THE INSTABILITY OF ROBUST AGGREGATION

Matthias Hild
California Institute of Technology &
Jet Propulsion Laboratory

The Instability of Robust Aggregation

Matthias Hild

Abstract

We discuss the feasibility of Levi's (1990) robust mode of aggregating individuals' evaluations of acts into a social choice function. We examine 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. We find that any stable robust social choice rule must collapse back into ex ante aggregation. We also provide sufficient conditions, such as Pareto optimality, under which robust aggregation leads to an infinite series of reversals of group choice. For ex ante aggregation, we find that stability is ubiquitous and that it follows from independence of irrelevant alternatives.

The Instability of Robust Aggregation

Matthias Hild[*]

1 Introduction

Levi (1990) proposes a novel mode for the aggregation of individuals' expected utilities into a social choice function. The idea behind this 'robust' mode of aggregation is to consider not only the individuals' actual evaluations of acts but also what we can call their 'empathetic' evaluations. Empathetic evaluations are the evaluations that real individuals would have if they were to keep their own utilities but adopted another individual's probabilities, or conversely. In the usual ex ante mode of aggregation, only the individuals' actual evaluations of acts enter the aggregation process. In Levi's robust mode, all empathetic evaluations enter into the aggregation process. Robust aggregation with I individuals is thus defined as ex ante aggregation of I^2 hypothetical individuals with all possible empathetic evaluations. For ease of presentation, we will restrict ourselves in this introduction to individuals who maximize expected utility. Later we will see that most of the presently available decision theories are covered by our theorems. Suppose individuals are expected utility maximizers characterized by pairs $\langle p_i, u_i \rangle$ of subjective probabilities and utilities. Ex ante aggregation aggregates the expected utilities U_i of these individuals. Robust aggregation forms all hypothetical individuals of the form $\langle p_i, u_j \rangle$ and then applies ex ante aggregation to the expected utilities $U_{i,j}$. Clearly, conventional ex ante aggregation is a special type of robust aggregation for which the aggregation rule S does not depend on evaluations the $U_{i,j}$ ($i \neq j$) of hypothetical individuals.

We intend to argue against the viability of robust aggregation and in favour of ex ante aggregation. To appreciate this argument, recall that Levi proposed the robust mode in reaction to the disturbing Ex Ante Homogeneity Theorem regarding the 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.[1]

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.[2] Then either $p_1, \ldots, p_I$ are linearly dependent or $u_1, \ldots, u_I$ are affinely dependent.*

Turning vice into virtue, robust aggregation evades the Ex Ante Homogeneity Theorem by embracing the homogeneity of hypothetical individuals. Trivially, the probabilities and utilities of the hypothetical individuals M_{ij} are linearly dependent even if the real individuals' probabilities and utilities are not, since each real individual's probability measure and utility function occurs at least twice among the models M_{ij}.

We examine the feasibility of robust aggregation from the viewpoint of decision theory. 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. For decision theory, it is therefore vital to consider the process of refining a given decision–theoretic model by taking previously neglected parameters into account (cf. Savage (1954) on 'small worlds'). We consider refinements $M_i = \langle p_i', u_i' \rangle$ of the individuals' decision–theoretic models $M_i = \langle p_i, u_i \rangle$. A decision–theoretic model M_i' is more detailed than model M_i iff, firstly, M_i' can describe all states and consequences of M_i and, secondly, M_i' individuates states and consequences between which the coarser model M_i 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 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_i' refines M_i exactly if M_i' is more detailed than M_i, M_i and M_i' agree on their probabilities for coarse–grained states and M_i and M_i' agree on their expected utilities for coarse–grained acts. All this will be made precise in what follows. We will argue that social choice

---

[1]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.

[2]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 $0 \leq i \leq I$.

2

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.

Our first set of findings is the following: Any robust social choice rule that is stable under refinements collapses back into ex ante aggregation. Moreover, we identify sufficient conditions, such as Pareto optimality, under which robust aggregation leads to repetitive violations of stability and an infinite oscillation of the group's choice. The existence of an infinite sequence of group choice reversal shows that instabilities are not guaranteed to disappear as we continue to refine the individuals' decision–theoretic models. We prove these results in a framework that follows Savage (1954) in representing acts as functions from states to consequences.[3] Since we make meek assumptions about the decision theory used to describe individual preferences or about the existence of group preferences, we are reassured 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 robust mode itself. Our second finding reaffirms this conclusion by demonstrating an attraction of ex ante aggregation. We adapt the condition of independence of irrelevant alternatives to the current framework with a variable set of possible, distinguishable, options. We find not only that stability under refinements follows from independence of irrelevant alternatives, but also that stability is ubiquitous (cf. Hild, 2001a). This result has a twofold significance. Firstly, it proves the feasibility of stable aggregation and, secondly, it provides new reasons for adopting independence of irrelevant alternatives as a desired property of ex ante social choice rules.

The literature distinguishes a third, ex post mode of aggregation that splits individuals' expected utilities into probabilities and utilities before aggregation takes place (cf. Hammond, 1981). The present instability results for robust aggregation have been replicated for the ex post mode. We find that there are no non–exceptional ex post social choice rules that are stable under refinements. What is more, we can show that, under some weak assumptions, ex post aggregation will also lead to an infinite series of group choice reversals (Hild, 2001a). In short, the combination of the Ex Ante Homogeneity Theorem and our instability theorems presents us with yet another difficulty in social choice theory: On the one hand, ex ante aggregation is troubled by the Homogeneity Theorem. On the other hand, ex ante aggregation seems to be the only reasonable way of satisfying stability under refinements.

---

[3]Our technical report Hild (2001c) discusses modifications of our definitions and theorems 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.

2 Decision–Theoretic Refinements

Fine–graining

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.* 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. Analogously, we will replace the worlds of a conventional model with the elements of a partition $\mathcal{W}$ (Savage's 'small 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.

We now give these ideas a formal expression. 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 $C \in \mathcal{C}$.[4] 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 note 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 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.[5] 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

---

[4]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.

[5]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.

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 introduction'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 at the end of this 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 acts are evaluated by the utilities of their consequences and by beliefs about the states in which the consequences are achieved. 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 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)[6] 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 (a) for every $p \in \mathbf{p}$ there is some Ω–partition $\mathcal{W}$ such that $[\mathcal{W}]$ is the domain of p, (b) for all Ω–partitions $\mathcal{W}, \mathcal{W}'$, every 1–1 mapping $\phi : \mathcal{W}' \to \mathcal{W}$ and every p with domain $[\mathcal{W}]$, if $p \in \mathbf{p}$, then $p \circ \bar{\phi} \in \mathbf{p}$ (where the extended mapping $\bar{\phi} : [\mathcal{W}'] \to [\mathcal{W}]$ is defined by $\bar{\phi}(A) := \bigcup \{ \phi(W) | W \in \mathcal{W}', W \subseteq A \}$ for all $A \in [\mathcal{W}']$). We refer to property (b) by saying that $\mathbf{p}$ is *closed under relabelling*. For any Ω–partition $\mathcal{W}$, we define $\mathbf{p}(\mathcal{W})$ as the set of all $p \in \mathbf{p}$ that have domain $[\mathcal{W}]$. We call any $p \in \mathbf{p}(\mathcal{W})$ a *($\mathbf{p}$–)belief measure on* $\mathcal{W}$. We make no assumptions about the range of belief measures and, for the sake of notational economy and without loss of generality, we choose $[\mathcal{W}]$ rather than $2^{\mathcal{W}}$ as the domain of belief measures. Examples of a belief type are set

---

[6]We can, however, allow state–dependent utilities; cf. our technical report Hild (2001c).

of all real valued set functions $p : [\mathcal{W}] \to \mathbb{R}$, the set of all capacities on any Ω–partitions,[7] or the set of all probabilities on any Ω–partitions.[8] Another example of a belief measure is a characteristic function k representing a *possibility set* $K \subseteq \mathcal{W}$ that dichotomizes the states in $\mathcal{W}$ into those considered subjectively possible and those considered subjectively impossible.[9] Suppose we have Ω–partitions $\mathcal{W}$, $\mathcal{W}'$ with $\mathcal{W} \subseteq [\mathcal{W}']$ and $p \in \mathbf{p}(\mathcal{W})$, $p' \in \mathbf{p}(\mathcal{W}')$. We then say that p' *refines* p if and only if $p(A) = p'(A)$ for all $A \in [\mathcal{W}]$. We say that a belief type is closed under cross–products if, for any two marginal belief measures, we can find a belief measure on a suitable two–dimensional partition with these marginal measures. Formally, a belief type $\mathbf{p}$ is *closed under cross–products* if and only if for all Ω–partitions $\mathcal{W}$, $\mathcal{W}^*$ and $\mathcal{W}'$ and for all $p \in \mathbf{p}(\mathcal{W})$, $p^* \in \mathbf{p}(\mathcal{W}^*)$, if $W \cap W^* \in \mathcal{W}'$ for all $W \in \mathcal{W}$, $W^* \in \mathcal{W}^*$,[10] then there exists a $q \in \mathbf{p}(\mathcal{W}')$ that refines both p and p^*. We note that probabilities, capacities (cf. Ghirardato, 1997) as well as characteristic functions of possibility sets are cross–product refinable.

$\mathbf{u}$ 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 particular example of a utility type is the class of real–valued one–dimensional utilities $u : \mathcal{C} \to \mathbb{R}$ with the canonical binary relation $\geq$ on $\mathbb{R}$. Since we make no assumptions about the set Z, 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). Moreover, since we do not assume that the scale relation defined on Z is an ordering, we can represent any binary preference relation on consequences, even relations that are not orderings, as a utility in our general sense. In particular, we can identify the set $\mathbf{r}$ of all r such that there is some Γ–partition $\mathcal{C}$ on which r is a binary relation with a utility type.

Our concept of a decision rule requires mainly that acts are evaluated by the utilities of their consequences and by beliefs about the states in which the consequences are achieved (properties (b) and (c)). For any belief type $\mathbf{p}$ and any utility type $\mathbf{u}$, $\langle G, \succeq \rangle$ 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 $f_1, f_2 \in \Phi_{\mathcal{W}, \mathcal{C}}$, the quadruple $\langle p, u, f_1, f_2 \rangle$ is in the domain of G and $\succeq$ is a reflexive binary relation on the range of G;

(b) for any grainings $\langle \mathcal{W}, \mathcal{C} \rangle$ and $\langle \mathcal{W}', \mathcal{C}' \rangle$ with $\mathcal{W} \subseteq [\mathcal{W}']$ and $\mathcal{C} \subseteq [\mathcal{C}']$, any $p \in \mathbf{p}(\mathcal{W})$, any $p' \in \mathbf{p}(\mathcal{W}')$, any $u \in \mathbf{u}(\mathcal{C})$, any $u' \in \mathbf{u}(\mathcal{C}')$, any $f_1, f_2 \in \Phi_{\mathcal{W}, \mathcal{C}} \cap \Phi_{\mathcal{W}', \mathcal{C}'}$: If p' refines p and $u \circ \mathcal{C} \circ f_k = u' \circ \mathcal{C}' \circ f_k$ $(k = 1, 2)$,[11] then $G(p, u, f_1, f_2) = G(p', u', f_1, f_2)$;

---

[7]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.

[8]p is a *probability on* $\mathcal{W}$ if and only if $p : [\mathcal{W}] \to [0, 1]$ $p(A) = \sum_{W \subseteq A} p(W)$ for any $A \in [\mathcal{W}]$ and $p(\Omega) = 1$.

[9]$k : [\mathcal{W}] \to \{0, 1\}$ *represents* $K \subseteq \mathcal{W}$ if and only if, for any $A \in [\mathcal{W}]$, $k(A) = 0$ if $A \subseteq -\bigcup K$ and $k(A) = 1$ else.

[10]By the definition of a partition (p. 5), this implies that $W \cap W^* \neq \varnothing$ for any $W \in \mathcal{W}$ and $W^* \in \mathcal{W}^*$.

[11]Recall that we write $\mathcal{C}(c)$ for the partition element of $\mathcal{C}$ containing $c \in \Gamma$ and that we can identify a partition with a particular type of function with domain Γ (for any Γ–partition $\mathcal{C}$).

(c) for any grainings $\langle \mathcal{W}, \mathcal{C} \rangle$ and $\langle \mathcal{W}', \mathcal{C}' \rangle$, any 1–1 mappings $\phi : \mathcal{W} \to \mathcal{W}'$ and $\psi : \mathcal{C} \to \mathcal{C}'$, any $p \in \mathbf{p}(\mathcal{W})$, any $p' \in \mathbf{p}(\mathcal{W}')$, any $u \in \mathbf{u}(\mathcal{C})$, any $u' \in \mathbf{u}(\mathcal{C}')$, any $f_1, f_2 \in \Phi_{\mathcal{W},\mathcal{C}}$ and any $f_1', f_2' \in \Phi_{\mathcal{W}',\mathcal{C}'}$: If $p = p' \circ \bar{\phi}$, if $u = u' \circ \psi$, if F_k is the $\langle \mathcal{W}, \mathcal{C} \rangle$–graining of f_k, and if $[\psi \circ F_k \circ \phi^{-1}]$ is the $\langle \mathcal{W}', \mathcal{C}' \rangle$–graining of f_k' (for $k = 1, 2$), then $G(p, u, f_1, f_2) = G(p', u', f_1', f_2')$;

(d) for any graining $\langle \mathcal{W}, \mathcal{C} \rangle$, any $p \in \mathbf{p}(\mathcal{W})$, any $f_1, f_2 \in \Phi_{\mathcal{W},\mathcal{C}}$ and any $C, D \in \mathcal{C}$ with $C \neq D$, $[\mathcal{C} \circ f_1](.) = C$ and $[\mathcal{C} \circ f_2](.) = D$, there are $u, u' \in \mathbf{u}(\mathcal{C})$ such that $G(p, u, f_1, f_2) \succ G(p, u, f_2, f_1)$ but $G(p, u', f_2, f_1) \succ G(p, u', f_1, f_2)$.

In a situation described by the antecedents of clause (b), the refined description of the decision situation is in a sense empty since the fine–grained descriptions of any consequences of f_1, f_2 do not differ from their coarse grained–descriptions in terms of utility. We therefore refer to property (b) by saying that G is *invariant under empty refinements*. We refer to property (c) by saying that G is *invariant under relabelling*. We interpret $G(p, u, f_1, f_2) \succeq G(p, u, f_2, f_1)$ to mean that f_1 is weakly preferred to f_2 by individuals with decision rule G, beliefs p and utilities u. We therefore refer to property (d) by saying the G is *non–trivial*. 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 $\langle p, u, f, g \rangle, \langle p, u, f, h \rangle$ 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 $f_1, f_2, f_1', f_2' \in \Phi_{\mathcal{W},\mathcal{C}}$: If $[u \circ \mathcal{C} \circ f_k](\omega) = [u' \circ \mathcal{C} \circ f_k'](\omega)$ for all $\omega \in -A$ and $k = 1, 2$, then $G(p, u, f_1, f_2) = G(p, u, f_1', f_2')$. Any A is p, G–*one* if and only if $-A$ is p, G–null.

Finally, $M = \langle \mathcal{W}, \mathbf{p}, p, \mathcal{C}, \mathbf{u}, u, \mathcal{F}, G, \succeq \rangle$ 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, $u \in \mathbf{u}(\mathcal{C})$, $\mathcal{F} \subseteq \Phi_{\mathcal{W},\mathcal{C}}$ is unambiguous w.r.t. $\langle \mathcal{W}, \mathcal{C} \rangle$, and $\langle G, \succeq \rangle$ is a decision rule for $\langle \mathbf{p}, \mathbf{u} \rangle$. We define $\mathcal{W}_M$, $\mathbf{p}_M$, p_M, $\mathcal{C}_M$, $\mathbf{u}_M$, u_M, $\mathcal{F}_M$, G_M and $\succeq_M$ to be the entities such that $M = \langle \mathcal{W}_M, \mathbf{p}_M, p_M, \mathcal{C}_M, \mathbf{u}_M, u_M, \mathcal{F}_M, G_M, \succeq_M \rangle$. 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*

8

evaluation function associated with M is the function $V_M : \mathcal{F}_M \to \mathrm{range}(G)$ such that $V_M(f) := V_M(f, f)$ for all $f \in \mathcal{F}_M$. The *preference relation associated with* M is the binary relation $\succeq_M \subseteq \mathcal{F}_M^2$ such that $f \succeq_M g$ if and only if $V_M(f, g) \succeq_M V_M(g, f)$ (for all $f, g \in \mathcal{F}_M$). We say that M *allows strong beliefs* if and only if there exists some Ω–partition $\mathcal{W} = \{W_1, W_2\}$ and some $p \in \mathbf{p}_M(\mathcal{W})$ such that W_1 is p, G_M–one.

We will have to assume that all individuals are described by the same decision–theoretic model because robust aggregation forms empathetic evaluations by mixing different individuals' beliefs and utilities. Let $I \geq 2$ ($I \in \mathbb{N}$) be a fixed number of individuals. Let $\mathbf{G}(I)$ be the set of all vectors $\langle M_i \rangle$ of decision–theoretic models such that $\mathcal{W}_{M_i} = \mathcal{W}_{M_j}$, $\mathbf{p}_{M_i} = \mathbf{p}_{M_j}$, $\mathcal{C}_{M_i} = \mathcal{C}_{M_j}$, $\mathbf{u}_{M_i} = \mathbf{u}_{M_j}$, $\mathcal{F}_{M_i} = \mathcal{F}_{M_j}$, $G_{M_i} = G_{M_j}$ and $\succeq_{M_i} = \succeq_{M_j}$ (for all $1 \leq i, j \leq I$). 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 $\mathcal{F}_{\langle M_i \rangle} := \mathcal{F}_{M_1}$. To simplify our notation, we henceforth hold the individuals' belief type, utility type and decision rule fixed and assume that, for any $\langle M_i \rangle \in \mathbf{G}(I)$, we have $\mathbf{p} = \mathbf{p}_{M_i}$, $\mathbf{u} = \mathbf{u}_{M_i}$, $G = G_{M_i}$ and $\succeq = \succeq_{M_i}$ (for each $1 \leq i \leq I$). In what follows, we therefore drop any reference to these components when we specify a vector of models in $\mathbf{G}(I)$.

Applications: Decision theory

We list several examples of decision–theoretic models in the sense of our definition. All of these models allow strong beliefs and possess belief types that are cross–product re-finable. 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 of $u_i \circ F$ w.r.t. p_i defined by $E(p_i, u_i \circ F) := \sum_{W \in \mathcal{W}} p_i(W) \cdot [u_i \circ F](W)$. Since our proofs construct only models with finite partitions of states, we avoid the issue of merely finitely vs. countably additive probabilities.[12]

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).

---

[12]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.

Choquet–expected utility. Beliefs are represented by capacities p_i on $\mathcal{W}$ and agents are endowed with real–valued one–dimensional utilities u_i on $\mathcal{C}$. Agents maximize the Choquet–expectation of $u \circ F$ w.r.t. p defined by $C(p, u \circ F) := \int_0^\infty p(u \circ F \geq x)\, dx + \int_{-\infty}^0 [p(u \circ F \geq x) - 1]\, dx$ (Gilboa, 1987, Schmeidler, 1989; cf. footnote **7**). 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 $\pi_i : [0,1] \rightarrow [0,1]$ 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 $\mathbf{1} : C \mapsto 1$, $\mathbf{1} : \mathcal{C} \rightarrow \mathbb{R}$ (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 $u_i, v_i : \mathcal{C} \rightarrow \mathbb{R}$. 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' $M_i : \mathbb{R}^2 \rightarrow \mathbb{R}$. 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 $V_i(f_1, f_2) \geq V_i(f_2, f_1)$.

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.**13** Any $A \in [\mathcal{W}]$ is null w.r.t. a possibility set K and decision–theoretic leximin (leximax) if and only if $A \subseteq - \bigcup K$.

13We 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 $\sigma, \tau : \{1, \ldots, M\} \rightarrow \mathcal{W}$ 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 $\sigma, \tau : \{1, \ldots, M\} \rightarrow \mathcal{W}$ 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 $1 \leq m \leq M$.

3 Ex Ante Aggregation

Ex ante aggregation rules

A social choice rule yields a social choice function based on the individuals' decision–theoretic models. Any $\mathcal{F} \subseteq \Phi$ is called a *set of possible distinguishable acts*. In our 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). 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 $M = \langle \mathcal{F}_M, V_M \rangle$. 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 $M = \langle \mathcal{F}_M, C_M \rangle$.

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).

Refinements

Consider an evaluation model M. Recall from our discussion of coarse–grained acts (p. 5) that fine–grained models can distinguish at least as many acts as coarse–grained models. In a refined model M' we therefore require that $\mathcal{F}_M \subseteq \mathcal{F}_{M'}$ and that evaluations of coarse–grained acts remain unaffected by the refinement. 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) $V = V'|\mathcal{F}$. 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 $X \subseteq \mathcal{F}$.

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 3.1 *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.

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 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.[14]

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.[15] 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.

---

[14]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.

[15]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 $C_{S(\langle V_i \rangle)}(X) = C_{S(\langle V_i' \rangle)}(X)$.

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:

- 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 3.2 *(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 robust social choice rules. Any non–trivial robust social choice rules is instable under refinements. Instabilities can therefore be counted as a defect of the robust 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.

4   Robust Aggregation

Robust social choice rules

The idea behind robust aggregation is to consider not only the individuals' actual evaluations but also what we have called their 'empathetic' evaluations. Empathetic evaluations are the evaluations that real individuals would have if they were to keep their own utilities but adopted another individual's beliefs. For $1 \leq k, l \leq I$ and $\langle M_i \rangle \in \mathbf{G}(I)$ we define M_{ij} as the evaluation model such that $\mathcal{F}_{M_{ij}} = \mathcal{F}_{\langle M_i \rangle}$ and $V_{M_{ij}}(f, g) = G(p_{M_i}, u_{M_j}, f, g)$ for all $f, g \in \mathcal{F}_{\langle M_i \rangle}$. Note that M_i denotes a decision–theoretic model while M_{ii} denotes the evaluation model with the evaluation function of a hypothetical individual with i's beliefs and j's utilities. There are I^2 such hypothetical individuals. For any $\langle M_i \rangle \in \mathbf{G}(I)$, we write $f \succeq_{M_{ij}} g$ if and only if $V_{M_{ij}}(f, g) \succeq V_{M_{ij}}(g, f)$ (for all $1 \leq i, j \leq I$ and $f, g \in \mathcal{F}_{\langle M_i \rangle}$).

Definition 4.1 *S is a* robust social choice rule *if and only if there is some non–empty set $\mathcal{G} \subseteq \mathbf{G}(I)$ such that $S : \mathcal{G} \to \mathbf{C}$ and there exists some ex ante social choice rule $S' : \mathbf{V}(I^2) \to \mathbf{C}$ such that $S(\langle M_i \rangle) = S'(\langle M_{ij} \rangle)$ for all $\langle M_i \rangle \in \mathcal{G}$.*

Any S' satisfying the above condition is *an ex ante rule associated with S*.[16] We now formulate several additional properties of a robust social choice rule S. We start with a non–triviality condition that requires a robust rule not collapse back into ex ante aggregation. S is *non–trivial* if and only if there are $\langle M_i \rangle, \langle M_i^* \rangle$ in the domain of S such that $\mathcal{W}_{\langle M_i \rangle} = \mathcal{W}_{\langle M_i^* \rangle}$, $\mathcal{C}_{\langle M_i \rangle} = \mathcal{C}_{\langle M_i^* \rangle}$, $\mathcal{F}_{\langle M_i \rangle} = \mathcal{F}_{\langle M_i^* \rangle}$, $\langle V_{M_i} \rangle = \langle V_{M_i^*} \rangle$ but $S(\langle M_i \rangle) \neq S(\langle M_i^* \rangle)$. S is *independent of irrelevant alternatives (IIA)* if and only if there exists an ex ante social choice rule associated with S that is IIA. S is *Pareto optimal* if and only if, for all $\langle M_i \rangle$ in the domain of S and all $f, g \in \mathcal{F}_{\langle M_i \rangle}$, if $f \succeq_{M_{ij}} g$ for all $1 \leq i, j \leq I$ but $f \succ_{M_{kl}} g$ for some $1 \leq k, l \leq I$, then $C_{S(\langle M_i \rangle)}(\{f, g\}) = \{f\}$. Clearly, if S is Pareto optimal, then it is non–trivial. *S has a wide domain* if and only if, for any graining $\langle \mathcal{W}, \mathcal{C} \rangle$, any $\langle p_i \rangle \in \mathbf{p}(\mathcal{W})^I$, any $\langle u_i \rangle \in \mathbf{u}(\mathcal{C})^I$ 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'. The final clause (*) of the definition allows the underlying structure of acts to force the inclusion of certain acts into the choice set $\mathcal{F}_{\langle M_i \rangle}$. A robust 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. Replacing condition (*) with (**) $\mathcal{F}_{\langle M_i \rangle} = \mathcal{F}$, we say that *$S$ has an unrestricted domain* if and only if $\mathbf{G}(I)$ is the domain of S. Finally, we ensure that S is invariant under relabelling the components of the individuals' decision theoretic models. So far, our definition of robust rules allows group preferences to depend not only on the values of the hypothetical individuals' evaluations of acts, but also on the acts themselves. We will assume that a simple relabelling, or renaming, of acts should not yield different social preferences. For any $\langle M_i \rangle, \langle M_i' \rangle \in \mathbf{G}(I)$, we say that

---

[16]There are generally several ex ante rule associated with S when S has a restricted domain.

$\langle M_i' \rangle$ is the ϕ, ψ, χ–*relabelling* of $\langle M_i \rangle$ if and only if ϕ, ψ and χ are 1–1 mappings with $\phi : \mathcal{W}_{\langle M_i \rangle} \to \mathcal{W}_{\langle M_i' \rangle}$, $\psi : \mathcal{C}_{\langle M_i \rangle} \to \mathcal{C}_{\langle M_i' \rangle}$, $\chi : \mathcal{F}_{\langle M_i \rangle} \to \mathcal{F}_{\langle M_i' \rangle}$ such that $\langle p_{M_i} \rangle = \langle p_{M_i'} \circ \bar{\phi} \rangle$, $\langle u_{M_i} \rangle = \langle u_{M_i'} \circ \psi \rangle$, $V_{M_{ij}}(f, g) = V_{M_{ij}'}(\chi(f), \chi(g))$ for all $f, g \in \mathcal{F}_{\langle M_i \rangle}$ and all $1 \leq i, j \leq I$, and, for any $f_1, f_2 \in \mathcal{F}_{\langle M_i \rangle}$ and any $f_1', f_2' \in \mathcal{F}_{\langle M_i' \rangle}$, if F_k is the $\langle \mathcal{W}_{\langle M_i \rangle}, \mathcal{C}_{\langle M_i \rangle} \rangle$–graining of f_k, and F_k' is the $\langle \mathcal{W}_{\langle M_i' \rangle}, \mathcal{C}_{\langle M_i \rangle} \rangle$–graining of $\chi(f_k)$, then $F_k' = \psi \circ F_k \circ \phi^{-1}$ (for $k = 1, 2$). S is *invariant under relabelling* if and only if, for any $\langle M_i \rangle, \langle M_i' \rangle \in \mathbf{G}(I)$, if $\langle M_i' \rangle$ is the ϕ, ψ, χ–relabelling of $\langle M_i \rangle$ for some ϕ, ψ, χ, then $\langle M_i \rangle$ is in the domain of S iff $\langle M_i' \rangle$ is in the domain of S and $f \in [S(\langle M_i \rangle)](X)$ iff $\chi(f) \in [S(\langle M_i' \rangle)](\{\chi(g) | g \in X\})$ for any $f \in \mathcal{F}_{\langle M_i \rangle}$ and $X \subseteq \mathcal{F}_{\langle M_i \rangle}$.

Instability

For the sake of a strong instability result, we now define a demanding notion of a refinement for individual models. For the group preference model, we use the same notion as in Section 3. If it were not for our interest in group preference reversals, we might wish to speak of an individual refinement already when individual preferences are left unchanged. To indulge robust aggregation, we ask moreover that both individual evaluation functions and beliefs remain unchanged in a refinement. For any vectors of decision–theoretic models $\langle M_i \rangle, \langle M_i' \rangle \in \mathbf{G}(I)$, we say that $\langle M_i' \rangle$ *refines* $\langle M_i \rangle$ if and only if, for each $1 \leq i \leq I$, (1) $\mathcal{W}_{\langle M_i \rangle} \subseteq [\mathcal{W}_{\langle M_i' \rangle}]$, (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) $p_{M_i'}$ refines 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.[17] A robust 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 $\langle M_i \rangle$.

Theorem 4.2 *Suppose that S is a robust social choice rule that is invariant under relabelling and has a wide domain and suppose that the belief types of the models in its domain are closed under cross–products.*

If S is non–trivial, then S cannot be stable under refinements.

In our view, this is a detrimental result for robust aggregation. We recall that we return to the ex ante mode of aggregation if our robust social choice rule is trivial. Hence, any stable robust social choice rule must collapse back into the ex ante mode. Ex ante social choice rules are stable whenever they are IIA. We can therefore conclude from Theorem 3.2 and Theorem 4.2 that the robust 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.

Theorem 4.3 *Suppose that S is a robust social choice rule that is invariant under relabelling and*

---

[17]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 $f, g \in \mathcal{F}_{\langle M_i \rangle}$.

(1) S is non–trivial and IIA, has a wide domain and the models in its domain have a monadic decision rule and a belief types closed under cross–products, or

(2) S is Pareto optimal and has a wide domain and the models its domain allow strong beliefs.

Then there is an infinite sequence $\langle M_i^n \rangle_{n \in \mathbb{N}}$ of vectors of decision–theoretic models in the domain of S such that $\langle M_i^{n+1} \rangle$ refines $\langle M_i^n \rangle$ (for each $n \in \mathbb{N}$) and S leads to a sequence of group models $\langle S(\langle M_i^n \rangle) \rangle_{n \in \mathbb{N}}$ that, for some $f \in \mathcal{F}_{\langle M_i^0 \rangle}$, oscillates between choosing and not choosing f from some choice set $X \subseteq \mathcal{F}_{\langle M_i^0 \rangle}$ (i.e., there is some $X \subseteq \mathcal{F}_{\langle M_i^0 \rangle}$ and $f \in X$ with $f \in C_{S(\langle M_i^{2n} \rangle)}(X)$ and $f \notin C_{S(\langle M_i^{2n+1} \rangle)}(X)$ for all $n \in \mathbb{N}$).

The infinite sequence of group preference reversals in Theorem 4.3 is particularly disturbing. We are confronted with the possibility of pervasive instabilities that refuse to disappear as we build increasingly fine–grained individual models. In the robust mode, instabilities are there to stay. We emphasize that instabilities are not the artifact of a narrow class of decision theories. All decision–theoretic models listed above allow strong beliefs and have belief types closed under cross–products and therefore satisfy all the decision–theoretic conditions of Theorem 4.2 and Theorem 4.3.1–2. Regret theory and decision–theoretic leximin/leximax are the only decision theories on our list that have a non–monadic decision rule. Hence, Theorem 4.3.3 applies to all our examples except regret theory and decision–theoretic leximin/leximax. The current version of Theorem 4.3.1 uses strong beliefs that make some worlds in $\mathcal{W}$ null. We can work with beliefs that are less strong if, instead, we make more assumptions about the decision rule used.

Appendix A   Proofs of Theorems

Theorem 3.2 *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 $C_{S(\langle \mathcal{F}, V_i \rangle)}(X) = C_{S(\langle X, V_i | X \rangle)}(X) = C_{S(\langle \mathcal{F}', V_i' \rangle)}(X)$. $\square$

Recall in what follows that we have fixed the individuals' belief type $\mathbf{p}$, utility type $\mathbf{u}$ and decision rule $\langle G, \succeq \rangle$.

Theorem 4.2 and Theorem 4.3.1 *Preparation:* Suppose that S is a robust social choice rule that is invariant under relabelling and the belief types of all models in its domain are closed under cross–products. Suppose S is non–trivial. Hence, there are $\langle M_i^* \rangle, \langle M_i^\dagger \rangle$ in the domain of S such that $\mathcal{W} := \mathcal{W}_{\langle M_i^* \rangle} = \mathcal{W}_{\langle M_i^\dagger \rangle}$, $\mathcal{C} := \mathcal{C}_{\langle M_i^* \rangle} = \mathcal{C}_{\langle M_i^\dagger \rangle}$, $\mathcal{F} := \mathcal{F}_{\langle M_i^* \rangle} = \mathcal{F}_{\langle M_i^\dagger \rangle}$, $\langle V_{M_i^*} \rangle = \langle V_{M_i^\dagger} \rangle$, and for some $X \subseteq \mathcal{F}$ and some $g \in X$, $(+)$ $g \in C_{S(\langle M_i^* \rangle)}(X)$ but $g \notin C_{S(\langle M_i^\dagger \rangle)}(X)$. Let $L := |\mathcal{W}|$ and $M := |\mathcal{C}|$. Then there is an Ω–partition $\mathcal{W}'$ with $|\mathcal{W}'| = L^2$ each of whose elements is at least countably infinite. We enumerate the elements of $\mathcal{W}$ and $\mathcal{W}'$ such that $\mathcal{W} = \{W_l | 1 \leq l \leq L\}$ and $\mathcal{W}' = \{W_{l_1, l_2} | 1 \leq l_1, l_2 \leq L\}$. We define $\mathcal{W}^* := \{\bigcup \{W_{l,1}, \ldots, W_{l,L}\} | 1 \leq l \leq L\}$ and $\mathcal{W}^\dagger := \{\bigcup \{W_{1,l}, \ldots, W_{L,l}\} | 1 \leq l \leq L\}$. Obviously, $\mathcal{W}^*$ and $\mathcal{W}^\dagger$ are Ω–partitions and subsets of $[\mathcal{W}']$. We define $\phi^* : \mathcal{W} \to \mathcal{W}^*$ by $\phi^*(W_l) := \bigcup \{W_{l,1}, \ldots, W_{l,L}\}$ and $\phi^\dagger : \mathcal{W} \to \mathcal{W}^\dagger$ by $\phi^\dagger(W_l) := \bigcup \{W_{1,l}, \ldots, W_{L,l}\}$ $(1 \leq l \leq L)$. Obviously, ϕ^* and $\phi^\dagger$ are 1–1. There also exists a Γ–partition $\mathcal{C}'$ with $|\mathcal{C}'| = M^2$ each of whose elements is at least countably infinite. We enumerate the elements of $\mathcal{C}$ and $\mathcal{C}'$ such that $\mathcal{C} = \{C_m | 1 \leq m \leq M\}$ and $\mathcal{C}' = \{C_{m_1, m_2} | 1 \leq m_1, m_2 \leq M\}$. We define $\mathcal{C}^* := \{\bigcup \{C_{m,1}, \ldots, C_{m,M}\} | 1 \leq m \leq M\}$ and $\mathcal{C}^\dagger := \{\bigcup \{C_{1,m}, \ldots, C_{M,m}\} | 1 \leq m \leq M\}$. Obviously, $\mathcal{C}^*$ and $\mathcal{C}^\dagger$ are Γ–partitions and subsets of $[\mathcal{C}']$. We define $\psi^* : \mathcal{C} \to \mathcal{C}^*$ by $\psi^*(C_m) := \bigcup \{C_{m,1}, \ldots, C_{m,M}\}$ and $\psi^\dagger : \mathcal{C} \to \mathcal{C}^\dagger$ by $\psi^\dagger(C_m) := \bigcup \{C_{1,m}, \ldots, C_{M,m}\}$ $(1 \leq m \leq M)$. Obviously, ψ^* and $\psi^\dagger$ are 1–1. Trivially, $W^* \cap W^\dagger \in \mathcal{W}'$ for every $W^* \in \mathcal{W}^*$, $W^\dagger \in \mathcal{W}^\dagger$ and $C^* \cap C^\dagger \in \mathcal{C}'$ for every $C^* \in \mathcal{C}^*$, $C^\dagger \in \mathcal{C}^\dagger$. We recall that every $C_{m_1, m_2} \in \mathcal{C}'$ is at least countably infinite $(1 \leq m_1, m_2 \leq M)$. We can hence enumerate the elements of some countably infinite subset of C_{m_1, m_2} $(1 \leq m_1, m_2 \leq M)$ in a sequence $c_{m_1, m_2}^0, c_{m_1, m_2}^1, \ldots, c_{m_1, m_2}^n, \ldots$ (where $c_{m_1, m_2}^n \neq c_{m_1, m_2}^{n'}$ for any $n, n' \in \mathbb{N}$ with $n \neq n'$). Let $\chi : \mathcal{F} \to \Phi_{\mathcal{W}', \mathcal{C}'}$ be the function such that, for any $f \in \mathcal{F}$, $f' = \chi(f)$ is the act defined by: For any $1 \leq l_1, l_2 \leq L$ and any $1 \leq m_1, m_2 \leq M$, if $f(\omega) \in C_{m_1}$ for all $\omega \in W_{l_1}$ and $f(\omega) \in C_{m_2}$ for all $\omega \in W_{l_2}$, then $f'(\omega) := c_{m_1, m_2}^0$ for all $\omega \in W_{l_1, l_2}$. Let $\mathcal{F}' := \chi(\mathcal{F})$ be the image of $\mathcal{F}$ under χ. Since $\mathcal{F}$ is unambiguous w.r.t. $\langle \mathcal{W}, \mathcal{C} \rangle$, $\chi : \mathcal{F} \to \mathcal{F}'$ is a 1–1 mapping. Note that $\mathcal{F}' \subseteq \Phi_{\mathcal{W}^*, \mathcal{C}^*}$ and $\mathcal{F}' \subseteq \Phi_{\mathcal{W}^\dagger, \mathcal{C}^\dagger}$. Moreover, for any $f \in \mathcal{F}$ with a $\langle \mathcal{W}, \mathcal{C} \rangle$–graining F, we find that $[\psi^* \circ F \circ (\phi^*)^{-1}]$ is the $\langle \mathcal{W}^*, \mathcal{C}^* \rangle$–graining of $\chi(f)$ and $[\psi^\dagger \circ F \circ (\phi^\dagger)^{-1}]$ is the $\langle \mathcal{W}^\dagger, \mathcal{C}^\dagger \rangle$–graining of $\chi(f)$.

Since $\mathbf{p}$ is closed under relabelling, $p_i^{**} := p_{M_i^*} \circ (\bar{\phi}^*)^{-1} \in \mathbf{p}(\mathcal{W}^*)$ and $p_i^{\dagger\dagger} := p_{M_i^\dagger} \circ (\bar{\phi}^\dagger)^{-1} \in \mathbf{p}(\mathcal{W}^\dagger)$ (where we use the definition of the extended mappings $(\bar{\phi}^*)$ and $(\bar{\phi}^\dagger)$ on p. 6). Let $u_i^{**} := u_{M_i^*} \circ (\psi^*)^{-1}$ and $u_i^{\dagger\dagger} := u_{M_i^\dagger} \circ (\psi^\dagger)^{-1}$. We write $M_i^{**} := \langle \mathcal{W}^*, p_i^{**}, \mathcal{C}^*, u_i^{**}, \mathcal{F}' \rangle$ and $M_i^{\dagger\dagger} := \langle \mathcal{W}^\dagger, p_i^{\dagger\dagger}, \mathcal{C}^\dagger, u_i^{\dagger\dagger}, \mathcal{F}' \rangle$ $(1 \leq i \leq I)$. We write $V_{ij}^*(f_1, f_2) := G(p_i^*, u_j^*, f_1, f_2)$, $V_{ij}^\dagger(f_1, f_2) := G(p_i^\dagger, u_j^\dagger, f_1, f_2)$ and $V_{ij}^{**}(f_1', f_2') := G(p_i^{**}, u_j^{**}, f_1', f_2')$, $V_{ij}^{\dagger\dagger}(f_1', f_2') := G(p_i^{\dagger\dagger}, u_j^{\dagger\dagger}, f_1', f_2')$ for all $f_1, f_2 \in \mathcal{F}$, $f_1', f_2' \in \mathcal{F}'$ and all $1 \leq i, j \leq I$. Since G is invariant under relabelling, we have $(*)$ $V_{ij}^*(f_1, f_2) = V_{ij}^{**}(\chi(f_1), \chi(f_2))$ and $V_{ij}^\dagger(f_1, f_2) = V_{ij}^{\dagger\dagger}(\chi(f_1), \chi(f_2))$ for all $f_1, f_2 \in \mathcal{F}$ and all $1 \leq i, j \leq I$. Since the S is invariant under relabelling, we then know that $\langle M_i^{**} \rangle, \langle M_i^{\dagger\dagger} \rangle$ are in the domain of S and there are $X' \subseteq \mathcal{F}'$ and $g' \in X'$ such that $(+')$ $g' \in C_{S(\langle M_i^{**} \rangle)}(X')$ but $g' \notin C_{S(\langle M_i^{\dagger\dagger} \rangle)}(X')$ (from $(+)$). Since $\mathbf{p}$ is closed under cross–products, there exist $p_i' \in \mathbf{p}(\mathcal{W}')$ such that p_i' refines p_i^{**} and $p_i^{\dagger\dagger}$ (for all $1 \leq i \leq I$). For every $1 \leq i \leq I$, let the functions $v_i, w_i \in \mathbf{u}(\mathcal{C}')$ be defined by $v_i(C_{m_1, m_2}) := u_i^{**}(\bigcup \{C_{m_1, 1}, \ldots, C_{m_1, M}\}) = u_{M_i^*}(C_{m_1})$ and $w_i(C_{m_1, m_2}) := u_i^{\dagger\dagger}(\bigcup \{C_{1, m_2}, \ldots, C_{M, m_2}\}) = u_{M_i^\dagger}(C_{m_2})$ for all $1 \leq m_1, m_2 \leq M$. We write $V_{ij}'(f_1, f_2) := G(p_i', v_j, f_1, f_2)$ and $V_{ij}''(f_1, f_2) := G(p_i', w_j, f_1, f_2)$ (for $f_1, f_2 \in \mathcal{F}_{\langle M_i' \rangle}$ and $1 \leq i, j \leq I$). Recall that we write $\mathcal{C}(c)$ for the partition element of $\mathcal{C}$ containing $c \in \Gamma$ and that we can identify a partition $\mathcal{C}$ with a

function from Γ onto $\mathcal{C}$ (for any Γ–partition $\mathcal{C}$). Since $v_j \circ \mathcal{C}' = u_j^{**} \circ \mathcal{C}^*$ and $w_j \circ \mathcal{C}' = u_j^{\dagger\dagger} \circ \mathcal{C}^\dagger$ and since G is invariant under empty refinements, we find, for any $f_1, f_2 \in \mathcal{F}'$, (**) $V'_{ij}(f_1, f_2) = V_{ij}^{**}(f_1, f_2)$ and $V''_{ij}(f_1, f_2) = V_{ij}^{\dagger\dagger}(f_1, f_2)$ ($1 \le i, j \le I$). We note that (***) $V'_{ii}(f_1, f_2) =_{\text{(by (**))}} V_{M_i^{**}}(f_1, f_2) =_{\text{(by (*))}} V_{M_i^*}(\chi^{-1}(f_1), \chi^{-1}(f_2)) =_{\text{(by assumption)}} V_{M_i^\dagger}(\chi^{-1}(f_1), \chi^{-1}(f_2)) =_{\text{(by (*))}} V_{M_i^{\dagger\dagger}}(f_1, f_2) =_{\text{(by (**))}} V''_{ii}(f_1, f_2)$ (for all $1 \le i \le I$ and all $f_1, f_2 \in \mathcal{F}'$). This completes our preparation and we now turn to finishing the proofs for our two separate propositions.

Proof of Theorem 4.2: If S has a wide domain, there exists $\langle M'_i \rangle$ in the domain of S such that (for all $1 \le i \le I$) $\mathcal{W}_{\langle M'_i \rangle} = \mathcal{W}'$, $p_{M'_i} = p'_i$, $\mathcal{C}_{\langle M'_i \rangle} = \mathcal{C}'$, $u_{M'_i} = v_i$ and $\mathcal{F}' \subseteq \mathcal{F}_{\langle M'_i \rangle}$. By (***), $\langle M'_i \rangle$ refines both $\langle M_i^{**} \rangle$ and $\langle M_i^{\dagger\dagger} \rangle$. Suppose S were stable and non–trivial. This would require that $g' \in C_{S(\langle M'_i \rangle)}(X')$ and $g' \notin C_{S(\langle M'_i \rangle)}(X')$ (by (+')). Contradiction! Hence, S is not stable if it is non–trivial. $\square$

Proof of Theorem 4.3.1: Suppose that S has a wide domain and all models in its domain have a monadic decision rule. We recall that we have enumerated the elements of some countably infinite subset of C_{m_1,m_2} ($1 \le m_1, m_2 \le M$) in a sequence $c_{m_1,m_2}^0, c_{m_1,m_2}^1, \ldots, c_{m_1,m_2}^n, \ldots$ (where $c_{m_1,m_2}^n \ne c_{m_1,m_2}^{n'}$ for any $n, n' \in \mathbb{N}$ with $n \ne n'$). We now define a sequence $\langle \mathcal{C}_{m_1,m_2}^n \rangle$ of partitions of C_{m_1,m_2} by $\mathcal{C}_{m_1,m_2}^n = \{\{c_{m_1,m_2}^1\}, \ldots, \{c_{m_1,m_2}^n\}, C_{m_1,m_2} - \{c_{m_1,m_2}^1, \ldots, c_{m_1,m_2}^n\}\}$ for all $n \in \mathbb{N}^+$. We write $C_{m_1,m_2}^0 := C_{m_1,m_2}$ and $C_{m_1,m_2}^n := C_{m_1,m_2} - \{c_{m_1,m_2}^1, \ldots, c_{m_1,m_2}^n\}$ for all $n \in \mathbb{N}^+$. Let $\mathcal{C}^0 := \mathcal{C}' = \bigcup\{\mathcal{C}_{m_1,m_2} | 1 \le m_1, m_2 \le M\}$ and $\mathcal{C}^n := \bigcup\{\mathcal{C}_{m_1,m_2}^n | 1 \le m_1, m_2 \le M\}$ for all $n \in \mathbb{N}^+$. We note that $c_{m_1,m_2}^n \in C_{m_1,m_2}^n$ ($1 \le m_1, m_2 \le M$, $n \in \mathbb{N}$). For any $n \in \mathbb{N}$, let $\psi^n : \mathcal{C}^n \to \mathcal{C}^{n+1}$ be defined by $\psi^n(C_{m_1,m_2}^n) := \{c_{m_1,m_2}^{n+1}\}$ and by $\psi^n(\{c_{m_1,m_2}^{n'}\}) := \{c_{m_1,m_2}^{n'}\}$ for all $1 \le n' \le n$ and all $1 \le m_1, m_2 \le M$. 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 $W \in \mathcal{W}'$ there exists some $1 \le m_1, m_2 \le M$ and some $1 \le n' \le n+1$ such that $h(\omega) = c_{m_1,m_2}^{n'}$ for all $\omega \in W$. Since $\mathcal{F}_\dagger^n$ is unambiguous w.r.t. $\langle \mathcal{W}', \mathcal{C}^n \rangle$, we know that, for every $f' \in \mathcal{F}'$, there is exactly one act $f_\dagger^n \in \mathcal{F}_\dagger^n$ that has the same $\langle \mathcal{W}', \mathcal{C}^n \rangle$–graining as f'. Let $\mathcal{F}^n$ be the set that consists of all acts in $\mathcal{F}'$ and those acts in $\mathcal{F}_\dagger^n$ that do not have the same $\langle \mathcal{W}', \mathcal{C}^n \rangle$–graining as some act in $\mathcal{F}'$ (for any $n \in \mathbb{N}$). Note that $\mathcal{F}^n \subseteq \Phi_{\mathcal{W}', \mathcal{C}^{n+1}}$ for any $n \in \mathbb{N}$. 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. $\langle \mathcal{W}', \mathcal{C}^n \rangle$ (for any $n \in \mathbb{N}$). We now define a sequence $\langle u_1^n, \ldots, u_I^n \rangle_{n \in \mathbb{N}}$ of vectors of individual utility functions such that $u_i^n \in \mathbf{u}(\mathcal{C}^n)$ for each $n \in \mathbb{N}$ and $1 \le i \le I$. For all even $n \in \mathbb{N}$, let $u_i^n(C_{m_1,m_2}^n) := v_i(C_{m_1,m_2})$ and, for all odd $n \in \mathbb{N}$, let $u_i^n(C_{m_1,m_2}^n) := w_i(C_{m_1,m_2})$ (for all $1 \le i \le I$, $1 \le m_1, m_2 \le M$). For any $n \in \mathbb{N}^+$ and any $1 \le n' \le n$, if n' is even, let $u_i^n(\{c_{m_1,m_2}^{n'}\}) := w_i(C_{m_1,m_2})$ and, if n' is odd, let $u_i^n(\{c_{m_1,m_2}^{n'}\}) := v_i(C_{m_1,m_2})$ (for all $1 \le i \le I$, $1 \le m_1, m_2 \le M$). Note that $u_i^n = u_i^{n+1} \circ \psi^n$ ($1 \le i \le I$, $n \in \mathbb{N}$). 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 $1 \le i \le I$, $\mathcal{W}_{\langle M_i^n \rangle} = \mathcal{W}'$, $p_{M_i^n} = p'_i$, $\mathcal{C}_{\langle M_i^n \rangle} = \mathcal{C}^n$, $u_{M_i^n} = u_i^n$ and $\mathcal{F}_{\langle M_i^n \rangle} = \mathcal{F}^n$. Assuming that G is monadic, we write $V_{ij}^n(f) := G(p'_i, u_j^n, f, f)$ for all $f \in \mathcal{F}^n$, $1 \le i, j \le I$ and $n \in \mathbb{N}$.

Note that for any $f \in \mathcal{F}'$ and any $\omega \in \Omega$ there exist $1 \le m_1, m_2 \le M$ such that $f(\omega) = c_{m_1,m_2}^0$. Since $u_j^{2n} \circ \mathcal{C}^{2n} \circ f = v_j \circ \mathcal{C}' \circ f$ and $u_j^{2n+1} \circ \mathcal{C}^{2n+1} \circ f = w_j \circ \mathcal{C}' \circ f$, the fact that G is invariant under empty refinements implies that (****) $V_{M_{ij}^{2n}}(f) = V'_{ij}(f, f) =_{\text{(by (**))}} V_{M_{ij}^{**}}(f, f)$ and $V_{M_{ij}^{2n+1}}(f) = V''_{ij}(f, f) =_{\text{(by (**))}} V_{M_{ij}^{\dagger\dagger}}(f, f)$ (for any $f \in \mathcal{F}'$, $1 \le i, j \le I$, $n \in \mathbb{N}$). We now show that M_i^{n+1} refines M_i^n, for all $n \in \mathbb{N}$ and $1 \le i \le I$. Conditions 1–4 of the definition of a refinement (p. 16) are trivially satisfied. We need to show that $V_{M_i^n}(f) = V_{M_i^{n+1}}(f)$ for all $f \in \mathcal{F}^n$, $1 \le i \le I$, $n \in \mathbb{N}$. We distinguish two cases. Case 1: $f \in \mathcal{F}'$. By (****), $V_{M_i^{2n}}(f) = V''_{ii}(f, f) =_{\text{(by (***))}} V''_{ii}(f, f) = V_{M_i^{2n+1}}(f)$ (for any $1 \le i \le I$, $n \in \mathbb{N}$). Case 2: $f \in \mathcal{F}^n - \mathcal{F}'$. Then $u_i^n \circ \mathcal{C}^n \circ f = u_i^{n+1} \circ \psi^n \circ \mathcal{C}^n \circ f = u_i^{n+1} \circ \mathcal{C}^{n+1} \circ f$ which yields the claim since G is invariant under empty refinements (for any $1 \le i, j \le I$, $n \in \mathbb{N}$). Hence, all conditions for a refinement are satisfied. If S is IIA and non–trivial, we finally obtain (by (****)) $C_{S(\langle M_i^{2n} \rangle)}(Y) = C_{S(\langle M_i^{**} \rangle)}(Y)$ and $C_{S(\langle M_i^{2n+1} \rangle)}(Y) = C_{S(\langle M_i^{\dagger\dagger} \rangle)}(Y)$ (for all $Y \subseteq F'$ and $n \in \mathbb{N}$). By (+'), we then have $g' \in C_{S(\langle M_i^{2n} \rangle)}(X')$ but $g' \notin C_{S(\langle M_i^{2n+1} \rangle)}(X')$ (for all $n \in \mathbb{N}$). $\square$

Lemma A.1 *Suppose $\langle G, \succeq \rangle$ 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 $k = 1, 2$. We note that $f_1, f_2 \in \Phi_{\mathcal{W}, \mathcal{C}}$. 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 $k = 1, 2$. Since G is invariant under empty refinements, we then have $G(p, v, f_1, f_2) = G(p, u, f_1, f_2)$. Since A is p, G–one and since $[v \circ \mathcal{C}' \circ f_k](\omega) = [u' \circ \mathcal{C}' \circ f_k](\omega)$ for all $\omega \in A$ ($k = 1, 2$), we also have $G(p, v, f_1, f_2) = G(p, u', f_1, f_2)$. Hence, $G(p, u, f_1, f_2) = G(p, u', f_1, f_2)$. $\square$

Theorem 4.3.2 *Proof:* Suppose that S is a robust social choice rule with a wide domain and that the models in its domain allow strong beliefs. Suppose furthermore that S is Pareto optimal. Since the models in the domain of S allow strong beliefs, there is some Ω–partition $\mathcal{W} = \{W_1, W_2\}$ and some $p \in \mathbf{p}(\mathcal{W})$ such that W_1 is p, G–one. Since $\mathbf{p}$ is closed under relabelling and G is invariant under relabelling, there also exists some $q \in \mathbf{p}(\mathcal{W})$ such that W_2 is q, G–one. Define $\langle p_i \rangle \in \mathbf{p}(\mathcal{W})^I$ by $p_1 := p$ and $p_j := q$ ($2 \leq j \leq I$). We partition Γ into two at least countably infinite sets C_k ($k = 1, 2$). For each $k = 1, 2$, we enumerate the elements of some countably infinite subset of C_k in a sequence $c_k^0, c_k^1, \ldots, c_k^n, \ldots$ (where $c_k^n \neq c_k^{n'}$ for any $n, n' \in \mathbb{N}$ with $n \neq n'$). We then define a sequence $\langle \mathcal{C}_k^n \rangle$ of partitions of C_k by $\mathcal{C}_k^n = \{\{c_k^1\}, \ldots, \{c_k^{n+1}\}, C_k - \{c_k^1, \ldots, c_k^{n+1}\}\}$ for all $n \in \mathbb{N}$. We write $C_k^n := C_k - \{c_k^1, \ldots, c_k^{n+1}\}$ and $\mathcal{C}^n := \mathcal{C}_1^n \cup \mathcal{C}_2^n$ for all $n \in \mathbb{N}$. Note that $c_k^0 \in C_k^n$ for all $n \in \mathbb{N}$ and $k = 1, 2$. For any $n \in \mathbb{N}$, let $\psi^n : \mathcal{C}^n \to \mathcal{C}^{n+1}$ be defined by $\psi^n(C_k^n) := \{c_k^{n+2}\}$ and $\psi^n(\{c_k^{n'}\}) := \{c_k^{n'}\}$ for all $1 \leq n' \leq n + 1$ ($k = 1, 2$). Define $f \in \Phi$ as the act such that, for any $\omega \in W_1$, $f(\omega) = c_1^1$ and, for any $\omega \in W_2$, $f(\omega) = c_1^0$. Define $g \in \Phi$ as the act such that, for any $\omega \in W_1$, $g(\omega) = c_2^1$ and, for any $\omega \in W_2$, $g(\omega) = c_2^0$. 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 $l = 1, 2$ there exists some $k = 1, 2$ and some $1 \leq m \leq n + 2$ such that $h(\omega) = c_k^m$ for all $\omega \in W_l$. 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^n, h_2^n 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. $\langle \mathcal{W}, \mathcal{C}^n \rangle$.

Let $\mathcal{C}' := \{C_1, C_2\}$. We note that $f, g \in \Phi_{\mathcal{W}, \mathcal{C}'}$ and that $[\mathcal{C}' \circ f](.) = C_1$ while $[\mathcal{C}' \circ g](.) = C_2$. Since G is non–trivial, there are $u, u' \in \mathbf{u}(\mathcal{C}')$ such that (*) $G(q, u, f, g) \succ G(q, u, g, f)$ but $G(q, u', g, f) \succ G(q, u', f, g)$. Define $a := u(C_1)$, $b := u(C_2)$ and $a' := u'(C_1)$, $b' := u'(C_2)$. We now 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 $1 \leq i \leq I$, $u_i^n \in \mathbf{u}(\mathcal{C}^n)$. Let $u_j^n(.) := a$ for all $2 \leq j \leq I$, $n \in \mathbb{N}$. For all even $n \in \mathbb{N}$, let $u_1^n(C_1^n) := a$ and $u_1^n(C_2^n) := b$. For all odd $n \in \mathbb{N}^+$, let $u_1^n(C_1^n) := a'$ and $u_1^n(C_2^n) := b'$. Finally, let $u_1^0(\{c_1^1\}) := u_1^0(\{c_2^1\}) := a$ and, for any $n, n' \in \mathbb{N}$ with $1 \leq n' \leq n + 1$, let $u_1^{n+1}(\{c_k^{n'}\}) := u_1^n(\{c_k^{n'}\})$ and $u_1^{n+1}(\{c_k^{n+2}\}) := u_1^n(C_k^n)$ (for $k = 1, 2$). We have (**) $u_i^n = u_i^{n+1} \circ \psi^n$ 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. 16) are trivially satisfied. On the one hand, (**) yields for any $h \in \mathcal{F}^n$ with $h \neq f, g$ and for any $n \in \mathbb{N}$: $u_i^n \circ \mathcal{C}^n \circ h = u_i^{n+1} \circ \psi^n \circ \mathcal{C}^n \circ h = u_i^{n+1} \circ \mathcal{C}^{n+1} \circ h$ for all $1 \leq i \leq I$. On the other hand, we have $u_1^n(\{c_k^1\}) = u_1^{n'}(\{c_k^1\})$ for any $n, n' \in \mathbb{N}$ and $k = 1, 2$. Hence, $[u_1^n \circ \mathcal{C}^n \circ f](\omega) = [u_1^{n+1} \circ \mathcal{C}^{n+1} \circ f](\omega)$ for all $\omega \in W_1$ and all $n \in \mathbb{N}$. Trivially, $u_j^n \circ \mathcal{C}^n \circ f = u_j^{n+1} \circ \mathcal{C}^{n+1} \circ f$ for all $2 \leq j \leq I$ and all $n \in \mathbb{N}$. Similarly for g. Since W_1 is p_1, G–one, Lemma A.1 yields $G(p_i, u_i^n, h_1, h_2) = G(p_i, u_i^{n+1}, h_1, h_2)$ for all $h_1, h_2 \in \mathcal{F}^n$, $1 \leq i \leq I$, $n \in \mathbb{N}$. Hence, all conditions for a refinement are satisfied.

Recall that $\langle G, \succeq \rangle$ is the individuals' shared decision rule. We write $f_1 \succeq_{ij}^n f_2$:iff $G(p_i, u_j^n, f_1, f_2) \succeq G(p_i, u_j^n, f_2, f_1)$ (for all $n \in \mathbb{N}^+$, $1 \leq i, j \leq I$ and $f_1, f_2 \in \mathcal{F}^n$). For all $n \in \mathbb{N}$, $1 \leq i \leq I$ and $2 \leq j \leq I$,

we have $G(p_i, u_j^n, f, g) = G(p_i, u_j^n, g, f)$ because G is invariant under relabelling. Define $f' \in \Phi$ as the act such that $f'(\omega) = c_1^1$ for any $\omega \in \Omega$ and define $g' \in \Phi$ as the act such that $g'(\omega) = c_2^1$ for any $\omega \in \Omega$. Since G is invariant under relabelling, we have $G(p_1, u_1^n, f', g') = G(p_1, u_1^n, g', f')$ $(n \in \mathbb{N})$. For any $n \in \mathbb{N}$, we also have $[u_1^n \circ \mathcal{C}^n \circ f](\omega) = [u_1^n \circ \mathcal{C}^n \circ f'](\omega)$ and $[u_1^n \circ \mathcal{C}^n \circ g](\omega) = [u_1^n \circ \mathcal{C}^n \circ g'](\omega)$ for all $\omega \in W_1$ because $u_1^n(\{c_1^1\}) = u_1^n(\{c_2^1\})$. Since W_1 is p_1, G–one, this implies $G(p_1, u_1^n, f, g) = G(p_1, u_1^n, f', g') = G(p_1, u_1^n, g', f') = G(p_1, u_1^n, g, f)$ (for all $n \in \mathbb{N}$). Since $\succeq$ is reflexive, we obtain $f \sim_{ij}^n g$ and $f \sim_{11}^n g$ for all $n \in \mathbb{N}$, $1 \leq i \leq I$ and $2 \leq j \leq I$. Moreover, $[u_1^{2n} \circ \mathcal{C}^{2n} \circ f](\omega) = a$ and $[u_1^{2n} \circ \mathcal{C}^{2n} \circ g](\omega) = b$ while $[u_1^{2n+1} \circ \mathcal{C}^{2n+1} \circ f](\omega) = a'$ and $[u_1^{2n+1} \circ \mathcal{C}^{2n+1} \circ f](\omega) = b'$ for all $\omega \in W_2$ and all $n \in \mathbb{N}$. By Lemma A.1 and (*), we therefore obtain $f \succ_{j1}^{2n} g$ and $g \succ_{j1}^{2n+1} f$ for all $n \in \mathbb{N}$ and all $2 \leq j \leq I$. If S is Pareto optimal, it follows that $C_{S(\langle M_i^{2n} \rangle)}(\{f, g\}) = \{f\}$ but $C_{S(\langle M_i^{2n+1} \rangle)}(\{f, g\}) = \{g\}$ (for all $n \in \mathbb{N}$). $\square$

References

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

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

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

[4] 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.

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

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

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

[8] P. Ghirardato. On independence for non–additive measures, with a Fubini theorem. *Journal of Economic Theory*, 73:261–291, 1997.

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

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

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

[12] M. Hild. Stable aggregation of preferences. *California Institute of Technology, Social Science Working Paper*, 1112, 2001a. Available at www.hild.org.

[13] 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.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

[28] 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.

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