DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

STABLE AGGREGATION OF PREFERENCES

Matthias Hild
California Institute of Technology &
Jet Propulsion Laboratory

SOCIAL SCIENCE WORKING PAPER 1112

March 2001

Stable Aggregation of Preferences

Matthias Hild

Abstract

We arrive at new conclusions for social choice theory by considering the process in which we refine decision–theoretic models and account for previously irrelevant parameters of a decision situation (cf. Savage's 'small worlds'). Suppose that, for each individual, we consider a coarse–grained and a fine–grained decision–theoretic model, both of which are consistent with each other in a sense to be defined. We desire any social choice rule to be stable under refinements in the sense that the group choice based on fine–grained individual models and the group choice based on coarse–grained individual models agree for choices among coarse–grained alternatives. For ex ante aggregation, we find that stability is ubiquitous since it follows from independence of irrelevant alternatives. In ex post aggregation, individuals' utilities are pooled separately from their beliefs before the group's choice function is constructed. We find that any 'non–exceptional' rule (e.g., any Pareto optimal rule) for ex post aggregation must be unstable. If the rule is, in addition, independent of irrelevant alternatives, we find an infinite series of reversals of binary group choices. We consider applications to risk management and the theory of consensus formation.

JEL classification numbers: D63, D71, D81

Key words: consensus formation, ex ante aggregation, ex post aggregation, independence of irrelevant alternatives, refinements, risk management, small worlds, social choice theory

Stable Aggregation of Preferences

Matthias Hild*

1 Introduction and Motivation

Stability under refinements

When decision theorists specify a model to describe an individual's preferences in a particular decision situation, they are faced with an infinity of potentially relevant details, or parameters. Fortunately, most of these details do not matter for decision situations in the real world. Nonetheless, the particular set of relevant details varies across different decision situations. Parameters that were irrelevant for some choices may well become relevant for other choices. It is therefore vital to consider the process of refining a given decision–theoretic model by taking previously neglected parameters into account. The *locus classicus* for such considerations is Savage (1954). In the present paper, we shall consider refinements of individual decision–theoretic models in the context of social choice theory. We will produce a new argument for adopting the condition of independence of irrelevant alternatives for ex ante social choice rules. Moreover, we state a new and general impossibility result for the ex post aggregation.

For a brief sketch of our framework, we restrict ourselves to individuals who maximize expected utility. Later we will see that the majority of all presently available decision theories is covered by our general result. For each individual $1, \ldots, I$, let $M_i = \langle P_i, u_i \rangle$ be a decision–theoretic model, where P_i is i's subjective probability measure and u_i is i's subjective utility function. A decision–theoretic model M' is more detailed than model M iff, firstly, M' can describe all consequences of M and, secondly, M' individuates consequences between which the coarser model M is unable to discriminate. As an example, think of the coarsely described consequence of a rise in the Euro/USD exchange rate and the more finely described consequences of a rise of at least 10% and a rise of less than 10%. We view the factorization of the individuals' evaluations of uncertain

*Address for correspondence: California Institute of Technology, Mailcode 228–77, Pasadena, CA 91125 (USA). Email: matthias@hild.org. Website: http://www.hild.org. I am grateful to Philippe Mongin for his intuitive explanation of the Ex Ante Homogeneity Theorem and for extensive help with an early draft. I am much indebted to Isaac Levi and Teddy Seidenfeld for their continued help and interest in this paper. I also thank Richard Jeffrey and Mathias Risse. I wish to thank Christ's College, Cambridge for their generous financial and academic support.

prospects into beliefs and utilities as an iterative, fractal–like process. What is a utility on one level of analysis is a compound of beliefs and utilities on another level of analysis. We say that M' refines M exactly if M' is more detailed than M and the probabilities and expected utilities in M' are consistent with those in M. We will argue that social choice rules should be stable under decision–theoretic refinements in the following sense: If the individual models $M'_1, \ldots, M'_I$ refine $M_1, \ldots, M_I$ respectively, then the group's resulting choice functions C' and C should be consistent with each other.

We consider two alternative modes of aggregating individual models into group choice. The first, ex ante mode aggregates the individuals' expected utilities. We adapt the notion of independence of irrelevant alternatives to our framework that allows variable sets of possible, distinguishable options. We then find that independence of irrelevant alternatives implies stability under refinements. The converse is true under a closure condition on the domain of the social choice rule. This first result provides new support for adopting independence of irrelevant alternatives as a desired property of ex ante social choice rules. The second, ex post mode splits individual expected utilities into probabilities and utilities before aggregation takes place (cf. Hammond, 1981). We find that any 'non–exceptional' rule (e.g., any Pareto optimal rule) for ex post aggregation must be unstable. If the rule is, in addition, independent of irrelevant alternatives, we find an infinite series of reversals of binary group choices.

Hence, we cannot be sure that instabilities will not eventually disappear as we continue to refine the individuals' decision–theoretic models. Hild (2001b) provides corresponding impossibility results for a hybrid form of aggregation that Levi (1990) calls 'robust aggregation'.[1] We follow Savage (1954) and represent acts as functions from states to consequences.[2] Our results make few presuppositions about the decision theory used to describe individual preferences. As far as the group is concerned, we make no assumptions about the existence of group beliefs or group preferences. We merely assume that an ex post social choice rule aggregates the individuals' utilities and that, in binary choices, it chooses absolutely dominant acts in a sense that is weaker than the sure–thing principle or related dominance conditions. These meek assumptions reassure us that group choice reversals are not the artifact of a particular decision theory or a narrow class of aggregation rules but a troubling feature of the ex post mode itself.

Risk management

The burgeoning discipline of risk management provides a concrete example of an ex post social choice rule (cf. Glickman/Gough, 1990). The decision–theoretic approach

[1]Levi proposed robust aggregation in reaction to Seidenfeld's version of the Ex Ante Homogeneity Theorem (cf. below). Robust aggregation forms hypothetical individuals $M_{ij} := \langle P_i, u_j \rangle$ $(1 \leq i, j \leq I)$ and proceeds by ex ante aggregation of these hypothetical individuals.

[2]Our technical report Hild (2001c) discusses modifications of our definitions and theorem in a framework that does not presuppose the separation of states and consequences and which contains no explicit representation of the underlying (causal) structure of acts. This framework allows us to simplify our notation and proofs in return for some stronger postulates.

to risk management structures the process of social decision making into two carefully separated parts and then uses concepts from decision analysis to arrive at a policy recommendation. In the first part, the highly professionalized discipline of probabilistic risk assessment procures probabilities for social decision making (Henley/Kumamoto, 1992). In the second part, a social utility function is constructed. In applications such as health care economics or environmental cost–benefit analysis, a social utility is obtained on the basis of individuals' subjective utilities for the outcomes of a risky policy measure. In these applications, social utility is typically defined as the average of the individuals' utilities (Haddix et al., 1996). Decision–theoretic risk management then proceeds by maximizing the expected social utility relative to the estimated probabilities. In the applications mentioned, the structure of the risk management process fits our definition of an ex post social choice rule. Since we make no assumptions about the existence of group probabilities or the manner in which they are constructed, our results apply even when individuals' probabilities are replaced by some expert's risk assessment.

The motivation for a factorization of ex ante evaluations into probabilities and utilities is fundamentally a philosophical one. According to a common normative argument, risk managers ought to separate questions of fact (probabilities) and questions of value (utilities). Implicit is the view that beliefs enjoy a special status vis-a-vis values that allows us to apply the norms of rational discourse and of scientific method to disagreeing beliefs. The U.S. Nuclear Regulatory Commission (1975) was the first to put this conviction into practice, thereby creating the discipline of probabilistic risk assessment (PRA). Disagreements about the social value of possible consequences, on the other hand, are settled by aggregating individuals' subjective values (at least in the applications that we mentioned).

For this reason, the current practice of risk management is closely bound to the ex post mode of aggregation and is, thus, liable to produce unstable social choices. Pursuing an empirical line of reasoning, we may conclude that the choice of different levels of detail may crucially influence the risk managers' decision. In its current structure, the risk management process therefore bestows potentially significant political power on whoever controls the choice of the degree of detail with which a risk management study is conducted. Pursuing a normative line of reasoning, we argue that the current practice of risk management is deficient. Our positive result that stability is ubiquitous among ex ante social choice rules suggests a concrete and practical alternative to the current decision–theoretic approach to risk management. Briefly, our solution retains PRA's maxim to be rational about probabilities but abandons the maxim that ex ante evaluations should be factorized into beliefs and values. The result is a two–stage process in which individuals first update their ex ante evaluations with the results of PRA and in which we then aggregate the individuals' ex ante evaluations by some ex ante social choice rule, such as voting or auctioning. This approach brings the large literature on the implementation of ex ante social choice rules to bear on risk management decisions (cf. Hild (2001d) for a more extensive discussion and organizational conclusions).

The possibility of consensus

Stability comes close to a conditio sine qua non for any self–contained and complete theory of consensus. If a rule for consensus formation depends on the choice of a level of detail, this choice will itself become the topic of disagreement among individuals. Any complete theory of consensus must therefore be either stable under refinements or explicitly determine a relevant level of detail for any given decision situation. If we do not share the hope that we can produce a theory that explicitly determines a relevant level of detail, it becomes interesting to consider our instability theorem in the context of recent work on ex ante aggregation of Bayesian preferences. Assuming that individuals and the group are expected utility maximizers, Goodman (1988), Seidenfeld/Kadane/Schervish (1989), Broome (1990), Schervish/Seidenfeld/Kadane (1991), and Mongin (1995, 1998) show in various decision–theoretic frameworks and in various degrees of generality that Pareto optimal ex ante aggregation already presupposes a disturbingly high degree of homogeneity of either individual probabilities or utilities. In the case of two individuals ($I = 2$), a high degree of homogeneity means that either the individuals' probabilities are identical or their utilities are identical up to affine transformations.[3]

Ex Ante Homogeneity Theorem (Mongin, 1995). *Assume that individuals and group maximize expected utility. Let $M_1, \ldots, M_I$ the individuals' decision–theoretic models and M_0 the group's decision–theoretic model such that ex ante Pareto optimality and a non–triviality condition are satisfied.*[4] *Then either $p_1, \ldots, p_I$ are linearly dependent or $u_1, \ldots, u_I$ are affinely dependent.*

In reaction to this impossibility theorem, it might at first seem tempting to abandon the ex ante mode and to shift either to the ex post mode or, as Levi (1990) advocates, to the robust mode. The reasons for suggesting a shift to the ex post mode could be as follows: The proof of the Ex Ante Homogeneity Theorem relies on the phenomenon of individuals who have identical preferences for 'different reasons'. More precisely, the identity of two individuals' expected utilities does not imply the identity of either their probabilities or their utilities. One individual might judge a certain consequence of an act unlikely but highly desirable, another might judge the consequence more likely but less desirable, and yet both may agree on their overall evaluation of the act. Since this phenomenon leads to the Ex Ante Homogeneity Theorem, it may seem opportune not to aggregate on the level of preferences but to proceed in the ex post mode on a deeper

[3] As the number I of individuals increases beyond 2, this conclusion becomes less severe because the probabilities and utilities of $I > 2$ individuals can be linearly or affinely dependent while not being identical. Note, however, that linear/affine independence of individual probabilities/utilities is a sufficient but not a necessary condition for the impossibility of Pareto optimal aggregation into Bayesian group preferences. Using the proof strategy of Seidenfeld/Kadane/Schervish (1989), Goodman (1988) discusses a case involving $N = 3$ individuals with linearly/affinely dependent probabilities/utilities and provides a necessary and sufficient condition for the impossibility of Pareto optimal Bayesian aggregation in this case.

[4] The required Pareto condition is: If $f \succeq_i g$ for all $1 \leq i \leq I$ but $f \succ_j g$ for some $1 \leq j \leq I$, then $f \succ_0 g$. The non–triviality condition is: There are consequences c, c' such that $u_i(c) > u_i(c')$ for all $0 \leq i \leq I$.

level where we can account for the individuals' 'reasons' for their preferences, namely their probabilities and utilities (cf. Mongin, 1998). Our instability theorem frustrates this proposal. Instabilities theorem arise, ironically, from the very same phenomenon of 'agreement for different reasons' that we have already identified as the source of the Ex Ante Homogeneity Theorem. Incidentally, Hylland/Zeckhauser (1979) also draw on 'agreement for different reasons' when they show that ex post aggregation may violate ex ante Pareto optimality. The importance of Hylland and Zeckhauser's result is, however, limited since there is no reason why we should continue to require ex ante Pareto optimality once we adopt the ex post mode of aggregation. In short, the combination of the Ex Ante Homogeneity Theorem and our instability theorem presents us with yet another difficulty in social choice theory: On the one hand, ex ante aggregation is troubled by the Ex Ante Homogeneity Theorem. On the other hand, ex ante aggregation seems to be the only available way of satisfying stability under refinements.

What is to be done? I cannot address this question in the present paper, but I am inclined to reconsider the current paradigm of social choice theory. I submit that this paradigm suffers from a lack of concern for the dynamic aspects of group formation. This trend started with Arrow's individualistic outlook under the guise of liberalism. Arrow (1951) writes: *"[...] we will [...] also assume in the present study that individual values are taken as data and are not capable of being altered by the nature of the decision process itself. This, of course, is the standard view in economic theory [...] and also in the classical liberal creed"* (p.7s). An alternative view stresses the dynamic aspects of group formation. On this view, being a member of a group and participating in social choice has itself an impact on the formation of individual preferences. We will study this feedback of aggregation processes onto individual preferences elsewhere. Perhaps we should interpret the troublesome ex ante Pareto condition as a solution concept for the aggregation process rather than a procedural condition. We would then say that an (ideal) group has been formed exactly when Pareto optimal ex ante aggregation has become possible. Since ex ante Pareto optimality implies a high degree of homogeneity, the successful completion of the dynamic aggregation process could, therefore, be characterized by a high degree of homogeneity.

2 Ex Ante Aggregation

Refinements

Let Γ be a non–empty set which we will call a *frame of reference for consequences*. Let Φ be a non–empty set which we will call a *frame of reference for acts*. It is relative to such frames of reference that we will compare the degrees of detail with which different decision–theoretic models describe the same decision situation. A frame of reference can be chosen arbitrarily as long as it is fine–grained enough to capture all the parameters of the most detailed model that we wish to consider. As far as consequences are concerned, we represent a model's degree of detail by some partition $\mathcal{C}$ of Γ, our frame of reference for consequences. The elements of this partition take the place of consequences in a conventional framework. We will eventually also replace the worlds of a conventional model with the elements of a partition $\mathcal{W}$ (Savage's 'small worlds') and explicitly consider the underlying structure of acts. Partitions of worlds will take centrestage in the following section. For now, we keep them backstage and assume that acts induce probabilities (or some other belief measure) for consequences given some suitably defined partition of worlds.

Illustration: Adapting Savage's (1954) example, let the frame of reference consist of the points in the real plane, i.e. pairs $\langle x, y \rangle$ of real numbers. The most fine–grained model relative to this frame of reference represents consequences as points in the real plane. A coarser model may ignore the second parameter y and represent consequences as lines in the real plane parallel to the y–axis.

Illustration: Consider the decision problem of a group of directors who consider building a production plant in Europe. We start with a model that accounts only for the individuals' preferences for building or not building the production plant. A second more detailed model accounts, in addition, for the individuals' beliefs and utilities concerning an upward or downward change in the Euro/USD exchange rate. The second model thus distinguishes the outcome of owning a production plant in Europe in a climate of an increasing exchange rate and the outcome of owning a production plant in Europe in a climate of a non–increasing exchange rate. A third yet more detailed model may account for the individuals' preferences over additional features, such as the magnitude of changes in the exchange rate or the future of the European stock market index.

Since we evaluate actions by their potential consequences, the degree of detail with which we describe consequence also affects how we individuate actions. Choosing, for example, a coarse–grained consequence C with certainty amounts to an uncertain prospects of more fine–grained descriptions of C. More generally, an action induces a certain prospect of coarse–grained consequences and another prospect of fine–grained consequences.

Illustration: We return to our managerial decision problem. In the first model, the agents can choose (at least) between building and not building the production plant. From the viewpoint of the second model, building the production plant amounts to choosing an uncertain prospect depending on the rise or fall of the exchange rate. In addition to the actions available in the first model, the second model allows the agent to make choices that were not available in the first model. For instance, the second model allows the agent to choose a hedging strategy for the risk of a rising exchange rate. In a more detailed model, the number of possible acts therefore increases.

In the next section, we will discuss the representation of acts as functions from states to consequences and formulate the requirements under which a prospect of coarse–grained consequences and a prospect of fine–grained consequences is induced by the same action. For now, we only use the fact that the more fine–grained a model becomes, the more acts it can distinguish. $M = \langle \mathcal{F}, R \rangle$ is a *preference model* if and only if $\mathcal{F} \subseteq \Phi$ and R is a relation on $\mathcal{F}$ (i.e., $R \subseteq \mathcal{F} \times \mathcal{F}$). Any $\mathcal{F} \subseteq \Phi$ is called a *set of possible distinguishable acts*. In this simple framework, the set $\mathcal{F}$ represents the degree of detail with which the model describes the decision situation at hand. The set $\mathcal{F}$ corresponds to Arrow's (1951) set of 'possible alternatives'. While Arrow kept this set fixed, we will study the behaviour of social choice rules under variations of the set of possible distinguishable acts $\mathcal{F}$ (cf. Laslier (2000) for a related approach). A preference model $M' = \langle \mathcal{F}', R' \rangle$ is *more detailed than* a preference model $M = \langle \mathcal{F}, R \rangle$ exactly when $\mathcal{F} \subseteq \mathcal{F}'$. For any binary relation R on $\mathcal{F}$ and any $X \subseteq \mathcal{F}$, let $R|X := R \cap (X \times X)$ be the restriction of R to X. For any preference models $M = \langle \mathcal{F}, R \rangle$ and $M' = \langle \mathcal{F}', R' \rangle$, we say that M' *refines* M if and only if (i) $\mathcal{F} \subseteq \mathcal{F}'$ and (ii) $R = R'|\mathcal{F}$.

For any $\mathcal{F} \subseteq \Phi$, we call V a (binary) *evaluation function* on $\mathcal{F}$ if and only if there is some non–empty set Z such that $V : \mathcal{F} \times \mathcal{F} \to Z$. Commonly, we choose $Z = \mathbb{R}$. We choose to work with the concept of binary evaluation functions because of its generality. Loomes/Sugden's (1982) regret theory provides an example of a genuinely binary evaluation function in which a preference $f \succ g$ is defined by $V(f,g) > V(g,f)$. Moreover, any binary relation R on $\mathcal{F}$ can be represented by a binary decision rule V_R defined as the characteristic function of R (i.e., $V_R : \mathcal{F}^2 \to \{0,1\}$ is defined by $V_R(f,g) := 1$ if $\langle f,g \rangle \in R$ and $V_R(f,g) := 0$ else). Finally, binary evaluation functions allow monadic evaluations as a special case. An evaluation function V on $\mathcal{F}$ is *monadic* if and only if $V(f,g) = V(f,f)$ for all $f,g \in \mathcal{F}$. For a monadic evaluation function V on $\mathcal{F}$, we define $V(f) := V(f,f)$ (for all $f \in \mathcal{F}$). Expected utility maximization is an example of a monadic decision rule in which a preference $f \succ g$ may be defined by $V(f) > V(g)$ or by $V(f) - V(g) > \alpha$ for some threshold $\alpha \in \mathbb{R}^+$ (cf. Fishburn, 1988). $M = \langle \mathcal{F}, V \rangle$ is an *evaluation model* if and only if $\mathcal{F} \subseteq \Phi$ and V is an evaluation function on $\mathcal{F}$. We define $\mathcal{F}_M$ and V_M as the entities such that $M = \langle \mathcal{F}_M, V_M \rangle$. For any evaluation function $V : \mathcal{F} \to Z$ and any $X \subseteq \mathcal{F}$, let $V|X := V \cap (X \times X \times Z)$ be the restriction of R to X. For arbitrary evaluation models $M = \langle \mathcal{F}, V \rangle$ and $M' = \langle \mathcal{F}', V' \rangle$, we say that M' *refines* M if and only if (i) $\mathcal{F} \subseteq \mathcal{F}'$ and (ii) $V = V'|\mathcal{F}$.

Recall that Arrow (1951) held the set $\mathcal{F}$ of possible distinguishable acts fixed, while we are concerned with variations of this set. In our framework, a set $X \subseteq \mathcal{F}$ is a *choice set for* $\mathcal{F}$ and corresponds to what Arrow (1951) calls an 'environment'. For any $\mathcal{F} \subseteq \Phi$, C is a *choice function for* $\mathcal{F}$ if and only if $C : (2^{\mathcal{F}} - \{\varnothing\}) \to (2^{\mathcal{F}} - \{\varnothing\})$ and $C(X) \subseteq X$ for any $X \subseteq \mathcal{F}$. $M = \langle \mathcal{F}, C \rangle$ is a *choice model* if and only if $\mathcal{F} \subseteq \Phi$ and C is a choice function on $\mathcal{F}$. We define $\mathcal{F}_M$ and C_M as the entities such that $M = \langle \mathcal{F}_M, C_M \rangle$. For arbitrary choice models $M = \langle \mathcal{F}, C \rangle$ and $M' = \langle \mathcal{F}', C' \rangle$, we say that M' *refines* M if and only if (i) $\mathcal{F} \subseteq \mathcal{F}'$ and (ii) $C(X) = C'(X)$ for all $X \subseteq \mathcal{F}$.

Stability

Let $I \geq 2$ ($I \in \mathbb{N}$) be a fixed number of individuals. Let $\mathbf{V}$ be the set of all evaluation models. Let $\mathbf{V}(I)$ be the set of all vectors $\langle \mathcal{F}_i, V_i \rangle$ of evaluation models with $\mathcal{F}_i = \mathcal{F}_j$ ($1 \leq i,j \leq I$). For $\langle M_i \rangle \in \mathcal{R}(I)$ or $\langle M_i \rangle \in \mathcal{V}(I)$, let $\mathcal{F}_{\langle M_i \rangle} := \mathcal{F}_{M_1}$. Let $\mathbf{C}$ be the set of all choice models. A social choice rule yields a choice function for the group as a function of individual evaluation models. Avoiding the assumption of an unrestricted domain, we say that S is an *ex ante social choice rule* if and only if there is some non–empty set $\mathcal{V} \subseteq \mathbf{V}(I)$ such that $S : \mathcal{V} \to \mathbf{C}$ and $\mathcal{F}_{S(\langle M_i \rangle)} = \mathcal{F}_{\langle M_i \rangle}$ for any $\langle M_i \rangle \in \mathcal{V}$. As special cases, we mention rules that yield a social choice function for any vector of individual preference orderings (cf. Arrow, 1951) and rules that yield a social choice function for any vector of individual monadic evaluations (cf. Sen, 1970). We now introduce the central notion of stability for social choice rules. Stability under refinements guarantees that social preferences based on refined individual models are compatible with coarser social preferences based on coarser individual models. We say that the vector $\langle M_i' \rangle$ *refines* the vector $\langle M_i \rangle$ if and only if M_i' refines M_i for all $1 \leq i \leq I$.

Definition 2.1 *A social choice rule S is* stable under refinements *if and only if, for all $\langle M_i \rangle$, $\langle M_i' \rangle$ in the domain of S, $S(\langle M_i' \rangle)$ refines $S(\langle M_i \rangle)$ whenever $\langle M_i' \rangle$ refines $\langle M_i \rangle$.*

Illustration 1: Stability is the property that makes opinion polls useful and feasible. If we use a stable social choice rule, opinion polls only need to elicit relatively coarse preferences about the options at stake. It then becomes unnecessary — as it is in practice impossible — to include all conceivable details in a poll's questionnaire.

Illustration 2: A CEO (taking the place of the group) commissions expert reports (taking the place of individuals). What information may a secretary neglect when summarizing the reports without manipulating the CEO's decision? What degree of detail in the expert reports is relevant? If the CEO's aggregation rule is stable, the secretary may choose any level of detail on which to summarize the expert reports. If, however, the CEO's aggregation rule is not stable, it becomes crucial for the outcome of his decision which details the secretary's summary includes.

Illustration 3: A government commissions a risk management study that is divided into a probabilistic risk assessment and an estimation of the effect of policy outcomes on the average citizen's welfare (cf. Section 1).

Is stability a desirable and important property of social choice rules? We offer four reasons for an affirmative answer. These reasons concern the existence of a buffering partition, i.e., a partition of consequences in whose refinements no changes in the group's preferences occur. First, *metaphysical*, reason: There is no reason to believe that an ultimate, maximally fine–grained, partition of reality exists. There is furthermore no reason to believe that a buffering partition exists. Second, *epistemic*, reason: Even if a buffering partition existed, we could never know when it has been reached. Third, *pragmatic*, reason: Even if we knew how to construct a buffering partition, the complexity of this partition would exceed our computational and other capacities. Fourth, *political reason*: Since the choice of the partition to be used in the analysis can significantly influence the recommendation of an unstable social choice rule, there is room for political manipulation through a clever choice of a favourable partition. This problem is mitigated only by the difficulty to foresee what level of detail will yield which social choice. Among the first three reasons, the pragmatic reason makes the least contentious philosophical assumption. Yet we believe that it is strong enough to motivate the desire for stability in view of our illustrations.

Isaac Levi (in private correspondence) suggests that the choice of an appropriate partition could be an ethico–political judgement and, therefore, an additional parameter in the modelling of a social decision problem. We would then add a new factor to the framework of social choice theory, namely the partition that is judged to be relevant. Our illustrations sketch scenarios in which it would be desirable to avoid ethico–political judgments of relevance all together. On the one hand, the addition of this new factor would subvert the initial project of a theory of consensus formation among disagreeing individuals since the choice of a graining would itself become a topic of disagreement. On the other hand, it is not clear how an explicit theory of social welfare could decide which details are relevant for determining the optimal social welfare and thus, implicitly, the appropriate tradeoffs between different individuals' welfare. The appeal of stability derives from the difficulty of choosing a reference partition. It is as coherent to choose the coarsest partition that allows to distinguish all feasible acts as it is to choose the most fine–grained partition available. No choice of a reference partition can claim to be self–evident.[5]

Independence of irrelevant alternatives

Arrow's (1951) formulation of independence of irrelevant alternatives keeps the set $\mathcal{F}$ of possible distinguishable actions fixed and does not make it explicit in the notation. Since we are concerned with changes in the set $\mathcal{F}$, we formulate a version of the independence condition with explicit reference to the set $\mathcal{F}$. Arrow's condition requires that, for fixed $\mathcal{F}$, social choice among the acts in a choice set $X \subseteq \mathcal{F}$, must not depend on individual preferences for acts outside of X.[6] Allowing $\mathcal{F}$ to vary, we require, in addition, that social choice must not depend on the choice of the set $\mathcal{F}$ of distinguishable possible acts. An ex ante social choice rule $S : \mathcal{V} \to \mathbf{C}$ is *independent of irrelevant alternatives (IIA)* if and only if, for all $\langle \mathcal{F}, V_i \rangle, \langle \mathcal{F}', V_i' \rangle \in \mathcal{V}$ and all $X \subseteq \mathcal{F} \cap \mathcal{F}'$, if $V_i|X = V_i'|X$ for all $1 \leq i \leq I$, then $C_{S(\langle M_i \rangle)}(X) = C_{S(\langle M_i' \rangle)}(X)$. Clearly, we obtain Arrow's version of the condition if we keep $\mathcal{F}$ fixed. Note that we do not presuppose the existence of a group preference or any rationality axioms like the weak axiom of revealed preference theory. As Plott (1976) emphasizes, IIA for choice functions is an extremely weak condition that is satisfied by any implementable social choice rules, including the following examples:

[5] Levi (in private correspondence) suggests to define the relevant level of detail as the coarsest common refinement of the individuals' 'basic partitions' (cf. Levi, 1986). Even if we agreed (which we do not) that there were some privileged 'basic partition' for each individual, the choice of the coarsest common refinement is itself a value judgement far from self–evident. Moreover, the individuals may have reasons, including economic incentives, to disagree with this value judgement.

[6] In the more general terms of evaluation functions, Arrow's condition requires that, for any evaluation functions V_i, V_i' on $\mathcal{F}$ $(1 \leq i \leq I)$ and any $X \subseteq \mathcal{F}$, if $V_i|X = V_i'|X$ for all $1 \leq i \leq I$, then $C_{S(\langle V_i \rangle)}(X) = C_{S(\langle V_i' \rangle)}(X)$.

- utilitarian rules with weights λ_i that do not depend on $\mathcal{F}$
- allocative leximin and leximax
- de Borda count
- Nash equilibrium
- auctioning

We now formulate a closure condition on the domain of social choice rules. It requires that if a social choice rule can be applied to some vector $\langle M_i \rangle$ of individual models, then it can also be applied to the restrictions of these models to any set of actions $X \subseteq \mathcal{F}_{\langle M_i \rangle}$. Formally, any set $\mathcal{V} \subseteq \mathbf{V}(I)$ is *closed under restrictions* if and only if, for any $\langle \mathcal{F}, V_i \rangle \in \mathcal{V}$ and any $X \in \mathcal{F}$, we also have $\langle X, V_i | X \rangle \in \mathcal{V}$. This condition is trivially satisfied by social choice rules with an unrestricted domain. An easy theorem now records the close connection between stability under refinements and the more familiar condition of independence of irrelevant alternatives.

Theorem 2.2 *(1) If an ex ante social choice rule is IIA, then it is also stable under refinements.*

(2) If an ex ante social choice rule is stable under refinements and its domain is closed under restrictions, then it is also IIA.

All proofs are collected in the appendix.

As a first consequence of this theorem, we find that stability is a ubiquitous property of ex ante social choice rules, even weaker than IIA. By asking for stability, we are clearly not asking for too much. In the next section where we will find that the situation is very different for ex post social choice rules. All non–exceptional ex post social choice rules, whether they are IIA or not, are instable under refinements. Instabilities can therefore be counted as a defect of the ex post mode that is easily avoided by the ex ante mode. As a second consequence of the theorem, we obtain additional reasons why we should require social choice rules to be IIA. These are the same three reasons that we have offered for requiring social choice rules to be stable under refinements. If the domain of social choice rules is closed under restrictions, the theorem shows that we have to require IIA if we desire stability under refinements.

3 Ex Post Aggregation

Numerical example

We start with an illustration of instability theorem for individuals who maximize expected utility. We will see that this restriction, along with other simplifications, is not essential and that the result generalizes to most known decision theories and to almost any conceivable ex post aggregation rules. We impose no restrictions on group choices other than an extremely weak dominance condition (absolute dominance) for binary choices. On the other hand, we work with a very stringent notion of refinements of individual decision–theoretic models. Despite this strong constraint on individuals' models, we will discover a severe ex post instability in which the group oscillates ad infinitum between absolute dominance of f over g and absolute dominance of g over f.[7]

We follow Savage (1954) and represent acts as functions from states to consequences.[8] We will eventually replace the worlds of a conventional model with the elements of a partition $\mathcal{W}$ (Savage's 'small worlds') and replace conventional consequences with the elements of a partition $\mathcal{C}$ (our 'small consequences'). Expected utility maximizers evaluate these acts by the probabilities of states in $\mathcal{W}$ and the utilities of consequences in $\mathcal{C}$. For simplicity's sake, our verbal exposition of the numerical example continues to consider only partitions of consequences and assumes that probabilities for consequences are induced by some suitably defined act and some suitably defined partition of worlds (cf. the appendix for details). As before, let Γ be a non–empty set, called a *frame of reference for consequences*. We assume that Γ is at least countably infinite. We say that $\mathcal{C}$ is a (finite) Γ*–partition* if and only if $\mathcal{C}$ is a finite collection of non–empty and mutually disjoint sets the union of which is Γ. A Γ–partition $\mathcal{C}$ represents the degree of detail with which the model describes the decision situation at hand. We say that p is a *probability on* $\mathcal{C}$ exactly when $p : \mathcal{C} \to [0, 1]$ and $\sum_{C \in \mathcal{C}} p(c) = 1$. Acts now induce probabilities for the consequences in $\mathcal{C}$. We again simplify our discussion by considering only two acts f and g that induce probabilities p_f and p_g on $\mathcal{C}$. A *(real–valued one–dimensional) utility* on $\mathcal{C}$ is a function $u : \mathcal{C} \to \mathbb{R}$. An expected utility maximizer's model is therefore characterized by the quadruple $M = \langle \mathcal{C}, p_f, p_g, u \rangle$. Define $U_M(f) := \sum_{C \in \mathcal{C}} u(C) \cdot p_f(C)$ as the expected utility of f in model M. Define $U_M(g)$ analogously.

The instability of group choices under ex post aggregation is a fairly obvious phenomenon if an individual's detailed model disagrees with the individual's coarse model on probabilities for coarse events. Savage (1954) points out that the agreement of preferences on coarse acts does not in itself guarantee the agreement of probabilities on coarse events ('small–world problem'). Such probability agreement imposes an additional constraint which we will include in our definition of a refinement of an individual's model.[9] A partition $\mathcal{C}'$ *details* a partition $\mathcal{C}$ exactly when each partition cell of $\mathcal{C}$ is itself partitioned by some more fine–grained partition cells in $\mathcal{C}'$. When $\mathcal{C}'$ details $\mathcal{C}$, when p is a probability on $\mathcal{C}$ and when p' is a probability on $\mathcal{C}'$, we say that p' *refines* p if and only

[7] Other authors have provides some examples of a single instability under the assumption that the group maximizes expected utility. Leeds (1990) provides what is effectively such an example in an attack on Levi's decision theory with indeterminate probabilities. Seidenfeld (1993) uses an example of this sort (involving state–dependent utilities) in a critique of Levi's robust aggregation. A numerical example with an infinity of group preference reversals is reported in Hild/Jeffrey/Risse (2001).

[8] Cf. Hild 2001c for a version of the theorem that does not presuppose a separation of states and consequences.

[9] We thus presuppose that the elicitation of probabilities does not draw on preferences alone but also uses some additional factor, such as external randomization by roulette–lotteries or a primitive qualitative probability relation (cf. Schervish/Seidenfeld/Kadane, 1990).

if p and p' agree on all coarse–grained consequences, i.e., $p(C) = p'(C)$ for all $C \in \mathcal{C}$. In addition to the invariance of probabilities on coarse–grained consequences, we will require that refinements of the individuals' models leave not only the individuals' preferences but also their expected utilities unchanged. If it were not for our interest in the instability theorem, we might wish to speak of an individual refinement already when individual preferences are left unchanged. With our stringent definition of individual refinements, however, we will obtain a strong instability result because our proofs will construct a group choice reversal even for individual refinements in which the evaluations of acts remain unchanged. To summarize, we say that an EU–model $M' = \langle \mathcal{C}', p'_f, p'_g, u' \rangle$ *refines* an EU–model $M = \langle \mathcal{C}', p'_f, p'_g, u' \rangle$ if and only if $\mathcal{C}'$ details $\mathcal{C}$, p'_f refines p_f, p'_g refines p_g, $U_M(f) = U_{M'}(f)$ and $U_M(g) = U_{M'}(g)$.

For concreteness, we consider two individuals whose utilities are aggregated by taking their average $u_0 = \frac{u_1+u_2}{2}$. We construct a sequence of individual refinements with two acts f and g such that, for any model in the sequence, both individuals' expected utility for f is 1 while their expected utility for g is 0. We here provide an intuitive sketch of the construction. The formal definition of this example is provided in appendix A. Starting from the Γ–partition $\{C, D\}$, we construct an infinite sequence of increasingly detailed Γ–partitions. In each of these partitions, we continue to partition C into an increasing number of sets each of which is, in turn, subdivided in the subsequent partition. The construction and indexing of the partition elements is illustrated in Figure 1.

We define g as an act that yields the consequence D with necessity, or in all possible states of the world. We define f as an act that, in all possible states of the world, yields some consequence from the C partition. The individuals differ in their probability judgments for the consequences of f. For the first four partitions, the upper part of Figure 1 shows the individual probabilities for the consequences if f is performed (i.e., $p^n_{1,f}$ and $p^n_{2,f}$ for $n = 0, 1, 2, 3$). The middle part of the same figure shows the individual utilities for consequences in the first four partitions. In the second partition, for example, the consequence C_1 has utility 3 for individual 1 and utility -5 for individual 2. The consequence D always has utility 0 for both individuals. The essence of this example lies in the individuals' strong disagreements about probabilities and utilities. With these values fixed, we find that the sequence of individual models $\langle M^n_i \rangle_{n\in\mathbb{N}}$ $(i = 1, 2)$ is indeed a sequence of refinements ($U^n_1(f) = U^n_2(f) = 1$ and $U^n_1(g) = U^n_2(g) = 0$, for all $n \in \mathbb{N}$).**[10]** We also find severe violations of stability.

Applying averaging to the individuals' utilities, we find that, in even numbered models, the group values all possible consequences of f at $+1$, while, in odd numbered models, the group values all possible consequences of f at -1 (i.e., $u^n_0(C_{k_1,\dots,k_n}) = (-1)^n$ for all $n \in \mathbb{N}$). At the same time, the only possible consequence of g is always valued at 0

[10] We remark that the elements in this sequence of individual models not only refine each other in the sense that we have define. They are moreover *local refinements* of each other in the sense that M^{n+1}_i refines M^n_i *and* $u^n_i(X) = \sum_{C\in\mathcal{C}^{n+1}} u^{n+1}_i(C) \cdot P^{n+1}_{1,f}(C|X)$ for all $X \in \mathcal{C}^n$ with $P^{n+1}_{1,f}(X) > 0$ and all $f \in \mathcal{F}$. In a model M^{n+1} that locally refines M^n, the consequences $X \in \mathcal{C}^n$ of the coarse model are treated as uncertain prospects with the distribution $P^{n+1}_f(.|X)$ and their expected utility equals the (risk–free) utility of X in the coarse model.

n — individual probabilities p_1^n, p_2^n for consequences of f

n		
0	C	D
	$1, 1$	$0, 0$

n			
1	C_1	C_2	D
	$\frac{3}{4}, \frac{1}{4}$	$\frac{1}{4}, \frac{3}{4}$	$0, 0$

n					
2	C_{11}	C_{12}	C_{21}	C_{22}	D
	$\frac{9}{16}, \frac{1}{16}$	$\frac{3}{16}, \frac{3}{16}$	$\frac{3}{16}, \frac{3}{16}$	$\frac{1}{16}, \frac{9}{16}$	$0, 0$

n									
3	C_{111}	C_{112}	C_{121}	C_{122}	C_{211}	C_{212}	C_{221}	C_{222}	D
	$\frac{27}{64}, \frac{1}{64}$	$\frac{9}{64}, \frac{3}{64}$	$\frac{9}{64}, \frac{3}{64}$	$\frac{3}{64}, \frac{9}{64}$	$\frac{9}{64}, \frac{3}{64}$	$\frac{3}{64}, \frac{9}{64}$	$\frac{3}{64}, \frac{9}{64}$	$\frac{1}{64}, \frac{27}{64}$	$0, 0$

n — individual utilities u_1^n, u_2^n

n		
0	C	D
	$1, 1$	$0, 0$

n			
1	C_1	C_2	D
	$3, -5$	$-5, 3$	$0, 0$

n					
2	C_{11}	C_{12}	C_{21}	C_{22}	D
	$1, 1$	$9, -7$	$-7, 9$	$1, 1$	$0, 0$

n									
3	C_{111}	C_{112}	C_{121}	C_{122}	C_{211}	C_{212}	C_{221}	C_{222}	D
	$3, -5$	$-5, 3$	$11, -13$	$3, -5$	$-5, 3$	$-13, 11$	$3, -5$	$-5, 3$	$0, 0$

n — group utilities u_0^n

n		
0	C	D
	1	0

n			
1	C_1	C_2	D
	-1	-1	0

n					
2	C_{11}	C_{12}	C_{21}	C_{22}	D
	1	1	1	1	0

n									
3	C_{111}	C_{112}	C_{121}	C_{122}	C_{211}	C_{212}	C_{221}	C_{222}	D
	-1	-1	-1	-1	-1	-1	-1	-1	0

Figure 1: Individual probabilities, individual utilities and group utilities for $n = 0, 1, 2, 3$.

(i.e., $u_0^n(D) = 0$ for all $n \in \mathbb{N}$). For the first four partitions, the lower part of Figure 1 tabulates these group utilities. We express the relationship between f and g by saying that f absolutely dominates g according to the group's utilities in any even numbered model while g absolutely dominates f according to the group's utilities in any odd numbered model. It therefore matters little what decision rule the group uses as long as the group strongly prefers strongly dominating acts. To mention but one example of a group decision rule, the group could aggregate the individuals' probabilities into a group probability and then maximize expected utility relative to the group probability and group utility. Notice, however, that our result makes no assumptions about the existence or properties of group beliefs. Another strength of our numerical example is the repetition of this strong instability ad infinitum as we refine individual models to higher and higher degrees.

Fine–graining

We turn to Savage's representation of acts and explicitly accommodate it in our notion of a refinement. Let Ω be a non–empty set, called a *frame of reference for states* and let Γ be a non–empty set, called a *frame of reference for consequences.* We assume that both Ω and Γ are at least countably infinite. Throughout, we will only construct models with a finite number of states $W \in \mathcal{W}$ and consequences $C \in \mathcal{C}$.[11] With this presupposition, we say that $\mathcal{W}$ is a (finite) Ω*–partition* if and only if $\mathcal{W}$ is a finite collection of non–empty and mutually disjoint sets the union of which is Ω. We use the analogous definition of a Γ–partition. We say that $\langle \mathcal{W}, \mathcal{C} \rangle$ is a *graining* if and only if $\mathcal{W}$ is an Ω–partition and $\mathcal{C}$ is a Γ–partition. Let $[\mathcal{W}] := \{\bigcup X | X \subseteq \mathcal{W}\}$ be the set of all events expressible in $\mathcal{W}$. Analogously, let $[\mathcal{C}] := \{\bigcup X | X \subseteq \mathcal{C}\}$. A graining $\langle \mathcal{W}', \mathcal{C}' \rangle$ *details* a graining $\langle \mathcal{W}, \mathcal{C} \rangle$ exactly when $\mathcal{W} \subseteq [\mathcal{W}']$ and $\mathcal{C} \subseteq [\mathcal{C}']$. For any Γ–partition $\mathcal{C}$, let $\mathcal{C}(c)$ be the partition cell of $\mathcal{C}$ that contains $c \in \Gamma$. We thus take the liberty to identify a partition $\mathcal{C}$ with a particular function from Γ onto $\mathcal{C}$.

The objects of our primary interest are $\langle \mathcal{W}, \mathcal{C} \rangle$*–grained acts* represented as functions $F : \mathcal{W} \to \mathcal{C}$ that map $\mathcal{W}$–grained states into $\mathcal{C}$–grained consequences. We then want to consider more finely grained representations of these acts. We already pointed out in Section 2 that the number of possible $\langle \mathcal{W}, \mathcal{C} \rangle$–grained acts increases as the graining $\langle \mathcal{W}, \mathcal{C} \rangle$ is replaced by a more detailed graining $\langle \mathcal{W}', \mathcal{C}' \rangle$ (i.e., $|\mathcal{C}^{\mathcal{W}}| \leq |(\mathcal{C}')^{(\mathcal{W}')}|$). What is more, that there are generally several alternative ways of adding details to the coarse description of an act. In other words, if a graining $\langle \mathcal{W}', \mathcal{C}' \rangle$ details a graining $\langle \mathcal{W}, \mathcal{C} \rangle$, there

[11] From the axiomatic viewpoint of Savage's approach, our restriction to models with a finite number of small worlds is not without problems (cf. Savage's postulate (P6)). Note, however, that all our definitions and proofs go through for partitions $\mathcal{W}$ with infinitely many partition cells and only finitely additive probability measures on $\mathcal{W}$. Infinite partitions $\mathcal{W}$ lead to complications, however, once we consider σ–additive probability measures defined on some σ–algebra. We then have to ensure the measurability of the acts which we construct. Although this can be done, restrictions of the sort we would lead to a loss of generality in other places. Finally, we can allow $\mathcal{C}$ to contain infinitely many partition cells as long as we are prepared to use the axiom of choice.

are several $\langle \mathcal{W}', \mathcal{C}' \rangle$–grained acts $F' : \mathcal{W}' \to \mathcal{C}'$ that we could consider as a more detailed description of a coarse–grained act $F : \mathcal{W} \to \mathcal{C}$. Thus, we need a means of identifying those fine–grained acts that count as refined descriptions of a given coarse–grained act. To achieve this end, we define a *reference act* as a function $f : \Omega \to \Gamma$ from the frame of reference for states to the frame of reference for consequences and then define the family of $\langle \mathcal{W}, \mathcal{C} \rangle$–grained acts induced by f.[12] Let $\Phi := \Gamma^{\Omega}$ be the frame of reference for acts. For any graining $\langle \mathcal{W}, \mathcal{C} \rangle$, we say that an act $f : \Omega \to \Gamma$ is *compatible with* $\langle \mathcal{W}, \mathcal{C} \rangle$ if and only if $\mathcal{C}(f(\omega)) = \mathcal{C}(f(\omega'))$ for all $W \in \mathcal{W}$ and all $\omega, \omega' \in W$. Hence, points $\omega, \omega' \in \Omega$ in the same partition cell of $\mathcal{W}$ must lead to points $c = f(\omega)$, $c' = f(\omega')$ in the same partition cell of $\mathcal{C}$. Let $\Phi_{\mathcal{W},\mathcal{C}}$ be the set of acts that are compatible with $\langle \mathcal{W}, \mathcal{C} \rangle$. For any act $f \in \Phi_{\mathcal{W},\mathcal{C}}$, we can then define *the* $\langle \mathcal{W}, \mathcal{C} \rangle$*–graining of* f as the function $F : \mathcal{W} \to \mathcal{C}$ such that $F(W) := \mathcal{C}(f(\omega))$ for any $\omega \in W$. A decision–theoretic model with a graining $\langle \mathcal{W}, \mathcal{C} \rangle$ will contain only acts that are compatible with $\langle \mathcal{W}, \mathcal{C} \rangle$. Moreover, we repeat that our primary interest is not in reference acts but in the $\langle \mathcal{W}, \mathcal{C} \rangle$–graining of reference acts. The purpose of a reference act $f : \Omega \to \Gamma$ is to perform as an identifier of those fine–grained acts that count as more detailed descriptions of a coarse–grained act. We therefore require the choice set of a decision–theoretic model to contain at most one reference act $f : \Omega \to \Gamma$ for each possible $\langle \mathcal{W}, \mathcal{C} \rangle$–grained act $F : \mathcal{W} \to \mathcal{C}$. For any graining $\langle \mathcal{W}, \mathcal{C} \rangle$, we say that a set $\mathcal{F} \subseteq \Phi_{\mathcal{W},\mathcal{C}}$ is *unambiguous w.r.t.* $\langle \mathcal{W}, \mathcal{C} \rangle$ if and only if no two reference acts in $\mathcal{F}$ have the same $\langle \mathcal{W}, \mathcal{C} \rangle$–graining. Using these terms, we will require that the choice set in a decision–theoretic model is unambiguous.

Generalized decision–theory

We now remove the example's narrow assumptions about the types of decision models used to describe the individuals. The discussion in this subsection will, therefore, be considerably more abstract. The reader may wish to consult the applications in the following section in order to appreciate the purpose of studying this generalized framework. As far as decision theory is concerned, we will mainly assume that individual utilities depend only on consequences and that consequences enter the evaluation of acts only through their utilities. We allow utilities to take values on some arbitrary scale and thus subsume real–valued one–dimensional utilities, real–valued multi–dimensional utilities and any ordinal preference relation over consequences including relations that are not orderings. We also make a generalized notion of a decision rule available that subsumes decision rules based on monadic evaluations (e.g., expected utility maximization), decision rules based on binary evaluations (e.g., regret theory) and any ordinal decision rule under which the relative ranking of any two acts depends only on those acts and on no other act, i.e., any ordinal decision rule that is independent of irrelevant alternatives (e.g., decision–theoretic leximin or leximax). Thus, the only types of decision theories not

[12] Savage takes 'small–world acts' to be functions from $\mathcal{W}$ to the set Γ^{Ω}, i.e. function from small worlds to what are acts in the basic framework. Given one model with small worlds and another with bigger worlds, it is therefore always clear which bigger–world act is detailed by a smaller–world act. I differ from Savage by introducing the notion of 'small consequences' and by considering sequences of increasingly fine–grained descriptions of an act.

covered by our framework are decision theories in which the utilities of consequences do not depend on consequences alone (e.g., Becker/Sarin's (1987) lottery–dependent utility theory where the utility of consequences depends on the gamble in which they occur)[13] and decision theories in which the relative ranking of two acts may depend on other acts in the choice set.

We say that $\mathbf{p}$ is a *belief type* if and only if $\mathbf{p}$ maps any Ω–partition $\mathcal{W}$ into a (possibly empty) set $\mathbf{p}(\mathcal{W})$. For any Ω–partition, we call any $p \in \mathbf{p}(\mathcal{W})$ a $\mathbf{p}$*–belief measure on* $\mathcal{W}$. Examples of a belief type are set of all real valued set functions $p : [\mathcal{W}] \to \mathbb{R}$, or the set of all probabilities on $\mathcal{W}$. Another example of a belief measure is any set $K \subseteq \mathcal{W}$ that dichotomizes the states in $\mathcal{W}$ into those considered subjectively possible and those considered subjectively impossible. An additional non–triviality constraint on the belief types that occur in the individual's decision–theoretic models will be introduced below (disagreement condition). All three examples satisfy this constraint.

$\mathbf{u}$ is a *utility type* if and only if there is some (non–empty) set Z such that $u \in \mathbf{u}$ iff there is some Γ–partition $\mathcal{C}$ with $u : \mathcal{C} \to Z$. $\mathbf{u}(\mathcal{C})$ is the set of all functions in $\mathbf{u}$ with domain $\mathcal{C}$; we call any $u \in \mathbf{u}(\mathcal{C})$ a $\mathbf{u}$*–utility on* $\mathcal{C}$. A *scale* is a pair $\langle Z, \trianglerighteq \rangle$ where Z is some non–empty set and $\trianglerighteq$ is a binary relation on Z with $z_1, z_2 \in Z$ such that $\langle z_1, z_2 \rangle \in \trianglerighteq$ but $\langle z_2, z_1 \rangle \notin \trianglerighteq$ (in which case we write '$z_1 \triangleright z_2$'). We say that $\trianglerighteq$ is a *scale relation for* the utility type $\mathbf{u}$ if and only if $\langle Z, \trianglerighteq \rangle$ is a scale and Z is the union of the domain of all functions in $\mathbf{u}$. We earlier considered the special class of real–valued one–dimensional utilities $u : \mathcal{C} \to \mathbb{R}$ (with the canonical scale relation $\geq$ on $\mathbb{R}$). Our general definition also admits real–valued multi–dimensional utilities $u : \mathcal{C} \to \mathbb{R}^L$ (for some $L \in \mathbb{N}^+$) and thus allows us to accommodate models like that of Machina (1982). For a multi-dimensional utility $u : \mathcal{C} \to \mathbb{R}^L$, we define the following scale relation on $\mathbb{R}^L$: For $\langle a_l \rangle, \langle b_l \rangle \in \mathbb{R}^L$, let $\langle a_l \rangle \trianglerighteq \langle b_l \rangle$:iff either $a_l = b_l$ for all $1 \leq l \leq L$, or $a_l > b_l$ for all $1 \leq l \leq L$. Let $\mathbf{r}$ be the set of all r such that there is some Γ–partition $\mathcal{C}$ on which r is a binary relation. Let $\mathbf{r}(\mathcal{C})$ be the set of all binary relations on $\mathcal{C}$. For any $u \in \mathbf{u}(\mathcal{C})$ and $r \in \mathbf{r}(\mathcal{C})$, we say that u $\trianglerighteq$*–represents* r if and only if $C\ r\ D$ iff $u(C) \trianglerighteq u(D)$ for all $C, D \in \mathcal{C}$. There exists a utility type $\mathbf{u}$ with a scale relation $\trianglerighteq$ such that for any Γ–partition $\mathcal{C}$, (1) any $r \in \mathbf{r}(\mathcal{C})$ is $\trianglerighteq$–represented by some $u \in \mathbf{u}(\mathcal{C})$ and (2) any $u \in \mathbf{u}(\mathcal{C})$ represents some $r \in \mathbf{r}(\mathcal{C})$. We can thus identify $\mathbf{r}$ with a particular utility type.

Our concept of a decision rule requires that consequences enter into the evaluation of an act only via their utilities (property (b)). We make no assumptions about the scale on which decision rules evaluate acts. For any belief type $\mathbf{p}$ and any utility type $\mathbf{u}$, G is a (binary) *decision rule for* $\langle \mathbf{p}, \mathbf{u} \rangle$ if and only if (a) G is a function such that, for any graining $\langle \mathcal{W}, \mathcal{C} \rangle$, any $p \in \mathbf{p}(\mathcal{W})$, any $u \in \mathbf{u}(\mathcal{C})$ and any $f, g \in \Phi_{\mathcal{W},\mathcal{C}}$, the quadruple $\langle p, u, f, g \rangle$ is in the domain of G, and (b) for any grainings $\langle \mathcal{W}, \mathcal{C} \rangle$ and $\langle \mathcal{W}, \mathcal{C}' \rangle$ with $\mathcal{C} \subseteq [\mathcal{C}']$, any $p \in \mathbf{p}(\mathcal{W})$, any $u \in \mathbf{u}(\mathcal{C})$, any $u' \in \mathbf{u}(\mathcal{C}')$ and any $f, g \in \Phi_{\mathcal{W},\mathcal{C}'}$:[14] If $u(\mathcal{C}(f(.))) = u'(\mathcal{C}'(f(.)))$ and $u(\mathcal{C}(g(.))) = u'(\mathcal{C}'(g(.)))$, then $G(p, u, f, g) = G(p, u', f, g)$. We say that a decision rule G is *monadic* if and only if $G(p, u, f, g) = G(p, u, f, h)$ for any $\langle p, u, f, g \rangle, \langle p, u, f, h \rangle$

[13]We can, however, allow state–dependent utilities; cf. Hild (2001c).

[14]N.b.: If $\mathcal{C} \subseteq [\mathcal{C}']$, then $\Phi_{\mathcal{W},\mathcal{C}'} \subseteq \Phi_{\mathcal{W},\mathcal{C}}$.

in the domain of G. For monadic decision rules, we define $G(p, u, f) := G(p, u, f, f)$ for any $\langle p, u, f, f\rangle$ in the domain of G. In the following examples, we keep p, u fixed and write $V(f, g) := G(p, u, f, g)$ and $V(f) := G(p, u, f, f)$. Expected utility maximization is an example of a monadic decision rule in which a preference $f \succ g$ may be defined by $V(f) > V(g)$ or by $V(f) - V(g) > \alpha$ for some threshold $\alpha \in \mathbb{R}^+$ (cf. Fishburn, 1988). Loomes/Sugden's (1982) regret theory provides an example of a genuinely binary decision rule in which a preference $f \succ g$ is defined by $V(f, g) > V(g, f)$. Finally, any ordinal decision rules, such as leximin or leximax, in which the relative ranking R of any two acts depends only on those two acts can also be represented by a binary decision rule $V_R(f, g)$ defined as the characteristic function of the rule's ordinal preference R over f, g (i.e., $V_R : \mathcal{F}^2 \to \{0, 1\}$ is defined by $V_R(f, g) := 1$ if $\langle f, g\rangle \in R$ and $V_R(f, g) := 0$ else).

We define a set of consequences to be null relative to a belief measure and a decision rule exactly when the utility assignments to the consequences in the set have no influence on the evaluation of acts. For any Ω–partition $\mathcal{W}$, any belief type $\mathbf{p}$, any utility type $\mathbf{u}$, any binary decision rule G for $\langle \mathbf{p}, \mathbf{u}\rangle$, any $p \in \mathbf{p}(\mathcal{W})$ and any $A \in [\mathcal{W}]$, we say that A is *p, G–null* if and only if for any Γ–partition $\mathcal{C}$, any $u, u' \in \mathbf{u}(\mathcal{C})$ and any $f, f', g, g' \in \Phi_{\mathcal{W},\mathcal{C}}$: If $u(\mathcal{C}(f(\omega))) = u'(\mathcal{C}(f'(\omega)))$ and $u(\mathcal{C}(g(\omega))) = u'(\mathcal{C}(g'(\omega)))$ for all $\omega \in -A$, then $G(p, u, f, g) = G(p, u, f', g')$. Any A is *p, G–one* if and only if $-A$ is p, G–null.

Finally, $M = \langle \mathcal{W}, \mathbf{p}, p, \mathcal{C}, \mathbf{u}, \trianglerighteq, u, \mathcal{F}, G\rangle$ is a (generalized) *decision–theoretic model* if and only if $\mathcal{W}$ is an Ω–partition, $\mathbf{p}$ is a belief type, $p \in \mathbf{p}(\mathcal{W})$, $\mathcal{C}$ is a Γ–partition, $\mathbf{u}$ is a utility type, $\trianglerighteq$ is a scale relation for $\mathbf{u}$, $u \in \mathbf{u}(\mathcal{C})$, $\mathcal{F} \subseteq \Phi_{\mathcal{W},\mathcal{C}}$ is unambiguous w.r.t. $\langle \mathcal{W}, \mathcal{C}\rangle$, and G is a decision rule for $\langle \mathbf{p}, \mathbf{u}\rangle$. We define $\mathcal{W}_M$, $\mathbf{p}_M$, p_M, $\mathcal{C}_M$, $\mathbf{u}_M$, $\trianglerighteq_M$, u_M, $\mathcal{F}_M$, and G_M to be the entities such that $M = \langle \mathcal{W}_M, \mathbf{p}_M, p_M, \mathcal{C}_M, \mathbf{u}_M, \trianglerighteq_M, u_M, \mathcal{F}_M, G_M\rangle$. The *binary evaluation function associated with* M is the function $V_M : \mathcal{F}_M^2 \to \text{range}(G)$ such that $V_M(f, g) := G_M(p_M, u_M, f, g)$ for all $f, g \in \mathcal{F}_M$. If G_M is monadic, then the *monadic evaluation function associated with* M is the function $V_M : \mathcal{F}_M \to \text{range}(G)$ such that $V_M(f) := V_M(f, f)$ for all $f \in \mathcal{F}_M$.

Let $\mathbf{G}(I)$ be the set of all vectors $\langle M_i\rangle$ of decision–theoretic models such that (1) $\mathcal{W}_{M_i} = \mathcal{W}_{M_j}$, $\mathcal{C}_{M_i} = \mathcal{C}_{M_j}$ and $\mathcal{F}_{M_i} = \mathcal{F}_{M_j}$ (for all $1 \le i, j \le I$) and (2) there exists some Ω–partition $\mathcal{W} = \{W_1, \ldots, W_I\}$ and some $\langle p_i\rangle \in \prod_i \mathbf{p}_{M_i}(\mathcal{W})$ such that, for each $1 \le i \le I$, W_i is p_{M_i}, G_{M_i}–one. For any $\langle M_i\rangle \in \mathbf{G}(I)$, let $\mathcal{W}_{\langle M_i\rangle} := \mathcal{W}_{M_1}$, $\mathcal{C}_{\langle M_i\rangle} := \mathcal{C}_{M_1}$, and $\mathcal{F}_{\langle M_i\rangle} := \mathcal{F}_{M_1}$. Condition (1) allows the models in a vector $\langle M_i\rangle \in \mathbf{G}(I)$ to have different belief types, utility types, utility scales and decision rules. To simplify our notation, we henceforth hold each individual's belief type, utility type, utility scale and decision rule fixed and assume that, for any $\langle M_i\rangle \in \mathbf{G}(I)$, we have $\mathbf{p}_i = \mathbf{p}_{M_i}$, $\mathbf{u}_i = \mathbf{u}_{M_i}$, $\trianglerighteq_i = \trianglerighteq_{M_i}$ and $G_i = G_{M_i}$ (for each $1 \le i \le I$). In what follows, we therefore drop any reference to these components when we specify a vector of models in $\mathbf{G}(I)$. Condition (2) imposes our only constraint on the individuals' belief types. We refer to condition (2) as *the disagreement condition.*

Ex post social choice rules

An ex post social choice rule yields a choice model $M = \langle \mathcal{F}, C\rangle$ for the group as a function of the individuals' decision–theoretic models. We make no assumptions about the existence of a group preference. Our definition merely requires ex post social choice rules to yield choice functions over the same set of acts that is evaluated by the individuals (clause 2), to aggregate individual utilities in a way that is not contaminated by beliefs or evaluations of acts and, in binary choices, to choose absolutely dominant acts in a sense that is much weaker than the sure–thing principle or related dominance conditions (clause 3). Under Savage's representation of acts, nothing needs to be said about the existence of group beliefs or the aggregation of individual beliefs. We call s a *utility aggregation rule* if and only if, for any Γ–partition $\mathcal{C}$, s maps any vector $\langle u_i\rangle \in \prod_i \mathbf{u}_i(\mathcal{C})$ to a choice function c on $\mathcal{C}$ (i.e., a function $c : (2^{\mathcal{C}} - \{\varnothing\}) \to (2^{\mathcal{C}} - \{\varnothing\})$ with $c(X) \subseteq X$ for any $X \subseteq \mathcal{C}$). We emphasize the generality of this concept. It subsumes the situation of our numerical example where real–valued one–dimensional individual utilities were aggregated into a real–valued one–dimensional group utility. Trivially, any group utility and any acyclical group preference generates a group choice function over consequences. Our concept of a utility aggregation rule avoids any rationality assumptions like the weak axiom of revealed preference. An act f absolutely dominates an act g relative to the group's choice function on consequences exactly when any $\mathcal{C}$–consequence of f is preferred to any $\mathcal{C}$–consequence of g in a binary choice. For any graining $\langle \mathcal{W}, \mathcal{C}\rangle$, any choice function c on $\mathcal{C}$ and any $f, g \in \Phi_{\mathcal{W},\mathcal{C}}$, we say that f *absolutely dominates* g w.r.t. c if and only if $c(\{\mathcal{C}(f(\omega)), \mathcal{C}(g(\omega'))\}) = \{\mathcal{C}(f(\omega))\}$ for any $\omega, \omega' \in \Omega$. The two acts in the numerical example display this property (cf. Figure 1).

Definition 3.1 *S is an* ex post *social choice rule if and only if (1) there is some non–empty set $\mathcal{G} \subseteq \mathbf{G}(I)$ such that $S : \mathcal{G} \to \mathbf{C}$, (2) for all $\langle M_i\rangle \in \mathcal{G}$, we have $\mathcal{F}_{S(\langle M_i\rangle)} = \mathcal{F}_{\langle M_i\rangle}$ and (3) there is some utility aggregation rule s such that for all $\langle M_i\rangle \in \mathcal{G}$ and all $f, g \in \mathcal{F}_{\langle M_i\rangle}$: If f absolutely dominates g, then $C_{S(\langle M_i\rangle)}(\{f, g\}) = \{f\}$.*

We call any s that satisfies condition (3) *a utility aggregation rule associated with* S. An ex post social choice rule S *has a wide domain* if and only if, for any graining $\langle \mathcal{W}, \mathcal{C}\rangle$, any $\langle p_i\rangle \in \prod_i \mathbf{p}_i(\mathcal{W})$, any $\langle u_i\rangle \in \prod_i \mathbf{u}_i(\mathcal{C})$ and any $\mathcal{F} \subseteq \Phi_{\mathcal{W},\mathcal{C}}$, there exists some $\langle M_i\rangle$ in the domain of S such that $\mathcal{W}_{\langle M_i\rangle} = \mathcal{W}$, $\langle p_{M_i}\rangle = \langle p_i\rangle$, $\mathcal{C}_{\langle M_i\rangle} = \mathcal{C}$, $\langle u_{M_i}\rangle = \langle u_i\rangle$ and (*) $\mathcal{F} \subseteq \mathcal{F}_{\langle M_i\rangle}$. Our definition of a 'wide domain' is in an important sense not a condition of a 'universal domain'. Firstly, the set of belief measures and utility measures that occur in models in the domain of S can be severely restricted (as long as they respect the definition of utility types and the disagreement condition). More importantly, the final clause (*) of the definition allows the underlying structure of acts to force the inclusion of certain acts into the choice set $\mathcal{F}_{\langle M_i\rangle}$. An ex post social choice rule S with a wide domain may, for example, be restricted to only those vectors $\langle M_i\rangle$ in which $\mathcal{F}_{\langle M_i\rangle}$ contains (a reference act for) each possible coarse–grained act $F : \mathcal{W}_{\langle M_i\rangle} \to \mathcal{C}_{\langle M_i\rangle}$. Domains thus restricted introduce considerable technical complications during the proofs.

We now formulate several properties of utility aggregation rules. A Γ–partition $\mathcal{C}$ is *refinable* if and only if every $C \in \mathcal{C}$ has at least two elements. Suppose s is a cardinal–

ordinal utility aggregation rule. s is *non–exceptional* if and only if there exists some refinable Γ–partition $\mathcal{C}$, (possibly identical) consequences $C_1, \ldots, C_I, D_1, \ldots, D_I \in \mathcal{C}$ and $\langle u_i \rangle, \langle u'_i \rangle \in \prod_i \mathbf{u}_i(\mathcal{C})$ such that for any $1 \leq i, j, k \leq I$) $u_i(C_i) = u'_i(C_i)$ and $u_i(D_i) = u'_i(D_i)$, but $[s(\langle u_i \rangle)](\{C_j, D_k\}) = \{C_j\}$ and $[s(\langle u'_i \rangle)](\{C_j, D_k\}) = \{D_k\}$. For concreteness, we mention already the special case in which s' aggregates the individuals' real–valued one–dimensional utilities on $\mathcal{C}$ into a real–valued one–dimensional group utility on $\mathcal{C}$ and where s is the utility aggregation rule generated by s' (i.e., s yields the choice function generated by the group utility aggregated by s'). If s' is a utilitarian rule with weights λ_i, then s is non–exceptional if and only if the weights of at least two individuals are non–zero. By Observation A.2 in the appendix, a utility aggregation rule must be non–exceptional if it is Pareto optimal (cf. below). s is *independent of irrelevant alternatives (IIA)* if and only if, if $u_i(C) = u'_i(C)$ for all $1 \leq i \leq I$ and all $C \in X$, then $[s(\langle u_i \rangle)](X) = [s(\langle u'_i \rangle)](X)$ (for any Γ–partitions $\mathcal{C}$ and $\mathcal{C}'$, any $X \subseteq \mathcal{C} \cap \mathcal{C}'$ and any $\langle u_i \rangle \in \prod_i \mathbf{u}_i(\mathcal{C})$ and $\langle u'_i \rangle \in \prod_i \mathbf{u}_i(\mathcal{C}')$). Note that IIA forces the aggregation of utilities to be independent of the fine–graining of the consequence partition. When, as is usual in the literature, the consequence partition is held fixed, this aspect of IIA cannot be expressed. Note, moreover, that IIA for utility aggregation rules (which deliver a group choice function over consequences) is an extremely weak and ubiquitous condition since we do not presuppose the existence of a group preference or an axiom of revealed preference (cf. the applications in the following section and Plott, 1976). Our final property involves the individuals' utility scales $\trianglerighteq_i$. s is (ex post) *Pareto optimal* if and only if $[s(\langle u_i \rangle)](\{C, D\}) = \{C\}$ when $u_i(C) \trianglerighteq_i u_i(D)$ for all $1 \leq i \leq I$ but $u_j(C) \triangleright_j u_j(D)$ for some $1 \leq j \leq I$ (for any Γ–partition $\mathcal{C}$, any $C, D \in \mathcal{C}$ and any $\langle u_i \rangle \in \prod_i \mathbf{u}_i(\mathcal{C})$).

Instability theorem

For the sake of a strong instability result, we now rig our setup against instabilities and formulate an excessively stringent notion of refinement for individual models. For the group choice model, we use the same notion as in Section 2. We indulge ex post aggregation in two excesses. On the one hand, we ask both that a refining model M' use the same state partition as the model M being refined and that M' have exactly the same belief measure as M. If it were not for our focus on instabilities, we would be satisfied if the state partition in M' detailed the state partition in M (i.e., $\mathcal{W}_M \subseteq [\mathcal{W}_{M'}]$) and if the belief measure of M' refined the belief measure of M in a suitable sense (cf. our definition on p. 11 in case p_M and $p_{M'}$ are probabilities). On the other hand, we ask that not only individual preferences but also individual evaluation functions remain unchanged in a refinement. Again, if it were not for our interest in group choice reversals, we might wish to speak of an individual refinement already when individual preferences are left unchanged. With our excessively stringent definition, however, we obtain a stronger instability result. Our proofs will construct blatant violations of a weak stability condition for group choice based on individual models in which both beliefs and evaluations remain untouched. For any vectors of decision–theoretic models $\langle M_i \rangle, \langle M'_i \rangle \in \mathbf{G}(I)$, we say that $\langle M'_i \rangle$ *stringently refines* $\langle M_i \rangle$ if and only if, for each $1 \leq i \leq I$, (1) $\mathcal{W}_{\langle M_i \rangle} = \mathcal{W}_{\langle M'_i \rangle}$, (2) $\mathcal{C}_{\langle M_i \rangle} \subseteq [\mathcal{C}_{\langle M'_i \rangle}]$, (3) $\mathcal{F}_{\langle M_i \rangle} \subseteq \mathcal{F}_{\langle M'_i \rangle}$, (4) $p_{M_i} = p_{M'_i}$ and (5) $V_{M_i}(f, g) = V_{M'_i}(f, g)$ for all

$f, g \in \mathcal{F}_{\langle M_i \rangle}$. Condition (5) implies that the individuals' preferences over acts in $\mathcal{F}_{\langle M_i \rangle}$ remain the same.[15] An ex post social choice rule S is *weakly stable under refinements* if and only if, for all $\langle M_i \rangle$, $\langle M_i' \rangle$ in the domain of S, $S(\langle M_i' \rangle)$ refines $S(\langle M_i \rangle)$ whenever $\langle M_i' \rangle$ stringently refines $\langle M_i \rangle$.

Theorem 3.2 *Suppose S is an ex post social choice rule with a wide domain and an associated utility aggregation rule that is (1) IIA and non–exceptional, or (2) Pareto optimal.*

Then there is an infinite sequence $\langle M_i^n \rangle_{n \in \mathbb{N}}$ of vectors of decision–theoretic models in $\mathcal{G}_S$ such that $\langle M_i^{n+1} \rangle$ stringently refines $\langle M_i^n \rangle$ (for all $n \in \mathbb{N}$) and S leads to a sequence of group models $\langle S(\langle M_i^n \rangle) \rangle_{n \in \mathbb{N}}$ that oscillates between absolute dominance of f over g and absolute dominance of g over f (for some $f, g \in \mathcal{F}_{\langle M_i^0 \rangle}$). Hence, S is not weakly stable under refinements.

In our view, these results devastate the hope that ex post aggregation could offer a viable alternative to ex ante aggregation. We recall that stability is a not an elusive property. For ex ante social choice rules it is a ubiquitous property and even weaker than IIA. Therefore, the combination of Theorem 2.2 and Theorem 3.2 supports the conclusion that the ex post mode of aggregation cannot yield a normative theory of fairness or of consensus formation. The only way to avoid this conclusion is to challenge the importance of stability or to accept the choice of a graining for the individual models as an additional factor in social choice. Moreover, Theorem 3.2 shows that instabilities may persist even though we build increasingly fine–grained individual models. In the ex post mode, instabilities are there to stay.

The proof of this theorem makes use of extreme disagreements of belief measures where one individual regards a consequence as certain, or one, while another regards it as impossible, or null (cf. the disagreement condition). The proof exploits the freedom to change utilities of null events without affecting individuals' preferences. If we make stronger assumption about the individuals' decision rule, for instance, that individuals maximize expected utility, we can construct refinements suitable for the proof of the theorem without resorting to such extreme disagreements of belief measures, as was the case in the numerical example. The virtue of the current proof technique based on the disagreement condition is its generality. We now consider applications of our theorem to concrete models. We emphasize that instabilities are not the artifact of a narrow class of decision theories. The theorem applies to all decision–theoretic models listed in the following section.

[15] In case G_i is monadic, clause (5) reduces to $V_{M_i}(f) = V_{M_i'}(f)$ for all $f \in \mathcal{F}_{\langle M_i \rangle}$. In case G_i is an ordinal decision rule, this clause reduces to $f \; R_{M_i} \; g$ iff $f \; R_{M_i'} \; g$ for all $f, g \in \mathcal{F}_{\langle M_i \rangle}$.

4 Applications

We now present an exemplary list of individual decision–theoretic models and utility aggregation rules to which our results apply. Different individuals may be described by any one of the models listed. Any of these models is a decision–theoretic models in the sense of our definition and any of their combinations satisfies the disagreement condition (p. 17). The group may or may not have any group beliefs and may or may not use a belief aggregation rule. In a special case, the group aggregates individual utilities into a group utility, aggregates individual beliefs into a group belief and then uses one of the below decision–theoretic models to construct a preference on the basis of these aggregated beliefs and aggregated utilities. Our results cover this special case since any of the models listed satisfies absolute dominance.[16]

Utility aggregation

1.) Cardinal aggregation of real–valued one–dimensional utilities. Consider individuals with real–valued one–dimensional utilities $u_i : \mathcal{C} \to \mathbb{R}$ and a utility aggregation rule s' that aggregates these utilities into a one–dimensional utility $u_0 : \mathcal{C} \to \mathbb{R}$. Any aggregation rule s' of this class generates a utility aggregation rule s in the sense of the definition. A smaller subclass consists of those utility aggregation rules that are generated by functions of the form $t' : \mathbb{R}^I \to \mathbb{R}$ and for which the aggregated utility $s'(\langle u_i \rangle)$ on $\mathcal{C}$ is defined by $s'(\langle u_i \rangle)(C) := t'(\langle u_i(C) \rangle)$ for all $C \in \mathcal{C}$. These aggregation rules are IIA iff the same function t is used for the aggregation of utilities on any partition. *Utilitarian* aggregation rules are generated by a linear function $t' : \mathbb{R}^I \to \mathbb{R}$. Utilitarian rules are non–exceptional iff at least two individuals' weights are non–zero. If all weights of a utilitarian rule are positive, then the rule is Pareto–optimal and satisfies the preconditions of Theorem 3.2 even if different weights are used for different partitions. If weights are independent of the partition on which utilities are defined and if at least two weights are non–zero, the rule also satisfies the preconditions of Theorem 3.2.

2.) Cardinal aggregation of real–valued multi–dimensional utilities. Restricting ourselves to $L = 2$, we represent any two–dimensional utility $u^* : \mathcal{C} \to \mathbb{R}^2$ by two one–dimensional utilities $u, v : \mathcal{C} \to \mathbb{R}$. A special class of utility aggregation rules for two–dimensional utilities is generated by pairs $\langle s'_1, s'_2 \rangle$ of cardinal–cardinal utility aggregation rules for one–dimensional utilities by using s'_1 and s'_2 to aggregate the two components of the individuals' two–dimensional utility. Thus, the rule generated by $\langle s'_1, s'_2 \rangle$ aggregates vectors of two–dimensional utilities (represented by) $\langle u_i, v_i \rangle$ into the pair $\langle s'_1(\langle u_i \rangle), s'_2(\langle v_i \rangle) \rangle$. Given our definition of the canonical scale relation on $\mathbb{R}^L$ (p. 16), both s'_1 and s'_2 have any of the properties of non–exceptionality, IIA or Pareto optimality

[16]In the case of regret theory, the group's modification function must be regular. In the case of expected utility with threshold $\alpha_0 \in \mathbb{R}$ and an aggregated group utility u_0, we define the group's choice function c on $\mathcal{C}$ such that, for any binary choices, $c(\{C, D\}) = \{C\}$ iff $u_0(C) - u_0(D) > \alpha_0$ (for any $C, D \in \mathcal{C}$). Then clause 3 of Definition 3.1 is satisfied.

exactly when $\langle s_1', s_2' \rangle$ has any of the corresponding properties. The reader may therefore refer back to the cardinal aggregation of one–dimensional utilities.

3.) Ordinal aggregation. A utility aggregation rule that only yields an ordinal preference over consequences does not furnish the cardinal information required by most of the decision theories on our list. Under ordinal aggregation, the group will therefore have to make use of an ordinal decision theory, such as decision–theoretic leximin or leximax (cf. below). *Allocative leximin* constructs a group preference over consequences by maximizing, first, the utility enjoyed by the worst–off individual, then the utility enjoyed by the second–worst–off individual etc. (cf. Sen, 1970). *Allocative leximax* is an opposing rule that constructs a group preference by maximizing, first, the utility enjoyed by the best–off individual, then the utility enjoyed by the second–best–off individual etc. (cf. Sen, 1970). The *de Borda count* constructs a group choice by assigning (a positive number of) points to consequences according to the rank that these consequences hold in each individuals' preference (or, utility) ranking and then choosing the consequences with the greatest number of points summed across individuals. Since we do not assume that utility aggregation rules yield a group preference over consequences (but merely a choice function), the de Borda count is IIA in the sense of the definition in Section 3. All three examples are, therefore, non–exceptional, IIA and Pareto optimal and satisfy the preconditions of Theorem 3.2. We end with a forth example that highlights non–exceptionality. For any $1 \leq \alpha \leq I$, *α–Pareto preference* constructs a group preference $\succeq_0$ over consequences such that $C \succ_0 D$ iff $u_i(C) \trianglerighteq_i u_i(D)$ for all $1 \leq i \leq I$ but $u_j(C) \triangleright_j u_j(D)$ for at least α different j $(1 \leq j \leq I)$. α–Pareto preference is IIA for any α. However, unless $\alpha = 1$ (Sen's (1970) Pareto extension rule), the rule is not Pareto optimal in the sense defined. Since the rule is non–exceptional iff $1 \leq \alpha \leq I - 1$, only I–Pareto preference satisfies none of the preconditions of Theorem 3.2.

Decision theories

Some of the following models have originally been proposed for agents' choices among roulette lotteries with fixed objective probabilities (especially, weighted utility and Machina). To explore the generality of our result, we extend this interpretation and allow individuals to use their own subjective probabilities. In what follows, we hold the graining $\langle \mathcal{W}, \mathcal{C} \rangle$ fixed and assume that $F, F_1, F_2 : \mathcal{W} \to \mathcal{C}$ are the $\langle \mathcal{W}, \mathcal{C} \rangle$–grainings of $f, f_1, f_2 \in \Phi_{\mathcal{W},\mathcal{C}}$. Under any of the following decision rules using probabilities, an event $A \in [\mathcal{W}]$ is null w.r.t. a probability p and the decision rule in question if and only if $p(A) = 0$.

Expected utility. Beliefs are represented by probabilities p_i on $\mathcal{W}$ and agents possess real–valued one–dimensional utilities u_i on $\mathcal{C}$. Agents maximize the expectation $E(p_i, u_i \circ F)$ of $u_i \circ F$ w.r.t. p_i. Since our proofs construct only models with finite partitions of states, we avoid the issue of merely finitely vs. countably additive probabilities.[17]

[17] We can modify the proofs for state partitions with infinitely many partition cells by introducing a suitable σ–algebra on which σ–additive probabilities are defined in a way that makes the acts constructed in the proof measurable.

Expected utility with threshold. Beliefs are represented by probabilities p_i on $\mathcal{W}$ and agents possess real–valued one–dimensional utilities u_i on $\mathcal{C}$ and a threshold $\alpha_i \in \mathbb{R}$. Agents' strongly prefer f_1 to f_2 exactly when $E(p_i, u_i \circ F_1) - E(p_i, u_i \circ F_2) > \alpha_i$ (Fishburn, 1988).

Choquet–expected utility. Beliefs are represented by capacities p_i on $\mathcal{W}$ and agents maximize the Choquet–expectation of real–valued one–dimensional utilities u_i on $\mathcal{C}$ (Gilboa, 1987, Schmeidler, 1989). p is a *capacity on* $\mathcal{W}$ iff $p : [\mathcal{W}] \to \mathbb{R}$, $p(\varnothing) = 0$, $p(\Omega) = 1$ and p is monotonic w.r.t. set inclusion. Agents maximize the Choquet–expectation of $u \circ F$ w.r.t. p defined by $C(p, u \circ F) := \int_0^\infty p(u \circ F \geq x)\ dx + \int_{-\infty}^0 [p(u \circ F \geq x) - 1]\ dx$. In the present setting, Choquet–expected utility theory subsumes Quiggin's (1982) rank dependent utility theory. In rank dependent utility theory, an event $A \in [\mathcal{W}]$ is null w.r.t. p if and only if $p(A) = 0$. In contrast, an event $A \in [\mathcal{W}]$ is null w.r.t. a capacity p and Choquet–expectation if and only if $p(B) = p(B \cap A)$ for all $B \in [\mathcal{W}]$.

Probability transforms. Beliefs are represented by probabilities p_i on $\mathcal{W}$, agents possess real–valued one–dimensional utilities u_i on $\mathcal{C}$ and, in addition, a probability transformation function $\pi_i : [0,1] \to [0,1]$ that is normalized to $\pi_i(0) = 0$ and $\pi_i(1) = 1$. This transformation function may differ across agents. Agents maximize $E(\pi_i \circ p_i, u_i \circ F)$ (Edwards, 1955, Kahneman/Tversky, 1979) or $E(\pi_i \circ p_i, u_i \circ F)/E(\pi_i \circ p_i, \mathbf{1})$ where $\mathbf{1} : C \mapsto 1$, $\mathbf{1} : \mathcal{C} \to \mathbb{R}$ (Karmarkar, 1978).

Weighted utility theory. Agents posses subjective probabilities p_i on $\mathcal{W}$ and two real–valued one–dimensional utilities u_i, v_i on $\mathcal{C}$. Agents maximize the function $E(p_i, [u_i \circ F])/E(p_i, [v_i \circ F])$ (Chew 1983, Fishburn, 1983).

Machina. Beliefs are represented by probabilities p_i on $\mathcal{W}$ and agents are endowed with a two–dimensional utility or, equivalently, with two one–dimensional utilities $u_i, v_i : \mathcal{C} \to \mathbb{R}$. Agents maximize Machina's (1982) functional defined by $V_i(f) := E(p_i, u_i \circ F) + \frac{1}{2} E(p_i, v_i \circ F)^2$.

Regret. Agents are characterized by subjective probabilities p_i on $\mathcal{W}$, real–valued one–dimensional utilities u_i on $\mathcal{C}$ and a 'modification function' $M_i : \mathbb{R}^2 \to \mathbb{R}$. Define the regret functional by $V_i(f_1, f_2) := E(p_i, M_i(u_i \circ F_1, u_i \circ F_2)$ (Loomes/Sugden, 1982). Agents hold preferences such that $f_1 \succeq_i f_2$ iff $V_i(f_1, f_2) \geq V_i(f_2, f_1)$. A modification function M is *regular* if and only if $M(x, x) = x$ (for any $x \in \mathbb{R}$), $M(., y)$ is strictly increasing increasing (for any $y \in \mathbb{R}$) and $M(x, .)$ is non–increasing (for any $x \in \mathbb{R}$).

Leximin or leximax. Agents are equipped with real–valued one–dimensional utilities u_i or orderings r_i on $\mathcal{C}$. Beliefs are represented by a non–empty 'possibility set' $K_i \subseteq \mathcal{W}$ (e.g., the support of a probability measure p_i on $\mathcal{W}$) which dichotomizes states into those considered possible and those considered impossible by the agent's lights. Decision–theoretic leximin w.r.t. all consequences in the possibility set K_i maximizes the worst–case outcome in $\{F(W) | W \in K_i\}$ and, in case of a tie, the second–worst outcome etc. Decision–theoretic leximax w.r.t. all consequences in the possibility set K_i maximizes the best–case outcome in $\{F(W) | W \in K_i\}$ and, in case of a tie, the second–best outcome etc.[18] Any $A \in [\mathcal{W}]$ is null w.r.t. a possibility set K and decision–theoretic leximin (leximax) if and only if $A \subseteq -\bigcup K$.

[18]We assume that $\mathcal{W}$ has $M \in \mathbb{N}^+$ elements. Agents following decision–theoretic leximin hold preferences such that (1) $f_1 \succ_i f_2$ iff there are bijections $\sigma, \tau : \{1, \ldots, M\} \to \mathcal{W}$ with $[u_i \circ F_1](\sigma(1)) \leq \ldots \leq [u_i \circ F_1](\sigma(M))$ and $[u_i \circ F_2](\tau(1)) \leq \ldots \leq [u_i \circ F_2](\tau(M))$ and there exists $1 \leq m \leq M$ such that $[u_i \circ F_1](\sigma(m)) > [u_i \circ F_2](\sigma(m))$ while, for all $1 \leq l < m$, $[u_i \circ F_1](\sigma(l)) = [u_i \circ F_2](\sigma(l))$; and (2) $f_1 \sim_i f_2$ iff there are bijections $\sigma, \tau : \{1, \ldots, M\} \to \mathcal{W}$ with $[u_i \circ F_1](\sigma(1)) \leq \ldots \leq [u_i \circ F_1](\sigma(M))$ and $[u_i \circ F_2](\tau(1)) \leq \ldots \leq [u_i \circ F_2](\tau(M))$ such that $[u_i \circ F_1](\sigma(m)) = [u_i \circ F_2](\sigma(m))$ for all $1 \leq m \leq M$.

Appendix A Proofs

Ex ante aggregation

Theorem 2.2 *Proof:* 1.) Trivial. 2.) Suppose S is an ex ante social choice rule that is IIA. Suppose, moreover, that $\mathcal{V}$ is the domain of S and $\langle \mathcal{F}, V_i \rangle, \langle \mathcal{F}', V_i' \rangle \in \mathcal{V}$. Suppose that $X \subseteq \mathcal{F} \cap \mathcal{F}'$. If $\mathcal{V}$ is closed under restrictions, then $\langle X, V_i|X \rangle \in \mathcal{V}$. Moreover, $\langle X, V_i|X \rangle$ refines both $\langle \mathcal{F}, V_i \rangle$ and $\langle \mathcal{F}', V_i' \rangle$. If S is stable under refinements, we then obtain $C_{S(\langle \mathcal{F}, V_i \rangle)}(X) = C_{S(\langle X, V_i|X \rangle)}(X) = C_{S(\langle \mathcal{F}', V_i' \rangle)}(X)$. $\square$

Numerical Example

The following definition provides the construction rule behind this example. We borrow definitions and notation from later parts of Section 3. Assume that both Ω and Γ are countably infinite. Take some $d \in \Gamma$ and define $C := \Gamma - \{d\}$ and $D := \{d\}$. Now take a sequence $\langle \mathcal{C}^n \rangle$ of partitions of C with $\mathcal{C}^0 = \{C\}$ and $\mathcal{C}^n = \{C_{k_1,\ldots,k_n} | k_1 = 1,2; \ldots; k_n = 1,2\}$ such that $C_{k_1,\ldots,k_n,1}$ and $C_{k_1,\ldots,k_n,2}$ are two non–empty sets that partition $C_{k_1,\ldots,k_n}$ (for all $n \in \mathbb{N}$ and all $k_1 = 1,2; \ldots; k_n = 1,2$). By assumption, there is a bijection $f : \Omega \to \Gamma - \{d\}$. Define the sequence $\langle \mathcal{W}^n \rangle$ of partitions of Ω by $\mathcal{W}^n := \{f^{-1}(X) | X \in \mathcal{C}^n\}$ where $f^{-1}(X)$ is the image of X under f^{-1} (for any $n \in \mathbb{N}$). Hence, the $\langle \mathcal{W}^n, \mathcal{C}^n \cup \{D\} \rangle$–graining of f is a function $F : \mathcal{W}^n \to \mathcal{C}^n \cup \{D\}$ such that $F(W_{k_1,\ldots,k_n}) = C_{k_1,\ldots,k_n}$ (for all $n \in \mathbb{N}$ and all $k_1 = 1,2; \ldots; k_n = 1,2$). Define $g(\omega) := d$ for all $\omega \in \Omega$. Hence, the $\langle \mathcal{W}^n, \mathcal{C}^n \cup \{D\} \rangle$–graining of g is a function $F : \mathcal{W}^n \to \mathcal{C}^n \cup \{D\}$ such that $F(W_{k_1,\ldots,k_n}) = D$ (for all $n \in \mathbb{N}$ and all $k_1 = 1,2; \ldots; k_n = 1,2$). Let $M_i^n := \langle \mathcal{W}^n, p_i^n, \mathcal{C}^n \cup \{D\}, u_i^n, \{f, g\}, U_i^n \rangle$ where p_i^n and u_i^n are as defined below and U_i^n evaluates acts in terms of their expected utility w.r.t. p_i^n and u_i^n (for $i = 1,2$, $n \in \mathbb{N}$). For all $n \in \mathbb{N}$ and all $k_1 = 1,2; \ldots; k_n = 1,2$, let:

$$\begin{aligned} p_1^n(W_{k_1,\ldots,k_n}) &:= (\frac{3}{4})^{\text{one}(k_1,\ldots,k_n)} \cdot (\frac{1}{4})^{\text{two}(k_1,\ldots,k_n)} \\ p_2^n(W_{k_1,\ldots,k_n}) &:= (\frac{1}{4})^{\text{one}(k_1,\ldots,k_n)} \cdot (\frac{3}{4})^{\text{two}(k_1,\ldots,k_n)} \end{aligned}$$

where one$(k_1,\ldots,k_n)$ and two$(k_1,\ldots,k_n)$ are the number of 1's and 2's contained in the string $k_1,\ldots,k_n$, respectively. For the empty string ($n = 0$), we have one() = two() = 0. Next, we define $u_1^n(.)$ recursively by $u_1^0(C) := 1$ and, for all $n \in \mathbb{N}^+$:

$$\begin{aligned} u_1^n(C_{k_1,\ldots,k_{n-1},1}) &:= u_1^{n-1}(C_{k_1,\ldots,k_{n-1}}) + (-1)^{n+1}2 \\ u_1^n(C_{k_1,\ldots,k_{n-1},2}) &:= u_1^{n-1}(C_{k_1,\ldots,k_{n-1}}) + (-1)^n 6 \end{aligned}$$

Finally, set $u_2^n(C_{k_1,\ldots,k_n}) := (-1)^n 2 - u_1^n(C_{k_1,\ldots,k_n})$ and $u_1^n(D) := u_2^n(D) := 0$ (for all $n \in \mathbb{N}$). For all $n \in \mathbb{N}$, we obtain the group utilities by averaging ($u_0^n = \frac{u_1^n + u_2^n}{2}$): $u_0^n(D) = 0$ and $u_0^n(C_{k_1,\ldots,k_n}) = \frac{1}{2}(u_1^n(C_{k_1,\ldots,k_n}) + (-1)^n 2 - u_1^n(C_{k_1,\ldots,k_n})) = (-1)^n$. We note that for all $n \in \mathbb{N}$:

$$u_1^n(C_{k_1,\ldots,k_n}) = \frac{3}{4}u_1^{n+1}(C_{k_1,\ldots,k_n,1}) + \frac{1}{4}u_1^{n+1}(C_{k_1,\ldots,k_n,2}) \tag{1}$$

$$u_2^n(C_{k_1,\ldots,k_n}) = \frac{1}{4}u_2^{n+1}(C_{k_1,\ldots,k_n,1}) + \frac{3}{4}u_2^{n+1}(C_{k_1,\ldots,k_n,2}) \tag{2}$$

It is then easy to see that $U_i^n(f) = 1$ and $U_i^n(g) = 0$ and, hence, M_i^{n+1} refines M_i^n (for $i = 1,2$ and any $n \in \mathbb{N}$). Moreover, Equations (1), (2) show that M_i^{n+1} locally refines M_i^n (for $i = 1,2$ and any $n \in \mathbb{N}$) in the sense defined in footnote **10**. Finally, we find that $u_0^n(D) = 0$, $u_0^{2n}(C_{k_1,\ldots,k_{2n}}) = 1$ and $u_0^{2n+1}(C_{k_1,\ldots,k_{2n+1}}) = -1$ (for any $n \in \mathbb{N}$ and all $k_1 = 1, 2; \ldots; k_{2n+1} = 1, 2$). Hence, f absolutely dominates g w.r.t. the group's utility in even numbered models while g absolutely dominates f w.r.t. the group's utility in odd numbered models. □

Ex post aggregation

Lemma A.1 *Suppose G is a decision rule for a belief type $\mathbf{p}$ and a utility type $\mathbf{u}$. Suppose $\langle \mathcal{W}, \mathcal{C} \rangle$ and $\langle \mathcal{W}, \mathcal{C}' \rangle$ are grainings with $\mathcal{C} \subseteq [\mathcal{C}']$, $f_1, f_2 \in \Phi_{\mathcal{W},\mathcal{C}'}$, $u \in \mathbf{u}(\mathcal{C})$, $u' \in \mathbf{u}(\mathcal{C}')$, $p \in \mathbf{p}(\mathcal{W})$ and $A \in [\mathcal{W}]$. Then If A is p, G–one and $[u \circ \mathcal{C} \circ f_k](\omega) = [u' \circ \mathcal{C}' \circ f_k](\omega)$ for all $\omega \in A$ and $k = 1, 2$, then $G(p, u, f_1, f_2) = G(p, u', f_1, f_2)$.*

Proof: Suppose that A is p, G–one and $[u \circ \mathcal{C} \circ f_k](\omega) = [u' \circ \mathcal{C}' \circ f_k](\omega)$ for all $\omega \in A$ and $k = 1, 2$. Let $v \in \mathbf{u}(\mathcal{C}')$ be the function such that $v(X) := u(C)$ for all $C \in \mathcal{C}$ and $X \in \mathcal{C}'$ with $X \subseteq C$. Hence, $[v \circ \mathcal{C}' \circ f_k](\omega) = [u \circ \mathcal{C} \circ f_k](\omega)$ for all $\omega \in \Omega$ and $k = 1, 2$. By property (b) of a decision rule, we then have $G(p, v, f_1, f_2) = G(p, u, f_1, f_2)$. Since A is p, G–one, we also have $G(p, v, f_1, f_2) = G(p, u', f_1, f_2)$. Hence, $G(p, u, f_1, f_2) = G(p, u', f_1, f_2)$. □

Recall in what follows that we have fixed each individual's belief type $\mathbf{p}_i$, utility type $\mathbf{u}_i$, scale $\langle Z_i, \trianglerighteq_i \rangle$ and decision rule G_i (for each $1 \le i \le I$).

Lemma A.2 *Suppose s is a cardinal–ordinal utility aggregation rule that is (1) IIA and non–exceptional, or (2) Pareto optimal. Then there exist $v_j^i, w_j^i \in Z_i$ with $v_i^i = w_i^i$ and $v_{i+I}^i = w_{i+I}^i$ ($1 \le i \le I$, $1 \le j \le 2I$) such that for all Γ–partitions $\mathcal{C}$ with $2 \cdot I$ different consequences $Y_1, \ldots, Y_{2I} \in \mathcal{C}$ and for all $\langle u_i \rangle, \langle u_i' \rangle \in \prod_i \mathbf{u}_i(\mathcal{C})$, $u_i(Y_j) = v_j^i$ and $u_i'(Y_j) = w_j^i$ (for all $1 \le i \le I$, $1 \le j \le 2I$), we have $[s(\langle u_i \rangle)](\{Y_l, Y_m\}) = \{Y_l\}$ but $[s(\langle u_i' \rangle)](\{Y_l, Y_m\}) = \{Y_m\}$ (for all $1 \le l \le I$ and $I + 1 \le m \le 2I$).*

Proof: Recall that $\langle Z_i, \trianglerighteq_i \rangle$ is individual i's scale for her utility type $\mathbf{u}_i$. 1.) Trivial. 2.) Suppose s is Pareto optimal. For any $1 \le i \le I$, let $\mathbf{0}_i, \mathbf{1}_i \in Z_i$ such that $\mathbf{1}_i \triangleright_i \mathbf{0}_i$. Let $\mathcal{C}$ be a partition with mutually non–identical consequences $Y_1, \ldots, Y_{2I} \in \mathcal{C}$ and let $u_i, u_i' \in \mathbf{u}_i(\mathcal{C})$ such that $u_i(Y_i) = u_i(Y_{I+i}) = \mathbf{0}_i$, $u_i(Y_j) = \mathbf{1}_i$, $u_i(Y_{I+j}) = \mathbf{0}_i$ and $u_i'(Y_i) = u_i'(Y_{I+i}) = \mathbf{0}_i$, $u_i'(Y_j) = \mathbf{0}_i$, $u_i'(Y_{I+j}) = \mathbf{1}_i$ (for all $1 \le i, j \le I$ with $i \ne j$). Hence, $u_i(Y_i) = u_i'(Y_i)$ and $u_i(Y_{i+I}) = u_i'(Y_{i+I})$ (for all $1 \le i \le I$). If s is Pareto optimal, we obtain $[s(\langle u_i \rangle)](\{Y_l, Y_m\}) = \{Y_l\}$ but $[s(\langle u_i' \rangle)](\{Y_l, Y_m\}) = \{Y_m\}$ (for all $1 \le l \le I$ and $I + 1 \le m \le 2I$). □

Theorem 3.2 *Proof:* Suppose S is an ex post social choice rule with a wide domain and s is a utility aggregation rule associated with S. Suppose furthermore that (1) s is IIA and non–exceptional, or (2) s is Pareto optimal. By assumption, Γ is at least countably infinite and, thus, we can partition it into $2 \cdot I$ different, at least countably infinite, sets $X_1, \ldots, X_{2I}$. For each $1 \le k \le 2I$, we can hence enumerate the elements of some countably infinite subset of X_k in a sequence $x_k^0, x_k^1, x_k^2, \ldots, x_k^n, \ldots$ (where $x_k^n \ne x_k^{n'}$ for any $n, n' \in \mathbb{N}$ with $n \ne n'$). We then define a sequence $\langle \mathcal{C}_k^n \rangle$ of partitions of X_k by $\mathcal{C}_k^n = \{\{x_k^1\}, \ldots, \{x_k^n\}, X_k - \{x_k^1, \ldots, x_k^n\}\}$ for all $n \in \mathbb{N}^+$. We write $X_k^0 := X_k$ and $X_k^n := X_k - \{x_k^1, \ldots, x_k^n\}$ for all $n \in \mathbb{N}^+$. Note that $x_k^0 \in X_k^n$ for all $n \in \mathbb{N}$. Let $\mathcal{C}^0 := \{X_1, \ldots, X_{2I}\}$ and $\mathcal{C}^n := \mathcal{C}_1^n \cup \ldots \cup \mathcal{C}_{2I}^n$ for all $n \in \mathbb{N}^+$. Let $\phi^0 : \mathcal{C}^0 \to \mathcal{C}^1$ and $\psi^0 : \mathcal{C}^0 \to \mathcal{C}^1$ be defined by $\phi^0(X_k^0) := X_k^1$ and $\psi^0(X_k^0) := \{x_k^1\}$ (for all $1 \le k \le 2I$). For any $n \in \mathbb{N}^+$, let $\phi^n : \mathcal{C}^n \to \mathcal{C}^{n+1}$ and $\psi^n : \mathcal{C}^n \to \mathcal{C}^{n+1}$ be defined by $\phi^n(\{x_k^{n'}\}) := \{x_k^{n'}\}$ and $\psi^n(\{x_k^{n'}\}) := \{x_k^{n'}\}$ for all $1 \le n' \le n$ ($n' \in \mathbb{N}^+$) and by $\phi^n(X_k^n) := X_k^{n+1}$ and $\psi^n(X_k^n) := \{x_k^{n+1}\}$ (for all $1 \le k \le 2I$). By the disagreement condition, there is some Ω–partition $\mathcal{W} = \{W_1, \ldots, W_I\}$ and some $p_i \in \mathbf{p}_i(\mathcal{W})$ such that W_i is p_i, G_i–one (for any $1 \le i \le I$). Define $f \in \Phi$ as the act such that, for any $1 \le i \le I$ and any $\omega \in W_i$, $f(\omega) = x_i^0$. Define $g \in \Phi$ as the act such that, for any $1 \le i \le I$ and any $\omega \in W_i$, $g(\omega) = x_{I+i}^0$. For any $n \in \mathbb{N}$, let $\mathcal{F}_\dagger^n \subseteq \Phi_{\mathcal{W},\mathcal{C}^n}$ be the set of acts h such that for any $1 \le i \le I$ there exists some $1 \le k \le 2 \cdot I$ and some $1 \le m \le n + 1$ such that $h(\omega) = x_k^m$ for all $\omega \in W_i$.

For any $n \in \mathbb{N}$, since $\mathcal{F}_\dagger^n$ is unambiguous w.r.t. $\langle \mathcal{W}, \mathcal{C}^n \rangle$, there is exactly one act $h_1^n \in \mathcal{F}_\dagger^n$ that has the same $\langle \mathcal{W}, \mathcal{C}^n \rangle$–graining as f and there is exactly one act $h_2^n \in \mathcal{F}_\dagger^n$ that has the same $\langle \mathcal{W}, \mathcal{C}^n \rangle$–graining as g. Replacing h_1, h_2 with f, g, we define $\mathcal{F}^n := (\mathcal{F}_\dagger^n - \{h_1^n, h_2^n\}) \cup \{f, g\}$ for any $n \in \mathbb{N}$. Note that $\mathcal{F}^n \subseteq \Phi_{\mathcal{W}, \mathcal{C}^{n+1}}$ ($n \in \mathbb{N}$). Moreover, $\mathcal{F}^n$ is unambiguous w.r.t. $\langle \mathcal{W}, \mathcal{C}^n \rangle$ and there is no set $\mathcal{F}'$ with $\mathcal{F}^n \subseteq \mathcal{F}' \subseteq \Phi_{\mathcal{W}, \mathcal{C}^n}$ that is unambiguous w.r.t. $\langle \mathcal{W}, \mathcal{C}^n \rangle$.

By Lemma A.2, there exist $v_j^i, w_j^i \in \prod_i Z_i$ with $v_i^i = w_i^i$ and $v_{I+i}^i = w_{I+i}^i$ ($1 \leq i \leq I$, $1 \leq j \leq 2I$) such that for all Γ–partitions $\mathcal{C}$ with $2 \cdot I$ different consequences $Y_1, \ldots, Y_{2I} \in \mathcal{C}$ and for all $\langle u_i \rangle, \langle u_i' \rangle \in \mathbf{u}(\mathcal{C})^I$ with $u_i(Y_j) = v_j^i$, and $u_i'(Y_j) = w_j^i$ (for all $1 \leq i \leq I$, $1 \leq j \leq 2I$), we have $[s(\langle u_i \rangle)](\{Y_l, Y_m\}) = \{Y_l\}$ but $[s(\langle u_i' \rangle)](\{Y_l, Y_m\}) = \{Y_m\}$ (for all $1 \leq l \leq I$ and $I+1 \leq m \leq 2I$). We recursively define a sequence $\langle u_1^n, \ldots, u_I^n \rangle_{n \in \mathbb{N}}$ of vectors of individual utility functions such that, for each $n \in \mathbb{N}$ and $1 \leq i \leq I$, $u_i^n \in \mathbf{u}_i(\mathcal{C}^n)$. For all even $n \in \mathbb{N}$, let $u_i^n(X_k^n) := v_k^i$ and, for all odd $n \in \mathbb{N}$, let $u_i^n(X_k^n) := w_k^i$ (for all $1 \leq i \leq I$ and all $1 \leq k \leq 2I$). For any $n, n' \in \mathbb{N}^+$ with $n' \leq n$, let $u_i^{n+1}(\{x_k^{n'}\}) := u_i^n(\{x_k^{n'}\})$ and let $u_i^{n+1}(\{x_k^{n+1}\}) := u_i^n(X_k^n)$ (for all $1 \leq i \leq I$, and $1 \leq k \leq 2I$). It follows from Lemma A.2 that, for all $n \in \mathbb{N}$, (*) $[s(\langle u_i^{2n} \rangle)](\{X_l^{2n}, X_m^{2n}\}) = \{X_l^{2n}\}$ and $[s(\langle u_i^{2n+1} \rangle)](\{X_l^{2n+1}, X_m^{2n+1}\}) = \{X_m^{2n+1}\}$ (for all $1 \leq l \leq I$ and $I+1 \leq m \leq 2I$). Moreover, (**) $u_i^n = u_i^{n+1} \circ \psi^n$ for any $n \in \mathbb{N}$ and $1 \leq i \leq I$. Since S has a wide domain, there is a sequence $\langle M_1^n, \ldots, M_I^n \rangle_{n \in \mathbb{N}}$ of vectors of individual decision–theoretic models in the domain of S such that, for all $n \in \mathbb{N}$ and all $1 \leq i \leq I$, $\mathcal{W}_{\langle M_i^n \rangle} = \mathcal{W}$, $p_{M_i^n} = p_i$, $\mathcal{C}_{\langle M_i^n \rangle} = \mathcal{C}^n$, $u_{M_i^n} = u_i^n$ and $\mathcal{F}_{\langle M_i^n \rangle} = \mathcal{F}^n$. We now show that $\langle M_i^{n+1} \rangle$ refines $\langle M_i^n \rangle$ (for all $n \in \mathbb{N}$). Conditions 1–4 of the definition of a refinement (p. 19) are trivially satisfied. On the one hand, (**) yields for any $h \in \mathcal{F}^n$ with $h \neq f, g$ and for any $n \in \mathbb{N}$: $u_i^n(\mathcal{C}^n(h(\omega))) = [u_i^{n+1} \circ \psi^n](\mathcal{C}^n(h(\omega))) = u_i^{n+1}(\mathcal{C}^{n+1}(h(\omega)))$ for all $1 \leq i \leq I$ and all $\omega \in \Omega$. On the other hand, we have (for any $n \in \mathbb{N}$) $u_i^n(X_i^n) = v_i^i = w_i^i = u_i^{n+1}(X_i^{n+1}) = [u_i^{n+1} \circ \phi^n](X_i^n)$ and, hence, $u_i^n(\mathcal{C}^n(f(\omega))) = [u_i^{n+1} \circ \phi^n](\mathcal{C}^n(f(\omega))) = u_i^{n+1}(\mathcal{C}^{n+1}(f(\omega)))$ for all $1 \leq i \leq I$ and all $\omega \in W_i$. Similarly for g. Since W_i is p_i, G_i–one, Lemma A.1 yields $G_i(p_i, u_i^n, h_1, h_2) = G_i(p_i, u_i^{n+1}, h_1, h_2)$ for all $h_1, h_2 \in \mathcal{F}^n$, $1 \leq i \leq I$, $n \in \mathbb{N}$. Similarly for g. Hence, all conditions for a refinement are satisfied. Finally, (*) implies, by Definition 3.1, $C_{S(\langle M_i^{2n} \rangle)}(\{f, g\}) = \{f\}$ and $C_{S(\langle M_i^{2n+1} \rangle)}(\{f, g\}) = \{g\}$ (for all $n \in \mathbb{N}$). In other words, f absolutely dominates g w.r.t. $s(\langle u_i^{2n} \rangle)$ whereas g absolutely dominates f w.r.t. $s(\langle u_i^{2n+1} \rangle)$ (for all $n \in \mathbb{N}$). □

References

[1] R. Zeckhauser A. Hylland. The impossibility of Bayesian group decision making with separate aggregation of beliefs and values. *Econometrica*, 47:1321–1336, 1979.

[2] K. Arrow. *Social Choice and Individual Values*. Wiley, New York, 1951.

[3] J.L. Becker and R. Sarin. Gamble dependent utility. *Management Science*, 33:1367–82, 1987.

[4] J. Broome. Bolker–Jeffrey expected utility theory and axiomatic utilitarianism. *Review of Economic Studies*, 57:477–503, 1990.

[5] S.H. Chew. A generalization of the quasilinear mean with applications to the measurement of income inequality and decision theory resolving the Allais paradox. *Econometrica*, 51:1065–92, 1983.

[6] W. Edwards. The prediction of decisions among bets. *Journal of Experimental Psychology*, 50:201–214, 1955.

[7] Henley E.J. and Kumamoto H. *Probabilistic Risk Assessment: Reliability Engineering, Design, and Analysis*. Institute of Electrical and Electronics Engineers, Inc., Piscataway, 1992.

[8] P.C. Fishburn. Transitive measurable utility. *Journal of Economic Theory*, 31:293–317, 1983.

[9] P.C. Fishburn. *Nonlinear Preference and Utility Theory*. Johns Hopkins University, Baltimore, 1988.

[10] I. Gilboa. Expected utility with purely subjective non–additive probabilities. *Journal of Mathematical Economics*, 16:65–88, 1987.

[11] T.S. Glickman and M. Gough, editors. *Readings in Risk*, Washington, D.C., 1990. Resources for the Future.

[12] J.H. Goodman. *Existence of Compromises in Simple Group Decisions*. PhD thesis, Department of Statistics, Carnegie Mellon University, 1988.

[13] A.C. Haddix, S.M. Teutsch, P.A. Shaffer, and D.O. Dunet, editors. *Prevention Effectiveness: A Guide to Decision Analysis and Economic Evaluation*, Oxford, 1996. OUP.

[14] P. Hammond. Ex–ante and ex–post welfare optimality under uncertainty. *Economica*, 48:235–250, 1981.

[15] M. Hild. The instability of robust aggregation. *California Institute of Technology, Social Science Working Paper*, 1113, 2001b. Available at www.hild.org.

[16] M. Hild. The instability of ex post and robust aggregation without state–consequence separation. *California Institute of Technology, Social Science Working Paper*, 1114, 2001c. Available at www.hild.org.

[17] M. Hild. Policy reversals in risk management: The effect of refined analyses. *California Institute of Technology, Social Science Working Paper*, 1116, 2001d. Available at www.hild.org.

[18] M. Hild, R.C. Jeffrey, and M. Risse. Preference aggregation after Harsanyi. In M. Salles and J. Weymark, editors, *Justice, Political Liberalism, and Utilitarianism*. CUP, Cambridge, 2001.

[19] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47:263–91, 1979.

[20] U.S. Karmarkar. Subjectively weighted utility: A descriptive extension of the expected utility model. *Organizational Behavior and Human Performance*, 21:61–72, 1978.

[21] J.-F. Laslier. Aggregation of preferences with a variable set of alternatives. *Social Choice and Welfare*, 17:269–282, 2000.

[22] S. Leeds. Levi's decision theory. *Philosophy of Science*, 57:158–168, 1990.

[23] I. Levi. *Hard Choices*. CUP, Cambridge, 1986.

[24] I. Levi. Pareto unanimity and consensus. *Journal of Philosophy*, 87:481–492, 1990.

[25] G. Loomes and R. Sugden. Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, 92:805–824, 1982.

[26] M.J. Machina. "Expected utility" analysis without the independence axiom. *Econometrica*, 50:277–323, 1982b.

[27] P. Mongin. Consistent Bayesian aggregation. *Journal of Economic Theory*, 66:313–351, 1995.

[28] P. Mongin. The paradox of the Bayesian experts and state–dependent utility theory. *Journal of Mathematical Economics*, 29:331–361, 1998.

[29] C. Plott. Axiomatic social choice theory: An overview and interpretation. *American Journal of Political Science*, 20:511–596, 1976.

[30] J. Quiggin. A theory of anticipated utility. *Journal of Economic Behavior and Organization*, 3:323–43, 1982.

[31] L.J. Savage. *The Foundations of Statistics*. Dover, New York, 1954.

[32] M.J. Schervish, T. Seidenfeld, and J.B. Kadane. State–dependent utilities. *JASA*, 85:840–847, 1990.

[33] M.J. Schervish, T. Seidenfeld, and J.B. Kadane. Shared preferences and state–dependent utilities. *Management Science*, 37:1575–1589, 1991.

[34] D. Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57:571–587, 1989.

[35] T. Seidenfeld. Outline of a theory of partially ordered preferences. *Philosophical Topics*, 21:173–189, 1993.

[36] T. Seidenfeld, J.B. Kadane, and M.J. Schervish. On the shared preferences of two Bayesian decision makers. *Journal of Philosophy*, 86:225–244, 1989.

[37] A. Sen. *Collective Choice and Social Welfare*. Holden–Day, San Fransisco, 1970.

[38] U.S. Nuclear Regulatory Commission. *Reactor Safety Study : An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants*. Nuclear Regulatory Commission, Washington, D.C., 1975.