(\cdot|E)$. Q.e.d.

(b) Assume F is minimally Paretian. 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 values 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.

(c) 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 values 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 for F , and thus by Lemma 5 $p_{\succeq^k} \rightarrow p_{\succeq}$, i.e., $\pi_{\mathbf{u}}(\mathbf{p}^k) \rightarrow \pi_{\mathbf{u}}(\mathbf{p})$. ■

By this lemma, each of the rules $\pi_{\mathbf{u}}$ ($\mathbf{u} \in \mathcal{U}$) satisfies the premises of Theorem 2, proved in Appendix A. So, by this theorem, Fact 2 holds.

B.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 values profile. For each values profile $\mathbf{u} \in \mathcal{U}$, denote the corresponding normalised 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 values 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 normalisation 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 values 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, $\beta_{j,\mathbf{u}^k} \rightarrow \beta_{j,\mathbf{u}}$. ■

B.4 Proof of Fact 4

Suppose F is linear-geometric with values-dependent weights, say w.r.t. weights $(\alpha_{i,\mathbf{u}}, \beta_{i,\mathbf{u}})_{i \in N, \mathbf{u} \in \mathcal{U}}$. I prove the three axioms.

Claim 1: F is dynamically rational.

Assume $(\succeq'_i) \in \mathcal{D}$ arises from $(\succeq_i) \in \mathcal{D}$ by conditionalising all preferences on $E \subseteq S$, i.e., $\succeq'_i = \succeq_{i,E}$ for all i . We must show that $\succeq' = \succeq_E$. Note that $(u_{\succeq'_i}) = (u_{\succeq_i})$ and that $(p_{\succeq'_i})$ arises from (p_{\succeq_i}) by conditionalising all beliefs on E , i.e., $p_{\succeq_i}(E) \neq 0$ and $p_{\succeq'_i} = p_{\succeq_i}(\cdot|E)$ for all i (no $p_{\succeq_i}(E)$ is zero because otherwise some $\succeq'_i = \succeq_{i,E}$ would be the all-indifferent relation, hence would lie outside \mathcal{P} , contradicting that $(\succeq'_i) \in \mathcal{D}$). To show that $\succeq' = \succeq_E$, we need to prove that $u_{\succeq'} = u_{\succeq} (= u_{\succeq_E})$ and that $p_{\succeq'}$ is the conditionalisation of p_{\succeq} on E (i.e., that $p_{\succeq}(E) \neq 0$ and $p_{\succeq'} = p_{\succeq}(\cdot|E)$). First, $u_{\succeq} = u_{\succeq'}$ because $(u_{\succeq'_i}) = (u_{\succeq_i})$, and because $u_{\succeq} = \sum_i \alpha_{i,\mathbf{u}} u_{\succeq_i}$ and $u_{\succeq'} = \sum_i \alpha_{i,\mathbf{u}} u_{\succeq'_i}$, where $\mathbf{u} := (u_{\succeq_i}) = (u_{\succeq'_i})$. Second, $p_{\succeq'}$ is the conditionalisation of p_{\succeq} on E because $p_{\succeq'} = \pi((\succeq_{i,E}))$ and $p_{\succeq} = \pi((\succeq_i))$ where π is a geometric opinion pooling rule (the one with weights $\beta_{i,\mathbf{u}}$, $i \in N$), and because geometric opinion pooling is dynamically rational by Theorem 2, proved in Appendix A. Q.e.d.

Claim 2: F is minimally Paretian.

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

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

Claim 3: F is continuous.

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. ■

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
- Alchourron, C. E., Gärdenfors, P. Makinson, D. (1985) On the Logic of Theory Change: Partial Meet Contraction and Revision Functions, *Journal of Symbolic Logic* 50(2): 510-530
- Baccelli, J. (2019) The problem of state-dependent utility: a reappraisal, *British Journal for the Philosophy of Science*, forth.

- Bacchelli, 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
- Broome, J. (2013) *Rationality Through Reasoning*, Hoboken: Wiley
- 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
- Dietrich, F., List, C. (2016) Probabilistic Opinion Pooling. In: C. Hitchcock & A. Hajek (eds.), *Oxford Handbook of Probability and Philosophy*, Oxford University Press
- Dietrich, F., Staras, A., Sugden, R. (2019) A Broomean model of rationality and reasoning, *Journal of Philosophy* 116: 585–614
- 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
- Joyce, J. (2019) The development of subjective Bayesianism, *Handbook of the History of Logic* 10: 415–475
- 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. (1998) The paradox of the Bayesian expert and state-dependent utility theory, *Journal of Mathematical Economics* 29: 331–361
- 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* 70: 633–663
- 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