

Social choice under risk and uncertainty

Oxford Handbook on Well-Being and Public Policy

Workshop II, Princeton

Philippe Mongin and Marcus Pivato

Centre National de la Recherche Scientifique & HEC Paris

Department of Mathematics, Trent University

March 1, 2014

The outcomes of policies are almost always unpredictable.
How should policy-makers cope with such unpredictability?

Economists conventionally distinguish two forms of unpredictability: *risk* and *uncertainty*.

A decision involves *risk* if the possible outcomes can be assigned probabilities in an unambiguous, objective way (e.g. by measuring frequencies in historical data drawn from large, homogeneous population).

Examples:

- ▶ Weather forecasting.
- ▶ Insurance (property, automobile, workplace, crop, medical, etc.).
- ▶ Medical diagnosis, treatment, prognosis.

A decision involves *uncertainty* if the phenomena are *sui generis*, or so poorly understood that we cannot assign clear probabilities to outcomes.

Examples:

- ▶ Macroeconomic policy.
- ▶ Climate change.
- ▶ National security and foreign policy.

The outcomes of policies are almost always unpredictable.

How should policy-makers cope with such unpredictability?

Economists conventionally distinguish two forms of unpredictability: *risk* and *uncertainty*.

A decision involves *risk* if the possible outcomes can be assigned probabilities in an unambiguous, objective way (e.g. by measuring frequencies in historical data drawn from large, homogeneous population).

Examples:

- ▶ Weather forecasting.
- ▶ Insurance (property, automobile, workplace, crop, medical, etc.).
- ▶ Medical diagnosis, treatment, prognosis.

A decision involves *uncertainty* if the phenomena are *sui generis*, or so poorly understood that we cannot assign clear probabilities to outcomes.

Examples:

- ▶ Macroeconomic policy.
- ▶ Climate change.
- ▶ National security and foreign policy.

The outcomes of policies are almost always unpredictable.

How should policy-makers cope with such unpredictability?

Economists conventionally distinguish two forms of unpredictability: *risk* and *uncertainty*.

A decision involves *risk* if the possible outcomes can be assigned probabilities in an unambiguous, objective way (e.g. by measuring frequencies in historical data drawn from large, homogeneous population).

Examples:

- ▶ Weather forecasting.
- ▶ Insurance (property, automobile, workplace, crop, medical, etc.).
- ▶ Medical diagnosis, treatment, prognosis.

A decision involves *uncertainty* if the phenomena are *sui generis*, or so poorly understood that we cannot assign clear probabilities to outcomes.

Examples:

- ▶ Macroeconomic policy.
- ▶ Climate change.
- ▶ National security and foreign policy.

The outcomes of policies are almost always unpredictable.

How should policy-makers cope with such unpredictability?

Economists conventionally distinguish two forms of unpredictability: *risk* and *uncertainty*.

A decision involves *risk* if the possible outcomes can be assigned probabilities in an unambiguous, objective way (e.g. by measuring frequencies in historical data drawn from large, homogeneous population).

Examples:

- ▶ Weather forecasting.
- ▶ Insurance (property, automobile, workplace, crop, medical, etc.).
- ▶ Medical diagnosis, treatment, prognosis.

A decision involves *uncertainty* if the phenomena are *sui generis*, or so poorly understood that we cannot assign clear probabilities to outcomes.

Examples:

- ▶ Macroeconomic policy.
- ▶ Climate change.
- ▶ National security and foreign policy.

The outcomes of policies are almost always unpredictable.

How should policy-makers cope with such unpredictability?

Economists conventionally distinguish two forms of unpredictability: *risk* and *uncertainty*.

A decision involves *risk* if the possible outcomes can be assigned probabilities in an unambiguous, objective way (e.g. by measuring frequencies in historical data drawn from large, homogeneous population).

Examples:

- ▶ Weather forecasting.
- ▶ Insurance (property, automobile, workplace, crop, medical, etc.).
- ▶ Medical diagnosis, treatment, prognosis.

A decision involves *uncertainty* if the phenomena are *sui generis*, or so poorly understood that we cannot assign clear probabilities to outcomes.

Examples:

- ▶ Macroeconomic policy.
- ▶ Climate change.
- ▶ National security and foreign policy.

The outcomes of policies are almost always unpredictable.

How should policy-makers cope with such unpredictability?

Economists conventionally distinguish two forms of unpredictability: *risk* and *uncertainty*.

A decision involves *risk* if the possible outcomes can be assigned probabilities in an unambiguous, objective way (e.g. by measuring frequencies in historical data drawn from large, homogeneous population).

Examples:

- ▶ Weather forecasting.
- ▶ Insurance (property, automobile, workplace, crop, medical, etc.).
- ▶ Medical diagnosis, treatment, prognosis.

A decision involves *uncertainty* if the phenomena are *sui generis*, or so poorly understood that we cannot assign clear probabilities to outcomes.

Examples:

- ▶ Macroeconomic policy.
- ▶ Climate change.
- ▶ National security and foreign policy.

The outcomes of policies are almost always unpredictable.

How should policy-makers cope with such unpredictability?

Economists conventionally distinguish two forms of unpredictability: *risk* and *uncertainty*.

A decision involves *risk* if the possible outcomes can be assigned probabilities in an unambiguous, objective way (e.g. by measuring frequencies in historical data drawn from large, homogeneous population).

Examples:

- ▶ Weather forecasting.
- ▶ Insurance (property, automobile, workplace, crop, medical, etc.).
- ▶ Medical diagnosis, treatment, prognosis.

A decision involves *uncertainty* if the phenomena are *sui generis*, or so poorly understood that we cannot assign clear probabilities to outcomes.

Examples:

- ▶ Macroeconomic policy.
- ▶ Climate change.
- ▶ National security and foreign policy.

The outcomes of policies are almost always unpredictable.

How should policy-makers cope with such unpredictability?

Economists conventionally distinguish two forms of unpredictability: *risk* and *uncertainty*.

A decision involves *risk* if the possible outcomes can be assigned probabilities in an unambiguous, objective way (e.g. by measuring frequencies in historical data drawn from large, homogeneous population).

Examples:

- ▶ Weather forecasting.
- ▶ Insurance (property, automobile, workplace, crop, medical, etc.).
- ▶ Medical diagnosis, treatment, prognosis.

A decision involves *uncertainty* if the phenomena are *sui generis*, or so poorly understood that we cannot assign clear probabilities to outcomes.

Examples:

- ▶ Macroeconomic policy.
- ▶ Climate change.
- ▶ National security and foreign policy.

The outcomes of policies are almost always unpredictable.

How should policy-makers cope with such unpredictability?

Economists conventionally distinguish two forms of unpredictability: *risk* and *uncertainty*.

A decision involves *risk* if the possible outcomes can be assigned probabilities in an unambiguous, objective way (e.g. by measuring frequencies in historical data drawn from large, homogeneous population).

Examples:

- ▶ Weather forecasting.
- ▶ Insurance (property, automobile, workplace, crop, medical, etc.).
- ▶ Medical diagnosis, treatment, prognosis.

A decision involves *uncertainty* if the phenomena are *sui generis*, or so poorly understood that we cannot assign clear probabilities to outcomes.

Examples:

- ▶ Macroeconomic policy.
- ▶ Climate change.
- ▶ National security and foreign policy.

The outcomes of policies are almost always unpredictable.

How should policy-makers cope with such unpredictability?

Economists conventionally distinguish two forms of unpredictability: *risk* and *uncertainty*.

A decision involves *risk* if the possible outcomes can be assigned probabilities in an unambiguous, objective way (e.g. by measuring frequencies in historical data drawn from large, homogeneous population).

Examples:

- ▶ Weather forecasting.
- ▶ Insurance (property, automobile, workplace, crop, medical, etc.).
- ▶ Medical diagnosis, treatment, prognosis.

A decision involves *uncertainty* if the phenomena are *sui generis*, or so poorly understood that we cannot assign clear probabilities to outcomes.

Examples:

- ▶ Macroeconomic policy.
- ▶ Climate change.
- ▶ National security and foreign policy.

Decision Theory

(a crash course)

Let $\mathcal{A} := \{a, b, c, \dots\}$ be a set of outcomes.

A *lottery* is a device \mathbf{p} which assigns a probability to each outcome in \mathcal{A} .

Formally, \mathbf{p} is a probability distribution: $\mathbf{p} := (p_a, p_b, p_c, \dots)$.

Let $\mathcal{P} = \{\mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \dots\}$ be a set of lotteries (e.g. due to different policies).

The problem of *risk*: How to pick the “best” lottery in \mathcal{P} .

Let u be a *utility function*, which assigns a “utility” to each outcome in \mathcal{A} .

The *expected utility* of a lottery \mathbf{p} is then defined:

$$\mathbb{E}(u, \mathbf{p}) := p_a u(a) + p_b u(b) + p_c u(c) + \dots$$

The standard way to cope with risk is to choose the lottery which

maximizes expected utility.

Question: Is this “rational”?

von Neumann-Morgenstern Theorem. Let $\Delta(\mathcal{A})$ be the set of all lotteries over \mathcal{A} . Let \succeq be a preference order on $\Delta(\mathcal{A})$ which satisfies certain axioms (describing “consistency” or “rationality”). Then there is a utility function u on \mathcal{A} such that \succeq seeks to maximize the expected value of u . That is for any lotteries \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

$$(\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(u, \mathbf{p}) \geq \mathbb{E}(u, \mathbf{q}))$$

Let $\mathcal{A} := \{a, b, c, \dots\}$ be a set of outcomes.

A **lottery** is a device \mathbf{p} which assigns a probability to each outcome in \mathcal{A}

Formally, \mathbf{p} is a probability distribution: $\mathbf{p} := (p_a, p_b, p_c, \dots)$.

Let $\mathcal{P} = \{\mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \dots\}$ be a set of lotteries (e.g. due to different policies)

The problem of *risk*: How to pick the “best” lottery in \mathcal{P} .

Let u be a *utility function*, which assigns a “utility” to each outcome in \mathcal{A} .

The *expected utility* of a lottery \mathbf{p} is then defined:

$$\mathbb{E}(u, \mathbf{p}) := p_a u(a) + p_b u(b) + p_c u(c) + \dots$$

The standard way to cope with risk is to choose the lottery which

maximizes expected utility.

Question: Is this “rational”?

von Neumann-Morgenstern Theorem. Let $\Delta(\mathcal{A})$ be the set of all lotteries over \mathcal{A} . Let \succeq be a preference order on $\Delta(\mathcal{A})$ which satisfies certain axioms (describing “consistency” or “rationality”). Then there is a utility function u on \mathcal{A} such that \succeq seeks to maximize the expected value of u . That is for any lotteries \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

$$(\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(u, \mathbf{p}) \geq \mathbb{E}(u, \mathbf{q}))$$

Let $\mathcal{A} := \{a, b, c, \dots\}$ be a set of outcomes.

A *lottery* is a device \mathbf{p} which assigns a probability to each outcome in \mathcal{A}

Formally, \mathbf{p} is a probability distribution: $\mathbf{p} := (p_a, p_b, p_c, \dots)$.

Let $\mathcal{P} = \{\mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \dots\}$ be a set of lotteries (e.g. due to different policies).

The problem of *risk*: How to pick the “best” lottery in \mathcal{P} .

Let u be a *utility function*, which assigns a “utility” to each outcome in \mathcal{A} .

The *expected utility* of a lottery \mathbf{p} is then defined:

$$\mathbb{E}(u, \mathbf{p}) := p_a u(a) + p_b u(b) + p_c u(c) + \dots$$

The standard way to cope with risk is to choose the lottery which

maximizes expected utility.

Question: Is this “rational”?

von Neumann-Morgenstern Theorem. Let $\Delta(\mathcal{A})$ be the set of all lotteries over \mathcal{A} . Let \succeq be a preference order on $\Delta(\mathcal{A})$ which satisfies certain axioms (describing “consistency” or “rationality”). Then there is a utility function u on \mathcal{A} such that \succeq seeks to maximize the expected value of u . That is for any lotteries \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

$$(\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(u, \mathbf{p}) \geq \mathbb{E}(u, \mathbf{q}))$$

Let $\mathcal{A} := \{a, b, c, \dots\}$ be a set of outcomes.

A *lottery* is a device \mathbf{p} which assigns a probability to each outcome in \mathcal{A}

Formally, \mathbf{p} is a probability distribution: $\mathbf{p} := (p_a, p_b, p_c, \dots)$.

Let $\mathcal{P} = \{\mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \dots\}$ be a set of lotteries (e.g. due to different policies).

The problem of *risk*: How to pick the “best” lottery in \mathcal{P} .

Let u be a *utility function*, which assigns a “utility” to each outcome in \mathcal{A} .

The *expected utility* of a lottery \mathbf{p} is then defined:

$$\mathbb{E}(u, \mathbf{p}) := p_a u(a) + p_b u(b) + p_c u(c) + \dots$$

The standard way to cope with risk is to choose the lottery which

maximizes expected utility.

Question: Is this “rational”?

von Neumann-Morgenstern Theorem. Let $\Delta(\mathcal{A})$ be the set of all lotteries over \mathcal{A} . Let \succeq be a preference order on $\Delta(\mathcal{A})$ which satisfies certain axioms (describing “consistency” or “rationality”). Then there is a utility function u on \mathcal{A} such that \succeq seeks to maximize the expected value of u . That is for any lotteries \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

$$(\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(u, \mathbf{p}) \geq \mathbb{E}(u, \mathbf{q}))$$

Let $\mathcal{A} := \{a, b, c, \dots\}$ be a set of outcomes.

A *lottery* is a device \mathbf{p} which assigns a probability to each outcome in \mathcal{A}

Formally, \mathbf{p} is a probability distribution: $\mathbf{p} := (p_a, p_b, p_c, \dots)$.

Let $\mathcal{P} = \{\mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \dots\}$ be a set of lotteries (e.g. due to different policies).

The problem of *risk*: How to pick the “best” lottery in \mathcal{P} .

Let u be a *utility function*, which assigns a “utility” to each outcome in \mathcal{A} .

The *expected utility* of a lottery \mathbf{p} is then defined:

$$\mathbb{E}(u, \mathbf{p}) := p_a u(a) + p_b u(b) + p_c u(c) + \dots$$

The standard way to cope with risk is to choose the lottery which

maximizes expected utility.

Question: Is this “rational”?

von Neumann-Morgenstern Theorem. Let $\Delta(\mathcal{A})$ be the set of all lotteries over \mathcal{A} . Let \succeq be a preference order on $\Delta(\mathcal{A})$ which satisfies certain axioms (describing “consistency” or “rationality”). Then there is a utility function u on \mathcal{A} such that \succeq seeks to maximize the expected value of u . That is for any lotteries \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

$$(\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(u, \mathbf{p}) \geq \mathbb{E}(u, \mathbf{q}))$$

Let $\mathcal{A} := \{a, b, c, \dots\}$ be a set of outcomes.

A *lottery* is a device \mathbf{p} which assigns a probability to each outcome in \mathcal{A}

Formally, \mathbf{p} is a probability distribution: $\mathbf{p} := (p_a, p_b, p_c, \dots)$.

Let $\mathcal{P} = \{\mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \dots\}$ be a set of lotteries (e.g. due to different policies).

The problem of *risk*: How to pick the “best” lottery in \mathcal{P} .

Let u be a *utility function*, which assigns a “utility” to each outcome in \mathcal{A} .

The *expected utility* of a lottery \mathbf{p} is then defined:

$$\mathbb{E}(u, \mathbf{p}) := p_a u(a) + p_b u(b) + p_c u(c) + \dots$$

The standard way to cope with risk is to choose the lottery which maximizes expected utility.

Question: Is this “rational”?

von Neumann-Morgenstern Theorem. Let $\Delta(\mathcal{A})$ be the set of all lotteries over \mathcal{A} . Let \succeq be a preference order on $\Delta(\mathcal{A})$ which satisfies certain axioms (describing “consistency” or “rationality”). Then there is a utility function u on \mathcal{A} such that \succeq seeks to maximize the expected value of u . That is for any lotteries \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

$$(\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(u, \mathbf{p}) \geq \mathbb{E}(u, \mathbf{q}))$$

Let $\mathcal{A} := \{a, b, c, \dots\}$ be a set of outcomes.

A *lottery* is a device \mathbf{p} which assigns a probability to each outcome in \mathcal{A}

Formally, \mathbf{p} is a probability distribution: $\mathbf{p} := (p_a, p_b, p_c, \dots)$.

Let $\mathcal{P} = \{\mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \dots\}$ be a set of lotteries (e.g. due to different policies).

The problem of *risk*: How to pick the “best” lottery in \mathcal{P} .

Let u be a *utility function*, which assigns a “utility” to each outcome in \mathcal{A} .

The *expected utility* of a lottery \mathbf{p} is then defined:

$$\mathbb{E}(u, \mathbf{p}) := p_a u(a) + p_b u(b) + p_c u(c) + \dots$$

The standard way to cope with risk is to choose the lottery which

maximizes expected utility.

Question: Is this “rational”?

von Neumann-Morgenstern Theorem. Let $\Delta(\mathcal{A})$ be the set of all lotteries over \mathcal{A} . Let \succeq be a preference order on $\Delta(\mathcal{A})$ which satisfies certain axioms (describing “consistency” or “rationality”). Then there is a utility function u on \mathcal{A} such that \succeq seeks to maximize the expected value of u . That is for any lotteries \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

$$(\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(u, \mathbf{p}) \geq \mathbb{E}(u, \mathbf{q}))$$

Let $\mathcal{A} := \{a, b, c, \dots\}$ be a set of outcomes.

A *lottery* is a device \mathbf{p} which assigns a probability to each outcome in \mathcal{A}

Formally, \mathbf{p} is a probability distribution: $\mathbf{p} := (p_a, p_b, p_c, \dots)$.

Let $\mathcal{P} = \{\mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \dots\}$ be a set of lotteries (e.g. due to different policies).

The problem of *risk*: How to pick the “best” lottery in \mathcal{P} .

Let u be a *utility function*, which assigns a “utility” to each outcome in \mathcal{A} .

The *expected utility* of a lottery \mathbf{p} is then defined:

$$\mathbb{E}(u, \mathbf{p}) := p_a u(a) + p_b u(b) + p_c u(c) + \dots$$

The standard way to cope with risk is to choose the lottery which

maximizes expected utility.

Question: Is this “rational”?

von Neumann-Morgenstern Theorem. Let $\Delta(\mathcal{A})$ be the set of all lotteries over \mathcal{A} . Let \succeq be a preference order on $\Delta(\mathcal{A})$ which satisfies certain axioms (describing “consistency” or “rationality”). Then there is a utility function u on \mathcal{A} such that \succeq seeks to maximize the expected value of u . That is for any lotteries \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

$$(\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(u, \mathbf{p}) \geq \mathbb{E}(u, \mathbf{q}))$$

Let $\mathcal{A} := \{a, b, c, \dots\}$ be a set of outcomes.

A *lottery* is a device \mathbf{p} which assigns a probability to each outcome in \mathcal{A}

Formally, \mathbf{p} is a probability distribution: $\mathbf{p} := (p_a, p_b, p_c, \dots)$.

Let $\mathcal{P} = \{\mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \dots\}$ be a set of lotteries (e.g. due to different policies).

The problem of *risk*: How to pick the “best” lottery in \mathcal{P} .

Let u be a *utility function*, which assigns a “utility” to each outcome in \mathcal{A} .

The *expected utility* of a lottery \mathbf{p} is then defined:

$$\mathbb{E}(u, \mathbf{p}) := p_a u(a) + p_b u(b) + p_c u(c) + \dots$$

The standard way to cope with risk is to choose the lottery which

maximizes expected utility.

Question: Is this “rational”?

von Neumann-Morgenstern Theorem. Let $\Delta(\mathcal{A})$ be the set of all lotteries over \mathcal{A} . Let \succeq be a preference order on $\Delta(\mathcal{A})$ which satisfies certain axioms (describing “consistency” or “rationality”). Then there is a utility function u on \mathcal{A} such that \succeq seeks to maximize the expected value of u . That is for any lotteries \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

$$(\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(u, \mathbf{p}) \geq \mathbb{E}(u, \mathbf{q}))$$

Let $\mathcal{A} := \{a, b, c, \dots\}$ be a set of outcomes.

A *lottery* is a device \mathbf{p} which assigns a probability to each outcome in \mathcal{A}

Formally, \mathbf{p} is a probability distribution: $\mathbf{p} := (p_a, p_b, p_c, \dots)$.

Let $\mathcal{P} = \{\mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \dots\}$ be a set of lotteries (e.g. due to different policies).

The problem of *risk*: How to pick the “best” lottery in \mathcal{P} .

Let u be a *utility function*, which assigns a “utility” to each outcome in \mathcal{A} .

The *expected utility* of a lottery \mathbf{p} is then defined:

$$\mathbb{E}(u, \mathbf{p}) := p_a u(a) + p_b u(b) + p_c u(c) + \dots$$

The standard way to cope with risk is to choose the lottery which

maximizes expected utility.

Question: Is this “rational”?

von Neumann-Morgenstern Theorem. Let $\Delta(\mathcal{A})$ be the set of all lotteries over \mathcal{A} . Let \succeq be a preference order on $\Delta(\mathcal{A})$ which satisfies certain axioms (describing “consistency” or “rationality”). Then there is a utility function u on \mathcal{A} such that \succeq seeks to maximize the expected value of u . That is for any lotteries \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

$$(\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(u, \mathbf{p}) \geq \mathbb{E}(u, \mathbf{q}))$$

Let $\mathcal{A} := \{a, b, c, \dots\}$ be a set of outcomes.

A *lottery* is a device \mathbf{p} which assigns a probability to each outcome in \mathcal{A}

Formally, \mathbf{p} is a probability distribution: $\mathbf{p} := (p_a, p_b, p_c, \dots)$.

Let $\mathcal{P} = \{\mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \dots\}$ be a set of lotteries (e.g. due to different policies).

The problem of *risk*: How to pick the “best” lottery in \mathcal{P} .

Let u be a *utility function*, which assigns a “utility” to each outcome in \mathcal{A} .

The *expected utility* of a lottery \mathbf{p} is then defined:

$$\mathbb{E}(u, \mathbf{p}) := p_a u(a) + p_b u(b) + p_c u(c) + \dots$$

The standard way to cope with risk is to choose the lottery which

maximizes expected utility.

Question: Is this “rational”?

von Neumann-Morgenstern Theorem. Let $\Delta(\mathcal{A})$ be the set of all lotteries over \mathcal{A} . Let \succeq be a preference order on $\Delta(\mathcal{A})$ which satisfies certain axioms (describing “consistency” or “rationality”). Then there is a utility function u on \mathcal{A} such that \succeq seeks to maximize the expected value of u . That is for any lotteries \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

$$(\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(u, \mathbf{p}) \geq \mathbb{E}(u, \mathbf{q})).$$

Let $\mathcal{S} = \{r, s, t, \dots\}$ be a set of possible “states of the world”.

We don’t know which is the true state, or even their probabilities.

A *prospect* is a device f which assigns an outcome in \mathcal{A} to each state in \mathcal{S} .
(Formally, f is a function from \mathcal{S} to \mathcal{A} .)

If $f(s) = a$, this means, “If the state of the world turns out to be s , then the prospect will yield the outcome a .”

Let $\mathcal{F} = \{f, g, h, \dots\}$ be a set of prospects (e.g. due to different policies).

The problem of *uncertainty*: How to pick the “best” prospect from \mathcal{F} .

A *subjective probability* is a probability distribution p on \mathcal{S} . It could represent the “beliefs” of an agent about \mathcal{S} .

For instance, if $p(s) > p(t)$, then the agent believes that state s is “more likely” than state t .

Let u be a utility function on the set of outcomes \mathcal{A} .

The *subjective expected utility* of prospect f relative to u and p is defined:

$$\mathbb{E}(f|u, p) := p(r) u[f(r)] + p(s) u[f(s)] + p(t) u[f(t)] + \dots$$

Let $\mathcal{S} = \{r, s, t, \dots\}$ be a set of possible “states of the world”.

We don’t know which is the true state, or even their probabilities.

A *prospect* is a device f which assigns an outcome in \mathcal{A} to each state in \mathcal{S} .
(Formally, f is a function from \mathcal{S} to \mathcal{A} .)

If $f(s) = a$, this means, “If the state of the world turns out to be s , then the prospect will yield the outcome a .”

Let $\mathcal{F} = \{f, g, h, \dots\}$ be a set of prospects (e.g. due to different policies).

The problem of *uncertainty*: How to pick the “best” prospect from \mathcal{F} .

A *subjective probability* is a probability distribution p on \mathcal{S} . It could represent the “beliefs” of an agent about \mathcal{S} .

For instance, if $p(s) > p(t)$, then the agent believes that state s is “more likely” than state t .

Let u be a utility function on the set of outcomes \mathcal{A} .

The *subjective expected utility* of prospect f relative to u and p is defined:

$$\mathbb{E}(f|u, p) := p(r) u[f(r)] + p(s) u[f(s)] + p(t) u[f(t)] + \dots$$

Let $\mathcal{S} = \{r, s, t, \dots\}$ be a set of possible “states of the world”.

We don’t know which is the true state, or even their probabilities.

A *prospect* is a device f which assigns an outcome in \mathcal{A} to each state in \mathcal{S} .
(Formally, f is a function from \mathcal{S} to \mathcal{A} .)

If $f(s) = a$, this means, “If the state of the world turns out to be s , then the prospect will yield the outcome a .”

Let $\mathcal{F} = \{f, g, h, \dots\}$ be a set of prospects (e.g. due to different policies).

The problem of *uncertainty*: How to pick the “best” prospect from \mathcal{F} .

A *subjective probability* is a probability distribution p on \mathcal{S} . It could represent the “beliefs” of an agent about \mathcal{S} .

For instance, if $p(s) > p(t)$, then the agent believes that state s is “more likely” than state t .

Let u be a utility function on the set of outcomes \mathcal{A} .

The *subjective expected utility* of prospect f relative to u and p is defined:

$$\mathbb{E}(f|u, p) := p(r) u[f(r)] + p(s) u[f(s)] + p(t) u[f(t)] + \dots$$

Let $\mathcal{S} = \{r, s, t, \dots\}$ be a set of possible “states of the world”.

We don’t know which is the true state, or even their probabilities.

A *prospect* is a device f which assigns an outcome in \mathcal{A} to each state in \mathcal{S} .
(Formally, f is a function from \mathcal{S} to \mathcal{A} .)

If $f(s) = a$, this means, “If the state of the world turns out to be s , then the prospect will yield the outcome a .”

Let $\mathcal{F} = \{f, g, h, \dots\}$ be a set of prospects (e.g. due to different policies).

The problem of *uncertainty*: How to pick the “best” prospect from \mathcal{F} .

A *subjective probability* is a probability distribution p on \mathcal{S} . It could represent the “beliefs” of an agent about \mathcal{S} .

For instance, if $p(s) > p(t)$, then the agent believes that state s is “more likely” than state t .

Let u be a utility function on the set of outcomes \mathcal{A} .

The *subjective expected utility* of prospect f relative to u and p is defined:

$$\mathbb{E}(f|u, p) := p(r) u[f(r)] + p(s) u[f(s)] + p(t) u[f(t)] + \dots$$

Let $\mathcal{S} = \{r, s, t, \dots\}$ be a set of possible “states of the world”.

We don’t know which is the true state, or even their probabilities.

A *prospect* is a device f which assigns an outcome in \mathcal{A} to each state in \mathcal{S} .
(Formally, f is a function from \mathcal{S} to \mathcal{A} .)

If $f(s) = a$, this means, “If the state of the world turns out to be s , then the prospect will yield the outcome a .”

Let $\mathcal{F} = \{f, g, h, \dots\}$ be a set of prospects (e.g. due to different policies).

The problem of *uncertainty*: How to pick the “best” prospect from \mathcal{F} .

A *subjective probability* is a probability distribution p on \mathcal{S} . It could represent the “beliefs” of an agent about \mathcal{S} .

For instance, if $p(s) > p(t)$, then the agent believes that state s is “more likely” than state t .

Let u be a utility function on the set of outcomes \mathcal{A} .

The *subjective expected utility* of prospect f relative to u and p is defined:

$$\mathbb{E}(f|u, p) := p(r) u[f(r)] + p(s) u[f(s)] + p(t) u[f(t)] + \dots$$

Let $\mathcal{S} = \{r, s, t, \dots\}$ be a set of possible “states of the world”.

We don’t know which is the true state, or even their probabilities.

A *prospect* is a device f which assigns an outcome in \mathcal{A} to each state in \mathcal{S} .
(Formally, f is a function from \mathcal{S} to \mathcal{A} .)

If $f(s) = a$, this means, “If the state of the world turns out to be s , then the prospect will yield the outcome a .”

Let $\mathcal{F} = \{f, g, h, \dots\}$ be a set of prospects (e.g. due to different policies).

The problem of *uncertainty*: How to pick the “best” prospect from \mathcal{F} .

A *subjective probability* is a probability distribution p on \mathcal{S} . It could represent the “beliefs” of an agent about \mathcal{S} .

For instance, if $p(s) > p(t)$, then the agent believes that state s is “more likely” than state t .

Let u be a utility function on the set of outcomes \mathcal{A} .

The *subjective expected utility* of prospect f relative to u and p is defined:

$$\mathbb{E}(f|u, p) := p(r) u[f(r)] + p(s) u[f(s)] + p(t) u[f(t)] + \dots$$

Let $\mathcal{S} = \{r, s, t, \dots\}$ be a set of possible “states of the world”.

We don’t know which is the true state, or even their probabilities.

A *prospect* is a device f which assigns an outcome in \mathcal{A} to each state in \mathcal{S} .
(Formally, f is a function from \mathcal{S} to \mathcal{A} .)

If $f(s) = a$, this means, “If the state of the world turns out to be s , then the prospect will yield the outcome a .”

Let $\mathcal{F} = \{f, g, h, \dots\}$ be a set of prospects (e.g. due to different policies).

The problem of *uncertainty*: How to pick the “best” prospect from \mathcal{F} .

A *subjective probability* is a probability distribution p on \mathcal{S} . It could represent the “beliefs” of an agent about \mathcal{S} .

For instance, if $p(s) > p(t)$, then the agent believes that state s is “more likely” than state t .

Let u be a utility function on the set of outcomes \mathcal{A} .

The *subjective expected utility* of prospect f relative to u and p is defined:

$$\mathbb{E}(f|u, p) := p(r) u[f(r)] + p(s) u[f(s)] + p(t) u[f(t)] + \dots$$

Let $\mathcal{S} = \{r, s, t, \dots\}$ be a set of possible “states of the world”.

We don’t know which is the true state, or even their probabilities.

A *prospect* is a device f which assigns an outcome in \mathcal{A} to each state in \mathcal{S} . (Formally, f is a function from \mathcal{S} to \mathcal{A} .)

If $f(s) = a$, this means, “If the state of the world turns out to be s , then the prospect will yield the outcome a .”

Let $\mathcal{F} = \{f, g, h, \dots\}$ be a set of prospects (e.g. due to different policies).

The problem of *uncertainty*: How to pick the “best” prospect from \mathcal{F} .

A *subjective probability* is a probability distribution p on \mathcal{S} . It could represent the “beliefs” of an agent about \mathcal{S} .

For instance, if $p(s) > p(t)$, then the agent believes that state s is “more likely” than state t .

Let u be a utility function on the set of outcomes \mathcal{A} .

The *subjective expected utility* of prospect f relative to u and p is defined:

$$\mathbb{E}(f|u, p) := p(r) u[f(r)] + p(s) u[f(s)] + p(t) u[f(t)] + \dots$$

The *subjective expected utility* of prospect f relative to u and p is defined:

$$\mathbb{E}(f|u, p) := p(r) u[f(r)] + p(s) u[f(s)] + p(t) u[f(t)] + \dots$$

The standard way to cope with uncertainty is to choose the prospect which *maximizes subjective expected utility*.

Question: Is this “rational”?

Savage’s Theorem. Suppose S is infinite. Let \mathcal{A}^S be the set of all prospects involving S and \mathcal{A} .

Let \succeq be a preference order on \mathcal{A}^S which satisfies certain axioms (describing “consistency” or “rationality”).

Then there is a utility function u on \mathcal{A} and a nonatomic* subjective probability p on S such that \succeq seeks to maximize subjective expected utility. That is: for any prospects f and g in \mathcal{A}^S ,

$$(f \succeq g) \iff (\mathbb{E}(f|u, p) \geq \mathbb{E}(g|u, p)).$$

Also, p is unique, and u is unique up to positive affine transformation.

(*) Nonatomic means $p(s) = 0$ for all s in S .

The *subjective expected utility* of prospect f relative to u and p is defined:

$$\mathbb{E}(f|u, p) := p(r) u[f(r)] + p(s) u[f(s)] + p(t) u[f(t)] + \dots$$

The standard way to cope with uncertainty is to choose the prospect which *maximizes subjective expected utility*.

Question: Is this “rational”?

Savage’s Theorem. Suppose S is infinite. Let \mathcal{A}^S be the set of all prospects involving S and \mathcal{A} .

Let \succeq be a preference order on \mathcal{A}^S which satisfies certain axioms (describing “consistency” or “rationality”).

Then there is a utility function u on \mathcal{A} and a nonatomic* subjective probability p on S such that \succeq seeks to maximize subjective expected utility. That is: for any prospects f and g in \mathcal{A}^S ,

$$(f \succeq g) \iff (\mathbb{E}(f|u, p) \geq \mathbb{E}(g|u, p)).$$

Also, p is unique, and u is unique up to positive affine transformation.

(*) Nonatomic means $p(s) = 0$ for all s in S .

The *subjective expected utility* of prospect f relative to u and p is defined:

$$\mathbb{E}(f|u, p) := p(r) u[f(r)] + p(s) u[f(s)] + p(t) u[f(t)] + \dots$$

The standard way to cope with uncertainty is to choose the prospect which *maximizes subjective expected utility*.

Question: Is this “rational”?

Savage's Theorem. Suppose S is infinite. Let \mathcal{A}^S be the set of all prospects involving S and \mathcal{A} .

Let \succeq be a preference order on \mathcal{A}^S which satisfies certain axioms (describing “consistency” or “rationality”).

Then there is a utility function u on \mathcal{A} and a nonatomic* subjective probability p on S such that \succeq seeks to maximize subjective expected utility. That is: for any prospects f and g in \mathcal{A}^S ,

$$(f \succeq g) \iff (\mathbb{E}(f|u, p) \geq \mathbb{E}(g|u, p)).$$

Also, p is unique, and u is unique up to positive affine transformation.

(*) Nonatomic means $p(s) = 0$ for all s in S .

The *subjective expected utility* of prospect f relative to u and p is defined:

$$\mathbb{E}(f|u, p) := p(r) u[f(r)] + p(s) u[f(s)] + p(t) u[f(t)] + \dots$$

The standard way to cope with uncertainty is to choose the prospect which *maximizes subjective expected utility*.

Question: Is this “rational”?

Savage's Theorem. Suppose S is infinite. Let \mathcal{A}^S be the set of all prospects involving S and \mathcal{A} .

Let \succeq be a preference order on \mathcal{A}^S which satisfies certain axioms (describing “consistency” or “rationality”).

Then there is a utility function u on \mathcal{A} and a nonatomic* subjective probability p on S such that \succeq seeks to maximize subjective expected utility. That is: for any prospects f and g in \mathcal{A}^S ,

$$(f \succeq g) \iff (\mathbb{E}(f|u, p) \geq \mathbb{E}(g|u, p)).$$

Also, p is unique, and u is unique up to positive affine transformation.

(*) Nonatomic means $p(s) = 0$ for all s in S .

The *subjective expected utility* of prospect f relative to u and p is defined:

$$\mathbb{E}(f|u, p) := p(r) u[f(r)] + p(s) u[f(s)] + p(t) u[f(t)] + \dots$$

The standard way to cope with uncertainty is to choose the prospect which *maximizes subjective expected utility*.

Question: Is this “rational”?

Savage's Theorem. Suppose S is infinite. Let \mathcal{A}^S be the set of all prospects involving S and \mathcal{A} .

Let \succeq be a preference order on \mathcal{A}^S which satisfies certain axioms (describing “consistency” or “rationality”).

Then there is a utility function u on \mathcal{A} and a nonatomic* subjective probability p on S such that \succeq seeks to maximize subjective expected utility. That is: for any prospects f and g in \mathcal{A}^S ,

$$(f \succeq g) \iff (\mathbb{E}(f|u, p) \geq \mathbb{E}(g|u, p)).$$

Also, p is unique, and u is unique up to positive affine transformation.

(*) Nonatomic means $p(s) = 0$ for all s in S .

The *subjective expected utility* of prospect f relative to u and p is defined:

$$\mathbb{E}(f|u, p) := p(r) u[f(r)] + p(s) u[f(s)] + p(t) u[f(t)] + \dots$$

The standard way to cope with uncertainty is to choose the prospect which *maximizes subjective expected utility*.

Question: Is this “rational”?

Savage's Theorem. Suppose S is infinite. Let \mathcal{A}^S be the set of all prospects involving S and \mathcal{A} .

Let \succeq be a preference order on \mathcal{A}^S which satisfies certain axioms (describing “consistency” or “rationality”).

Then there is a utility function u on \mathcal{A} and a nonatomic* subjective probability p on S such that \succeq seeks to maximize subjective expected utility. That is: for any prospects f and g in \mathcal{A}^S ,

$$(f \succeq g) \iff (\mathbb{E}(f|u, p) \geq \mathbb{E}(g|u, p)).$$

Also, p is unique, and u is unique up to positive affine transformation.

(*) Nonatomic means $p(s) = 0$ for all s in S .

Question: How can we extend the vNM-Savage theory to decision-making for *groups* of people? In effect, this involves two kinds of aggregation:

- ▶ Aggregation across different states of nature (i.e. uncertainty).
- ▶ Aggregation across different people (i.e. social choice).

Furthermore, social aggregation could be applied either *ex ante* (i.e. *before* the uncertainty is resolved) or *ex post* (i.e. *after* the uncertainty is resolved). The problem is that these two forms of aggregation often clash.

Consider three natural guidelines for social choice under uncertainty:

- ▶ **Statewise Dominance principle:** If the lottery/prospect **P** produces a better *ex post* social outcome than **Q** in every state of nature, then **P** is *ex ante* better than **Q**. (Rationality; intertemporal consistency.)
- ▶ **Ex ante Pareto principle:** If everyone prefers lottery/prospect **P** over **Q**, then society should prefer **P** over **Q**. (Nonpaternalism; contractarianism.)
- ▶ **Egalitarianism:** *Ceteris paribus* (or even *ceteris not-so-paribus*), it is better to minimize inequality between people, both at the *ex ante* stage and the *ex post* stage.

We shall see that these three principles are often in conflict.

Question: How can we extend the vNM-Savage theory to decision-making for *groups* of people? In effect, this involves two kinds of aggregation:

- ▶ Aggregation across different states of nature (i.e. **uncertainty**).
- ▶ Aggregation across different people (i.e. social choice).

Furthermore, social aggregation could be applied either *ex ante* (i.e. *before* the uncertainty is resolved) or *ex post* (i.e. *after* the uncertainty is resolved).

The problem is that these two forms of aggregation often clash.

Consider three natural guidelines for social choice under uncertainty:

- ▶ **Statewise Dominance principle:** If the lottery/prospect **P** produces a better *ex post* social outcome than **Q** in every state of nature, then **P** is *ex ante* better than **Q**. (Rationality; intertemporal consistency.)
- ▶ **Ex ante Pareto principle:** If everyone prefers lottery/prospect **P** over **Q**, then society should prefer **P** over **Q**. (Nonpaternalism; contractarianism.)
- ▶ **Egalitarianism:** *Ceteris paribus* (or even *ceteris not-so-paribus*), it is better to minimize inequality between people, both at the *ex ante* stage and the *ex post* stage.

We shall see that these three principles are often in conflict.

Question: How can we extend the vNM-Savage theory to decision-making for *groups* of people? In effect, this involves two kinds of aggregation:

- ▶ Aggregation across different states of nature (i.e. uncertainty).
- ▶ Aggregation across different people (i.e. **social choice**).

Furthermore, social aggregation could be applied either *ex ante* (i.e. *before* the uncertainty is resolved) or *ex post* (i.e. *after* the uncertainty is resolved).

The problem is that these two forms of aggregation often clash.

Consider three natural guidelines for social choice under uncertainty:

- ▶ **Statewise Dominance principle:** If the lottery/prospect **P** produces a better *ex post* social outcome than **Q** in every state of nature, then **P** is *ex ante* better than **Q**. (Rationality; intertemporal consistency.)
- ▶ **Ex ante Pareto principle:** If everyone prefers lottery/prospect **P** over **Q**, then society should prefer **P** over **Q**. (Nonpaternalism; contractarianism.)
- ▶ **Egalitarianism:** *Ceteris paribus* (or even *ceteris not-so-paribus*), it is better to minimize inequality between people, both at the *ex ante* stage and the *ex post* stage.

We shall see that these three principles are often in conflict.

Question: How can we extend the vNM-Savage theory to decision-making for *groups* of people? In effect, this involves two kinds of aggregation:

- ▶ Aggregation across different states of nature (i.e. uncertainty).
- ▶ Aggregation across different people (i.e. social choice).

Furthermore, social aggregation could be applied either *ex ante* (i.e. *before* the uncertainty is resolved) or *ex post* (i.e. *after* the uncertainty is resolved).

The problem is that these two forms of aggregation often clash.

Consider three natural guidelines for social choice under uncertainty:

- ▶ **Statewise Dominance principle:** If the lottery/prospect **P** produces a better *ex post* social outcome than **Q** in every state of nature, then **P** is *ex ante* better than **Q**. (Rationality; intertemporal consistency.)
- ▶ **Ex ante Pareto principle:** If everyone prefers lottery/prospect **P** over **Q**, then society should prefer **P** over **Q**. (Nonpaternalism; contractarianism.)
- ▶ **Egalitarianism:** *Ceteris paribus* (or even *ceteris not-so-paribus*), it is better to minimize inequality between people, both at the *ex ante* stage and the *ex post* stage.

We shall see that these three principles are often in conflict.

Question: How can we extend the vNM-Savage theory to decision-making for *groups* of people? In effect, this involves two kinds of aggregation:

- ▶ Aggregation across different states of nature (i.e. uncertainty).
- ▶ Aggregation across different people (i.e. social choice).

Furthermore, social aggregation could be applied either *ex ante* (i.e. *before* the uncertainty is resolved) or *ex post* (i.e. *after* the uncertainty is resolved). The problem is that these two forms of aggregation often clash.

Consider three natural guidelines for social choice under uncertainty:

- ▶ **Statewise Dominance principle:** If the lottery/prospect **P** produces a better *ex post* social outcome than **Q** in every state of nature, then **P** is *ex ante* better than **Q**. (Rationality; intertemporal consistency.)
- ▶ **Ex ante Pareto principle:** If everyone prefers lottery/prospect **P** over **Q**, then society should prefer **P** over **Q**. (Nonpaternalism; contractarianism.)
- ▶ **Egalitarianism:** *Ceteris paribus* (or even *ceteris not-so-paribus*), it is better to minimize inequality between people, both at the *ex ante* stage and the *ex post* stage.

We shall see that these three principles are often in conflict.

Question: How can we extend the vNM-Savage theory to decision-making for *groups* of people? In effect, this involves two kinds of aggregation:

- ▶ Aggregation across different states of nature (i.e. **uncertainty**).
- ▶ Aggregation across different people (i.e. social choice).

Furthermore, social aggregation could be applied either *ex ante* (i.e. *before* the uncertainty is resolved) or *ex post* (i.e. *after* the uncertainty is resolved). The problem is that these two forms of aggregation often clash.

Consider three natural guidelines for social choice under uncertainty:

- ▶ **Statewise Dominance principle:** If the lottery/prospect **P** produces a better *ex post* social outcome than **Q** in every state of nature, then **P** is *ex ante* better than **Q**. (Rationality; intertemporal consistency.)
- ▶ *Ex ante Pareto principle:* If everyone prefers lottery/prospect **P** over **Q**, then society should prefer **P** over **Q**. (Nonpaternalism; contractarianism.)
- ▶ **Egalitarianism:** *Ceteris paribus* (or even *ceteris not-so-paribus*), it is better to minimize inequality between people, both at the *ex ante* stage and the *ex post* stage.

We shall see that these three principles are often in conflict.

Question: How can we extend the vNM-Savage theory to decision-making for *groups* of people? In effect, this involves two kinds of aggregation:

- ▶ Aggregation across different states of nature (i.e. uncertainty).
- ▶ Aggregation across different people (i.e. **social choice**).

Furthermore, social aggregation could be applied either *ex ante* (i.e. *before* the uncertainty is resolved) or *ex post* (i.e. *after* the uncertainty is resolved). The problem is that these two forms of aggregation often clash.

Consider three natural guidelines for social choice under uncertainty:

- ▶ **Statewise Dominance principle:** If the lottery/prospect **P** produces a better *ex post* social outcome than **Q** in every state of nature, then **P** is *ex ante* better than **Q**. (Rationality; intertemporal consistency.)
- ▶ **Ex ante Pareto principle:** If everyone prefers lottery/prospect **P** over **Q**, then society should prefer **P** over **Q**. (Nonpaternalism; contractarianism.)
- ▶ **Egalitarianism:** *Ceteris paribus* (or even *ceteris not-so-paribus*), it is better to minimize inequality between people, both at the *ex ante* stage and the *ex post* stage.

We shall see that these three principles are often in conflict.

Question: How can we extend the vNM-Savage theory to decision-making for *groups* of people? In effect, this involves two kinds of aggregation:

- ▶ Aggregation across different states of nature (i.e. uncertainty).
- ▶ Aggregation across different people (i.e. **social choice**).

Furthermore, social aggregation could be applied either *ex ante* (i.e. *before* the uncertainty is resolved) or *ex post* (i.e. *after* the uncertainty is resolved). The problem is that these two forms of aggregation often clash.

Consider three natural guidelines for social choice under uncertainty:

- ▶ **Statewise Dominance principle:** If the lottery/prospect **P** produces a better *ex post* social outcome than **Q** in every state of nature, then **P** is *ex ante* better than **Q**. (Rationality; intertemporal consistency.)
- ▶ *Ex ante* **Pareto principle:** If everyone prefers lottery/prospect **P** over **Q**, then society should prefer **P** over **Q**. (Nonpaternalism; contractarianism.)
- ▶ **Egalitarianism:** *Ceteris paribus* (or even *ceteris not-so-paribus*), it is better to minimize inequality between people, both at the *ex ante* stage and the *ex post* stage.

Question: How can we extend the vNM-Savage theory to decision-making for *groups* of people? In effect, this involves two kinds of aggregation:

- ▶ Aggregation across different states of nature (i.e. uncertainty).
- ▶ Aggregation across different people (i.e. social choice).

Furthermore, social aggregation could be applied either *ex ante* (i.e. *before* the uncertainty is resolved) or *ex post* (i.e. *after* the uncertainty is resolved). The problem is that these two forms of aggregation often clash.

Consider three natural guidelines for social choice under uncertainty:

- ▶ **Statewise Dominance principle:** If the lottery/prospect **P** produces a better *ex post* social outcome than **Q** in every state of nature, then **P** is *ex ante* better than **Q**. (Rationality; intertemporal consistency.)
- ▶ **Ex ante Pareto principle:** If everyone prefers lottery/prospect **P** over **Q**, then society should prefer **P** over **Q**. (Nonpaternalism; contractarianism.)
- ▶ **Egalitarianism:** *Ceteris paribus* (or even *ceteris not-so-paribus*), it is better to minimize inequality between people, both at the *ex ante* stage and the *ex post* stage.

We shall see that these three principles are often in conflict.

Prolog Crash course in decision theory.

- I Risk: Harsanyi's Social Aggregation Theorem.
- II Uncertainty: Heterogeneity and spurious unanimity.
- III Equality: The Diamond paradox.

Risk:

Harsanyi's Social Aggregation Theorem

Harsanyi's social aggregation theorem: hypotheses

(10/47)

Let $\mathcal{I} := \{\text{Ann, Bob, Carl, } \dots\}$ be a set of people.

Let \mathcal{A} be a set of social outcomes. Each social outcome in \mathcal{A} yields an outcome for each person in \mathcal{I} , determining her consumption bundle, health, autonomy and any other factors relevant to her well-being.

Let $\Delta(\mathcal{A})$ be the space of all lotteries over \mathcal{A} (i.e. "social lotteries").

Each policy corresponds to some element of $\Delta(\mathcal{A})$.

To choose the "best" policy, we need a *social preference order* \succeq on $\Delta(\mathcal{A})$.

If \succeq satisfies the vNM axioms of "rationality", then there is some *ex ante* social welfare function (SWF) W on \mathcal{A} such that \succeq maximizes the expected value of W :

$$\text{For any } \mathbf{p} \text{ and } \mathbf{q} \text{ in } \Delta(\mathcal{A}), \quad (\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(W, \mathbf{p}) \geq \mathbb{E}(W, \mathbf{q})).$$

Every individual i in \mathcal{I} has a preference order \succeq_i over $\Delta(\mathcal{A})$.

If \succeq_i satisfies the vNM axioms of "rationality", then there is some utility function u_i on \mathcal{A} such that \succeq_i maximizes the expected value of u_i :

$$\text{For any } \mathbf{p} \text{ and } \mathbf{q} \text{ in } \Delta(\mathcal{A}), \quad (\mathbf{p} \succeq_i \mathbf{q}) \iff (\mathbb{E}(u_i, \mathbf{p}) \geq \mathbb{E}(u_i, \mathbf{q}))$$

Harsanyi's social aggregation theorem: hypotheses

(10/47)

Let $\mathcal{I} := \{\text{Ann, Bob, Carl, } \dots\}$ be a set of people.

Let \mathcal{A} be a set of **social outcomes**. Each social outcome in \mathcal{A} yields an outcome for each person in \mathcal{I} , determining her consumption bundle, health, autonomy and any other factors relevant to her well-being.

Let $\Delta(\mathcal{A})$ be the space of all lotteries over \mathcal{A} (i.e. "social lotteries").

Each policy corresponds to some element of $\Delta(\mathcal{A})$.

To choose the "best" policy, we need a *social preference order* \succeq on $\Delta(\mathcal{A})$.

If \succeq satisfies the vNM axioms of "rationality", then there is some *ex ante* social welfare function (SWF) W on \mathcal{A} such that \succeq maximizes the expected value of W :

$$\text{For any } \mathbf{p} \text{ and } \mathbf{q} \text{ in } \Delta(\mathcal{A}), \quad (\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(W, \mathbf{p}) \geq \mathbb{E}(W, \mathbf{q})).$$

Every individual i in \mathcal{I} has a preference order \succeq_i over $\Delta(\mathcal{A})$.

If \succeq_i satisfies the vNM axioms of "rationality", then there is some utility function u_i on \mathcal{A} such that \succeq_i maximizes the expected value of u_i :

$$\text{For any } \mathbf{p} \text{ and } \mathbf{q} \text{ in } \Delta(\mathcal{A}), \quad (\mathbf{p} \succeq_i \mathbf{q}) \iff (\mathbb{E}(u_i, \mathbf{p}) \geq \mathbb{E}(u_i, \mathbf{q}))$$

Harsanyi's social aggregation theorem: hypotheses

(10/47)

Let $\mathcal{I} := \{\text{Ann, Bob, Carl, } \dots\}$ be a set of people.

Let \mathcal{A} be a set of social outcomes. Each social outcome in \mathcal{A} yields an outcome for each person in \mathcal{I} , determining her consumption bundle, health, autonomy and any other factors relevant to her well-being.

Let $\Delta(\mathcal{A})$ be the space of all lotteries over \mathcal{A} (i.e. “social lotteries”).

Each policy corresponds to some element of $\Delta(\mathcal{A})$.

To choose the “best” policy, we need a *social preference order* \succeq on $\Delta(\mathcal{A})$.

If \succeq satisfies the vNM axioms of “rationality”, then there is some *ex ante* social welfare function (SWF) W on \mathcal{A} such that \succeq maximizes the expected value of W :

$$\text{For any } \mathbf{p} \text{ and } \mathbf{q} \text{ in } \Delta(\mathcal{A}), \quad (\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(W, \mathbf{p}) \geq \mathbb{E}(W, \mathbf{q})).$$

Every individual i in \mathcal{I} has a preference order \succeq_i over $\Delta(\mathcal{A})$.

If \succeq_i satisfies the vNM axioms of “rationality”, then there is some utility function u_i on \mathcal{A} such that \succeq_i maximizes the expected value of u_i :

$$\text{For any } \mathbf{p} \text{ and } \mathbf{q} \text{ in } \Delta(\mathcal{A}), \quad (\mathbf{p} \succeq_i \mathbf{q}) \iff (\mathbb{E}(u_i, \mathbf{p}) \geq \mathbb{E}(u_i, \mathbf{q}))$$

Harsanyi's social aggregation theorem: hypotheses

(10/47)

Let $\mathcal{I} := \{\text{Ann, Bob, Carl, } \dots\}$ be a set of people.

Let \mathcal{A} be a set of social outcomes. Each social outcome in \mathcal{A} yields an outcome for each person in \mathcal{I} , determining her consumption bundle, health, autonomy and any other factors relevant to her well-being.

Let $\Delta(\mathcal{A})$ be the space of all lotteries over \mathcal{A} (i.e. “social lotteries”).

Each policy corresponds to some element of $\Delta(\mathcal{A})$.

To choose the “best” policy, we need a *social preference order* \succeq on $\Delta(\mathcal{A})$.

If \succeq satisfies the vNM axioms of “rationality”, then there is some *ex ante* social welfare function (SWF) W on \mathcal{A} such that \succeq maximizes the expected value of W :

$$\text{For any } \mathbf{p} \text{ and } \mathbf{q} \text{ in } \Delta(\mathcal{A}), \quad (\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(W, \mathbf{p}) \geq \mathbb{E}(W, \mathbf{q})).$$

Every individual i in \mathcal{I} has a preference order \succeq_i over $\Delta(\mathcal{A})$.

If \succeq_i satisfies the vNM axioms of “rationality”, then there is some utility function u_i on \mathcal{A} such that \succeq_i maximizes the expected value of u_i :

$$\text{For any } \mathbf{p} \text{ and } \mathbf{q} \text{ in } \Delta(\mathcal{A}), \quad (\mathbf{p} \succeq_i \mathbf{q}) \iff (\mathbb{E}(u_i, \mathbf{p}) \geq \mathbb{E}(u_i, \mathbf{q}))$$

Harsanyi's social aggregation theorem: hypotheses

(10/47)

Let $\mathcal{I} := \{\text{Ann, Bob, Carl, } \dots\}$ be a set of people.

Let \mathcal{A} be a set of social outcomes. Each social outcome in \mathcal{A} yields an outcome for each person in \mathcal{I} , determining her consumption bundle, health, autonomy and any other factors relevant to her well-being.

Let $\Delta(\mathcal{A})$ be the space of all lotteries over \mathcal{A} (i.e. “social lotteries”).

Each policy corresponds to some element of $\Delta(\mathcal{A})$.

To choose the “best” policy, we need a *social preference order* \succeq on $\Delta(\mathcal{A})$.

If \succeq satisfies the vNM axioms of “rationality”, then there is some **ex ante social welfare function** (SWF) W on \mathcal{A} such that \succeq maximizes the expected value of W :

$$\text{For any } \mathbf{p} \text{ and } \mathbf{q} \text{ in } \Delta(\mathcal{A}), \quad (\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(W, \mathbf{p}) \geq \mathbb{E}(W, \mathbf{q})).$$

Every individual i in \mathcal{I} has a preference order \succeq_i over $\Delta(\mathcal{A})$.

If \succeq_i satisfies the vNM axioms of “rationality”, then there is some utility function u_i on \mathcal{A} such that \succeq_i maximizes the expected value of u_i :

$$\text{For any } \mathbf{p} \text{ and } \mathbf{q} \text{ in } \Delta(\mathcal{A}), \quad (\mathbf{p} \succeq_i \mathbf{q}) \iff (\mathbb{E}(u_i, \mathbf{p}) \geq \mathbb{E}(u_i, \mathbf{q}))$$

Harsanyi's social aggregation theorem: hypotheses

(10/47)

Let $\mathcal{I} := \{\text{Ann, Bob, Carl, } \dots\}$ be a set of people.

Let \mathcal{A} be a set of social outcomes. Each social outcome in \mathcal{A} yields an outcome for each person in \mathcal{I} , determining her consumption bundle, health, autonomy and any other factors relevant to her well-being.

Let $\Delta(\mathcal{A})$ be the space of all lotteries over \mathcal{A} (i.e. “social lotteries”).

Each policy corresponds to some element of $\Delta(\mathcal{A})$.

To choose the “best” policy, we need a *social preference order* \succeq on $\Delta(\mathcal{A})$.

If \succeq satisfies the vNM axioms of “rationality”, then there is some *ex ante* social welfare function (SWF) W on \mathcal{A} such that \succeq maximizes the expected value of W :

$$\text{For any } \mathbf{p} \text{ and } \mathbf{q} \text{ in } \Delta(\mathcal{A}), \quad (\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(W, \mathbf{p}) \geq \mathbb{E}(W, \mathbf{q})).$$

Every individual i in \mathcal{I} has a preference order \succeq_i over $\Delta(\mathcal{A})$.

If \succeq_i satisfies the vNM axioms of “rationality”, then there is some utility function u_i on \mathcal{A} such that \succeq_i maximizes the expected value of u_i :

$$\text{For any } \mathbf{p} \text{ and } \mathbf{q} \text{ in } \Delta(\mathcal{A}), \quad (\mathbf{p} \succeq_i \mathbf{q}) \iff (\mathbb{E}(u_i, \mathbf{p}) \geq \mathbb{E}(u_i, \mathbf{q}))$$

Harsanyi's social aggregation theorem: hypotheses

(10/47)

Let $\mathcal{I} := \{\text{Ann, Bob, Carl, } \dots\}$ be a set of people.

Let \mathcal{A} be a set of social outcomes. Each social outcome in \mathcal{A} yields an outcome for each person in \mathcal{I} , determining her consumption bundle, health, autonomy and any other factors relevant to her well-being.

Let $\Delta(\mathcal{A})$ be the space of all lotteries over \mathcal{A} (i.e. “social lotteries”).

Each policy corresponds to some element of $\Delta(\mathcal{A})$.

To choose the “best” policy, we need a *social preference order* \succeq on $\Delta(\mathcal{A})$.

If \succeq satisfies the vNM axioms of “rationality”, then there is some *ex ante* social welfare function (SWF) W on \mathcal{A} such that \succeq maximizes the expected value of W :

$$\text{For any } \mathbf{p} \text{ and } \mathbf{q} \text{ in } \Delta(\mathcal{A}), \quad (\mathbf{p} \succeq \mathbf{q}) \iff (\mathbb{E}(W, \mathbf{p}) \geq \mathbb{E}(W, \mathbf{q})).$$

Every individual i in \mathcal{I} has a preference order \succeq_i over $\Delta(\mathcal{A})$.

If \succeq_i satisfies the vNM axioms of “rationality”, then there is some utility function u_i on \mathcal{A} such that \succeq_i maximizes the expected value of u_i :

$$\text{For any } \mathbf{p} \text{ and } \mathbf{q} \text{ in } \Delta(\mathcal{A}), \quad (\mathbf{p} \succeq_i \mathbf{q}) \iff (\mathbb{E}(u_i, \mathbf{p}) \geq \mathbb{E}(u_i, \mathbf{q})).$$

For any \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

Individual vNM: For all i in \mathcal{I} , $(\mathbf{p} \succeq_i \mathbf{q}) \Leftrightarrow (\mathbb{E}(u_i, \mathbf{p}) \geq \mathbb{E}(u_i, \mathbf{q}))$.

Group vNM: $(\mathbf{p} \succeq \mathbf{q}) \Leftrightarrow (\mathbb{E}(W, \mathbf{p}) \geq \mathbb{E}(W, \mathbf{q}))$.

Finally, we suppose \succeq satisfies the *ex ante* Pareto principle:

XAP: For any \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$, if $\mathbf{p} \succeq_i \mathbf{q}$ for all i in \mathcal{I} , then $\mathbf{p} \succeq \mathbf{q}$.
If, furthermore, $\mathbf{p} \succ_i \mathbf{q}$ for some i in \mathcal{I} , then $\mathbf{p} \succ \mathbf{q}$.

Harsanyi's Social Aggregation Theorem (1955) Suppose that the individual preferences \succeq_i all satisfy **Individual vNM** and the social preference \succeq satisfies **Group vNM** and **XAP**.

Then the *ex ante* social welfare function W is weighted utilitarian. That is, there exist weights $c_i > 0$ for all individuals i in \mathcal{I} , and some constant K , such that, for any outcome a in \mathcal{A} , we have

$$W(a) = K + c_{\text{Ann}} u_{\text{Ann}}(a) + c_{\text{Bob}} u_{\text{Bob}}(a) + c_{\text{Carl}} u_{\text{Carl}}(a) + \cdots$$

For any \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

Individual vNM: For all i in \mathcal{I} , $(\mathbf{p} \succeq_i \mathbf{q}) \Leftrightarrow (\mathbb{E}(u_i, \mathbf{p}) \geq \mathbb{E}(u_i, \mathbf{q}))$.

Group vNM: $(\mathbf{p} \succeq \mathbf{q}) \Leftrightarrow (\mathbb{E}(W, \mathbf{p}) \geq \mathbb{E}(W, \mathbf{q}))$.

Finally, we suppose \succeq satisfies the **ex ante Pareto principle**:

XAP: For any \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$, if $\mathbf{p} \succeq_i \mathbf{q}$ for all i in \mathcal{I} , then $\mathbf{p} \succeq \mathbf{q}$.
If, furthermore, $\mathbf{p} \succ_i \mathbf{q}$ for some i in \mathcal{I} , then $\mathbf{p} \succ \mathbf{q}$.

Harsanyi's Social Aggregation Theorem (1955) Suppose that the individual preferences \succeq_i all satisfy **Individual vNM** and the social preference \succeq satisfies **Group vNM** and **XAP**.

Then the ex ante social welfare function W is weighted utilitarian. That is, there exist weights $c_i > 0$ for all individuals i in \mathcal{I} , and some constant K , such that, for any outcome a in \mathcal{A} , we have

$$W(a) = K + c_{\text{Ann}} u_{\text{Ann}}(a) + c_{\text{Bob}} u_{\text{Bob}}(a) + c_{\text{Carl}} u_{\text{Carl}}(a) + \dots$$

For any \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

Individual vNM: For all i in \mathcal{I} , $(\mathbf{p} \succeq_i \mathbf{q}) \Leftrightarrow (\mathbb{E}(u_i, \mathbf{p}) \geq \mathbb{E}(u_i, \mathbf{q}))$.

Group vNM: $(\mathbf{p} \succeq \mathbf{q}) \Leftrightarrow (\mathbb{E}(W, \mathbf{p}) \geq \mathbb{E}(W, \mathbf{q}))$.

Finally, we suppose \succeq satisfies the *ex ante* Pareto principle:

XAP: For any \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$, if $\mathbf{p} \succeq_i \mathbf{q}$ for all i in \mathcal{I} , then $\mathbf{p} \succeq \mathbf{q}$.
If, furthermore, $\mathbf{p} \succ_i \mathbf{q}$ for some i in \mathcal{I} , then $\mathbf{p} \succ \mathbf{q}$.

Harsanyi's Social Aggregation Theorem (1955) Suppose that the individual preferences \succeq_i all satisfy **Individual vNM** and the social preference \succeq satisfies **Group vNM** and **XAP**.

Then the *ex ante* social welfare function W is weighted utilitarian. That is, there exist weights $c_i > 0$ for all individuals i in \mathcal{I} , and some constant K , such that, for any outcome a in \mathcal{A} , we have

$$W(a) = K + c_{\text{Ann}} u_{\text{Ann}}(a) + c_{\text{Bob}} u_{\text{Bob}}(a) + c_{\text{Carl}} u_{\text{Carl}}(a) + \dots$$

For any \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$,

Individual vNM: For all i in \mathcal{I} , $(\mathbf{p} \succeq_i \mathbf{q}) \Leftrightarrow (\mathbb{E}(u_i, \mathbf{p}) \geq \mathbb{E}(u_i, \mathbf{q}))$.

Group vNM: $(\mathbf{p} \succeq \mathbf{q}) \Leftrightarrow (\mathbb{E}(W, \mathbf{p}) \geq \mathbb{E}(W, \mathbf{q}))$.

Finally, we suppose \succeq satisfies the *ex ante* Pareto principle:

XAP: For any \mathbf{p} and \mathbf{q} in $\Delta(\mathcal{A})$, if $\mathbf{p} \succeq_i \mathbf{q}$ for all i in \mathcal{I} , then $\mathbf{p} \succeq \mathbf{q}$.
If, furthermore, $\mathbf{p} \succ_i \mathbf{q}$ for some i in \mathcal{I} , then $\mathbf{p} \succ \mathbf{q}$.

Harsanyi's Social Aggregation Theorem (1955) Suppose that the individual preferences \succeq_i all satisfy **Individual vNM** and the social preference \succeq satisfies **Group vNM** and **XAP**.

Then the *ex ante* social welfare function W is **weighted utilitarian**. That is, there exist weights $c_i > 0$ for all individuals i in \mathcal{I} , and some constant K , such that, for any outcome a in \mathcal{A} , we have

$$W(a) = K + c_{\text{Ann}} u_{\text{Ann}}(a) + c_{\text{Bob}} u_{\text{Bob}}(a) + c_{\text{Carl}} u_{\text{Carl}}(a) + \dots$$

- ▶ Harsanyi interpreted his SAT as an argument for utilitarianism.
- ▶ But Sen argued that Harsanyi's SAT was merely a hollow formal "representation" result, without real ethical content.
- ▶ In particular, Sen argued that a vNM utility function conflated an individual's risk preferences with her well-being; hence an aggregate of vNM utilites is not really a measure of social welfare.
- ▶ But the SAT is also vulnerable to another serious criticism.
- ▶ It assumes a setting of *risk*, where probabilities are *objective*, and known to everyone.
- ▶ But many policy problems involve *uncertainty*.
- ▶ Each individual might have her own *subjective* probabilities (*à la* Savage), but there is no reason for these subjective probabilities to *agree*.
- ▶ It turns out that Harsanyi's theorem is fundamentally incompatible with such heterogeneity of beliefs.

- ▶ Harsanyi interpreted his SAT as an argument for utilitarianism.
- ▶ But Sen argued that Harsanyi's SAT was merely a hollow formal “representation” result, without real ethical content.
- ▶ In particular, Sen argued that a vNM utility function conflated an individual's risk preferences with her well-being; hence an aggregate of vNM utilities is not really a measure of social welfare.
- ▶ But the SAT is also vulnerable to another serious criticism.
- ▶ It assumes a setting of *risk*, where probabilities are *objective*, and known to everyone.
- ▶ But many policy problems involve *uncertainty*.
- ▶ Each individual might have her own *subjective* probabilities (*à la* Savage), but there is no reason for these subjective probabilities to *agree*.
- ▶ It turns out that Harsanyi's theorem is fundamentally incompatible with such heterogeneity of beliefs.

- ▶ Harsanyi interpreted his SAT as an argument for utilitarianism.
- ▶ But Sen argued that Harsanyi's SAT was merely a hollow formal “representation” result, without real ethical content.
- ▶ In particular, Sen argued that a vNM utility function conflated an individual's risk preferences with her well-being; hence an aggregate of vNM utilities is not really a measure of social welfare.
- ▶ But the SAT is also vulnerable to another serious criticism.
- ▶ It assumes a setting of *risk*, where probabilities are *objective*, and known to everyone.
- ▶ But many policy problems involve *uncertainty*.
- ▶ Each individual might have her own *subjective* probabilities (*à la* Savage), but there is no reason for these subjective probabilities to *agree*.
- ▶ It turns out that Harsanyi's theorem is fundamentally incompatible with such heterogeneity of beliefs.

- ▶ Harsanyi interpreted his SAT as an argument for utilitarianism.
- ▶ But Sen argued that Harsanyi's SAT was merely a hollow formal “representation” result, without real ethical content.
- ▶ In particular, Sen argued that a vNM utility function conflated an individual's risk preferences with her well-being; hence an aggregate of vNM utilities is not really a measure of social welfare.
- ▶ But the SAT is also vulnerable to another serious criticism.
- ▶ It assumes a setting of *risk*, where probabilities are *objective*, and known to everyone.
- ▶ But many policy problems involve *uncertainty*.
- ▶ Each individual might have her own *subjective* probabilities (*à la* Savage), but there is no reason for these subjective probabilities to *agree*.
- ▶ It turns out that Harsanyi's theorem is fundamentally incompatible with such heterogeneity of beliefs.

- ▶ Harsanyi interpreted his SAT as an argument for utilitarianism.
- ▶ But Sen argued that Harsanyi's SAT was merely a hollow formal “representation” result, without real ethical content.
- ▶ In particular, Sen argued that a vNM utility function conflated an individual's risk preferences with her well-being; hence an aggregate of vNM utilities is not really a measure of social welfare.
- ▶ But the SAT is also vulnerable to another serious criticism.
- ▶ It assumes a setting of *risk*, where probabilities are *objective*, and known to everyone.
- ▶ But many policy problems involve *uncertainty*.
- ▶ Each individual might have her own *subjective* probabilities (*à la* Savage), but there is no reason for these subjective probabilities to agree.
- ▶ It turns out that Harsanyi's theorem is fundamentally incompatible with such heterogeneity of beliefs.

- ▶ Harsanyi interpreted his SAT as an argument for utilitarianism.
- ▶ But Sen argued that Harsanyi's SAT was merely a hollow formal “representation” result, without real ethical content.
- ▶ In particular, Sen argued that a vNM utility function conflated an individual's risk preferences with her well-being; hence an aggregate of vNM utilities is not really a measure of social welfare.
- ▶ But the SAT is also vulnerable to another serious criticism.
- ▶ It assumes a setting of *risk*, where probabilities are *objective*, and known to everyone.
- ▶ But many policy problems involve *uncertainty*.
- ▶ Each individual might have her own *subjective* probabilities (*à la* Savage), but there is no reason for these subjective probabilities to agree.
- ▶ It turns out that Harsanyi's theorem is fundamentally incompatible with such heterogeneity of beliefs.

- ▶ Harsanyi interpreted his SAT as an argument for utilitarianism.
- ▶ But Sen argued that Harsanyi's SAT was merely a hollow formal “representation” result, without real ethical content.
- ▶ In particular, Sen argued that a vNM utility function conflated an individual's risk preferences with her well-being; hence an aggregate of vNM utilities is not really a measure of social welfare.
- ▶ But the SAT is also vulnerable to another serious criticism.
- ▶ It assumes a setting of *risk*, where probabilities are *objective*, and known to everyone.
- ▶ But many policy problems involve *uncertainty*.
- ▶ Each individual might have her own *subjective* probabilities (*à la* Savage), but there is no reason for these subjective probabilities to *agree*.
- ▶ It turns out that Harsanyi's theorem is fundamentally incompatible with such heterogeneity of beliefs.

- ▶ Harsanyi interpreted his SAT as an argument for utilitarianism.
- ▶ But Sen argued that Harsanyi's SAT was merely a hollow formal “representation” result, without real ethical content.
- ▶ In particular, Sen argued that a vNM utility function conflated an individual's risk preferences with her well-being; hence an aggregate of vNM utilities is not really a measure of social welfare.
- ▶ But the SAT is also vulnerable to another serious criticism.
- ▶ It assumes a setting of *risk*, where probabilities are *objective*, and known to everyone.
- ▶ But many policy problems involve *uncertainty*.
- ▶ Each individual might have her own *subjective* probabilities (*à la* Savage), but there is no reason for these subjective probabilities to *agree*.
- ▶ It turns out that Harsanyi's theorem is fundamentally incompatible with such heterogeneity of beliefs.

Uncertainty: Heterogeneity and spurious unanimity

Let \mathcal{S} be an (infinite) set of states of the world. Let \mathcal{A} be a set of outcomes. Let $\mathcal{A}^{\mathcal{S}}$ be the set of all social prospects (mapping states to outcomes). Let \mathcal{I} be a set of people. For all i in \mathcal{I} , let \succeq_i be a preference order on $\mathcal{A}^{\mathcal{S}}$. If \succeq_i satisfies the Savage axioms of “rationality”, then there is a utility function u_i on \mathcal{A} , and a probability distribution p_i on \mathcal{S} yielding an SEU representation for \succeq_i :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq_i g) \iff (\mathbb{E}(f|u_i, p_i) \geq \mathbb{E}(g|u_i, p_i)).$$

Let \succeq be a social preference order over $\mathcal{A}^{\mathcal{S}}$.

If \succeq satisfies the Savage axioms of “rationality”, then there is an *ex ante* social welfare function W on \mathcal{A} , and a probability distribution P on \mathcal{S} , yielding an SEU representation for \succeq :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq g) \iff (\mathbb{E}(f|W, P) \geq \mathbb{E}(g|W, P)).$$

(For technical reasons, we also assume that \mathcal{S} is infinite, and p_i and P are *nonatomic* —i.e. $p_i(s) = 0 = P(s)$ for all s in \mathcal{S} .)

Let \mathcal{S} be an (infinite) set of states of the world. Let \mathcal{A} be a set of outcomes. Let $\mathcal{A}^{\mathcal{S}}$ be the set of all social prospects (mapping states to outcomes). Let \mathcal{I} be a set of people. For all i in \mathcal{I} , let \succeq_i be a preference order on $\mathcal{A}^{\mathcal{S}}$. If \succeq_i satisfies the Savage axioms of “rationality”, then there is a utility function u_i on \mathcal{A} , and a probability distribution p_i on \mathcal{S} yielding an SEU representation for \succeq_i :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq_i g) \iff (\mathbb{E}(f|u_i, p_i) \geq \mathbb{E}(g|u_i, p_i)).$$

Let \succeq be a social preference order over $\mathcal{A}^{\mathcal{S}}$.

If \succeq satisfies the Savage axioms of “rationality”, then there is an *ex ante* social welfare function W on \mathcal{A} , and a probability distribution P on \mathcal{S} , yielding an SEU representation for \succeq :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq g) \iff (\mathbb{E}(f|W, P) \geq \mathbb{E}(g|W, P)).$$

(For technical reasons, we also assume that \mathcal{S} is infinite, and p_i and P are *nonatomic* —i.e. $p_i(s) = 0 = P(s)$ for all s in \mathcal{S} .)

Let \mathcal{S} be an (infinite) set of states of the world. Let \mathcal{A} be a set of outcomes. Let $\mathcal{A}^{\mathcal{S}}$ be the set of all **social prospects** (mapping states to outcomes).

Let \mathcal{I} be a set of people. For all i in \mathcal{I} , let \succeq_i be a preference order on $\mathcal{A}^{\mathcal{S}}$. If \succeq_i satisfies the Savage axioms of “rationality”, then there is a utility function u_i on \mathcal{A} , and a probability distribution p_i on \mathcal{S} yielding an SEU representation for \succeq_i :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq_i g) \iff (\mathbb{E}(f|u_i, p_i) \geq \mathbb{E}(g|u_i, p_i)).$$

Let \succeq be a social preference order over $\mathcal{A}^{\mathcal{S}}$.

If \succeq satisfies the Savage axioms of “rationality”, then there is an *ex ante* social welfare function W on \mathcal{A} , and a probability distribution P on \mathcal{S} , yielding an SEU representation for \succeq :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq g) \iff (\mathbb{E}(f|W, P) \geq \mathbb{E}(g|W, P)).$$

(For technical reasons, we also assume that \mathcal{S} is infinite, and p_i and P are *nonatomic* —i.e. $p_i(s) = 0 = P(s)$ for all s in \mathcal{S} .)

Let \mathcal{S} be an (infinite) set of states of the world. Let \mathcal{A} be a set of outcomes. Let $\mathcal{A}^{\mathcal{S}}$ be the set of all social prospects (mapping states to outcomes). Let \mathcal{I} be a set of people. For all i in \mathcal{I} , let \succeq_i be a preference order on $\mathcal{A}^{\mathcal{S}}$. If \succeq_i satisfies the Savage axioms of “rationality”, then there is a utility function u_i on \mathcal{A} , and a probability distribution p_i on \mathcal{S} yielding an SEU representation for \succeq_i :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq_i g) \iff (\mathbb{E}(f|u_i, p_i) \geq \mathbb{E}(g|u_i, p_i)).$$

Let \succeq be a social preference order over $\mathcal{A}^{\mathcal{S}}$.

If \succeq satisfies the Savage axioms of “rationality”, then there is an *ex ante* social welfare function W on \mathcal{A} , and a probability distribution P on \mathcal{S} , yielding an SEU representation for \succeq :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq g) \iff (\mathbb{E}(f|W, P) \geq \mathbb{E}(g|W, P)).$$

(For technical reasons, we also assume that \mathcal{S} is infinite, and p_i and P are *nonatomic* —i.e. $p_i(s) = 0 = P(s)$ for all s in \mathcal{S} .)

Let \mathcal{S} be an (infinite) set of states of the world. Let \mathcal{A} be a set of outcomes. Let $\mathcal{A}^{\mathcal{S}}$ be the set of all social prospects (mapping states to outcomes). Let \mathcal{I} be a set of people. For all i in \mathcal{I} , let \succeq_i be a preference order on $\mathcal{A}^{\mathcal{S}}$. If \succeq_i satisfies the Savage axioms of “rationality”, then there is a utility function u_i on \mathcal{A} , and a probability distribution p_i on \mathcal{S} yielding an SEU representation for \succeq_i :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq_i g) \iff (\mathbb{E}(f|u_i, p_i) \geq \mathbb{E}(g|u_i, p_i)).$$

Let \succeq be a social preference order over $\mathcal{A}^{\mathcal{S}}$.

If \succeq satisfies the Savage axioms of “rationality”, then there is an *ex ante* social welfare function W on \mathcal{A} , and a probability distribution P on \mathcal{S} , yielding an SEU representation for \succeq :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq g) \iff (\mathbb{E}(f|W, P) \geq \mathbb{E}(g|W, P)).$$

(For technical reasons, we also assume that \mathcal{S} is infinite, and p_i and P are *nonatomic* —i.e. $p_i(s) = 0 = P(s)$ for all s in \mathcal{S} .)

Let \mathcal{S} be an (infinite) set of states of the world. Let \mathcal{A} be a set of outcomes. Let $\mathcal{A}^{\mathcal{S}}$ be the set of all social prospects (mapping states to outcomes). Let \mathcal{I} be a set of people. For all i in \mathcal{I} , let \succeq_i be a preference order on $\mathcal{A}^{\mathcal{S}}$. If \succeq_i satisfies the Savage axioms of “rationality”, then there is a utility function u_i on \mathcal{A} , and a probability distribution p_i on \mathcal{S} yielding an SEU representation for \succeq_i :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq_i g) \iff (\mathbb{E}(f|u_i, p_i) \geq \mathbb{E}(g|u_i, p_i)).$$

Let \succeq be a social preference order over $\mathcal{A}^{\mathcal{S}}$.

If \succeq satisfies the Savage axioms of “rationality”, then there is an *ex ante* social welfare function W on \mathcal{A} , and a probability distribution P on \mathcal{S} , yielding an SEU representation for \succeq :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq g) \iff (\mathbb{E}(f|W, P) \geq \mathbb{E}(g|W, P)).$$

(For technical reasons, we also assume that \mathcal{S} is infinite, and p_i and P are *nonatomic* —i.e. $p_i(s) = 0 = P(s)$ for all s in \mathcal{S} .)

Let \mathcal{S} be an (infinite) set of states of the world. Let \mathcal{A} be a set of outcomes. Let $\mathcal{A}^{\mathcal{S}}$ be the set of all social prospects (mapping states to outcomes). Let \mathcal{I} be a set of people. For all i in \mathcal{I} , let \succeq_i be a preference order on $\mathcal{A}^{\mathcal{S}}$. If \succeq_i satisfies the Savage axioms of “rationality”, then there is a utility function u_i on \mathcal{A} , and a probability distribution p_i on \mathcal{S} yielding an SEU representation for \succeq_i :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq_i g) \iff (\mathbb{E}(f|u_i, p_i) \geq \mathbb{E}(g|u_i, p_i)).$$

Let \succeq be a social preference order over $\mathcal{A}^{\mathcal{S}}$.

If \succeq satisfies the Savage axioms of “rationality”, then there is an *ex ante* social welfare function W on \mathcal{A} , and a probability distribution P on \mathcal{S} , yielding an SEU representation for \succeq :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq g) \iff (\mathbb{E}(f|W, P) \geq \mathbb{E}(g|W, P)).$$

(For technical reasons, we also assume that \mathcal{S} is infinite, and p_i and P are *nonatomic* —i.e. $p_i(s) = 0 = P(s)$ for all s in \mathcal{S} .)

Let \mathcal{S} be an (infinite) set of states of the world. Let \mathcal{A} be a set of outcomes. Let $\mathcal{A}^{\mathcal{S}}$ be the set of all social prospects (mapping states to outcomes). Let \mathcal{I} be a set of people. For all i in \mathcal{I} , let \succeq_i be a preference order on $\mathcal{A}^{\mathcal{S}}$. If \succeq_i satisfies the Savage axioms of “rationality”, then there is a utility function u_i on \mathcal{A} , and a probability distribution p_i on \mathcal{S} yielding an SEU representation for \succeq_i :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq_i g) \iff (\mathbb{E}(f|u_i, p_i) \geq \mathbb{E}(g|u_i, p_i)).$$

Let \succeq be a social preference order over $\mathcal{A}^{\mathcal{S}}$.

If \succeq satisfies the Savage axioms of “rationality”, then there is an *ex ante* social welfare function W on \mathcal{A} , and a probability distribution P on \mathcal{S} , yielding an SEU representation for \succeq :

$$\text{For all } f \text{ and } g \text{ in } \mathcal{A}^{\mathcal{S}}, \quad (f \succeq g) \iff (\mathbb{E}(f|W, P) \geq \mathbb{E}(g|W, P)).$$

(For technical reasons, we also assume that \mathcal{S} is infinite, and p_i and P are *nonatomic* —i.e. $p_i(s) = 0 = P(s)$ for all s in \mathcal{S} .)

Individual SEU: For all i in \mathcal{I} , and all f and g in \mathcal{A}^S ,

$$(f \succeq_i g) \iff (\mathbb{E}(f|u_i, p_i) \geq \mathbb{E}(g|u_i, p_i)).$$

Group SEU: $(f \succeq g) \iff (\mathbb{E}(f|W, P) \geq \mathbb{E}(g|W, P)).$

Finally, we suppose \succeq satisfies the *ex ante* Pareto principle:

XAP: For any f and g in \mathcal{A}^S , if $f \succeq_i g$ for all i in \mathcal{I} , then $f \succeq g$.
If, furthermore, $f \succ_i g$ for some i in \mathcal{I} , then $f \succ g$.

Bayesian Social Aggregation Theorem: Suppose that \succeq_i all satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**.

Then the social welfare function W is weighted utilitarian: there exist weights $c_i > 0$ for all individuals i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

However, if the individual utility functions are sufficiently diverse, then this can only happen if everyone has the *same beliefs*:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then **Group SEU** and **XAP** are *incompatible*.

Individual SEU: For all i in \mathcal{I} , and all f and g in \mathcal{A}^S ,

$$(f \succeq_i g) \iff (\mathbb{E}(f|u_i, p_i) \geq \mathbb{E}(g|u_i, p_i)).$$

Group SEU: $(f \succeq g) \iff (\mathbb{E}(f|W, P) \geq \mathbb{E}(g|W, P)).$

Finally, we suppose \succeq satisfies the **ex ante Pareto principle**:

XAP: For any f and g in \mathcal{A}^S , if $f \succeq_i g$ for all i in \mathcal{I} , then $f \succeq g$.
If, furthermore, $f \succ_i g$ for some i in \mathcal{I} , then $f \succ g$.

Bayesian Social Aggregation Theorem: Suppose that \succeq_i all satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**.

Then the social welfare function W is weighted utilitarian: there exist weights $c_i > 0$ for all individuals i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

However, if the individual utility functions are sufficiently diverse, then this can only happen if everyone has the **same beliefs**:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then **Group SEU** and **XAP** are **incompatible**.

Individual SEU: For all i in \mathcal{I} , and all f and g in \mathcal{A}^S ,

$$(f \succeq_i g) \iff (\mathbb{E}(f|u_i, p_i) \geq \mathbb{E}(g|u_i, p_i)).$$

Group SEU: $(f \succeq g) \iff (\mathbb{E}(f|W, P) \geq \mathbb{E}(g|W, P)).$

Finally, we suppose \succeq satisfies the *ex ante* Pareto principle:

XAP: For any f and g in \mathcal{A}^S , if $f \succeq_i g$ for all i in \mathcal{I} , then $f \succeq g$.
If, furthermore, $f \succ_i g$ for some i in \mathcal{I} , then $f \succ g$.

Bayesian Social Aggregation Theorem: Suppose that \succeq_i all satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**.

Then the social welfare function W is weighted utilitarian: there exist weights $c_i > 0$ for all individuals i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

However, if the individual utility functions are sufficiently diverse, then this can only happen if everyone has the *same beliefs*:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then Group SEU and XAP are *incompatible*.

Individual SEU: For all i in \mathcal{I} , and all f and g in \mathcal{A}^S ,

$$(f \succeq_i g) \iff (\mathbb{E}(f|u_i, p_i) \geq \mathbb{E}(g|u_i, p_i)).$$

Group SEU: $(f \succeq g) \iff (\mathbb{E}(f|W, P) \geq \mathbb{E}(g|W, P)).$

Finally, we suppose \succeq satisfies the *ex ante* Pareto principle:

XAP: For any f and g in \mathcal{A}^S , if $f \succeq_i g$ for all i in \mathcal{I} , then $f \succeq g$.
If, furthermore, $f \succ_i g$ for some i in \mathcal{I} , then $f \succ g$.

Bayesian Social Aggregation Theorem: Suppose that \succeq_i all satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**.
Then the social welfare function W is *weighted utilitarian*: there exist weights $c_i > 0$ for all individuals i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

However, if the individual utility functions are sufficiently diverse, then this can only happen if everyone has the *same beliefs*:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then Group SEU and XAP are *incompatible*.

Individual SEU: For all i in \mathcal{I} , and all f and g in \mathcal{A}^S ,

$$(f \succeq_i g) \iff (\mathbb{E}(f|u_i, p_i) \geq \mathbb{E}(g|u_i, p_i)).$$

Group SEU: $(f \succeq g) \iff (\mathbb{E}(f|W, P) \geq \mathbb{E}(g|W, P)).$

Finally, we suppose \succeq satisfies the *ex ante* Pareto principle:

XAP: For any f and g in \mathcal{A}^S , if $f \succeq_i g$ for all i in \mathcal{I} , then $f \succeq g$.
If, furthermore, $f \succ_i g$ for some i in \mathcal{I} , then $f \succ g$.

Bayesian Social Aggregation Theorem: Suppose that \succeq_i all satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**.
Then the social welfare function W is weighted utilitarian: there exist weights $c_i > 0$ for all individuals i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

However, if the individual utility functions are sufficiently diverse, then this can only happen if everyone has the **same beliefs**:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then Group SEU and XAP are **incompatible**.

Individual SEU: For all i in \mathcal{I} , and all f and g in \mathcal{A}^S ,

$$(f \succeq_i g) \iff (\mathbb{E}(f|u_i, p_i) \geq \mathbb{E}(g|u_i, p_i)).$$

Group SEU: $(f \succeq g) \iff (\mathbb{E}(f|W, P) \geq \mathbb{E}(g|W, P)).$

Finally, we suppose \succeq satisfies the *ex ante* Pareto principle:

XAP: For any f and g in \mathcal{A}^S , if $f \succeq_i g$ for all i in \mathcal{I} , then $f \succeq g$.
If, furthermore, $f \succ_i g$ for some i in \mathcal{I} , then $f \succ g$.

Bayesian Social Aggregation Theorem: Suppose that \succeq_i all satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**.

Then the social welfare function W is weighted utilitarian: there exist weights $c_i > 0$ for all individuals i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

However, if the individual utility functions are sufficiently diverse, then this can only happen if everyone has the **same beliefs**:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then Group SEU and XAP are **incompatible**.

Bayesian Social Aggregation Theorem: Suppose all the \succeq_i satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**. Then there exist weights $c_i > 0$ for all i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

However, if $\{u_i\}_{i \in \mathcal{I}}$ are sufficiently diverse, this can only happen if everyone has the same beliefs:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then Group SEU and XAP are incompatible.

This “probability agreement” result was intuited by Raiffa (1968). Early versions were obtained by Starr (1973), Hylland & Zeckhauser (1979), Hammond (1981), Broome (1991), and others.

The version we present here is due to Mongin (1995).

It is highly unlikely that all agents will agree about subjective probabilities.

Thus, this result is interpreted as an *impossibility theorem*, saying, in effect,

“Consistent Bayesian aggregation is impossible if agents have heterogeneous beliefs.”

Bayesian Social Aggregation Theorem: Suppose all the \succeq_i satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**. Then there exist weights $c_i > 0$ for all i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

However, if $\{u_i\}_{i \in \mathcal{I}}$ are sufficiently diverse, this can only happen if everyone has the **same beliefs**:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then **Group SEU** and **XAP** are incompatible.

This “probability agreement” result was intuited by Raiffa (1968). Early versions were obtained by Starr (1973), Hylland & Zeckhauser (1979), Hammond (1981), Broome (1991), and others.

The version we present here is due to Mongin (1995).

It is highly unlikely that all agents will **agree** about subjective probabilities.

Thus, this result is interpreted as an *impossibility theorem*, saying, in effect,

“Consistent Bayesian aggregation is impossible if agents have heterogeneous beliefs.”

Bayesian Social Aggregation Theorem: Suppose all the \succeq_i satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**. Then there exist weights $c_i > 0$ for all i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

However, if $\{u_i\}_{i \in \mathcal{I}}$ are sufficiently diverse, this can only happen if everyone has the same beliefs:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then Group SEU and XAP are *incompatible*.

This “probability agreement” result was intuited by Raiffa (1968). Early versions were obtained by Starr (1973), Hylland & Zeckhauser (1979), Hammond (1981), Broome (1991), and others.

The version we present here is due to Mongin (1995).

It is highly unlikely that all agents will agree about subjective probabilities.

Thus, this result is interpreted as an *impossibility theorem*, saying, in effect,

“Consistent Bayesian aggregation is impossible if agents have heterogeneous beliefs.”

Bayesian Social Aggregation Theorem: Suppose all the \succeq_i satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**. Then there exist weights $c_i > 0$ for all i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

If $\{u_i\}_{i \in \mathcal{I}}$ are sufficiently diverse, this can only happen if everyone has the same beliefs:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then **Group SEU** and **XAP** are incompatible.

Possible escapes: Perhaps **Individual SEU** and **Group SEU** demand too much “rationality” from the agents.

Can we escape by requiring a weaker degree of “rationality”?

Also, we assumed \mathcal{S} was infinite and p_i and P were nonatomic.

Are these “technical” assumptions driving the result?

Finally, we assumed that \succeq_i and \succeq were defined for *all* prospects in $\mathcal{A}^{\mathcal{S}}$. But many prospects in $\mathcal{A}^{\mathcal{S}}$ may be unfeasible, or even logically impossible.

Can we escape by weakening this “universal domain” assumption?

Sadly, no. The conclusions of the BSAT (i.e. utilitarianism + “probability agreement”) remain in force **even** if we relax the “universal domain” assumption, allow \mathcal{S} to be finite, and only require the agents to satisfy *Statewise Dominance* (“minimal rationality”), rather than being SEU maximizers (M&P, 2013, extending Blackorby, Donaldson & Mongin, 2004).

Bayesian Social Aggregation Theorem: Suppose all the \succeq_i satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**. Then there exist weights $c_i > 0$ for all i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

If $\{u_i\}_{i \in \mathcal{I}}$ are sufficiently diverse, this can only happen if everyone has the **same beliefs**:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then **Group SEU** and **XAP** are incompatible.

Possible escapes: Perhaps **Individual SEU** and **Group SEU** demand too much “rationality” from the agents.

Can we escape by requiring a weaker degree of “rationality”?

Also, we assumed \mathcal{S} was infinite and p_i and P were nonatomic.

Are these “technical” assumptions driving the result?

Finally, we assumed that \succeq_i and \succeq were defined for *all* prospects in $\mathcal{A}^{\mathcal{S}}$. But many prospects in $\mathcal{A}^{\mathcal{S}}$ may be unfeasible, or even logically impossible.

Can we escape by weakening this “universal domain” assumption?

Sadly, no. The conclusions of the BSAT (i.e. utilitarianism + “probability agreement”) remain in force **even** if we relax the “universal domain” assumption, allow \mathcal{S} to be finite, and only require the agents to satisfy *Statewise Dominance* (“minimal rationality”), rather than being SEU maximizers (M&P, 2013, extending Blackorby, Donaldson & Mongin, 2004).

Bayesian Social Aggregation Theorem: Suppose all the \succeq_i satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**. Then there exist weights $c_i > 0$ for all i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

If $\{u_i\}_{i \in \mathcal{I}}$ are sufficiently diverse, this can only happen if everyone has the same beliefs:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then **Group SEU** and **XAP** are **incompatible**.

Possible escapes: Perhaps **Individual SEU** and **Group SEU** demand too much “rationality” from the agents.

Can we escape by requiring a weaker degree of “rationality”?

Also, we assumed \mathcal{S} was infinite and p_i and P were nonatomic.

Are these “technical” assumptions driving the result?

Finally, we assumed that \succeq_i and \succeq were defined for *all* prospects in $\mathcal{A}^{\mathcal{S}}$. But many prospects in $\mathcal{A}^{\mathcal{S}}$ may be unfeasible, or even logically impossible.

Can we escape by weakening this “universal domain” assumption?

Sadly, no. The conclusions of the BSAT (i.e. utilitarianism + “probability agreement”) remain in force **even** if we relax the “universal domain” assumption, allow \mathcal{S} to be finite, and only require the agents to satisfy *Statewise Dominance* (“minimal rationality”), rather than being SEU maximizers (M&P, 2013, extending Blackorby, Donaldson & Mongin, 2004).

Bayesian Social Aggregation Theorem: Suppose all the \succeq_i satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**. Then there exist weights $c_i > 0$ for all i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

If $\{u_i\}_{i \in \mathcal{I}}$ are sufficiently diverse, this can only happen if everyone has the same beliefs:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then **Group SEU** and **XAP** are incompatible.

Possible escapes: Perhaps **Individual SEU** and **Group SEU** demand too much “rationality” from the agents.

Can we escape by requiring a weaker degree of “rationality”?

Also, we assumed \mathcal{S} was infinite and p_i and P were nonatomic.

Are these “technical” assumptions driving the result?

Finally, we assumed that \succeq_i and \succeq were defined for *all* prospects in $\mathcal{A}^{\mathcal{S}}$. But many prospects in $\mathcal{A}^{\mathcal{S}}$ may be unfeasible, or even logically impossible.

Can we escape by weakening this “universal domain” assumption?

Sadly, no. The conclusions of the BSAT (i.e. utilitarianism + “probability agreement”) remain in force **even** if we relax the “universal domain” assumption, allow \mathcal{S} to be finite, and only require the agents to satisfy *Statewise Dominance* (“minimal rationality”), rather than being SEU maximizers (M&P, 2013, extending Blackorby, Donaldson & Mongin, 2004).

Bayesian Social Aggregation Theorem: Suppose all the \succeq_i satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**. Then there exist weights $c_i > 0$ for all i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

If $\{u_i\}_{i \in \mathcal{I}}$ are sufficiently diverse, this can only happen if everyone has the same beliefs:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then **Group SEU** and **XAP** are incompatible.

Possible escapes: Perhaps **Individual SEU** and **Group SEU** demand too much “rationality” from the agents.

Can we escape by requiring a weaker degree of “rationality”?

Also, we assumed \mathcal{S} was infinite and p_i and P were nonatomic.

Are these “technical” assumptions driving the result?

Finally, we assumed that \succeq_i and \succeq were defined for *all* prospects in $\mathcal{A}^{\mathcal{S}}$. But many prospects in $\mathcal{A}^{\mathcal{S}}$ may be unfeasible, or even logically impossible.

Can we escape by weakening this “universal domain” assumption?

Sadly, no. The conclusions of the BSAT (i.e. utilitarianism + “probability agreement”) remain in force **even** if we relax the “universal domain” assumption, allow \mathcal{S} to be finite, and only require the agents to satisfy *Statewise Dominance* (“minimal rationality”), rather than being SEU maximizers (M&P, 2013, extending Blackorby, Donaldson & Mongin, 2004).

Bayesian Social Aggregation Theorem: Suppose all the \succeq_i satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**. Then there exist weights $c_i > 0$ for all i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

If $\{u_i\}_{i \in \mathcal{I}}$ are sufficiently diverse, this can only happen if everyone has the same beliefs:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then **Group SEU** and **XAP** are incompatible.

Possible escapes: Perhaps **Individual SEU** and **Group SEU** demand too much “rationality” from the agents.

Can we escape by requiring a weaker degree of “rationality”?

Also, we assumed \mathcal{S} was infinite and p_i and P were nonatomic.

Are these “technical” assumptions driving the result?

Finally, we assumed that \succeq_i and \succeq were defined for *all* prospects in $\mathcal{A}^{\mathcal{S}}$. But many prospects in $\mathcal{A}^{\mathcal{S}}$ may be unfeasible, or even logically impossible.

Can we escape by weakening this “universal domain” assumption?

Sadly, no. The conclusions of the BSAT (i.e. utilitarianism + “probability agreement”) remain in force **even** if we relax the “universal domain” assumption, allow \mathcal{S} to be finite, and only require the agents to satisfy *Statewise Dominance* (“minimal rationality”), rather than being SEU maximizers (M&P, 2013, extending Blackorby, Donaldson & Mongin, 2004).

Bayesian Social Aggregation Theorem: Suppose all the \succeq_i satisfy **Individual SEU**, and suppose \succeq satisfies **Group SEU** and **XAP**. Then there exist weights $c_i > 0$ for all i in \mathcal{I} , and some constant K , such that

$$W = K + c_{\text{Ann}} u_{\text{Ann}} + c_{\text{Bob}} u_{\text{Bob}} + c_{\text{Carl}} u_{\text{Carl}} + \dots$$

If $\{u_i\}_{i \in \mathcal{I}}$ are sufficiently diverse, this can only happen if everyone has the same beliefs:

$$p_{\text{Ann}} = p_{\text{Bob}} = p_{\text{Carl}} = \dots = P.$$

If agents have different beliefs, then **Group SEU** and **XAP** are incompatible.

Possible escapes: Perhaps **Individual SEU** and **Group SEU** demand too much “rationality” from the agents.

Can we escape by requiring a weaker degree of “rationality”?

Also, we assumed \mathcal{S} was infinite and p_i and P were nonatomic.

Are these “technical” assumptions driving the result?

Finally, we assumed that \succeq_i and \succeq were defined for *all* prospects in $\mathcal{A}^{\mathcal{S}}$. But many prospects in $\mathcal{A}^{\mathcal{S}}$ may be unfeasible, or even logically impossible.

Can we escape by weakening this “universal domain” assumption?

Sadly, no. The conclusions of the BSAT (i.e. utilitarianism + “probability agreement”) remain in force **even** if we relax the “universal domain” assumption, allow \mathcal{S} to be finite, and only require the agents to satisfy *Statewise Dominance* (“minimal rationality”), rather than being SEU maximizers (M&P, 2013, extending Blackorby, Donaldson & Mongin, 2004).

Statewise Dominance seems non-negotiable. Is *Ex ante* Pareto the culprit?

Indeed, *Ex ante* Pareto is already suspect, for other reasons.

To see this, suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann, Bob}\}$, with the beliefs:

	h	t
Ann's probability	0.9	0.1
Bob's probability	0.1	0.9

(i.e. $p_{\text{Ann}}(h) = 0.9$, etc.)

Consider two social prospects \mathbf{X} and \mathbf{Y} , with payoffs defined as follows:

$\mathbf{X} :=$		h	t
Ann	10	-20	
Bob	-20	10	

$\mathbf{Y} :=$		h	t
Ann	0	0	
Bob	0	0	

$\mathbf{X} \succ_A \mathbf{Y}$, because $\mathbb{E}(\mathbf{X} | u_A, p_A) = 7 > 0 = \mathbb{E}(\mathbf{Y} | u_A, p_A)$. Likewise, $\mathbf{X} \succ_B \mathbf{Y}$.

Thus, *Ex ante* Pareto dictates that $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.

But A&B's *ex ante* unanimity is "spurious", arising from different beliefs.

At least one of Ann or Bob must be *wrong*.

Indeed, if the *ex post* social preference \succ_{xp} is utilitarian, then $\mathbf{x}_h \prec_{\text{xp}} \mathbf{y}_h$ and

$\mathbf{x}_t \prec_{\text{xp}} \mathbf{y}_t$. Thus, Group Statewise Dominance dictates that $\mathbf{X} \prec \mathbf{Y}$.

Statewise Dominance seems non-negotiable. Is *Ex ante* Pareto the culprit? Indeed, *Ex ante* Pareto is already suspect, for other reasons.

To see this, suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann, Bob}\}$, with the beliefs:

	h	t
Ann's probability	0.9	0.1
Bob's probability	0.1	0.9

(i.e. $p_{\text{Ann}}(h) = 0.9$, etc.)

Consider two social prospects \mathbf{X} and \mathbf{Y} , with payoffs defined as follows:

$$\mathbf{X} :=$$

	h	t
Ann	10	-20
Bob	-20	10

$$\mathbf{Y} :=$$

	h	t
Ann	0	0
Bob	0	0

$\mathbf{X} \succ_A \mathbf{Y}$, because $\mathbb{E}(\mathbf{X} | u_A, p_A) = 7 > 0 = \mathbb{E}(\mathbf{Y} | u_A, p_A)$. Likewise, $\mathbf{X} \succ_B \mathbf{Y}$. Thus, *Ex ante* Pareto dictates that $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.

But A&B's *ex ante* unanimity is "spurious", arising from different beliefs. At least one of Ann or Bob must be *wrong*.

Indeed, if the *ex post* social preference \succ_{xp} is utilitarian, then $x_h \prec_{\text{xp}} y_h$ and $x_t \prec_{\text{xp}} y_t$. Thus, Group Statewise Dominance dictates that $\mathbf{X} \prec \mathbf{Y}$.

Statewise Dominance seems non-negotiable. Is *Ex ante* Pareto the culprit? Indeed, *Ex ante* Pareto is already suspect, for other reasons.

To see this, suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann, Bob}\}$, with the beliefs:

	h	t
Ann's probability	0.9	0.1
Bob's probability	0.1	0.9

(i.e. $p_{\text{Ann}}(h) = 0.9$, etc.)

Consider two social prospects \mathbf{X} and \mathbf{Y} , with payoffs defined as follows:

	h	t
$\mathbf{X} :=$ Ann	10	-20
Bob	-20	10

	h	t
$\mathbf{Y} :=$ Ann	0	0
Bob	0	0

$\mathbf{X} \succ_A \mathbf{Y}$, because $\mathbb{E}(\mathbf{X} | u_A, p_A) = 7 > 0 = \mathbb{E}(\mathbf{Y} | u_A, p_A)$. Likewise, $\mathbf{X} \succ_B \mathbf{Y}$. Thus, *Ex ante* Pareto dictates that $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.

But A&B's *ex ante* unanimity is "spurious", arising from different beliefs. At least one of Ann or Bob must be *wrong*.

Indeed, if the *ex post* social preference \succ_{xp} is utilitarian, then $x_h \prec_{\text{xp}} y_h$ and $x_t \prec_{\text{xp}} y_t$. Thus, Group Statewise Dominance dictates that $\mathbf{X} \prec \mathbf{Y}$.

Statewise Dominance seems non-negotiable. Is *Ex ante* Pareto the culprit? Indeed, *Ex ante* Pareto is already suspect, for other reasons.

To see this, suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann, Bob}\}$, with the beliefs:

	h	t
Ann's probability	0.9	0.1
Bob's probability	0.1	0.9

(i.e. $p_{\text{Ann}}(h) = 0.9$, etc.)

Consider two social prospects \mathbf{X} and \mathbf{Y} , with payoffs defined as follows:

		h	t
$\mathbf{X} :=$	Ann	10	-20
	Bob	-20	10

		h	t
$\mathbf{Y} :=$	Ann	0	0
	Bob	0	0

$\mathbf{X} \succ_A \mathbf{Y}$, because $\mathbb{E}(\mathbf{X} | u_A, p_A) = 7 > 0 = \mathbb{E}(\mathbf{Y} | u_A, p_A)$. Likewise, $\mathbf{X} \succ_B \mathbf{Y}$.

Thus, *Ex ante* Pareto dictates that $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.

But A&B's *ex ante* unanimity is "spurious", arising from different beliefs.

At least one of Ann or Bob must be *wrong*.

Indeed, if the *ex post* social preference \succ_{xp} is utilitarian, then $x_h \prec_{\text{xp}} y_h$ and

$x_t \prec_{\text{xp}} y_t$. Thus, Group Statewise Dominance dictates that $\mathbf{X} \prec \mathbf{Y}$.

Statewise Dominance seems non-negotiable. Is *Ex ante* Pareto the culprit? Indeed, *Ex ante* Pareto is already suspect, for other reasons.

To see this, suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann, Bob}\}$, with the beliefs:

	h	t	
Ann's probability	0.9	0.1	(i.e. $p_{\text{Ann}}(h) = 0.9$, etc.)
Bob's probability	0.1	0.9	

Consider two social prospects \mathbf{X} and \mathbf{Y} , with payoffs defined as follows:

$\mathbf{X} :=$		h	t	$\mathbf{Y} :=$		h	t
	Ann	10	-20		Ann	0	0
	Bob	-20	10		Bob	0	0

$\mathbf{X} \succ_A \mathbf{Y}$, because $\mathbb{E}(\mathbf{X} | u_A, p_A) = 7 > 0 = \mathbb{E}(\mathbf{Y} | u_A, p_A)$. Likewise, $\mathbf{X} \succ_B \mathbf{Y}$.

Thus, *Ex ante* Pareto dictates that $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.

But A&B's *ex ante* unanimity is "spurious", arising from different beliefs.

At least one of Ann or Bob must be *wrong*.

Indeed, if the *ex post* social preference \succ_{xp} is utilitarian, then $x_h \prec_{\text{xp}} y_h$ and

$x_t \prec_{\text{xp}} y_t$. Thus, Group Statewise Dominance dictates that $\mathbf{X} \prec \mathbf{Y}$.

Statewise Dominance seems non-negotiable. Is *Ex ante* Pareto the culprit? Indeed, *Ex ante* Pareto is already suspect, for other reasons.

To see this, suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann, Bob}\}$, with the beliefs:

	h	t
Ann's probability	0.9	0.1
Bob's probability	0.1	0.9

(i.e. $p_{\text{Ann}}(h) = 0.9$, etc.)

Consider two social prospects \mathbf{X} and \mathbf{Y} , with payoffs defined as follows:

		h	t
$\mathbf{X} :=$	Ann	10	-20
	Bob	-20	10

		h	t
$\mathbf{Y} :=$	Ann	0	0
	Bob	0	0

$\mathbf{X} \succ_A \mathbf{Y}$, because $\mathbb{E}(\mathbf{X} | u_A, p_A) = 7 > 0 = \mathbb{E}(\mathbf{Y} | u_A, p_A)$. Likewise, $\mathbf{X} \succ_B \mathbf{Y}$. Thus, *Ex ante* Pareto dictates that $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.

But A&B's *ex ante* unanimity is "spurious", arising from different beliefs. At least one of Ann or Bob must be *wrong*.

Indeed, if the *ex post* social preference \succ_{xp} is utilitarian, then $x_h \prec_{\text{xp}} y_h$ and $x_t \prec_{\text{xp}} y_t$. Thus, Group Statewise Dominance dictates that $\mathbf{X} \prec \mathbf{Y}$.

Statewise Dominance seems non-negotiable. Is *Ex ante* Pareto the culprit? Indeed, *Ex ante* Pareto is already suspect, for other reasons.

To see this, suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann, Bob}\}$, with the beliefs:

	h	t
Ann's probability	0.9	0.1
Bob's probability	0.1	0.9

(i.e. $p_{\text{Ann}}(h) = 0.9$, etc.)

Consider two social prospects \mathbf{X} and \mathbf{Y} , with payoffs defined as follows:

		h	t
$\mathbf{X} :=$	Ann	10	-20
	Bob	-20	10

		h	t
$\mathbf{Y} :=$	Ann	0	0
	Bob	0	0

$\mathbf{X} \succ_A \mathbf{Y}$, because $\mathbb{E}(\mathbf{X} | u_A, p_A) = 7 > 0 = \mathbb{E}(\mathbf{Y} | u_A, p_A)$. Likewise, $\mathbf{X} \succ_B \mathbf{Y}$. Thus, *Ex ante* Pareto dictates that $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.

But A&B's *ex ante* unanimity is "spurious", arising from different beliefs. At least one of Ann or Bob must be *wrong*.

Indeed, if the *ex post* social preference \succeq_{xp} is utilitarian, then $x_h \prec_{\text{xp}} y_h$ and $x_t \prec_{\text{xp}} y_t$. Thus, Group Statewise Dominance dictates that $\mathbf{X} \prec \mathbf{Y}$.

Statewise Dominance seems non-negotiable. Is *Ex ante* Pareto the culprit? Indeed, *Ex ante* Pareto is already suspect, for other reasons.

To see this, suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann, Bob}\}$, with the beliefs:

	h	t
Ann's probability	0.9	0.1
Bob's probability	0.1	0.9

(i.e. $p_{\text{Ann}}(h) = 0.9$, etc.)

Consider two social prospects \mathbf{X} and \mathbf{Y} , with payoffs defined as follows:

		h	t
$\mathbf{X} :=$	Ann	10	-20
	Bob	-20	10

		h	t
$\mathbf{Y} :=$	Ann	0	0
	Bob	0	0

$\mathbf{X} \succ_A \mathbf{Y}$, because $\mathbb{E}(\mathbf{X} | u_A, p_A) = 7 > 0 = \mathbb{E}(\mathbf{Y} | u_A, p_A)$. Likewise, $\mathbf{X} \succ_B \mathbf{Y}$. Thus, *Ex ante* Pareto dictates that $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.

But A&B's *ex ante* unanimity is "spurious", arising from different beliefs. At least one of Ann or Bob must be *wrong*.

Indeed, if the *ex post* social preference \succ_{xp} is utilitarian, then $\mathbf{x}_h \prec_{\text{xp}} \mathbf{y}_h$ and

$\mathbf{x}_t \prec_{\text{xp}} \mathbf{y}_t$.

Thus, Group Statewise Dominance dictates that $\mathbf{X} \succ \mathbf{Y}$.

Statewise Dominance seems non-negotiable. Is *Ex ante* Pareto the culprit? Indeed, *Ex ante* Pareto is already suspect, for other reasons.

To see this, suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann, Bob}\}$, with the beliefs:

	h	t
Ann's probability	0.9	0.1
Bob's probability	0.1	0.9

(i.e. $p_{\text{Ann}}(h) = 0.9$, etc.)

Consider two social prospects \mathbf{X} and \mathbf{Y} , with payoffs defined as follows:

		h	t
$\mathbf{X} :=$	Ann	10	-20
	Bob	-20	10

		h	t
$\mathbf{Y} :=$	Ann	0	0
	Bob	0	0

$\mathbf{X} \succ_A \mathbf{Y}$, because $\mathbb{E}(\mathbf{X} | u_A, p_A) = 7 > 0 = \mathbb{E}(\mathbf{Y} | u_A, p_A)$. Likewise, $\mathbf{X} \succ_B \mathbf{Y}$. Thus, *Ex ante* Pareto dictates that $\mathbf{X} \succ_{xa} \mathbf{Y}$.

But A&B's *ex ante* unanimity is "spurious", arising from different beliefs. At least one of Ann or Bob must be *wrong*.

Indeed, if the *ex post* social preference \succ_{xp} is utilitarian, then $\mathbf{x}_h \prec_{xp} \mathbf{y}_h$ and $\mathbf{x}_t \prec_{xp} \mathbf{y}_t$. Thus, Group Statewise Dominance dictates that $\mathbf{X} \prec_{gsd} \mathbf{Y}$.

Idea: Weaken *ex ante* Pareto to avoid such cases of “spurious unanimity”.

Gilboa, Samet, and Schmeidler (2004) suppose each individual i is an SEU-maximizer with a utility function u_i and probabilistic beliefs p_i on an infinite set \mathcal{S} of states of nature.

Let \mathfrak{B} be the set of events on whose probabilities all agents agree.
(Formally $\mathfrak{B} := \{\mathcal{E} \subseteq \mathcal{S}; p_i[\mathcal{E}] = p_j[\mathcal{E}], \text{ for all } i \text{ and } j \text{ in } \mathcal{I}\}.$)

A prospect f in $\mathcal{A}^{\mathcal{S}}$ is *admissible* if it only depends on information in \mathfrak{B} .
(Formally, this means f is \mathfrak{B} -measurable: $f^{-1}(\mathcal{E}) \in \mathfrak{B}$ for any measurable $\mathcal{E} \subseteq \mathcal{A}$.)

GSS restrict the *ex ante* Pareto condition to apply *only* to comparisons between admissible prospects (thereby excluding spurious unanimity.)

Theorem. (GSS) Let W be an ex post social welfare function on \mathcal{A} , let P be a probability distribution on \mathcal{S} , and let \succeq be the ex ante social preference relation on $\mathcal{A}^{\mathcal{S}}$ which maximizes the P -expected value of W . Then \succeq satisfies the ex ante Pareto condition restricted to admissible prospects if and only if W is a weighted utilitarian sum of the individual utilities u_i , and P is a weighted average of the individual probabilities p_i .

Idea: Weaken *ex ante* Pareto to avoid such cases of “spurious unanimity”.

Gilboa, Samet, and Schmeidler (2004) suppose each individual i is an SEU-maximizer with a utility function u_i and probabilistic beliefs p_i on an infinite set \mathcal{S} of states of nature.

Let \mathfrak{B} be the set of events on whose probabilities all agents agree.

(Formally $\mathfrak{B} := \{\mathcal{E} \subseteq \mathcal{S}; p_i[\mathcal{E}] = p_j[\mathcal{E}], \text{ for all } i \text{ and } j \text{ in } \mathcal{I}\}.$)

A prospect f in $\mathcal{A}^{\mathcal{S}}$ is *admissible* if it only depends on information in \mathfrak{B} .

(Formally, this means f is \mathfrak{B} -measurable: $f^{-1}(\mathcal{E}) \in \mathfrak{B}$ for any measurable $\mathcal{E} \subseteq \mathcal{A}$.)

GSS restrict the *ex ante* Pareto condition to apply *only* to comparisons between admissible prospects (thereby excluding spurious unanimity.)

Theorem. (GSS) Let W be an ex post social welfare function on \mathcal{A} , let P be a probability distribution on \mathcal{S} , and let \succeq be the ex ante social preference relation on $\mathcal{A}^{\mathcal{S}}$ which maximizes the P -expected value of W .

Then \succeq satisfies the ex ante Pareto condition restricted to admissible prospects if and only if W is a weighted utilitarian sum of the individual utilities u_i , and P is a weighted average of the individual probabilities p_i .

Idea: Weaken *ex ante* Pareto to avoid such cases of “spurious unanimity”.

Gilboa, Samet, and Schmeidler (2004) suppose each individual i is an SEU-maximizer with a utility function u_i and probabilistic beliefs p_i on an infinite set \mathcal{S} of states of nature.

Let \mathfrak{B} be the set of events on whose probabilities all agents agree.
(Formally $\mathfrak{B} := \{\mathcal{E} \subseteq \mathcal{S}; p_i[\mathcal{E}] = p_j[\mathcal{E}], \text{ for all } i \text{ and } j \text{ in } \mathcal{I}\}.$)

A prospect f in $\mathcal{A}^{\mathcal{S}}$ is *admissible* if it only depends on information in \mathfrak{B} .
(Formally, this means f is \mathfrak{B} -measurable: $f^{-1}(\mathcal{E}) \in \mathfrak{B}$ for any measurable $\mathcal{E} \subseteq \mathcal{A}$.)

GSS restrict the *ex ante* Pareto condition to apply *only* to comparisons between admissible prospects (thereby excluding spurious unanimity.)

Theorem. (GSS) Let W be an ex post social welfare function on \mathcal{A} , let P be a probability distribution on \mathcal{S} , and let \succeq be the ex ante social preference relation on $\mathcal{A}^{\mathcal{S}}$ which maximizes the P -expected value of W . Then \succeq satisfies the ex ante Pareto condition restricted to admissible prospects if and only if W is a weighted utilitarian sum of the individual utilities u_i , and P is a weighted average of the individual probabilities p_i .

Idea: Weaken *ex ante* Pareto to avoid such cases of “spurious unanimity”.

Gilboa, Samet, and Schmeidler (2004) suppose each individual i is an SEU-maximizer with a utility function u_i and probabilistic beliefs p_i on an infinite set \mathcal{S} of states of nature.

Let \mathfrak{B} be the set of events on whose probabilities all agents agree.

(Formally $\mathfrak{B} := \{\mathcal{E} \subseteq \mathcal{S}; p_i[\mathcal{E}] = p_j[\mathcal{E}], \text{ for all } i \text{ and } j \text{ in } \mathcal{I}\}.$)

A prospect f in $\mathcal{A}^{\mathcal{S}}$ is *admissible* if it only depends on information in \mathfrak{B} .

(Formally, this means f is \mathfrak{B} -measurable: $f^{-1}(\mathcal{E}) \in \mathfrak{B}$ for any measurable $\mathcal{E} \subseteq \mathcal{A}$.)

GSS restrict the *ex ante* Pareto condition to apply *only* to comparisons between admissible prospects (thereby excluding spurious unanimity.)

Theorem. (GSS) Let W be an ex post social welfare function on \mathcal{A} , let P be a probability distribution on \mathcal{S} , and let \succeq be the ex ante social preference relation on $\mathcal{A}^{\mathcal{S}}$ which maximizes the P -expected value of W .

Then \succeq satisfies the ex ante Pareto condition restricted to admissible prospects if and only if W is a weighted utilitarian sum of the individual utilities u_i , and P is a weighted average of the individual probabilities p_i .

Idea: Weaken *ex ante* Pareto to avoid such cases of “spurious unanimity”.

Gilboa, Samet, and Schmeidler (2004) suppose each individual i is an SEU-maximizer with a utility function u_i and probabilistic beliefs p_i on an infinite set \mathcal{S} of states of nature.

Let \mathfrak{B} be the set of events on whose probabilities all agents agree.

(Formally $\mathfrak{B} := \{\mathcal{E} \subseteq \mathcal{S}; p_i[\mathcal{E}] = p_j[\mathcal{E}], \text{ for all } i \text{ and } j \text{ in } \mathcal{I}\}.$)

A prospect f in $\mathcal{A}^{\mathcal{S}}$ is *admissible* if it only depends on information in \mathfrak{B} .

(Formally, this means f is \mathfrak{B} -measurable: $f^{-1}(\mathcal{E}) \in \mathfrak{B}$ for any measurable $\mathcal{E} \subseteq \mathcal{A}$.)

GSS restrict the *ex ante* Pareto condition to apply *only* to comparisons between admissible prospects (thereby excluding spurious unanimity.)

Theorem. (GSS) Let W be an ex post social welfare function on \mathcal{A} , let P be a probability distribution on \mathcal{S} , and let \succeq be the ex ante social preference relation on $\mathcal{A}^{\mathcal{S}}$ which maximizes the P -expected value of W . Then \succeq satisfies the ex ante Pareto condition restricted to admissible prospects if and only if W is a weighted utilitarian sum of the individual utilities u_i , and P is a weighted average of the individual probabilities p_i .

Idea: Weaken *ex ante* Pareto to avoid such cases of “spurious unanimity”.

Gilboa, Samet, and Schmeidler (2004) suppose each individual i is an SEU-maximizer with a utility function u_i and probabilistic beliefs p_i on an infinite set \mathcal{S} of states of nature.

Let \mathfrak{B} be the set of events on whose probabilities all agents agree.

(Formally $\mathfrak{B} := \{\mathcal{E} \subseteq \mathcal{S}; p_i[\mathcal{E}] = p_j[\mathcal{E}], \text{ for all } i \text{ and } j \text{ in } \mathcal{I}\}.$)

A prospect f in $\mathcal{A}^{\mathcal{S}}$ is *admissible* if it only depends on information in \mathfrak{B} .

(Formally, this means f is \mathfrak{B} -measurable: $f^{-1}(\mathcal{E}) \in \mathfrak{B}$ for any measurable $\mathcal{E} \subseteq \mathcal{A}$.)

GSS restrict the *ex ante* Pareto condition to apply *only* to comparisons between admissible prospects (thereby excluding spurious unanimity.)

Theorem. (GSS) Let W be an ex post social welfare function on \mathcal{A} , let P be a probability distribution on \mathcal{S} , and let \succeq be the ex ante social preference relation on $\mathcal{A}^{\mathcal{S}}$ which maximizes the P -expected value of W .

Then \succeq satisfies the ex ante Pareto condition restricted to admissible prospects if and only if W is a weighted utilitarian sum of the individual utilities u_i , and P is a weighted average of the individual probabilities p_i .

Idea: Weaken *ex ante* Pareto to avoid such cases of “spurious unanimity”.

Gilboa, Samet, and Schmeidler (2004) suppose each individual i is an SEU-maximizer with a utility function u_i and probabilistic beliefs p_i on an infinite set \mathcal{S} of states of nature.

Let \mathfrak{B} be the set of events on whose probabilities all agents agree.

(Formally $\mathfrak{B} := \{\mathcal{E} \subseteq \mathcal{S}; p_i[\mathcal{E}] = p_j[\mathcal{E}], \text{ for all } i \text{ and } j \text{ in } \mathcal{I}\}.$)

A prospect f in $\mathcal{A}^{\mathcal{S}}$ is *admissible* if it only depends on information in \mathfrak{B} .

(Formally, this means f is \mathfrak{B} -measurable: $f^{-1}(\mathcal{E}) \in \mathfrak{B}$ for any measurable $\mathcal{E} \subseteq \mathcal{A}$.)

GSS restrict the *ex ante* Pareto condition to apply *only* to comparisons between admissible prospects (thereby excluding spurious unanimity.)

Theorem. (GSS) Let W be an ex post social welfare function on \mathcal{A} , let P be a probability distribution on \mathcal{S} , and let \succeq be the ex ante social preference relation on $\mathcal{A}^{\mathcal{S}}$ which maximizes the P -expected value of W . Then \succeq satisfies the ex ante Pareto condition restricted to admissible prospects if and only if W is a weighted utilitarian sum of the individual utilities u_i , and P is a weighted average of the individual probabilities p_i .

Theorem. (GSS) Let W be an ex post SWF on \mathcal{A} , let P be a probability on \mathcal{S} , and let \succeq be the ex ante preference relation on $\mathcal{A}^{\mathcal{S}}$ which maximizes the P -expected value of W .

Then \succeq satisfies the restricted ex ante Pareto condition $\iff W$ is a weighted utilitarian sum of the utilities $\{u_i\}$, and P is a weighted average of the probabilities $\{p_i\}$.

This seems like a perfect solution. It does **not** require probability agreement, and it is **not** susceptible to spurious unanimity. Or is it?

Suppose $\mathcal{S} = \{r, s, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$.

Assume two prospects, f and g , which yield the *same* payoff for both agents in each state of nature.

	r	s	t
f	100	0	100
g	0	100	0

Ann and Bob begin with the *same* prior probability p :

$$p(r) = 0.49, \quad p(s) = 0.02, \quad \text{and} \quad p(t) = 0.49.$$

Ann privately observes the event $\{r, s\}$, while Bob privately observes $\{s, t\}$.

After Bayesian updating, they have the following posterior probabilities:

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96

Theorem. (GSS) Let W be an ex post SWF on \mathcal{A} , let P be a probability on \mathcal{S} , and let \succeq be the ex ante preference relation on $\mathcal{A}^{\mathcal{S}}$ which maximizes the P -expected value of W .

Then \succeq satisfies the restricted ex ante Pareto condition $\iff W$ is a weighted utilitarian sum of the utilities $\{u_i\}$, and P is a weighted average of the probabilities $\{p_i\}$.

This seems like a perfect solution. It does *not* require probability agreement, and it is *not* susceptible to spurious unanimity. Or is it?

Suppose $\mathcal{S} = \{r, s, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$.

Assume two prospects, f and g , which yield the *same* payoff for both agents in each state of nature.

	r	s	t
f	100	0	100
g	0	100	0

Ann and Bob begin with the *same* prior probability p :

$$p(r) = 0.49, \quad p(s) = 0.02, \quad \text{and} \quad p(t) = 0.49.$$

Ann privately observes the event $\{r, s\}$, while Bob privately observes $\{s, t\}$.

After Bayesian updating, they have the following posterior probabilities:

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96

Theorem. (GSS) Let W be an ex post SWF on \mathcal{A} , let P be a probability on \mathcal{S} , and let \succeq be the ex ante preference relation on $\mathcal{A}^{\mathcal{S}}$ which maximizes the P -expected value of W .

Then \succeq satisfies the restricted ex ante Pareto condition $\iff W$ is a weighted utilitarian sum of the utilities $\{u_i\}$, and P is a weighted average of the probabilities $\{p_i\}$.

This seems like a perfect solution. It does *not* require probability agreement, and it is *not* susceptible to spurious unanimity. Or is it?

Suppose $\mathcal{S} = \{r, s, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$.

Assume two prospects, f and g , which yield the *same* payoff for both agents in each state of nature.

	r	s	t
f	100	0	100
g	0	100	0

Ann and Bob begin with the *same* prior probability p :

$$p(r) = 0.49, \quad p(s) = 0.02, \quad \text{and} \quad p(t) = 0.49.$$

Ann privately observes the event $\{r, s\}$, while Bob privately observes $\{s, t\}$.

After Bayesian updating, they have the following posterior probabilities:

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96

Theorem. (GSS) Let W be an ex post SWF on \mathcal{A} , let P be a probability on \mathcal{S} , and let \succeq be the ex ante preference relation on $\mathcal{A}^{\mathcal{S}}$ which maximizes the P -expected value of W . Then \succeq satisfies the restricted ex ante Pareto condition $\iff W$ is a weighted utilitarian sum of the utilities $\{u_i\}$, and P is a weighted average of the probabilities $\{p_i\}$.

This seems like a perfect solution. It does *not* require probability agreement, and it is *not* susceptible to spurious unanimity. Or is it?

Suppose $\mathcal{S} = \{r, s, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$.

Assume two prospects, f and g , which yield the *same* payoff for both agents in each state of nature.

	r	s	t
f	100	0	100
g	0	100	0

Ann and Bob begin with the *same* prior probability p :

$$p(r) = 0.49, \quad p(s) = 0.02, \quad \text{and} \quad p(t) = 0.49.$$

Ann privately observes the event $\{r, s\}$, while Bob privately observes $\{s, t\}$.

After Bayesian updating, they have the following posterior probabilities:

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96

Theorem. (GSS) Let W be an ex post SWF on \mathcal{A} , let P be a probability on \mathcal{S} , and let \succeq be the ex ante preference relation on $\mathcal{A}^{\mathcal{S}}$ which maximizes the P -expected value of W . Then \succeq satisfies the restricted ex ante Pareto condition $\iff W$ is a weighted utilitarian sum of the utilities $\{u_i\}$, and P is a weighted average of the probabilities $\{p_i\}$.

This seems like a perfect solution. It does *not* require probability agreement, and it is *not* susceptible to spurious unanimity. Or is it?

Suppose $\mathcal{S} = \{r, s, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$.

Assume two prospects, f and g , which yield the *same* payoff for both agents in each state of nature.

	r	s	t
f	100	0	100
g	0	100	0

Ann and Bob begin with the *same* prior probability p :

$$p(r) = 0.49, \quad p(s) = 0.02, \quad \text{and} \quad p(t) = 0.49.$$

Ann privately observes the event $\{r, s\}$, while Bob privately observes $\{s, t\}$.

After Bayesian updating, they have the following posterior probabilities:

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96

Theorem. (GSS) Let W be an ex post SWF on \mathcal{A} , let P be a probability on \mathcal{S} , and let \succeq be the ex ante preference relation on $\mathcal{A}^{\mathcal{S}}$ which maximizes the P -expected value of W . Then \succeq satisfies the restricted ex ante Pareto condition $\iff W$ is a weighted utilitarian sum of the utilities $\{u_i\}$, and P is a weighted average of the probabilities $\{p_i\}$.

This seems like a perfect solution. It does *not* require probability agreement, and it is *not* susceptible to spurious unanimity. Or is it?

Suppose $\mathcal{S} = \{r, s, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$.

Assume two prospects, f and g , which yield the *same* payoff for both agents in each state of nature.

	r	s	t
f	100	0	100
g	0	100	0

Ann and Bob begin with the *same* prior probability p :

$$p(r) = 0.49, \quad p(s) = 0.02, \quad \text{and} \quad p(t) = 0.49.$$

Ann privately observes the event $\{r, s\}$, while Bob privately observes $\{s, t\}$.

After Bayesian updating, they have the following posterior probabilities:

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96

Theorem. (GSS) Let W be an ex post SWF on \mathcal{A} , let P be a probability on \mathcal{S} , and let \succeq be the ex ante preference relation on $\mathcal{A}^{\mathcal{S}}$ which maximizes the P -expected value of W . Then \succeq satisfies the restricted ex ante Pareto condition $\iff W$ is a weighted utilitarian sum of the utilities $\{u_i\}$, and P is a weighted average of the probabilities $\{p_i\}$.

This seems like a perfect solution. It does *not* require probability agreement, and it is *not* susceptible to spurious unanimity. Or is it?

Suppose $\mathcal{S} = \{r, s, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$.

Assume two prospects, f and g , which yield the *same* payoff for both agents in each state of nature.

	r	s	t
f	100	0	100
g	0	100	0

Ann and Bob begin with the *same* prior probability p :

$$p(r) = 0.49, \quad p(s) = 0.02, \quad \text{and} \quad p(t) = 0.49.$$

Ann privately observes the event $\{r, s\}$, while Bob privately observes $\{s, t\}$.

After Bayesian updating, they have the following posterior probabilities:

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96

	r	s	t
f	100	0	100
g	0	100	0

Ann & Bob agree: Expected Utility(f) = 96, while Expected Utility(g) = 4.

Thus, $f \succ_{\text{Ann}} g$ and $f \succ_{\text{Bob}} g$.

Furthermore, $\mathfrak{B} = \{S, \{r, t\}, \{s\}, \emptyset\}$, so both f and g are admissible.

Thus, even GSS's restricted *ex ante* Pareto dictates that $f \succ_{\text{xa}} g$.

Indeed, if P is the average of Ann's and Bob's beliefs (as GSS recommend), then P also says Expected SWF(f) = 96, while Expected SWF(g) = 4.

However, clearly, the *true* state is s .

Thus, g is *actually* the better choice.

By ignoring private information, the GSS theorem gets the wrong answer.

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96

	r	s	t
f	100	0	100
g	0	100	0

Ann & Bob agree: Expected Utility(f) = 96, while Expected Utility(g) = 4.

Thus, $f \succ_{\text{Ann}} g$ and $f \succ_{\text{Bob}} g$.

Furthermore, $\mathfrak{B} = \{S, \{r, t\}, \{s\}, \emptyset\}$, so both f and g are admissible.

Thus, even GSS's restricted *ex ante* Pareto dictates that $f \succ_{\text{xa}} g$.

Indeed, if P is the average of Ann's and Bob's beliefs (as GSS recommend), then P also says Expected SWF(f) = 96, while Expected SWF(g) = 4.

However, clearly, the *true* state is s .

Thus, g is *actually* the better choice.

By ignoring private information, the GSS theorem gets the wrong answer.

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96

	r	s	t
f	100	0	100
g	0	100	0

Ann & Bob agree: Expected Utility(f) = 96, while Expected Utility(g) = 4.

Thus, $f \succ_{\text{Ann}} g$ and $f \succ_{\text{Bob}} g$.

Furthermore, $\mathfrak{B} = \{\mathcal{S}, \{r, t\}, \{s\}, \emptyset\}$, so both f and g are admissible.

Thus, even GSS's restricted *ex ante* Pareto dictates that $f \succ_{\text{xa}} g$.

Indeed, if P is the average of Ann's and Bob's beliefs (as GSS recommend), then P also says Expected SWF(f) = 96, while Expected SWF(g) = 4.

However, clearly, the *true* state is s .

Thus, g is *actually* the better choice.

By ignoring private information, the GSS theorem gets the wrong answer.

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96

	r	s	t
f	100	0	100
g	0	100	0

Ann & Bob agree: Expected Utility(f) = 96, while Expected Utility(g) = 4.

Thus, $f \succ_{\text{Ann}} g$ and $f \succ_{\text{Bob}} g$.

Furthermore, $\mathfrak{B} = \{\mathcal{S}, \{r, t\}, \{s\}, \emptyset\}$, so both f and g are admissible.

Thus, even GSS's restricted *ex ante* Pareto dictates that $f \succ_{\text{xa}} g$.

Indeed, if P is the average of Ann's and Bob's beliefs (as GSS recommend), then P also says Expected SWF(f) = 96, while Expected SWF(g) = 4.

However, clearly, the *true* state is s .

Thus, g is *actually* the better choice.

By ignoring private information, the GSS theorem gets the wrong answer.

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96

	r	s	t
f	100	0	100
g	0	100	0

Ann & Bob agree: Expected Utility(f) = 96, while Expected Utility(g) = 4.

Thus, $f \succ_{\text{Ann}} g$ and $f \succ_{\text{Bob}} g$.

Furthermore, $\mathfrak{B} = \{\mathcal{S}, \{r, t\}, \{s\}, \emptyset\}$, so both f and g are admissible.

Thus, even GSS's restricted *ex ante* Pareto dictates that $f \succ_{\text{xa}} g$.

Indeed, if P is the average of Ann's and Bob's beliefs (as GSS recommend), then P also says Expected SWF(f) = 96, while Expected SWF(g) = 4.

However, clearly, the *true* state is s .

Thus, g is *actually* the better choice.

By ignoring private information, the GSS theorem gets the wrong answer.

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96
Average		0.48	0.4	0.48

	r	s	t
f	100	0	100
g	0	100	0

Ann & Bob agree: Expected Utility(f) = 96, while Expected Utility(g) = 4.

Thus, $f \succ_{\text{Ann}} g$ and $f \succ_{\text{Bob}} g$.

Furthermore, $\mathfrak{B} = \{\mathcal{S}, \{r, t\}, \{s\}, \emptyset\}$, so both f and g are admissible.

Thus, even GSS's restricted *ex ante* Pareto dictates that $f \succ_{\text{xa}} g$.

Indeed, if P is the average of Ann's and Bob's beliefs (as GSS recommend), then P also says Expected SWF(f) = 96, while Expected SWF(g) = 4.

However, clearly, the *true* state is s .

Thus, g is *actually* the better choice.

By ignoring private information, the GSS theorem gets the wrong answer.

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96
Average		0.48	0.4	0.48

	r	s	t
f	100	0	100
g	0	100	0

Ann & Bob agree: Expected Utility(f) = 96, while Expected Utility(g) = 4.

Thus, $f \succ_{\text{Ann}} g$ and $f \succ_{\text{Bob}} g$.

Furthermore, $\mathfrak{B} = \{\mathcal{S}, \{r, t\}, \{s\}, \emptyset\}$, so both f and g are admissible.

Thus, even GSS's restricted *ex ante* Pareto dictates that $f \succ_{\text{xa}} g$.

Indeed, if P is the average of Ann's and Bob's beliefs (as GSS recommend), then P also says Expected SWF(f) = 96, while Expected SWF(g) = 4.

However, clearly, the *true* state is s .

Thus, g is *actually* the better choice.

By ignoring private information, the GSS theorem gets the wrong answer.

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96
Average		0.48	0.4	0.48

	r	s	t
f	100	0	100
g	0	100	0

Ann & Bob agree: Expected Utility(f) = 96, while Expected Utility(g) = 4.

Thus, $f \succ_{\text{Ann}} g$ and $f \succ_{\text{Bob}} g$.

Furthermore, $\mathfrak{B} = \{\mathcal{S}, \{r, t\}, \{s\}, \emptyset\}$, so both f and g are admissible.

Thus, even GSS's restricted *ex ante* Pareto dictates that $f \succ_{\text{xa}} g$.

Indeed, if P is the average of Ann's and Bob's beliefs (as GSS recommend), then P also says Expected SWF(f) = 96, while Expected SWF(g) = 4.

However, clearly, the *true* state is s .

Thus, g is *actually* the better choice.

By ignoring private information, the GSS theorem gets the wrong answer.

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96
Average		0.48	0.4	0.48

	r	s	t
f	100	0	100
g	0	100	0

Ann & Bob agree: Expected Utility(f) = 96, while Expected Utility(g) = 4.

Thus, $f \succ_{\text{Ann}} g$ and $f \succ_{\text{Bob}} g$.

Furthermore, $\mathfrak{B} = \{\mathcal{S}, \{r, t\}, \{s\}, \emptyset\}$, so both f and g are admissible.

Thus, even GSS's restricted *ex ante* Pareto dictates that $f \succ_{\text{xa}} g$.

Indeed, if P is the average of Ann's and Bob's beliefs (as GSS recommend), then P also says Expected SWF(f) = 96, while Expected SWF(g) = 4.

However, clearly, the *true* state is s .

Thus, g is *actually* the better choice.

By ignoring private information, the GSS theorem gets the wrong answer.

	Info	r	s	t
Prior		0.49	0.02	0.49
Ann	$\{r, s\}$	0.96	0.04	0
Bob	$\{s, t\}$	0	0.04	0.96
Average		0.48	0.4	0.48

	r	s	t
f	100	0	100
g	0	100	0

Ann & Bob agree: Expected Utility(f) = 96, while Expected Utility(g) = 4.

Thus, $f \succ_{\text{Ann}} g$ and $f \succ_{\text{Bob}} g$.

Furthermore, $\mathfrak{B} = \{\mathcal{S}, \{r, t\}, \{s\}, \emptyset\}$, so both f and g are admissible.

Thus, even GSS's restricted *ex ante* Pareto dictates that $f \succ_{\text{xa}} g$.

Indeed, if P is the average of Ann's and Bob's beliefs (as GSS recommend), then P also says Expected SWF(f) = 96, while Expected SWF(g) = 4.

However, clearly, the *true* state is s .

Thus, g is *actually* the better choice.

By ignoring private information, the GSS theorem gets the wrong answer. 

GSS attempt to distinguish between “legitimate” unanimity and “spurious” unanimity by an *endogenous* criterion: the agreement set \mathfrak{B} .

But this attempt fails. Maybe instead we should use an *exogenous* criterion.

Idea: We should distinguish between *objective* randomness (i.e. “risk”) and *subjective* randomness (arising from “uncertainty”).

- *Ex ante* Pareto only makes sense for *objective* randomness, where the agents can agree for legitimate reasons.
- *Ex ante* Pareto is never appropriate for *subjective* randomness, where “spurious unanimity” is possible.

M & P (2013) consider a model of social choice which with two independent sources of randomness: one objective and one subjective.

They apply *ex ante* Pareto *only* to the agents' preferences over *objective* randomness. This yields a new version of the Social Aggregation Theorem:

- *Ex ante* social preferences maximize expected value of a utilitarian SWF.
- All agents must have the same beliefs about the *objective* randomness.
- But they can have different beliefs about the *subjective* randomness.

GSS attempt to distinguish between “legitimate” unanimity and “spurious” unanimity by an *endogenous* criterion: the agreement set \mathfrak{B} . But this attempt fails. Maybe instead we should use an *exogenous* criterion.

Idea: We should distinguish between *objective* randomness (i.e. “risk”) and *subjective* randomness (arising from “uncertainty”).

- *Ex ante* Pareto only makes sense for *objective* randomness, where the agents can agree for legitimate reasons.
- *Ex ante* Pareto is never appropriate for *subjective* randomness, where “spurious unanimity” is possible.

M & P (2013) consider a model of social choice which with two independent sources of randomness: one objective and one subjective.

They apply *ex ante* Pareto *only* to the agents' preferences over *objective* randomness. This yields a new version of the Social Aggregation Theorem:

- *Ex ante* social preferences maximize expected value of a utilitarian SWF.
- All agents must have the same beliefs about the *objective* randomness.
- But they can have different beliefs about the *subjective* randomness.

GSS attempt to distinguish between “legitimate” unanimity and “spurious” unanimity by an *endogenous* criterion: the agreement set \mathfrak{B} .

But this attempt fails. Maybe instead we should use an *exogenous* criterion.

Idea: We should distinguish between *objective* randomness (i.e. “risk”) and *subjective* randomness (arising from “uncertainty”).

- *Ex ante* Pareto only makes sense for *objective* randomness, where the agents can agree for legitimate reasons.
- *Ex ante* Pareto is never appropriate for *subjective* randomness, where “spurious unanimity” is possible.

M & P (2013) consider a model of social choice which with two independent sources of randomness: one objective and one subjective.

They apply *ex ante* Pareto *only* to the agents' preferences over *objective* randomness. This yields a new version of the Social Aggregation Theorem:

- *Ex ante* social preferences maximize expected value of a utilitarian SWF.
- All agents must have the same beliefs about the *objective* randomness.
- But they can have different beliefs about the *subjective* randomness.

GSS attempt to distinguish between “legitimate” unanimity and “spurious” unanimity by an *endogenous* criterion: the agreement set \mathfrak{B} .

But this attempt fails. Maybe instead we should use an *exogenous* criterion.

Idea: We should distinguish between *objective* randomness (i.e. “risk”) and *subjective* randomness (arising from “uncertainty”).

- *Ex ante* Pareto **only** makes sense for **objective** randomness, where the agents can agree for legitimate reasons.
 - *Ex ante* Pareto is never appropriate for *subjective* randomness, where “spurious unanimity” is possible.

M & P (2013) consider a model of social choice which with two independent sources of randomness: one objective and one subjective.

They apply *ex ante* Pareto *only* to the agents' preferences over *objective* randomness. This yields a new version of the Social Aggregation Theorem:

- *Ex ante* social preferences maximize expected value of a utilitarian SWF.
- All agents must have the same beliefs about the *objective* randomness.
- But they can have different beliefs about the *subjective* randomness.

GSS attempt to distinguish between “legitimate” unanimity and “spurious” unanimity by an *endogenous* criterion: the agreement set \mathfrak{B} .

But this attempt fails. Maybe instead we should use an *exogenous* criterion.

Idea: We should distinguish between *objective* randomness (i.e. “risk”) and *subjective* randomness (arising from “uncertainty”).

- *Ex ante* Pareto only makes sense for *objective* randomness, where the agents can agree for legitimate reasons.
- *Ex ante* Pareto is **never** appropriate for **subjective** randomness, where “spurious unanimity” is possible.

M & P (2013) consider a model of social choice which with two independent sources of randomness: one objective and one subjective.

They apply *ex ante* Pareto *only* to the agents' preferences over *objective* randomness. This yields a new version of the Social Aggregation Theorem:

- *Ex ante* social preferences maximize expected value of a utilitarian SWF.
- All agents must have the same beliefs about the *objective* randomness.
- But they can have different beliefs about the *subjective* randomness.

GSS attempt to distinguish between “legitimate” unanimity and “spurious” unanimity by an *endogenous* criterion: the agreement set \mathfrak{B} .

But this attempt fails. Maybe instead we should use an *exogenous* criterion.

Idea: We should distinguish between *objective* randomness (i.e. “risk”) and *subjective* randomness (arising from “uncertainty”).

- *Ex ante* Pareto only makes sense for *objective* randomness, where the agents can agree for legitimate reasons.
- *Ex ante* Pareto is never appropriate for *subjective* randomness, where “spurious unanimity” is possible.

M & P (2013) consider a model of social choice which with two independent sources of randomness: one objective and one subjective.

They apply *ex ante* Pareto *only* to the agents' preferences over *objective* randomness. This yields a new version of the Social Aggregation Theorem:

- *Ex ante* social preferences maximize expected value of a utilitarian SWF.
- All agents must have the same beliefs about the *objective* randomness.
- But they can have different beliefs about the *subjective* randomness.

GSS attempt to distinguish between “legitimate” unanimity and “spurious” unanimity by an *endogenous* criterion: the agreement set \mathfrak{B} .

But this attempt fails. Maybe instead we should use an *exogenous* criterion.

Idea: We should distinguish between *objective* randomness (i.e. “risk”) and *subjective* randomness (arising from “uncertainty”).

- *Ex ante* Pareto only makes sense for *objective* randomness, where the agents can agree for legitimate reasons.
- *Ex ante* Pareto is never appropriate for *subjective* randomness, where “spurious unanimity” is possible.

M & P (2013) consider a model of social choice which with two independent sources of randomness: one objective and one subjective.

They apply *ex ante* Pareto *only* to the agents’ preferences over *objective* randomness. This yields a new version of the Social Aggregation Theorem:

- *Ex ante* social preferences maximize expected value of a utilitarian SWF.
- All agents must have the same beliefs about the *objective* randomness.
- But they can have different beliefs about the *subjective* randomness.

GSS attempt to distinguish between “legitimate” unanimity and “spurious” unanimity by an *endogenous* criterion: the agreement set \mathfrak{B} .

But this attempt fails. Maybe instead we should use an *exogenous* criterion.

Idea: We should distinguish between *objective* randomness (i.e. “risk”) and *subjective* randomness (arising from “uncertainty”).

- *Ex ante* Pareto only makes sense for *objective* randomness, where the agents can agree for legitimate reasons.
- *Ex ante* Pareto is never appropriate for *subjective* randomness, where “spurious unanimity” is possible.

M & P (2013) consider a model of social choice which with two independent sources of randomness: one objective and one subjective.

They apply *ex ante* Pareto *only* to the agents’ preferences over *objective* randomness. This yields a new version of the Social Aggregation Theorem:

- *Ex ante* social preferences maximize expected value of a utilitarian SWF.
- All agents must have the same beliefs about the *objective* randomness.
- But they can have different beliefs about the *subjective* randomness.

GSS attempt to distinguish between “legitimate” unanimity and “spurious” unanimity by an *endogenous* criterion: the agreement set \mathfrak{B} .

But this attempt fails. Maybe instead we should use an *exogenous* criterion.

Idea: We should distinguish between *objective* randomness (i.e. “risk”) and *subjective* randomness (arising from “uncertainty”).

- *Ex ante* Pareto only makes sense for *objective* randomness, where the agents can agree for legitimate reasons.
- *Ex ante* Pareto is never appropriate for *subjective* randomness, where “spurious unanimity” is possible.

M & P (2013) consider a model of social choice which with two independent sources of randomness: one objective and one subjective.

They apply *ex ante* Pareto *only* to the agents’ preferences over *objective* randomness. This yields a new version of the Social Aggregation Theorem:

- *Ex ante* social preferences maximize expected value of a utilitarian SWF.
- All agents must have the same beliefs about the *objective* randomness.
- But they can have different beliefs about the *subjective* randomness.

GSS attempt to distinguish between “legitimate” unanimity and “spurious” unanimity by an *endogenous* criterion: the agreement set \mathfrak{B} .

But this attempt fails. Maybe instead we should use an *exogenous* criterion.

Idea: We should distinguish between *objective* randomness (i.e. “risk”) and *subjective* randomness (arising from “uncertainty”).

- *Ex ante* Pareto only makes sense for *objective* randomness, where the agents can agree for legitimate reasons.
- *Ex ante* Pareto is never appropriate for *subjective* randomness, where “spurious unanimity” is possible.

M & P (2013) consider a model of social choice which with two independent sources of randomness: one objective and one subjective.

They apply *ex ante* Pareto *only* to the agents’ preferences over *objective* randomness. This yields a new version of the Social Aggregation Theorem:

- *Ex ante* social preferences maximize expected value of a utilitarian SWF.
- All agents must have the same beliefs about the *objective* randomness.
- But they can have different beliefs about the *subjective* randomness.

Equality:

The Diamond Paradox

The Diamond Paradox

(24/47)

Statewise Dominance is the “minimal” criterion for rational choice.

But Statewise Dominance can contradict our intuitions about fairness.

Suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$. Suppose F and G are social prospects that yield state-contingent payoffs as follows (Diamond, 1967):

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

(Canonical example: allocating a hard candy to one of two children.)

Let \succeq_{xp} be an impartial *ex post* social preference order. So $\begin{bmatrix} 1 \\ 0 \end{bmatrix} \approx_{\text{xp}} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$.

Thus, $F(h) \approx_{\text{xp}} G(h)$ and $F(t) \approx_{\text{xp}} G(t)$.

Thus, Statewise Dominance implies that $F \approx_{\text{xa}} G$.

But this contradicts our intuition: F is *fair*, while G is *unfair*.

(Note: This paradox does not depend on SEU or subjective beliefs.)

The Diamond Paradox

(24/47)

Statewise Dominance is the “minimal” criterion for rational choice. But Statewise Dominance can contradict our intuitions about fairness. Suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$. Suppose F and G are social prospects that yield state-contingent payoffs as follows (Diamond, 1967):

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

(Canonical example: allocating a hard candy to one of two children.)

Let \succeq_{xp} be an impartial *ex post* social preference order. So $\begin{bmatrix} 1 \\ 0 \end{bmatrix} \approx_{\text{xp}} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$.

Thus, $F(h) \approx_{\text{xp}} G(h)$ and $F(t) \approx_{\text{xp}} G(t)$.

Thus, Statewise Dominance implies that $F \approx_{\text{xa}} G$.

But this contradicts our intuition: F is *fair*, while G is *unfair*.

(Note: This paradox does not depend on SEU or subjective beliefs.)

The Diamond Paradox

(24/47)

Statewise Dominance is the “minimal” criterion for rational choice. But Statewise Dominance can contradict our intuitions about fairness. Suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$. Suppose F and G are social prospects that yield state-contingent payoffs as follows (Diamond, 1967):

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

(Canonical example: allocating a hard candy to one of two children.)

Let \succeq_{xp} be an impartial *ex post* social preference order. So $\begin{bmatrix} 1 \\ 0 \end{bmatrix} \approx_{\text{xp}} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$.

Thus, $F(h) \approx_{\text{xp}} G(h)$ and $F(t) \approx_{\text{xp}} G(t)$.

Thus, Statewise Dominance implies that $F \approx_{\text{xa}} G$.

But this contradicts our intuition: F is *fair*, while G is *unfair*.

(Note: This paradox does not depend on SEU or subjective beliefs.)

The Diamond Paradox

(24/47)

Statewise Dominance is the “minimal” criterion for rational choice. But Statewise Dominance can contradict our intuitions about fairness. Suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$. Suppose F and G are social prospects that yield state-contingent payoffs as follows (Diamond, 1967):

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

(Canonical example: allocating a hard candy to one of two children.)

Let \succeq_{xp} be an impartial *ex post* social preference order. So $\begin{bmatrix} 1 \\ 0 \end{bmatrix} \approx_{xp} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$.

Thus, $F(h) \approx_{xp} G(h)$ and $F(t) \approx_{xp} G(t)$.

Thus, Statewise Dominance implies that $F \approx_{xa} G$.

But this contradicts our intuition: F is *fair*, while G is *unfair*.

(Note: This paradox does not depend on SEU or subjective beliefs.)

The Diamond Paradox

(24/47)

Statewise Dominance is the “minimal” criterion for rational choice. But Statewise Dominance can contradict our intuitions about fairness. Suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$. Suppose F and G are social prospects that yield state-contingent payoffs as follows (Diamond, 1967):

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

(Canonical example: allocating a hard candy to one of two children.)

Let \succsim_{xp} be an impartial *ex post* social preference order. So $\begin{bmatrix} 1 \\ 0 \end{bmatrix} \approx_{xp} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$.

Thus, $F(h) \approx_{xp} G(h)$ and $F(t) \approx_{xp} G(t)$.

Thus, Statewise Dominance implies that $F \approx_{xa} G$.

But this contradicts our intuition: F is *fair*, while G is *unfair*.

(Note: This paradox does not depend on SEU or subjective beliefs.)

The Diamond Paradox

(24/47)

Statewise Dominance is the “minimal” criterion for rational choice. But Statewise Dominance can contradict our intuitions about fairness. Suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$. Suppose F and G are social prospects that yield state-contingent payoffs as follows (Diamond, 1967):

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

(Canonical example: allocating a hard candy to one of two children.)

Let \succeq_{xp} be an impartial *ex post* social preference order. So $\begin{bmatrix} 1 \\ 0 \end{bmatrix} \approx_{\text{xp}} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$.

Thus, $F(h) \approx_{\text{xp}} G(h)$ and $F(t) \approx_{\text{xp}} G(t)$.

Thus, Statewise Dominance implies that $F \approx_{\text{sa}} G$.

But this contradicts our intuition: F is *fair*, while G is *unfair*.

(Note: This paradox does not depend on SEU or subjective beliefs.)

The Diamond Paradox

(24/47)

Statewise Dominance is the “minimal” criterion for rational choice. But Statewise Dominance can contradict our intuitions about fairness. Suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$. Suppose F and G are social prospects that yield state-contingent payoffs as follows (Diamond, 1967):

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

(Canonical example: allocating a hard candy to one of two children.)

Let \succeq_{xp} be an impartial *ex post* social preference order. So $\begin{bmatrix} 1 \\ 0 \end{bmatrix} \approx_{\text{xp}} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$.

Thus, $F(h) \approx_{\text{xp}} G(h)$ and $F(t) \approx_{\text{xp}} G(t)$.

Thus, Statewise Dominance implies that $F \approx_{\text{xa}} G$.

But this contradicts our intuition: F is *fair*, while G is *unfair*.

(Note: This paradox does not depend on SEU or subjective beliefs.)

Statewise Dominance is the “minimal” criterion for rational choice. But Statewise Dominance can contradict our intuitions about fairness. Suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$. Suppose F and G are social prospects that yield state-contingent payoffs as follows (Diamond, 1967):

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

(Canonical example: allocating a hard candy to one of two children.)

Let \succ_{xp} be an impartial *ex post* social preference order. So $\begin{bmatrix} 1 \\ 0 \end{bmatrix} \approx_{\text{xp}} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$.

Thus, $F(h) \approx_{\text{xp}} G(h)$ and $F(t) \approx_{\text{xp}} G(t)$.

Thus, Statewise Dominance implies that $F \approx_{\text{xa}} G$.

But this contradicts our intuition: F is *fair*, while G is *unfair*.

(Note: This paradox does not depend on SEU or subjective beliefs.)

Statewise Dominance is the “minimal” criterion for rational choice. But Statewise Dominance can contradict our intuitions about fairness. Suppose $\mathcal{S} = \{h, t\}$ and $\mathcal{I} = \{\text{Ann}, \text{Bob}\}$. Suppose F and G are social prospects that yield state-contingent payoffs as follows (Diamond, 1967):

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

(Canonical example: allocating a hard candy to one of two children.)

Let \succsim_{xp} be an impartial *ex post* social preference order. So $\begin{bmatrix} 1 \\ 0 \end{bmatrix} \approx_{\text{xp}} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$.

Thus, $F(h) \approx_{\text{xp}} G(h)$ and $F(t) \approx_{\text{xp}} G(t)$.

Thus, Statewise Dominance implies that $F \approx_{\text{xa}} G$.

But this contradicts our intuition: F is *fair*, while G is *unfair*.

(*Note:* This paradox does not depend on SEU or subjective beliefs.)

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

Our ethical intuitions in Diamond's Paradox exemplify *ex ante egalitarianism* —that is, applying egalitarian principles to the *expected utilities* of the agents, *before* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex ante* social welfare function W_{xa} to the ensemble of expected utilities of the agents.

For simplicity, suppose $\text{Probability}(h) = \frac{1}{2} = \text{Probability}(t)$.

Let $EU_A(F) :=$ expected utility of Ann in prospect F , etc.

Thus, $EU_A(F) = \frac{1}{2} = EU_B(F)$, whereas $EU_A(G) = 1$ and $EU_B(G) = 0$.

Thus, $F \succ_{xa} G$ because

$$W_{xa}[EU_A(F), EU_B(F)] = W_{xa}\left(\frac{1}{2}, \frac{1}{2}\right) > W_{xa}(1, 0) = W_{xa}[EU_A(G), EU_B(G)]$$

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

Our ethical intuitions in Diamond's Paradox exemplify *ex ante* egalitarianism —that is, applying egalitarian principles to the *expected utilities* of the agents, *before* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex ante* social welfare function W_{xa} to the ensemble of expected utilities of the agents.

For simplicity, suppose $\text{Probability}(h) = \frac{1}{2} = \text{Probability}(t)$.

Let $EU_A(F) :=$ expected utility of Ann in prospect F , etc.

Thus, $EU_A(F) = \frac{1}{2} = EU_B(F)$, whereas $EU_A(G) = 1$ and $EU_B(G) = 0$.

Thus, $F \succ_{xa} G$ because

$$W_{xa}[EU_A(F), EU_B(F)] = W_{xa}\left(\frac{1}{2}, \frac{1}{2}\right) > W_{xa}(1, 0) = W_{xa}[EU_A(G), EU_B(G)]$$

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

Our ethical intuitions in Diamond's Paradox exemplify *ex ante* egalitarianism —that is, applying egalitarian principles to the *expected utilities* of the agents, *before* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex ante* social welfare function W_{xa} to the ensemble of expected utilities of the agents.

For simplicity, suppose $\text{Probability}(h) = \frac{1}{2} = \text{Probability}(t)$.

Let $EU_A(F) :=$ expected utility of Ann in prospect F , etc.

Thus, $EU_A(F) = \frac{1}{2} = EU_B(F)$, whereas $EU_A(G) = 1$ and $EU_B(G) = 0$.

Thus, $F \succ_{\text{xa}} G$ because

$$W_{\text{xa}}[EU_A(F), EU_B(F)] = W_{\text{xa}}\left(\frac{1}{2}, \frac{1}{2}\right) > W_{\text{xa}}(1, 0) = W_{\text{xa}}[EU_A(G), EU_B(G)]$$

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

Our ethical intuitions in Diamond's Paradox exemplify *ex ante* egalitarianism —that is, applying egalitarian principles to the *expected utilities* of the agents, *before* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex ante* social welfare function W_{xa} to the ensemble of expected utilities of the agents.

For simplicity, suppose $\text{Probability}(h) = \frac{1}{2} = \text{Probability}(t)$.

Let $EU_A(F) := \text{expected utility of Ann in prospect } F$, etc.

Thus, $EU_A(F) = \frac{1}{2} = EU_B(F)$, whereas $EU_A(G) = 1$ and $EU_B(G) = 0$.

Thus, $F \succ_{\text{xa}} G$ because

$$W_{\text{xa}}[EU_A(F), EU_B(F)] = W_{\text{xa}}\left(\frac{1}{2}, \frac{1}{2}\right) > W_{\text{xa}}(1, 0) = W_{\text{xa}}[EU_A(G), EU_B(G)]$$

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

Our ethical intuitions in Diamond's Paradox exemplify *ex ante* egalitarianism —that is, applying egalitarian principles to the *expected utilities* of the agents, *before* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex ante* social welfare function W_{xa} to the ensemble of expected utilities of the agents.

For simplicity, suppose $\text{Probability}(h) = \frac{1}{2} = \text{Probability}(t)$.

Let $EU_A(F) := \text{expected utility of Ann in prospect } F$, etc.

Thus, $EU_A(F) = \frac{1}{2} = EU_B(F)$, whereas $EU_A(G) = 1$ and $EU_B(G) = 0$.

Thus, $F \succ_{\text{xa}} G$ because

$$W_{\text{xa}}[EU_A(F), EU_B(F)] = W_{\text{xa}}\left(\frac{1}{2}, \frac{1}{2}\right) > W_{\text{xa}}(1, 0) = W_{\text{xa}}[EU_A(G), EU_B(G)]$$

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

Our ethical intuitions in Diamond's Paradox exemplify *ex ante* egalitarianism —that is, applying egalitarian principles to the *expected utilities* of the agents, *before* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex ante* social welfare function W_{xa} to the ensemble of expected utilities of the agents.

For simplicity, suppose $\text{Probability}(h) = \frac{1}{2} = \text{Probability}(t)$.

Let $EU_A(F) := \text{expected utility of Ann in prospect } F$, etc.

Thus, $EU_A(F) = \frac{1}{2} = EU_B(F)$, whereas $EU_A(G) = 1$ and $EU_B(G) = 0$.

Thus, $F \succ_{\text{xa}} G$ because

$$W_{\text{xa}}[EU_A(F), EU_B(F)] = W_{\text{xa}}\left(\frac{1}{2}, \frac{1}{2}\right) > W_{\text{xa}}(1, 0) = W_{\text{xa}}[EU_A(G), EU_B(G)]$$

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

Our ethical intuitions in Diamond's Paradox exemplify *ex ante* egalitarianism —that is, applying egalitarian principles to the *expected utilities* of the agents, *before* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex ante* social welfare function W_{xa} to the ensemble of expected utilities of the agents.

For simplicity, suppose $\text{Probability}(h) = \frac{1}{2} = \text{Probability}(t)$.

Let $EU_A(F) := \text{expected utility of Ann in prospect } F$, etc.

Thus, $EU_A(F) = \frac{1}{2} = EU_B(F)$, whereas $EU_A(G) = 1$ and $EU_B(G) = 0$.

Thus, $F \succ_{\text{xa}} G$ because

$$W_{\text{xa}}[EU_A(F), EU_B(F)] = W_{\text{xa}}\left(\frac{1}{2}, \frac{1}{2}\right) > W_{\text{xa}}(1, 0) = W_{\text{xa}}[EU_A(G), EU_B(G)]$$

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

Our ethical intuitions in Diamond's Paradox exemplify *ex ante* egalitarianism —that is, applying egalitarian principles to the *expected utilities* of the agents, *before* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex ante* social welfare function W_{xa} to the ensemble of expected utilities of the agents.

For simplicity, suppose $\text{Probability}(h) = \frac{1}{2} = \text{Probability}(t)$.

Let $EU_A(F) := \text{expected utility of Ann in prospect } F$, etc.

Thus, $EU_A(F) = \frac{1}{2} = EU_B(F)$, whereas $EU_A(G) = 1$ and $EU_B(G) = 0$.

Thus, $F \succ_{\text{xa}} G$ because

$$W_{\text{xa}}[EU_A(F), EU_B(F)] = W_{\text{xa}}\left(\frac{1}{2}, \frac{1}{2}\right) > W_{\text{xa}}(1, 0) = W_{\text{xa}}[EU_A(G), EU_B(G)]$$

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

Our ethical intuitions in Diamond's Paradox exemplify *ex ante* egalitarianism —that is, applying egalitarian principles to the *expected utilities* of the agents, *before* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex ante* social welfare function W_{xa} to the ensemble of expected utilities of the agents.

For simplicity, suppose $\text{Probability}(h) = \frac{1}{2} = \text{Probability}(t)$.

Let $EU_A(F) := \text{expected utility of Ann in prospect } F$, etc.

Thus, $EU_A(F) = \frac{1}{2} = EU_B(F)$, whereas $EU_A(G) = 1$ and $EU_B(G) = 0$.

Thus, $F \succ_{\text{xa}} G$ because

$$W_{\text{xa}}[EU_A(F), EU_B(F)] = W_{\text{xa}}\left(\frac{1}{2}, \frac{1}{2}\right) > W_{\text{xa}}(1, 0) = W_{\text{xa}}[EU_A(G), EU_B(G)]$$

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

Our ethical intuitions in Diamond's Paradox exemplify *ex ante* egalitarianism —that is, applying egalitarian principles to the *expected utilities* of the agents, *before* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex ante* social welfare function W_{xa} to the ensemble of expected utilities of the agents.

For simplicity, suppose $\text{Probability}(h) = \frac{1}{2} = \text{Probability}(t)$.

Let $EU_A(F) := \text{expected utility of Ann in prospect } F$, etc.

Thus, $EU_A(F) = \frac{1}{2} = EU_B(F)$, whereas $EU_A(G) = 1$ and $EU_B(G) = 0$.

Thus, $F \succ_{\text{xa}} G$ because

$$W_{\text{xa}}[EU_A(F), EU_B(F)] = W_{\text{xa}}\left(\frac{1}{2}, \frac{1}{2}\right) > W_{\text{xa}}(1, 0) = W_{\text{xa}}[EU_A(G), EU_B(G)].$$

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

This is to be contrasted with **ex post egalitarianism**, which applies egalitarian principles to social outcomes, *after* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex post* social welfare function W_{xp} to each social outcome, and then computing the expected value of W_{xp} .

We have $W_{xp} \left(\begin{pmatrix} 1 \\ 0 \end{pmatrix} \right) = C = W_{xp} \left(\begin{pmatrix} 0 \\ 1 \end{pmatrix} \right)$ for some constant C .

Thus, $F \approx_{xa} G$ because

$$\frac{1}{2} W_{xp} [F(h)] + \frac{1}{2} W_{xp} [F(t)] = \frac{1}{2} C + \frac{1}{2} C = \frac{1}{2} W_{xp} [G(h)] + \frac{1}{2} W_{xp} [G(t)].$$

Thus, Diamond's Paradox reveals a disagreement between *ex ante* egalitarianism and *ex post* egalitarianism.

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

This is to be contrasted with *ex post* egalitarianism, which applies egalitarian principles to social outcomes, *after* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex post* social welfare function W_{xp} to each social outcome, and then computing the expected value of W_{xp} .

We have $W_{xp} \left(\begin{pmatrix} 1 \\ 0 \end{pmatrix} \right) = C = W_{xp} \left(\begin{pmatrix} 0 \\ 1 \end{pmatrix} \right)$ for some constant C .

Thus, $F \approx_{xa} G$ because

$$\frac{1}{2} W_{xp} [F(h)] + \frac{1}{2} W_{xp} [F(t)] = \frac{1}{2} C + \frac{1}{2} C = \frac{1}{2} W_{xp} [G(h)] + \frac{1}{2} W_{xp} [G(t)].$$

Thus, Diamond's Paradox reveals a disagreement between *ex ante* egalitarianism and *ex post* egalitarianism.

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

This is to be contrasted with *ex post* egalitarianism, which applies egalitarian principles to social outcomes, *after* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex post* social welfare function W_{xp} to each social outcome, and then computing the expected value of W_{xp} .

We have $W_{xp} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = C = W_{xp} \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ for some constant C .

Thus, $F \approx_{xa} G$ because

$$\frac{1}{2} W_{xp} [F(h)] + \frac{1}{2} W_{xp} [F(t)] = \frac{1}{2} C + \frac{1}{2} C = \frac{1}{2} W_{xp} [G(h)] + \frac{1}{2} W_{xp} [G(t)].$$

Thus, Diamond's Paradox reveals a disagreement between *ex ante* egalitarianism and *ex post* egalitarianism.

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

This is to be contrasted with *ex post* egalitarianism, which applies egalitarian principles to social outcomes, *after* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex post* social welfare function W_{xp} to each social outcome, and then computing the expected value of W_{xp} .

We have $W_{xp} \left(\begin{pmatrix} 1 \\ 0 \end{pmatrix} \right) = C = W_{xp} \left(\begin{pmatrix} 0 \\ 1 \end{pmatrix} \right)$ for some constant C .

Thus, $F \approx_{xa} G$ because

$$\frac{1}{2} W_{xp} [F(h)] + \frac{1}{2} W_{xp} [F(t)] = \frac{1}{2} C + \frac{1}{2} C = \frac{1}{2} W_{xp} [G(h)] + \frac{1}{2} W_{xp} [G(t)].$$

Thus, Diamond's Paradox reveals a disagreement between *ex ante* egalitarianism and *ex post* egalitarianism.

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

This is to be contrasted with *ex post* egalitarianism, which applies egalitarian principles to social outcomes, *after* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex post* social welfare function W_{xp} to each social outcome, and then computing the expected value of W_{xp} .

We have $W_{xp} \left(\begin{pmatrix} 1 \\ 0 \end{pmatrix} \right) = C = W_{xp} \left(\begin{pmatrix} 0 \\ 1 \end{pmatrix} \right)$ for some constant C .

Thus, $F \approx_{xa} G$ because

$$\frac{1}{2} W_{xp} [F(h)] + \frac{1}{2} W_{xp} [F(t)] = \frac{1}{2} C + \frac{1}{2} C = \frac{1}{2} W_{xp} [G(h)] + \frac{1}{2} W_{xp} [G(t)].$$

Thus, Diamond's Paradox reveals a disagreement between *ex ante* egalitarianism and *ex post* egalitarianism.

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

This is to be contrasted with *ex post* egalitarianism, which applies egalitarian principles to social outcomes, *after* the resolution of uncertainty.

Formally, this means comparing social prospects by applying an (impartial, egalitarian) *ex post* social welfare function W_{xp} to each social outcome, and then computing the expected value of W_{xp} .

We have $W_{xp} \left(\begin{pmatrix} 1 \\ 0 \end{pmatrix} \right) = C = W_{xp} \left(\begin{pmatrix} 0 \\ 1 \end{pmatrix} \right)$ for some constant C .

Thus, $F \approx_{xa} G$ because

$$\frac{1}{2} W_{xp} [F(h)] + \frac{1}{2} W_{xp} [F(t)] = \frac{1}{2} C + \frac{1}{2} C = \frac{1}{2} W_{xp} [G(h)] + \frac{1}{2} W_{xp} [G(t)].$$

Thus, Diamond's Paradox reveals a disagreement between *ex ante* egalitarianism and *ex post* egalitarianism.

Ex ante egalitarianism vs. *ex post* egalitarianism

(27/47)

In Diamond's Paradox, *ex post* egalitarianism gets the “wrong” answer.

This suggests that *ex ante* egalitarianism is superior.

But consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

H	h	t
Ann	1	0
Bob	1	0

E and H yield more egalitarian outcomes than F in *both* states of nature. So *ex post* egalitarianism would prefer both E and H over F .

(Meanwhile, the choice between E and H might be based on social risk aversion.)

But all three prospects yield the same expected utility ($1/2$) for each agent. So *ex ante* egalitarianism would be indifferent between them.

In this case, it is *ex ante* egalitarianism which gets the “wrong” answer.

Furthermore, XAE and XPE each suffer from a critical flaw.

- ▶ XAE violates the Statewise Dominance axiom (so it is “irrational”).
- ▶ XPE violates the *ex ante* Pareto axiom (so it is “paternalistic”).

It seems that *neither* approach is really optimal. Instead, we should employ a “hybrid” of XAE and XPE (Ben Porath, Gilboa & Schmeidler; Gajdos & Maurin).

Ex ante egalitarianism vs. *ex post* egalitarianism

(27/47)

In Diamond's Paradox, *ex post* egalitarianism gets the “wrong” answer. This suggests that *ex ante* egalitarianism is superior.

But consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

H	h	t
Ann	1	0
Bob	1	0

E and H yield more egalitarian outcomes than F in *both* states of nature. So *ex post* egalitarianism would prefer both E and H over F .

(Meanwhile, the choice between E and H might be based on social risk aversion.)

But all three prospects yield the same expected utility ($1/2$) for each agent. So *ex ante* egalitarianism would be indifferent between them.

In this case, it is *ex ante* egalitarianism which gets the “wrong” answer.

Furthermore, XAE and XPE each suffer from a critical flaw.

- ▶ XAE violates the Statewise Dominance axiom (so it is “irrational”).
- ▶ XPE violates the *ex ante* Pareto axiom (so it is “paternalistic”).

It seems that *neither* approach is really optimal. Instead, we should employ a “hybrid” of XAE and XPE (Ben Porath, Gilboa & Schmeidler; Gajdos & Maurin).

Ex ante egalitarianism vs. *ex post* egalitarianism

(27/47)

In Diamond's Paradox, *ex post* egalitarianism gets the “wrong” answer. This suggests that *ex ante* egalitarianism is superior. But consider the following three social prospects:

<i>F</i>	<i>h</i>	<i>t</i>
Ann	1	0
Bob	0	1

<i>E</i>	<i>h</i>	<i>t</i>
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

<i>H</i>	<i>h</i>	<i>t</i>
Ann	1	0
Bob	1	0

E and *H* yield more egalitarian outcomes than *F* in *both* states of nature. So *ex post* egalitarianism would prefer both *E* and *H* over *F*.

(Meanwhile, the choice between *E* and *H* might be based on social risk aversion.)

But all three prospects yield the same expected utility ($1/2$) for each agent. So *ex ante* egalitarianism would be indifferent between them.

In this case, it is *ex ante* egalitarianism which gets the “wrong” answer.

Furthermore, XAE and XPE each suffer from a critical flaw.

- ▶ XAE violates the Statewise Dominance axiom (so it is “irrational”).
- ▶ XPE violates the *ex ante* Pareto axiom (so it is “paternalistic”).

It seems that *neither* approach is really optimal. Instead, we should employ a “hybrid” of XAE and XPE (Ben Porath, Gilboa & Schmeidler; Gajdos & Maurin).

Ex ante egalitarianism vs. *ex post* egalitarianism

(27/47)

In Diamond's Paradox, *ex post* egalitarianism gets the “wrong” answer. This suggests that *ex ante* egalitarianism is superior. But consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

H	h	t
Ann	1	0
Bob	1	0

E and H yield **more egalitarian** outcomes than F in **both** states of nature. So *ex post* egalitarianism would prefer both E and H over F .

(Meanwhile, the choice between E and H might be based on social risk aversion.)

But all three prospects yield the same expected utility ($1/2$) for each agent. So *ex ante* egalitarianism would be indifferent between them.

In this case, it is *ex ante* egalitarianism which gets the “wrong” answer. Furthermore, XAE and XPE each suffer from a critical flaw.

- ▶ XAE violates the Statewise Dominance axiom (so it is “irrational”).
- ▶ XPE violates the *ex ante* Pareto axiom (so it is “paternalistic”).

It seems that *neither* approach is really optimal. Instead, we should employ a “**hybrid**” of XAE and XPE (Ben Porath, Gilboa & Schmeidler; Gajdos & Maurin).

Ex ante egalitarianism vs. *ex post* egalitarianism

(27/47)

In Diamond's Paradox, *ex post* egalitarianism gets the “wrong” answer. This suggests that *ex ante* egalitarianism is superior. But consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

H	h	t
Ann	1	0
Bob	1	0

E and H yield more egalitarian outcomes than F in *both* states of nature. So *ex post* egalitarianism would prefer both E and H over F .

(Meanwhile, the choice between E and H might be based on social risk aversion.)

But all three prospects yield the same expected utility ($1/2$) for each agent. So *ex ante* egalitarianism would be indifferent between them.

In this case, it is *ex ante* egalitarianism which gets the “wrong” answer.

Furthermore, XAE and XPE each suffer from a critical flaw.

- ▶ XAE violates the Statewise Dominance axiom (so it is “irrational”).
- ▶ XPE violates the *ex ante* Pareto axiom (so it is “paternalistic”).

It seems that *neither* approach is really optimal. Instead, we should employ a “hybrid” of XAE and XPE (Ben Porath, Gilboa & Schmeidler; Gajdos & Maurin).

Ex ante egalitarianism vs. *ex post* egalitarianism

(27/47)

In Diamond's Paradox, *ex post* egalitarianism gets the “wrong” answer. This suggests that *ex ante* egalitarianism is superior. But consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

H	h	t
Ann	1	0
Bob	1	0

E and H yield more egalitarian outcomes than F in *both* states of nature. So *ex post* egalitarianism would prefer both E and H over F .

(Meanwhile, the choice between E and H might be based on social risk aversion.)

But all three prospects yield the **same** expected utility ($1/2$) for each agent.

So *ex ante* egalitarianism would be indifferent between them.

In this case, it is *ex ante* egalitarianism which gets the “wrong” answer.

Furthermore, XAE and XPE each suffer from a critical flaw.

- ▶ XAE violates the Statewise Dominance axiom (so it is “irrational”).
- ▶ XPE violates the *ex ante* Pareto axiom (so it is “paternalistic”).

It seems that *neither* approach is really optimal. Instead, we should employ

a “hybrid” of XAE and XPE (Ben Porath, Gilboa & Schmeidler; Gajdos & Maurin)

Ex ante egalitarianism vs. *ex post* egalitarianism

(27/47)

In Diamond's Paradox, *ex post* egalitarianism gets the “wrong” answer. This suggests that *ex ante* egalitarianism is superior. But consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

H	h	t
Ann	1	0
Bob	1	0

E and H yield more egalitarian outcomes than F in *both* states of nature. So *ex post* egalitarianism would prefer both E and H over F .

(Meanwhile, the choice between E and H might be based on social risk aversion.)

But all three prospects yield the same expected utility ($1/2$) for each agent. So *ex ante* egalitarianism would be indifferent between them.

In this case, it is *ex ante* egalitarianism which gets the “wrong” answer. Furthermore, XAE and XPE each suffer from a critical flaw.

- ▶ XAE violates the Statewise Dominance axiom (so it is “irrational”).
- ▶ XPE violates the *ex ante* Pareto axiom (so it is “paternalistic”).

It seems that *neither* approach is really optimal. Instead, we should employ a “hybrid” of XAE and XPE (Ben Porath, Gilboa & Schmeidler; Gajdos & Maurin).

Ex ante egalitarianism vs. *ex post* egalitarianism

(27/47)

In Diamond's Paradox, *ex post* egalitarianism gets the “wrong” answer. This suggests that *ex ante* egalitarianism is superior. But consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

H	h	t
Ann	1	0
Bob	1	0

E and H yield more egalitarian outcomes than F in *both* states of nature. So *ex post* egalitarianism would prefer both E and H over F .

(Meanwhile, the choice between E and H might be based on social risk aversion.)

But all three prospects yield the same expected utility ($1/2$) for each agent. So *ex ante* egalitarianism would be indifferent between them.

In this case, it is *ex ante* egalitarianism which gets the “wrong” answer.

Furthermore, XAE and XPE each suffer from a critical flaw.

- ▶ XAE violates the Statewise Dominance axiom (so it is “irrational”).
- ▶ XPE violates the *ex ante* Pareto axiom (so it is “paternalistic”).

It seems that *neither* approach is really optimal. Instead, we should employ a “hybrid” of XAE and XPE (Ben Porath, Gilboa & Schmeidler; Gajdos & Maurin).

Ex ante egalitarianism vs. *ex post* egalitarianism

(27/47)

In Diamond's Paradox, *ex post* egalitarianism gets the “wrong” answer. This suggests that *ex ante* egalitarianism is superior. But consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

H	h	t
Ann	1	0
Bob	1	0

E and H yield more egalitarian outcomes than F in *both* states of nature. So *ex post* egalitarianism would prefer both E and H over F .

(Meanwhile, the choice between E and H might be based on social risk aversion.)

But all three prospects yield the same expected utility ($1/2$) for each agent. So *ex ante* egalitarianism would be indifferent between them.

In this case, it is *ex ante* egalitarianism which gets the “wrong” answer. Furthermore, XAE and XPE each suffer from a critical flaw.

- ▶ XAE violates the Statewise Dominance axiom (so it is “irrational”).
- ▶ XPE violates the *ex ante* Pareto axiom (so it is “paternalistic”).

It seems that *neither* approach is really optimal. Instead, we should employ a “hybrid” of XAE and XPE (Ben Porath, Gilboa & Schmeidler; Gajdos & Maurin).

Ex ante egalitarianism vs. *ex post* egalitarianism

(27/47)

In Diamond's Paradox, *ex post* egalitarianism gets the “wrong” answer. This suggests that *ex ante* egalitarianism is superior. But consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

H	h	t
Ann	1	0
Bob	1	0

E and H yield more egalitarian outcomes than F in *both* states of nature. So *ex post* egalitarianism would prefer both E and H over F .

(Meanwhile, the choice between E and H might be based on social risk aversion.)

But all three prospects yield the same expected utility ($1/2$) for each agent. So *ex ante* egalitarianism would be indifferent between them.

In this case, it is *ex ante* egalitarianism which gets the “wrong” answer. Furthermore, XAE and XPE each suffer from a critical flaw.

- ▶ XAE violates the **Statewise Dominance** axiom (so it is “irrational”).
- ▶ XPE violates the *ex ante* Pareto axiom (so it is “paternalistic”).

It seems that *neither* approach is really optimal. Instead, we should employ a “hybrid” of XAE and XPE (Ben Porath, Gilboa & Schmeidler; Gajdos & Maurin).

In Diamond's Paradox, *ex post* egalitarianism gets the “wrong” answer. This suggests that *ex ante* egalitarianism is superior. But consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

H	h	t
Ann	1	0
Bob	1	0

E and H yield more egalitarian outcomes than F in *both* states of nature. So *ex post* egalitarianism would prefer both E and H over F .

(Meanwhile, the choice between E and H might be based on social risk aversion.)

But all three prospects yield the same expected utility ($1/2$) for each agent. So *ex ante* egalitarianism would be indifferent between them.

In this case, it is *ex ante* egalitarianism which gets the “wrong” answer. Furthermore, XAE and XPE each suffer from a critical flaw.

- ▶ XAE violates the Statewise Dominance axiom (so it is “irrational”).
- ▶ XPE violates the *ex ante* Pareto axiom (so it is “paternalistic”).

It seems that *neither* approach is really optimal. Instead, we should employ a “hybrid” of XAE and XPE (Ben Porath, Gilboa & Schmeidler; Gajdos & Maurin).

In Diamond's Paradox, *ex post* egalitarianism gets the “wrong” answer. This suggests that *ex ante* egalitarianism is superior. But consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

H	h	t
Ann	1	0
Bob	1	0

E and H yield more egalitarian outcomes than F in *both* states of nature. So *ex post* egalitarianism would prefer both E and H over F .

(Meanwhile, the choice between E and H might be based on social risk aversion.)

But all three prospects yield the same expected utility ($1/2$) for each agent. So *ex ante* egalitarianism would be indifferent between them.

In this case, it is *ex ante* egalitarianism which gets the “wrong” answer. Furthermore, XAE and XPE each suffer from a critical flaw.

- ▶ XAE violates the Statewise Dominance axiom (so it is “irrational”).
- ▶ XPE violates the *ex ante* Pareto axiom (so it is “paternalistic”).

It seems that *neither* approach is really optimal. Instead, we should employ a “**hybrid**” of XAE and XPE (Ben Porath, Gilboa & Schmeidler; Gajdos & Maurin).

A problem with these “hybrids” of XAE and XPE is that we might end up with the worst features of both.

Instead, Fleurbaey (2010) argued that we should try to combine the *best* features of XAE and XPE.

He proposed to rank such prospects by the expected value of an *ex post* SWF W_{xp} having the property: $W_{xp}(u, u, \dots, u) = u$, for any u in \mathbb{R} .

He called these *expected equally distributed equivalent* (EED) rankings.

An *egalitarian prospect* is one which yields a perfectly egalitarian social outcome in every state. For example:

F	s_1	s_2	s_3	s_4	s_5
Ann	3	0	5	2	-1
Bob	3	0	5	2	-1
Carol	3	0	5	2	-1

If F is egalitarian, then $\text{Expected Value}(W_{xp}, F) = \text{Expected Utility}_i(F)$ for every individual i in \mathcal{I} (assuming common probabilistic beliefs).

Thus, EED satisfies *ex ante* Pareto, when comparing egalitarian prospects.

A problem with these “hybrids” of XAE and XPE is that we might end up with the worst features of both.

Instead, Fleurbaey (2010) argued that we should try to combine the *best* features of XAE and XPE.

He proposed to rank such prospects by the expected value of an *ex post* SWF W_{xp} having the property: $W_{xp}(u, u, \dots, u) = u$, for any u in \mathbb{R} .

He called these *expected equally distributed equivalent* (EEDe) rankings.

An *egalitarian prospect* is one which yields a perfectly egalitarian social outcome in every state. For example:

F	s_1	s_2	s_3	s_4	s_5
Ann	3	0	5	2	-1
Bob	3	0	5	2	-1
Carol	3	0	5	2	-1

If F is egalitarian, then $\text{Expected Value}(W_{xp}, F) = \text{Expected Utility}_i(F)$ for every individual i in \mathcal{I} (assuming common probabilistic beliefs).

Thus, EEDe satisfies *ex ante* Pareto, when comparing egalitarian prospects.

A problem with these “hybrids” of XAE and XPE is that we might end up with the worst features of both.

Instead, Fleurbaey (2010) argued that we should try to combine the *best* features of XAE and XPE.

He proposed to rank such prospects by the expected value of an *ex post* SWF W_{xp} having the property: $W_{xp}(u, u, \dots, u) = u$, for any u in \mathbb{R} .

He called these *expected equally distributed equivalent* (EED) rankings.

An *egalitarian prospect* is one which yields a perfectly egalitarian social outcome in every state. For example:

F	s_1	s_2	s_3	s_4	s_5
Ann	3	0	5	2	-1
Bob	3	0	5	2	-1
Carol	3	0	5	2	-1

If F is egalitarian, then $\text{Expected Value}(W_{xp}, F) = \text{Expected Utility}_i(F)$ for every individual i in \mathcal{I} (assuming common probabilistic beliefs).

Thus, EED satisfies *ex ante* Pareto, when comparing egalitarian prospects.

A problem with these “hybrids” of XAE and XPE is that we might end up with the worst features of both.

Instead, Fleurbaey (2010) argued that we should try to combine the *best* features of XAE and XPE.

He proposed to rank such prospects by the expected value of an *ex post* SWF W_{xp} having the property: $W_{xp}(u, u, \dots, u) = u$, for any u in \mathbb{R} .

He called these *expected equally distributed equivalent* (EEDe) rankings.

An *egalitarian prospect* is one which yields a perfectly egalitarian social outcome in every state. For example:

F	s_1	s_2	s_3	s_4	s_5
Ann	3	0	5	2	-1
Bob	3	0	5	2	-1
Carol	3	0	5	2	-1

If F is egalitarian, then $\text{Expected Value}(W_{xp}, F) = \text{Expected Utility}_i(F)$ for every individual i in \mathcal{I} (assuming common probabilistic beliefs).

Thus, EEDe satisfies *ex ante* Pareto, when comparing egalitarian prospects.

A problem with these “hybrids” of XAE and XPE is that we might end up with the worst features of both.

Instead, Fleurbaey (2010) argued that we should try to combine the *best* features of XAE and XPE.

He proposed to rank such prospects by the expected value of an *ex post* SWF W_{xp} having the property: $W_{xp}(u, u, \dots, u) = u$, for any u in \mathbb{R} .

He called these *expected equally distributed equivalent* (EEDe) rankings.

An *egalitarian prospect* is one which yields a perfectly egalitarian social outcome in every state. For example:

F	s_1	s_2	s_3	s_4	s_5
Ann	3	0	5	2	-1
Bob	3	0	5	2	-1
Carol	3	0	5	2	-1

If F is egalitarian, then $\text{Expected Value}(W_{xp}, F) = \text{Expected Utility}_i(F)$ for every individual i in \mathcal{I} (assuming common probabilistic beliefs).

Thus, EEDe satisfies *ex ante* Pareto, when comparing egalitarian prospects.

A problem with these “hybrids” of XAE and XPE is that we might end up with the worst features of both.

Instead, Fleurbaey (2010) argued that we should try to combine the *best* features of XAE and XPE.

He proposed to rank such prospects by the expected value of an *ex post* SWF W_{xp} having the property: $W_{xp}(u, u, \dots, u) = u$, for any u in \mathbb{R} .

He called these *expected equally distributed equivalent* (EED) rankings.

An *egalitarian prospect* is one which yields a perfectly egalitarian social outcome in every state. For example:

F	s_1	s_2	s_3	s_4	s_5
Ann	3	0	5	2	-1
Bob	3	0	5	2	-1
Carol	3	0	5	2	-1

If F is egalitarian, then $\text{Expected Value}(W_{xp}, F) = \text{Expected Utility}_i(F)$ for every individual i in \mathcal{I} (assuming common probabilistic beliefs).

Thus, EED satisfies *ex ante* Pareto, when comparing egalitarian prospects.

A problem with these “hybrids” of XAE and XPE is that we might end up with the worst features of both.

Instead, Fleurbaey (2010) argued that we should try to combine the *best* features of XAE and XPE.

He proposed to rank such prospects by the expected value of an *ex post* SWF W_{xp} having the property: $W_{xp}(u, u, \dots, u) = u$, for any u in \mathbb{R} .

He called these *expected equally distributed equivalent* (EEDe) rankings.

An *egalitarian prospect* is one which yields a perfectly egalitarian social outcome in every state. For example:

F	s_1	s_2	s_3	s_4	s_5
Ann	3	0	5	2	-1
Bob	3	0	5	2	-1
Carol	3	0	5	2	-1

If F is egalitarian, then $\text{Expected Value}(W_{xp}, F) = \text{Expected Utility}_i(F)$ for every individual i in \mathcal{I} (assuming common probabilistic beliefs).

Thus, EEDe satisfies *ex ante* Pareto, when comparing egalitarian prospects.

A *riskless prospect* is one where each person gets the same outcome in every state.
For example:

F	s_1	s_2	s_3	s_4	s_5
Ann	3	3	3	3	3
Bob	5	5	5	5	5
Carol	1	1	1	1	1

Any XPE ranking satisfies *ex ante* Pareto if comparing riskless prospects.

And any XPE ranking satisfies Statewise Dominance.

In fact, these three properties characterise the EEDE rankings.

Theorem. (Fleurbaey, 2010) *Let \succeq be a ranking of all social prospects. Then \succeq satisfies Statewise Dominance, ex ante Pareto for riskless prospects, and ex ante Pareto for egalitarian prospects, if and only if \succeq is an EEDE ranking.*

By extending *ex ante* Pareto to “non-reranking prospects”, Fleurbaey also characterized two subclasses of EEDE rankings: *ex post generalized Gini* and *ex post leximin*.

A *riskless prospect* is one where each person gets the same outcome in every state.
For example:

F	s_1	s_2	s_3	s_4	s_5
Ann	3	3	3	3	3
Bob	5	5	5	5	5
Carol	1	1	1	1	1

Any XPE ranking satisfies *ex ante* Pareto if comparing riskless prospects.

And any XPE ranking satisfies Statewise Dominance.

In fact, these three properties characterise the EEDE rankings.

Theorem. (Fleurbaey, 2010) *Let \succeq be a ranking of all social prospects. Then \succeq satisfies Statewise Dominance, ex ante Pareto for riskless prospects, and ex ante Pareto for egalitarian prospects, if and only if \succeq is an EEDE ranking.*

By extending *ex ante* Pareto to “non-reranking prospects”, Fleurbaey also characterized two subclasses of EEDE rankings: *ex post generalized Gini* and *ex post leximin*.

A *riskless prospect* is one where each person gets the same outcome in every state.
For example:

F	s_1	s_2	s_3	s_4	s_5
Ann	3	3	3	3	3
Bob	5	5	5	5	5
Carol	1	1	1	1	1

Any XPE ranking satisfies *ex ante* Pareto if comparing riskless prospects.

And any XPE ranking satisfies Statewise Dominance.

In fact, these three properties characterise the EEDE rankings.

Theorem. (Fleurbaey, 2010) Let \succeq be a ranking of all social prospects. Then \succeq satisfies Statewise Dominance, ex ante Pareto for riskless prospects, and ex ante Pareto for egalitarian prospects, if and only if \succeq is an EEDE ranking.

By extending *ex ante* Pareto to “non-reranking prospects”, Fleurbaey also characterized two subclasses of EEDE rankings: *ex post generalized Gini* and *ex post leximin*.

A *riskless prospect* is one where each person gets the same outcome in every state.
For example:

F	s_1	s_2	s_3	s_4	s_5
Ann	3	3	3	3	3
Bob	5	5	5	5	5
Carol	1	1	1	1	1

Any XPE ranking satisfies *ex ante* Pareto if comparing riskless prospects.

And any XPE ranking satisfies Statewise Dominance.

In fact, these three properties characterise the EEDE rankings.

Theorem. (Fleurbaey, 2010) Let \succeq be a ranking of all social prospects. Then \succeq satisfies Statewise Dominance, ex ante Pareto for riskless prospects, and ex ante Pareto for egalitarian prospects, if and only if \succeq is an EEDE ranking.

By extending *ex ante* Pareto to “non-reranking prospects”, Fleurbaey also characterized two subclasses of EEDE rankings: *ex post generalized Gini* and *ex post leximin*.

A *riskless prospect* is one where each person gets the same outcome in every state.
For example:

F	s_1	s_2	s_3	s_4	s_5
Ann	3	3	3	3	3
Bob	5	5	5	5	5
Carol	1	1	1	1	1

Any XPE ranking satisfies *ex ante* Pareto if comparing riskless prospects.

And any XPE ranking satisfies Statewise Dominance.

In fact, these three properties characterise the EEDE rankings.

Theorem. (Fleurbaey, 2010) *Let \succeq be a ranking of all social prospects. Then \succeq satisfies Statewise Dominance, ex ante Pareto for riskless prospects, and ex ante Pareto for egalitarian prospects, if and only if \succeq is an EEDE ranking.*

By extending *ex ante* Pareto to “non-reranking prospects”, Fleurbaey also characterized two subclasses of EEDE rankings: *ex post generalized Gini* and *ex post leximin*.

A *riskless prospect* is one where each person gets the same outcome in every state.
For example:

F	s_1	s_2	s_3	s_4	s_5
Ann	3	3	3	3	3
Bob	5	5	5	5	5
Carol	1	1	1	1	1

Any XPE ranking satisfies *ex ante* Pareto if comparing riskless prospects.

And any XPE ranking satisfies Statewise Dominance.

In fact, these three properties characterise the EEDE rankings.

Theorem. (Fleurbaey, 2010) *Let \succeq be a ranking of all social prospects. Then \succeq satisfies Statewise Dominance, ex ante Pareto for riskless prospects, and ex ante Pareto for egalitarian prospects, if and only if \succeq is an EEDE ranking.*

By extending *ex ante* Pareto to “non-reranking prospects”, Fleurbaey also characterized two subclasses of EEDE rankings: *ex post generalized Gini* and *ex post leximin*.

A *riskless prospect* is one where each person gets the same outcome in every state.
For example:

F	s_1	s_2	s_3	s_4	s_5
Ann	3	3	3	3	3
Bob	5	5	5	5	5
Carol	1	1	1	1	1

Any XPE ranking satisfies *ex ante* Pareto if comparing riskless prospects.

And any XPE ranking satisfies Statewise Dominance.

In fact, these three properties characterise the EEDE rankings.

Theorem. (Fleurbaey, 2010) *Let \succeq be a ranking of all social prospects. Then \succeq satisfies Statewise Dominance, ex ante Pareto for riskless prospects, and ex ante Pareto for egalitarian prospects, if and only if \succeq is an EEDE ranking.*

By extending *ex ante* Pareto to “non-reranking prospects”, Fleurbaey also characterized two subclasses of EEDE rankings: *ex post generalized Gini* and *ex post leximin*.

Social choice under uncertainty becomes surprisingly complicated, once issues of fairness and heterogeneity come into play.

However, there are also other problems we haven't even discussed...

- ▶ Increasing *ex post* equality correlates individuals' risks, meaning greater collective risk, *ex ante*. So there is a tradeoff between egalitarianism and social risk-aversion. (e.g. Keeney, 1980; Fishburn & Straffin, 1989).
- ▶ Diamond argued that “fairness” requires a coin toss. But if it was necessary to flip the coin once, why not flip it again? To avoid falling into an infinite cycle of coin flips, it seems that our *ex post* SWF must be *history-dependent* (Machina, 1988; Hayashi, 2013).
- ▶ *Ex ante* Pareto assumes people face uncertainty “rationally” (e.g. are SEU maximizers), and that the utility they maximize *ex ante* accurately predicts their welfare *ex post*. But empirically, neither claim is true (e.g. Kahneman & Schkade, 1998).
- ▶ This suggests we turn away from *ex ante* and towards *ex post*. But in a sense, *there no such thing as ex post* (e.g. Hild, Jeffrey & Risse, 1997). There is no “end of history” when all uncertainty has been resolved.

Social choice under uncertainty becomes surprisingly complicated, once issues of fairness and heterogeneity come into play.

However, there are also other problems we haven't even discussed...

- ▶ Increasing *ex post* equality correlates individuals' risks, meaning greater collective risk, *ex ante*. So there is a tradeoff between egalitarianism and social risk-aversion. (e.g. Keeney, 1980; Fishburn & Straffin, 1989).
- ▶ Diamond argued that “fairness” requires a coin toss. But if it was necessary to flip the coin once, why not flip it again? To avoid falling into an infinite cycle of coin flips, it seems that our *ex post* SWF must be *history-dependent* (Machina, 1988; Hayashi, 2013).
- ▶ *Ex ante* Pareto assumes people face uncertainty “rationally” (e.g. are SEU maximizers), and that the utility they maximize *ex ante* accurately predicts their welfare *ex post*. But empirically, neither claim is true (e.g. Kahneman & Schkade, 1998).
- ▶ This suggests we turn away from *ex ante* and towards *ex post*. But in a sense, *there no such thing as ex post* (e.g. Hild, Jeffrey & Risse, 1997). There is no “end of history” when all uncertainty has been resolved.

Social choice under uncertainty becomes surprisingly complicated, once issues of fairness and heterogeneity come into play.

However, there are also other problems we haven't even discussed...

- ▶ Increasing *ex post* equality correlates individuals' risks, meaning greater collective risk, *ex ante*. So there is a tradeoff between egalitarianism and social risk-aversion. (e.g. Keeney, 1980; Fishburn & Straffin, 1989).
- ▶ Diamond argued that "fairness" requires a coin toss. But if it was necessary to flip the coin once, why not flip it again? To avoid falling into an infinite cycle of coin flips, it seems that our *ex post* SWF must be *history-dependent* (Machina, 1988; Hayashi, 2013).
- ▶ *Ex ante* Pareto assumes people face uncertainty "rationally" (e.g. are SEU maximizers), and that the utility they maximize *ex ante* accurately predicts their welfare *ex post*. But empirically, neither claim is true (e.g. Kahneman & Schkade, 1998).
- ▶ This suggests we turn away from *ex ante* and towards *ex post*. But in a sense, *there no such thing as ex post* (e.g. Hild, Jeffrey & Risse, 1997). There is no "end of history" when all uncertainty has been resolved.

Social choice under uncertainty becomes surprisingly complicated, once issues of fairness and heterogeneity come into play.

However, there are also other problems we haven't even discussed...

- ▶ Increasing *ex post* equality correlates individuals' risks, meaning greater collective risk, *ex ante*. So there is a tradeoff between egalitarianism and social risk-aversion. (e.g. Keeney, 1980; Fishburn & Straffin, 1989).
- ▶ Diamond argued that “fairness” requires a coin toss. But if it was necessary to flip the coin once, why not flip it again? To avoid falling into an infinite cycle of coin flips, it seems that our *ex post* SWF must be *history-dependent* (Machina, 1988; Hayashi, 2013).
- ▶ *Ex ante* Pareto assumes people face uncertainty “rationally” (e.g. are SEU maximizers), and that the utility they maximize *ex ante* accurately predicts their welfare *ex post*. But empirically, neither claim is true (e.g. Kahneman & Schkade, 1998).
- ▶ This suggests we turn away from *ex ante* and towards *ex post*. But in a sense, *there no such thing as ex post* (e.g. Hild, Jeffrey & Risse, 1997). There is no “end of history” when all uncertainty has been resolved.

Social choice under uncertainty becomes surprisingly complicated, once issues of fairness and heterogeneity come into play.

However, there are also other problems we haven't even discussed...

- ▶ Increasing *ex post* equality correlates individuals' risks, meaning greater collective risk, *ex ante*. So there is a tradeoff between egalitarianism and social risk-aversion. (e.g. Keeney, 1980; Fishburn & Straffin, 1989).
- ▶ Diamond argued that “fairness” requires a coin toss. But if it was necessary to flip the coin once, why not flip it again? To avoid falling into an infinite cycle of coin flips, it seems that our *ex post* SWF must be *history-dependent* (Machina, 1988; Hayashi, 2013).
- ▶ *Ex ante* Pareto assumes people face uncertainty “rationally” (e.g. are SEU maximizers), and that the utility they maximize *ex ante* accurately predicts their welfare *ex post*. But empirically, neither claim is true (e.g. Kahneman & Schkade, 1998).
- ▶ This suggests we turn away from *ex ante* and towards *ex post*. But in a sense, *there no such thing as ex post* (e.g. Hild, Jeffrey & Risse, 1997). There is no “end of history” when all uncertainty has been resolved.

Social choice under uncertainty becomes surprisingly complicated, once issues of fairness and heterogeneity come into play.

However, there are also other problems we haven't even discussed...

- ▶ Increasing *ex post* equality correlates individuals' risks, meaning greater collective risk, *ex ante*. So there is a tradeoff between egalitarianism and social risk-aversion. (e.g. Keeney, 1980; Fishburn & Straffin, 1989).
- ▶ Diamond argued that “fairness” requires a coin toss. But if it was necessary to flip the coin once, why not flip it again? To avoid falling into an infinite cycle of coin flips, it seems that our *ex post* SWF must be *history-dependent* (Machina, 1988; Hayashi, 2013).
- ▶ *Ex ante* Pareto assumes people face uncertainty “rationally” (e.g. are SEU maximizers), and that the utility they maximize *ex ante* accurately predicts their welfare *ex post*. But empirically, neither claim is true (e.g. Kahneman & Schkade, 1998).
- ▶ This suggests we turn away from *ex ante* and towards *ex post*. But in a sense, *there no such thing as ex post* (e.g. Hild, Jeffrey & Risse, 1997). There is no “end of history” when all uncertainty has been resolved.

Thank you.

Introduction

Crash course in decision theory

Risk: Lotteries and von Neumann and Morgenstern

Uncertainty: Prospects and Savage

Risk and uncertainty in public policy

Part I. Harsanyi's social aggregation theorem

Hypotheses

Theorem statement

Harsanyi and his discontents

Part II. Uncertainty: Heterogeneity and spurious unanimity

Social prospects

Bayesian Social Aggregation Theorem

"Consistent Bayesian aggregation is impossible"

Possible escapes?

Spurious Unanimity

Gilboa, Samet & Schmeidler

Restricted *ex ante* Pareto

Spurious unanimity returns

Objective vs. subjective uncertainty

Part III. Equality: The Diamond Paradox

The Diamond Paradox

Ex ante egalitarianism

Ex post egalitarianism

Ex ante egalitarianism vs. *ex post* egalitarianism

Fleurbaey: expected equally distributed equivalent rankings

Conclusion

Appendices

Appendix A: Sen's objections to Harsanyi

Appendix B: Bayesian Social Aggregation with minimal hypotheses

Appendix C: Complete specification in the Diamond Paradox

Appendix D: Hybrids of XAE and XPE

Appendix A

Sen's objections to Harsanyi's SAT

Harsanyi's Social Aggregation Theorem: If \succeq_i all satisfy **Individual vNM** and \succeq satisfies **Group vNM** and **XAP**, then there are weights $c_i > 0$ for all i in \mathcal{I} , and a constant K , such that, for any a in \mathcal{A} , we have

$$W(a) = K + c_{\text{Ann}} u_{\text{Ann}}(a) + c_{\text{Bob}} u_{\text{Bob}}(a) + c_{\text{Carl}} u_{\text{Carl}}(a) + \dots$$

Harsanyi interpreted his SAT as an argument for utilitarianism.

But Sen made three objections to this interpretation....

(S1) The SAT just says that any \succeq satisfying **Group vNM** and **XAP** can be interpreted *as if* it came from some weighted utilitarian social welfare function, for some suitable choice of weights $c_{\text{Ann}}, c_{\text{Bob}}, c_{\text{Carl}}, \dots$

But it doesn't tell us how to *obtain* these weights. Nor does it tell us how they should change if $\succeq_{\text{Ann}}, \succeq_{\text{Bob}}, \succeq_{\text{Carl}}, \dots$ change. It doesn't give us a “profile-independent” formula for W as a function of $u_{\text{Ann}}, u_{\text{Bob}}, u_{\text{Carl}}, \dots$

(S2) The vNM utility functions u_i are only defined up to multiplication by a constant: If u_i represents \succeq_i , then so does $3 \cdot u_i$.

Thus, the weights c_i are *also* only defined up to multiplication by a constant. (If we replace u_{Ann} with $3 u_{\text{Ann}}$, then we replace c_{Ann} with $c_{\text{Ann}}/3$.)

(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really “utilitarian” —it is an aggregate of “lottery preferences”, not an aggregate of well-being per se.

Harsanyi's Social Aggregation Theorem: If \succeq_i all satisfy **Individual vNM** and \succeq satisfies **Group vNM** and **XAP**, then there are weights $c_i > 0$ for all i in \mathcal{I} , and a constant K , such that, for any a in \mathcal{A} , we have

$$W(a) = K + c_{\text{Ann}} u_{\text{Ann}}(a) + c_{\text{Bob}} u_{\text{Bob}}(a) + c_{\text{Carl}} u_{\text{Carl}}(a) + \dots$$

Harsanyi interpreted his SAT as an argument for utilitarianism.

But Sen made three objections to this interpretation....

(S1) The SAT just says that any \succeq satisfying **Group vNM** and **XAP** can be interpreted *as if* it came from some weighted utilitarian social welfare function, for some suitable choice of weights $c_{\text{Ann}}, c_{\text{Bob}}, c_{\text{Carl}}, \dots$

But it doesn't tell us how to *obtain* these weights. Nor does it tell us how they should change if $\succeq_{\text{Ann}}, \succeq_{\text{Bob}}, \succeq_{\text{Carl}}, \dots$ change. It doesn't give us a "profile-independent" formula for W as a function of $u_{\text{Ann}}, u_{\text{Bob}}, u_{\text{Carl}}, \dots$

(S2) The vNM utility functions u_i are only defined up to multiplication by a constant: If u_i represents \succeq_i , then so does $3 \cdot u_i$.

Thus, the weights c_i are *also* only defined up to multiplication by a constant. (If we replace u_{Ann} with $3 u_{\text{Ann}}$, then we replace c_{Ann} with $c_{\text{Ann}}/3$.)

(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really "utilitarian" —it is an aggregate of "lottery preferences", not an aggregate of well-being per se.

Harsanyi's Social Aggregation Theorem: If \succeq_i all satisfy **Individual vNM** and \succeq satisfies **Group vNM** and **XAP**, then there are weights $c_i > 0$ for all i in \mathcal{I} , and a constant K , such that, for any a in \mathcal{A} , we have

$$W(a) = K + c_{\text{Ann}} u_{\text{Ann}}(a) + c_{\text{Bob}} u_{\text{Bob}}(a) + c_{\text{Carl}} u_{\text{Carl}}(a) + \dots$$

Harsanyi interpreted his SAT as an argument for utilitarianism.

But Sen made three objections to this interpretation....

(S1) The SAT just says that any \succeq satisfying **Group vNM** and **XAP** can be interpreted *as if* it came from some weighted utilitarian social welfare function, for some suitable choice of weights $c_{\text{Ann}}, c_{\text{Bob}}, c_{\text{Carl}}, \dots$

But it doesn't tell us how to *obtain* these weights. Nor does it tell us how they should change if $\succeq_{\text{Ann}}, \succeq_{\text{Bob}}, \succeq_{\text{Carl}}, \dots$ change. It doesn't give us a "profile-independent" formula for W as a function of $u_{\text{Ann}}, u_{\text{Bob}}, u_{\text{Carl}}, \dots$

(S2) The vNM utility functions u_i are only defined up to multiplication by a constant: If u_i represents \succeq_i , then so does $3 \cdot u_i$.

Thus, the weights c_i are *also* only defined up to multiplication by a constant. (If we replace u_{Ann} with $3 u_{\text{Ann}}$, then we replace c_{Ann} with $c_{\text{Ann}}/3$.)

(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really "utilitarian" —it is an aggregate of "lottery preferences", not an aggregate of well-being per se.

Harsanyi's Social Aggregation Theorem: If \succeq_i all satisfy **Individual vNM** and \succeq satisfies **Group vNM** and **XAP**, then there are weights $c_i > 0$ for all i in \mathcal{I} , and a constant K , such that, for any a in \mathcal{A} , we have

$$W(a) = K + c_{\text{Ann}} u_{\text{Ann}}(a) + c_{\text{Bob}} u_{\text{Bob}}(a) + c_{\text{Carl}} u_{\text{Carl}}(a) + \dots$$

Harsanyi interpreted his SAT as an argument for utilitarianism.

But Sen made three objections to this interpretation....

(S1) The SAT just says that any \succeq satisfying **Group vNM** and **XAP** can be interpreted as if it came from some weighted utilitarian social welfare function, for some suitable choice of weights $c_{\text{Ann}}, c_{\text{Bob}}, c_{\text{Carl}}, \dots$

But it doesn't tell us how to **obtain** these weights. Nor does it tell us how they should change if $\succeq_{\text{Ann}}, \succeq_{\text{Bob}}, \succeq_{\text{Carl}}, \dots$ change. It doesn't give us a "profile-independent" formula for W as a function of $u_{\text{Ann}}, u_{\text{Bob}}, u_{\text{Carl}}, \dots$

(S2) The vNM utility functions u_i are only defined up to multiplication by a constant: If u_i represents \succeq_i , then so does $3 \cdot u_i$.

Thus, the weights c_i are also only defined up to multiplication by a constant. (If we replace u_{Ann} with $3 u_{\text{Ann}}$, then we replace c_{Ann} with $c_{\text{Ann}}/3$.)

(S3) u_i represents person i 's preferences over lotteries. But it isn't clear u_i really measures i 's well-being. So W is not really "utilitarian" —it is an aggregate of "lottery preferences", not an aggregate of well-being per se.

Harsanyi's Social Aggregation Theorem: If \succeq_i all satisfy **Individual vNM** and \succeq satisfies **Group vNM** and **XAP**, then there are weights $c_i > 0$ for all i in \mathcal{I} , and a constant K , such that, for any a in \mathcal{A} , we have

$$W(a) = K + c_{\text{Ann}} u_{\text{Ann}}(a) + c_{\text{Bob}} u_{\text{Bob}}(a) + c_{\text{Carl}} u_{\text{Carl}}(a) + \dots$$

Harsanyi interpreted his SAT as an argument for utilitarianism.

But Sen made three objections to this interpretation....

(S1) The SAT just says that any \succeq satisfying **Group vNM** and **XAP** can be interpreted *as if* it came from some weighted utilitarian social welfare function, for some suitable choice of weights $c_{\text{Ann}}, c_{\text{Bob}}, c_{\text{Carl}}, \dots$

But it doesn't tell us how to *obtain* these weights. Nor does it tell us how they should change if $\succeq_{\text{Ann}}, \succeq_{\text{Bob}}, \succeq_{\text{Carl}}, \dots$ change. It doesn't give us a “**profile-independent**” formula for W as a function of $u_{\text{Ann}}, u_{\text{Bob}}, u_{\text{Carl}}, \dots$

(S2) The vNM utility functions u_i are only defined up to multiplication by a constant: If u_i represents \succeq_i , then so does $3 \cdot u_i$.

Thus, the weights c_i are *also* only defined up to multiplication by a constant. (If we replace u_{Ann} with $3 u_{\text{Ann}}$, then we replace c_{Ann} with $c_{\text{Ann}}/3$.)

(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really “utilitarian” —it is an aggregate of “lottery preferences”, not an aggregate of well-being per se.

Harsanyi's Social Aggregation Theorem: If \succeq_i all satisfy **Individual vNM** and \succeq satisfies **Group vNM** and **XAP**, then there are weights $c_i > 0$ for all i in \mathcal{I} , and a constant K , such that, for any a in \mathcal{A} , we have

$$W(a) = K + c_{\text{Ann}} u_{\text{Ann}}(a) + c_{\text{Bob}} u_{\text{Bob}}(a) + c_{\text{Carl}} u_{\text{Carl}}(a) + \dots$$

Harsanyi interpreted his SAT as an argument for utilitarianism.

But Sen made three objections to this interpretation....

(S1) The SAT just says that any \succeq satisfying **Group vNM** and **XAP** can be interpreted *as if* it came from some weighted utilitarian social welfare function, for some suitable choice of weights $c_{\text{Ann}}, c_{\text{Bob}}, c_{\text{Carl}}, \dots$

But it doesn't tell us how to *obtain* these weights. Nor does it tell us how they should change if $\succeq_{\text{Ann}}, \succeq_{\text{Bob}}, \succeq_{\text{Carl}}, \dots$ change. It doesn't give us a "profile-independent" formula for W as a function of $u_{\text{Ann}}, u_{\text{Bob}}, u_{\text{Carl}}, \dots$

(S2) The vNM utility functions u_i are only defined up to multiplication by a constant: If u_i represents \succeq_i , then so does $3 \cdot u_i$.

Thus, the weights c_i are *also* only defined up to multiplication by a constant. (If we replace u_{Ann} with $3 u_{\text{Ann}}$, then we replace c_{Ann} with $c_{\text{Ann}}/3$.)

(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really "utilitarian" —it is an aggregate of "lottery preferences", not an aggregate of well-being per se.

Harsanyi's Social Aggregation Theorem: If \succeq_i all satisfy **Individual vNM** and \succeq satisfies **Group vNM** and **XAP**, then there are weights $c_i > 0$ for all i in \mathcal{I} , and a constant K , such that, for any a in \mathcal{A} , we have

$$W(a) = K + c_{\text{Ann}} u_{\text{Ann}}(a) + c_{\text{Bob}} u_{\text{Bob}}(a) + c_{\text{Carl}} u_{\text{Carl}}(a) + \dots$$

Harsanyi interpreted his SAT as an argument for utilitarianism.

But Sen made three objections to this interpretation....

(S1) The SAT just says that any \succeq satisfying **Group vNM** and **XAP** can be interpreted *as if* it came from some weighted utilitarian social welfare function, for some suitable choice of weights $c_{\text{Ann}}, c_{\text{Bob}}, c_{\text{Carl}}, \dots$

But it doesn't tell us how to *obtain* these weights. Nor does it tell us how they should change if $\succeq_{\text{Ann}}, \succeq_{\text{Bob}}, \succeq_{\text{Carl}}, \dots$ change. It doesn't give us a "profile-independent" formula for W as a function of $u_{\text{Ann}}, u_{\text{Bob}}, u_{\text{Carl}}, \dots$

(S2) The vNM utility functions u_i are only defined up to multiplication by a constant: If u_i represents \succeq_i , then so does $3 \cdot u_i$.

Thus, the weights c_i are *also* only defined up to multiplication by a constant. (If we replace u_{Ann} with $3 u_{\text{Ann}}$, then we replace c_{Ann} with $c_{\text{Ann}}/3$.)

(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really "utilitarian" —it is an aggregate of "lottery preferences", not an aggregate of well-being *per se*.

Harsanyi's Social Aggregation Theorem: If \succeq_i all satisfy **Individual vNM** and \succeq satisfies **Group vNM** and **XAP**, then there are weights $c_i > 0$ for all i in \mathcal{I} , and a constant K , such that, for any a in \mathcal{A} , we have

$$W(a) = K + c_{\text{Ann}} u_{\text{Ann}}(a) + c_{\text{Bob}} u_{\text{Bob}}(a) + c_{\text{Carl}} u_{\text{Carl}}(a) + \dots$$

Harsanyi interpreted his SAT as an argument for utilitarianism.

But Sen made three objections to this interpretation....

(S1) The SAT just says that any \succeq satisfying **Group vNM** and **XAP** can be interpreted *as if* it came from some weighted utilitarian social welfare function, for some suitable choice of weights $c_{\text{Ann}}, c_{\text{Bob}}, c_{\text{Carl}}, \dots$

But it doesn't tell us how to *obtain* these weights. Nor does it tell us how they should change if $\succeq_{\text{Ann}}, \succeq_{\text{Bob}}, \succeq_{\text{Carl}}, \dots$ change. It doesn't give us a "profile-independent" formula for W as a function of $u_{\text{Ann}}, u_{\text{Bob}}, u_{\text{Carl}}, \dots$

(S2) The vNM utility functions u_i are only defined up to multiplication by a constant: If u_i represents \succeq_i , then so does $3 \cdot u_i$.

Thus, the weights c_i are *also* only defined up to multiplication by a constant. (If we replace u_{Ann} with $3 u_{\text{Ann}}$, then we replace c_{Ann} with $c_{\text{Ann}}/3$.)

(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really "utilitarian" —it is an aggregate of "lottery preferences", not an aggregate of well-being per se.

Harsanyi's Social Aggregation Theorem: If \succeq_i all satisfy **Individual vNM** and \succeq satisfies **Group vNM** and **XAP**, then there are weights $c_i > 0$ for all i in \mathcal{I} , and a constant K , such that, for any a in \mathcal{A} , we have

$$W(a) = K + c_{\text{Ann}} u_{\text{Ann}}(a) + c_{\text{Bob}} u_{\text{Bob}}(a) + c_{\text{Carl}} u_{\text{Carl}}(a) + \dots$$

Harsanyi interpreted his SAT as an argument for utilitarianism.

But Sen made three objections to this interpretation....

(S1) The SAT just says that any \succeq satisfying **Group vNM** and **XAP** can be interpreted *as if* it came from some weighted utilitarian social welfare function, for some suitable choice of weights $c_{\text{Ann}}, c_{\text{Bob}}, c_{\text{Carl}}, \dots$

But it doesn't tell us how to *obtain* these weights. Nor does it tell us how they should change if $\succeq_{\text{Ann}}, \succeq_{\text{Bob}}, \succeq_{\text{Carl}}, \dots$ change. It doesn't give us a "profile-independent" formula for W as a function of $u_{\text{Ann}}, u_{\text{Bob}}, u_{\text{Carl}}, \dots$

(S2) The vNM utility functions u_i are only defined up to multiplication by a constant: If u_i represents \succeq_i , then so does $3 \cdot u_i$.

Thus, the weights c_i are *also* only defined up to multiplication by a constant. (If we replace u_{Ann} with $3 u_{\text{Ann}}$, then we replace c_{Ann} with $c_{\text{Ann}}/3$.)

(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really "utilitarian" —it is an aggregate of "lottery preferences", not an aggregate of well-being *per se*.

“(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really 'utilitarian' —it is an aggregate of 'lottery preferences', not an aggregate of well-being *per se*.”

To appreciate this, suppose w_i is person i 's “true” welfare function. (So $w_i(a)$ = the welfare of person i in outcome a .)

Assuming i is self-interested, her vNM utility function u_i should have the form $u_i(a) = \phi_i[w_i(a)]$, where ϕ_i is some increasing function.

The shape of ϕ_i expresses i 's “risk attitudes”. For example:

- ▶ if i is *risk-averse*, then ϕ_i is concave.
- ▶ If i is *risk-prone* then ϕ_i is convex.

But *there is no reason to suppose that ϕ_i is linear*.

Thus, Harsanyi's SAT actually characterizes an *ex ante* SWF of the form:

$$W(a) = K + c_{\text{Ann}} \phi_{\text{Ann}}[w_{\text{Ann}}(a)] + c_{\text{Bob}} \phi_{\text{Bob}}[w_{\text{Bob}}(a)] + c_{\text{Carl}} \phi_{\text{Carl}}[w_{\text{Carl}}(a)] + \dots$$

This does not invalidate SAT's analysis of social choice under uncertainty. It merely invalidates Harsanyi's “utilitarian” interpretation of the SAT.

“(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really 'utilitarian' —it is an aggregate of 'lottery preferences', not an aggregate of well-being *per se*.”

To appreciate this, suppose w_i is person i 's “true” welfare function. (So $w_i(a)$ = the welfare of person i in outcome a .)

Assuming i is self-interested, her vNM utility function u_i should have the form $u_i(a) = \phi_i[w_i(a)]$, where ϕ_i is some increasing function.

The shape of ϕ_i expresses i 's “risk attitudes”. For example:

- ▶ if i is *risk-averse*, then ϕ_i is concave.
- ▶ If i is *risk-prone* then ϕ_i is convex.

But *there is no reason to suppose that ϕ_i is linear*.

Thus, Harsanyi's SAT actually characterizes an *ex ante* SWF of the form:

$$W(a) = K + c_{\text{Ann}} \phi_{\text{Ann}}[w_{\text{Ann}}(a)] + c_{\text{Bob}} \phi_{\text{Bob}}[w_{\text{Bob}}(a)] + c_{\text{Carl}} \phi_{\text{Carl}}[w_{\text{Carl}}(a)] + \dots$$

This does not invalidate SAT's analysis of social choice under uncertainty. It merely invalidates Harsanyi's “utilitarian” interpretation of the SAT.

“(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really 'utilitarian' —it is an aggregate of 'lottery preferences', not an aggregate of well-being *per se*.”

To appreciate this, suppose w_i is person i 's “true” welfare function. (So $w_i(a)$ = the welfare of person i in outcome a .)

Assuming i is self-interested, her vNM utility function u_i should have the form $u_i(a) = \phi_i[w_i(a)]$, where ϕ_i is some increasing function.

The shape of ϕ_i expresses i 's “risk attitudes”. For example:

- ▶ if i is *risk-averse*, then ϕ_i is concave.
- ▶ If i is *risk-prone* then ϕ_i is convex.

But *there is no reason to suppose that ϕ_i is linear*.

Thus, Harsanyi's SAT actually characterizes an *ex ante* SWF of the form:

$$W(a) = K + c_{\text{Ann}} \phi_{\text{Ann}}[w_{\text{Ann}}(a)] + c_{\text{Bob}} \phi_{\text{Bob}}[w_{\text{Bob}}(a)] + c_{\text{Carl}} \phi_{\text{Carl}}[w_{\text{Carl}}(a)] + \dots$$

This does not invalidate SAT's analysis of social choice under uncertainty. It merely invalidates Harsanyi's “utilitarian” interpretation of the SAT.

“(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really 'utilitarian' —it is an aggregate of 'lottery preferences', not an aggregate of well-being *per se*.”

To appreciate this, suppose w_i is person i 's “true” welfare function. (So $w_i(a)$ = the welfare of person i in outcome a .)

Assuming i is self-interested, her vNM utility function u_i should have the form $u_i(a) = \phi_i[w_i(a)]$, where ϕ_i is some increasing function.

The shape of ϕ_i expresses i 's “risk attitudes”. For example:

- ▶ if i is *risk-averse*, then ϕ_i is concave.
- ▶ If i is *risk-prone* then ϕ_i is convex.

But *there is no reason to suppose that ϕ_i is linear*.

Thus, Harsanyi's SAT actually characterizes an *ex ante* SWF of the form:

$$W(a) = K + c_{\text{Ann}} \phi_{\text{Ann}}[w_{\text{Ann}}(a)] + c_{\text{Bob}} \phi_{\text{Bob}}[w_{\text{Bob}}(a)] + c_{\text{Carl}} \phi_{\text{Carl}}[w_{\text{Carl}}(a)] + \dots$$

This does not invalidate SAT's analysis of social choice under uncertainty. It merely invalidates Harsanyi's “utilitarian” interpretation of the SAT.

“(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really 'utilitarian' —it is an aggregate of 'lottery preferences', not an aggregate of well-being *per se*.”

To appreciate this, suppose w_i is person i 's “true” welfare function. (So $w_i(a)$ = the welfare of person i in outcome a .)

Assuming i is self-interested, her vNM utility function u_i should have the form $u_i(a) = \phi_i[w_i(a)]$, where ϕ_i is some increasing function.

The shape of ϕ_i expresses i 's “risk attitudes”. For example:

- ▶ if i is *risk-averse*, then ϕ_i is concave.
- ▶ If i is *risk-prone* then ϕ_i is convex.

But *there is no reason to suppose that ϕ_i is linear*.

Thus, Harsanyi's SAT actually characterizes an *ex ante* SWF of the form:

$$W(a) = K + c_{\text{Ann}} \phi_{\text{Ann}}[w_{\text{Ann}}(a)] + c_{\text{Bob}} \phi_{\text{Bob}}[w_{\text{Bob}}(a)] + c_{\text{Carl}} \phi_{\text{Carl}}[w_{\text{Carl}}(a)] + \dots$$

This does not invalidate SAT's analysis of social choice under uncertainty. It merely invalidates Harsanyi's “utilitarian” interpretation of the SAT.

“(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really 'utilitarian' —it is an aggregate of 'lottery preferences', not an aggregate of well-being *per se*.”

To appreciate this, suppose w_i is person i 's “true” welfare function. (So $w_i(a)$ = the welfare of person i in outcome a .)

Assuming i is self-interested, her vNM utility function u_i should have the form $u_i(a) = \phi_i[w_i(a)]$, where ϕ_i is some increasing function.

The shape of ϕ_i expresses i 's “risk attitudes”. For example:

- ▶ if i is *risk-averse*, then ϕ_i is concave.
- ▶ If i is *risk-prone* then ϕ_i is convex.

But *there is no reason to suppose that ϕ_i is linear*.

Thus, Harsanyi's SAT actually characterizes an *ex ante* SWF of the form:

$$W(a) = K + c_{\text{Ann}} \phi_{\text{Ann}}[w_{\text{Ann}}(a)] + c_{\text{Bob}} \phi_{\text{Bob}}[w_{\text{Bob}}(a)] + c_{\text{Carl}} \phi_{\text{Carl}}[w_{\text{Carl}}(a)] + \dots$$

This does not invalidate SAT's analysis of social choice under uncertainty.
It merely invalidates Harsanyi's “utilitarian” interpretation of the SAT.

“(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really 'utilitarian' —it is an aggregate of 'lottery preferences', not an aggregate of well-being *per se*.”

To appreciate this, suppose w_i is person i 's “true” welfare function. (So $w_i(a)$ = the welfare of person i in outcome a .)

Assuming i is self-interested, her vNM utility function u_i should have the form $u_i(a) = \phi_i[w_i(a)]$, where ϕ_i is some increasing function.

The shape of ϕ_i expresses i 's “risk attitudes”. For example:

- ▶ if i is *risk-averse*, then ϕ_i is concave.
- ▶ If i is *risk-prone* then ϕ_i is convex.

But *there is no reason to suppose that ϕ_i is linear*.

Thus, Harsanyi's SAT actually characterizes an *ex ante* SWF of the form:

$$W(a) = K + c_{\text{Ann}} \phi_{\text{Ann}}[w_{\text{Ann}}(a)] + c_{\text{Bob}} \phi_{\text{Bob}}[w_{\text{Bob}}(a)] + c_{\text{Carl}} \phi_{\text{Carl}}[w_{\text{Carl}}(a)] + \dots$$

This does not invalidate SAT's analysis of social choice under uncertainty.
It merely invalidates Harsanyi's “utilitarian” interpretation of the SAT.

“(S3) u_i represents person i 's preferences over *lotteries*. But it isn't clear u_i really measures i 's *well-being*. So W is not really 'utilitarian' —it is an aggregate of 'lottery preferences', not an aggregate of well-being *per se*.”

To appreciate this, suppose w_i is person i 's “true” welfare function. (So $w_i(a)$ = the welfare of person i in outcome a .)

Assuming i is self-interested, her vNM utility function u_i should have the form $u_i(a) = \phi_i[w_i(a)]$, where ϕ_i is some increasing function.

The shape of ϕ_i expresses i 's “risk attitudes”. For example:

- ▶ if i is *risk-averse*, then ϕ_i is concave.
- ▶ If i is *risk-prone* then ϕ_i is convex.

But *there is no reason to suppose that ϕ_i is linear*.

Thus, Harsanyi's SAT actually characterizes an *ex ante* SWF of the form:

$$W(a) = K + c_{\text{Ann}} \phi_{\text{Ann}}[w_{\text{Ann}}(a)] + c_{\text{Bob}} \phi_{\text{Bob}}[w_{\text{Bob}}(a)] + c_{\text{Carl}} \phi_{\text{Carl}}[w_{\text{Carl}}(a)] + \dots$$

This does not invalidate SAT's analysis of social choice under uncertainty. It merely invalidates Harsanyi's “utilitarian” interpretation of the SAT.

“(S1) The SAT says \succeq can be interpreted as *if* it came from some weighted utilitarian SWF... But it doesn't tell us how to *obtain* these weights.

“(S2) The vNM utility functions u_i are only defined up to multiplication by a constant. Thus, the weights c_i are *also* only defined up to multiplication by a constant. ”

One solution is to supplement the SAT with a rule to “normalize” the vNM utility functions.

For example, Dhillon and Mertens axiomatically characterized the *relative utilitarian ex ante* social welfare function RU , defined

$$RU(a) := \frac{u_{\text{Ann}}(a)}{R_{\text{Ann}}} + \frac{u_{\text{Bob}}(a)}{R_{\text{Bob}}} + \frac{u_{\text{Carl}}(a)}{R_{\text{Carl}}} + \dots \quad \text{for all } a \text{ in } \mathcal{A}.$$

Here, for all i in \mathcal{I} , $R_i := \max\{u_i(a); a \in \mathcal{A}\} - \min\{u_i(a); a \in \mathcal{A}\}$.

Thus, for all i in \mathcal{I} , the function u_i/R_i ranges over an interval of length 1.

So everyone has equal “weight”, in some sense.

But it's not clear this renormalization is ethically appropriate.

For example, there may be a legitimate reason why Ann has much more intense preferences over \mathcal{A} than Bob does. Perhaps she has a much greater stake in the outcomes. Relative utilitarianism neglects these differences.

“(S1) The SAT says \succeq can be interpreted as *if* it came from some weighted utilitarian SWF... But it doesn't tell us how to *obtain* these weights.

“(S2) The vNM utility functions u_i are only defined up to multiplication by a constant. Thus, the weights c_i are *also* only defined up to multiplication by a constant. ”

One solution is to supplement the SAT with a rule to “normalize” the vNM utility functions.

For example, Dhillon and Mertens axiomatically characterized the *relative utilitarian ex ante* social welfare function RU , defined

$$RU(a) := \frac{u_{\text{Ann}}(a)}{R_{\text{Ann}}} + \frac{u_{\text{Bob}}(a)}{R_{\text{Bob}}} + \frac{u_{\text{Carl}}(a)}{R_{\text{Carl}}} + \dots \quad \text{for all } a \text{ in } \mathcal{A}.$$

Here, for all i in \mathcal{I} , $R_i := \max\{u_i(a); a \in \mathcal{A}\} - \min\{u_i(a); a \in \mathcal{A}\}$.

Thus, for all i in \mathcal{I} , the function u_i/R_i ranges over an interval of length 1.

So everyone has equal “weight”, in some sense.

But it's not clear this renormalization is ethically appropriate.

For example, there may be a legitimate reason why Ann has much more intense preferences over \mathcal{A} than Bob does. Perhaps she has a much greater stake in the outcomes. Relative utilitarianism neglects these differences.

“(S1) The SAT says \succeq can be interpreted *as if* it came from some weighted utilitarian SWF... But it doesn't tell us how to *obtain* these weights.

“(S2) The vNM utility functions u_i are only defined up to multiplication by a constant. Thus, the weights c_i are *also* only defined up to multiplication by a constant. ”

One solution is to supplement the SAT with a rule to “normalize” the vNM utility functions.

For example, Dhillon and Mertens axiomatically characterized the *relative utilitarian* ex ante social welfare function RU , defined

$$RU(a) := \frac{u_{\text{Ann}}(a)}{R_{\text{Ann}}} + \frac{u_{\text{Bob}}(a)}{R_{\text{Bob}}} + \frac{u_{\text{Carl}}(a)}{R_{\text{Carl}}} + \dots \quad \text{for all } a \text{ in } \mathcal{A}.$$

Here, for all i in \mathcal{I} , $R_i := \max\{u_i(a); a \in \mathcal{A}\} - \min\{u_i(a); a \in \mathcal{A}\}$.

Thus, for all i in \mathcal{I} , the function u_i/R_i ranges over an interval of length 1.

So everyone has equal “weight”, in some sense.

But it's not clear this renormalization is ethically appropriate.

For example, there may be a legitimate reason why Ann has much more intense preferences over \mathcal{A} than Bob does. Perhaps she has a much greater stake in the outcomes. Relative utilitarianism neglects these differences.

“(S1) The SAT says \succeq can be interpreted as if it came from some weighted utilitarian SWF... But it doesn't tell us how to *obtain* these weights.

“(S2) The vNM utility functions u_i are only defined up to multiplication by a constant. Thus, the weights c_i are *also* only defined up to multiplication by a constant.”

One solution is to supplement the SAT with a rule to “normalize” the vNM utility functions.

For example, Dhillon and Mertens axiomatically characterized the *relative utilitarian ex ante* social welfare function RU , defined

$$RU(a) := \frac{u_{\text{Ann}}(a)}{R_{\text{Ann}}} + \frac{u_{\text{Bob}}(a)}{R_{\text{Bob}}} + \frac{u_{\text{Carl}}(a)}{R_{\text{Carl}}} + \dots \quad \text{for all } a \text{ in } \mathcal{A}.$$

Here, for all i in \mathcal{I} , $R_i := \max\{u_i(a); a \in \mathcal{A}\} - \min\{u_i(a); a \in \mathcal{A}\}$.

Thus, for all i in \mathcal{I} , the function u_i/R_i ranges over an interval of length 1. So everyone has equal “weight”, in some sense.

But it's not clear this renormalization is ethically appropriate.

For example, there may be a legitimate reason why Ann has much more intense preferences over \mathcal{A} than Bob does. Perhaps she has a much greater stake in the outcomes. Relative utilitarianism neglects these differences.

“(S1) The SAT says \succeq can be interpreted as if it came from some weighted utilitarian SWF... But it doesn't tell us how to *obtain* these weights.

“(S2) The vNM utility functions u_i are only defined up to multiplication by a constant. Thus, the weights c_i are *also* only defined up to multiplication by a constant.”

One solution is to supplement the SAT with a rule to “normalize” the vNM utility functions.

For example, Dhillon and Mertens axiomatically characterized the *relative utilitarian ex ante* social welfare function RU , defined

$$RU(a) := \frac{u_{\text{Ann}}(a)}{R_{\text{Ann}}} + \frac{u_{\text{Bob}}(a)}{R_{\text{Bob}}} + \frac{u_{\text{Carl}}(a)}{R_{\text{Carl}}} + \dots \quad \text{for all } a \text{ in } \mathcal{A}.$$

Here, for all i in \mathcal{I} , $R_i := \max\{u_i(a); a \in \mathcal{A}\} - \min\{u_i(a); a \in \mathcal{A}\}$.

Thus, for all i in \mathcal{I} , the function u_i/R_i ranges over an interval of length 1.

So everyone has equal “weight”, in some sense.

But it's not clear this renormalization is ethically appropriate.

For example, there may be a legitimate reason why Ann has much more intense preferences over \mathcal{A} than Bob does. Perhaps she has a much greater stake in the outcomes.

Relative utilitarianism neglects these differences.

“(S1) The SAT says \succeq can be interpreted as if it came from some weighted utilitarian SWF... But it doesn't tell us how to *obtain* these weights.

“(S2) The vNM utility functions u_i are only defined up to multiplication by a constant. Thus, the weights c_i are *also* only defined up to multiplication by a constant.”

One solution is to supplement the SAT with a rule to “normalize” the vNM utility functions.

For example, Dhillon and Mertens axiomatically characterized the *relative utilitarian ex ante* social welfare function RU , defined

$$RU(a) := \frac{u_{\text{Ann}}(a)}{R_{\text{Ann}}} + \frac{u_{\text{Bob}}(a)}{R_{\text{Bob}}} + \frac{u_{\text{Carl}}(a)}{R_{\text{Carl}}} + \dots \quad \text{for all } a \text{ in } \mathcal{A}.$$

Here, for all i in \mathcal{I} , $R_i := \max\{u_i(a); a \in \mathcal{A}\} - \min\{u_i(a); a \in \mathcal{A}\}$.

Thus, for all i in \mathcal{I} , the function u_i/R_i ranges over an interval of length 1.

So everyone has equal “weight”, in some sense.

But it's not clear this renormalization is ethically appropriate.

For example, there may be a legitimate reason why Ann has much more intense preferences over \mathcal{A} than Bob does. Perhaps she has a much greater stake in the outcomes. Relative utilitarianism neglects these differences.

Appendix B

Bayesian social aggregation with minimal hypotheses

Let \mathcal{S} be a finite set of states.

We now define an *individual prospect* to be a device \mathbf{x} which assigns a real-valued “payoff” x_s (e.g. an income or utility level) to each state $s \in \mathcal{S}$. (Formally, this means that \mathbf{x} is an \mathcal{S} -dimensional vector of real numbers.)

Let \mathcal{I} be a set of individuals. For all $i \in \mathcal{I}$, let \mathcal{X}^i be the set of individual prospects which are feasible for i (a subset of $\mathbb{R}^{\mathcal{S}}$).

Let \succeq^i be individual i 's *ex ante* preference order on \mathcal{X}^i .

Instead of the Savage SEU axioms, we will now only require:

Individual statewise dominance: For all \mathbf{x} and \mathbf{y} in \mathcal{X}^i , if $x_s \geq y_s$ for all s in \mathcal{S} , then $\mathbf{x} \succeq^i \mathbf{y}$. If, also, $x_s > y_s$ for some s in \mathcal{S} , then $\mathbf{x} \succ^i \mathbf{y}$.

We define a *social outcome* to be a device \mathbf{x} which assigns a real-valued “payoff” x^i to each individual i in \mathcal{I} . (Formally, this means that \mathbf{x} is an \mathcal{I} -dimensional vector of real numbers.)

Let \mathcal{X}_{xp} be the set of feasible social outcomes (a subset of $\mathbb{R}^{\mathcal{I}}$).

Let \succeq_{xp} be an *ex post* social preference order on \mathcal{X}_{xp} . We will require:

Ex post Pareto: For all outcomes \mathbf{x} and \mathbf{y} in \mathcal{X}_{xp} , if $x^i \geq y^i$ for all i in \mathcal{I} , then $\mathbf{x} \succeq_{\text{xp}} \mathbf{y}$. If, also, $x^i > y^i$ for some i in \mathcal{I} , then $\mathbf{x} \succ_{\text{xp}} \mathbf{y}$.

Let \mathcal{S} be a finite set of states.

We now define an *individual prospect* to be a device \mathbf{x} which assigns a real-valued “payoff” x_s (e.g. an income or utility level) to each state $s \in \mathcal{S}$. (Formally, this means that \mathbf{x} is an \mathcal{S} -dimensional vector of real numbers.)

Let \mathcal{I} be a set of individuals. For all $i \in \mathcal{I}$, let \mathcal{X}^i be the set of individual prospects which are feasible for i (a subset of $\mathbb{R}^{\mathcal{S}}$).

Let \succeq^i be individual i 's *ex ante* preference order on \mathcal{X}^i .

Instead of the Savage SEU axioms, we will now only require:

Individual statewise dominance: For all \mathbf{x} and \mathbf{y} in \mathcal{X}^i , if $x_s \geq y_s$ for all s in \mathcal{S} , then $\mathbf{x} \succeq^i \mathbf{y}$. If, also, $x_s > y_s$ for some s in \mathcal{S} , then $\mathbf{x} \succ^i \mathbf{y}$.

We define a *social outcome* to be a device \mathbf{x} which assigns a real-valued “payoff” x^i to each individual i in \mathcal{I} .

(Formally, this means that \mathbf{x} is an \mathcal{I} -dimensional vector of real numbers.)

Let \mathcal{X}_{xp} be the set of feasible social outcomes (a subset of $\mathbb{R}^{\mathcal{I}}$).

Let \succeq_{xp} be an *ex post* social preference order on \mathcal{X}_{xp} . We will require:

Ex post Pareto: For all outcomes \mathbf{x} and \mathbf{y} in \mathcal{X}_{xp} , if $x^i \geq y^i$ for all i in \mathcal{I} , then $\mathbf{x} \succeq_{xp} \mathbf{y}$. If, also, $x^i > y^i$ for some i in \mathcal{I} , then $\mathbf{x} \succ_{xp} \mathbf{y}$.

Let \mathcal{S} be a finite set of states.

We now define an *individual prospect* to be a device \mathbf{x} which assigns a real-valued “payoff” x_s (e.g. an income or utility level) to each state $s \in \mathcal{S}$. (Formally, this means that \mathbf{x} is an \mathcal{S} -dimensional vector of real numbers.)

Let \mathcal{I} be a set of individuals. For all $i \in \mathcal{I}$, let \mathcal{X}^i be the set of individual prospects which are feasible for i (a subset of $\mathbb{R}^{\mathcal{S}}$).

Let \succeq^i be individual i 's *ex ante* preference order on \mathcal{X}^i .

Instead of the Savage SEU axioms, we will now only require:

Individual statewise dominance: For all \mathbf{x} and \mathbf{y} in \mathcal{X}^i , if $x_s \geq y_s$ for all s in \mathcal{S} , then $\mathbf{x} \succeq^i \mathbf{y}$. If, also, $x_s > y_s$ for some s in \mathcal{S} , then $\mathbf{x} \succ^i \mathbf{y}$.

We define a *social outcome* to be a device \mathbf{x} which assigns a real-valued “payoff” x^i to each individual i in \mathcal{I} .

(Formally, this means that \mathbf{x} is an \mathcal{I} -dimensional vector of real numbers.)

Let \mathcal{X}_{xp} be the set of feasible social outcomes (a subset of $\mathbb{R}^{\mathcal{I}}$).

Let \succeq_{xp} be an *ex post* social preference order on \mathcal{X}_{xp} . We will require:

Ex post Pareto: For all outcomes \mathbf{x} and \mathbf{y} in \mathcal{X}_{xp} , if $x^i \geq y^i$ for all i in \mathcal{I} , then $\mathbf{x} \succeq_{\text{xp}} \mathbf{y}$. If, also, $x^i > y^i$ for some i in \mathcal{I} , then $\mathbf{x} \succ_{\text{xp}} \mathbf{y}$.

Let \mathcal{S} be a finite set of states.

We now define an *individual prospect* to be a device \mathbf{x} which assigns a real-valued “payoff” x_s (e.g. an income or utility level) to each state $s \in \mathcal{S}$. (Formally, this means that \mathbf{x} is an \mathcal{S} -dimensional vector of real numbers.)

Let \mathcal{I} be a set of individuals. For all $i \in \mathcal{I}$, let \mathcal{X}^i be the set of individual prospects which are feasible for i (a subset of $\mathbb{R}^{\mathcal{S}}$).

Let \succeq^i be individual i 's *ex ante* preference order on \mathcal{X}^i .

Instead of the Savage SEU axioms, we will now only require:

Individual statewise dominance: For all \mathbf{x} and \mathbf{y} in \mathcal{X}^i , if $x_s \geq y_s$ for all s in \mathcal{S} , then $\mathbf{x} \succeq^i \mathbf{y}$. If, also, $x_s > y_s$ for some s in \mathcal{S} , then $\mathbf{x} \succ^i \mathbf{y}$.

We define a *social outcome* to be a device \mathbf{x} which assigns a real-valued “payoff” x^i to each individual i in \mathcal{I} .

(Formally, this means that \mathbf{x} is an \mathcal{I} -dimensional vector of real numbers.)

Let \mathcal{X}_{xp} be the set of feasible social outcomes (a subset of $\mathbb{R}^{\mathcal{I}}$).

Let \succeq_{xp} be an *ex post* social preference order on \mathcal{X}_{xp} . We will require:

Ex post Pareto: For all outcomes \mathbf{x} and \mathbf{y} in \mathcal{X}_{xp} , if $x^i \geq y^i$ for all i in \mathcal{I} , then $\mathbf{x} \succeq_{\text{xp}} \mathbf{y}$. If, also, $x^i > y^i$ for some i in \mathcal{I} , then $\mathbf{x} \succ_{\text{xp}} \mathbf{y}$.

Let \mathcal{S} be a finite set of states.

We now define an *individual prospect* to be a device \mathbf{x} which assigns a real-valued “payoff” x_s (e.g. an income or utility level) to each state $s \in \mathcal{S}$. (Formally, this means that \mathbf{x} is an \mathcal{S} -dimensional vector of real numbers.)

Let \mathcal{I} be a set of individuals. For all $i \in \mathcal{I}$, let \mathcal{X}^i be the set of individual prospects which are feasible for i (a subset of $\mathbb{R}^{\mathcal{S}}$).

Let \succeq^i be individual i 's *ex ante* preference order on \mathcal{X}^i .

Instead of the Savage SEU axioms, we will now only require:

Individual statewise dominance: For all \mathbf{x} and \mathbf{y} in \mathcal{X}^i , if $x_s \geq y_s$ for all s in \mathcal{S} , then $\mathbf{x} \succeq^i \mathbf{y}$. If, also, $x_s > y_s$ for some s in \mathcal{S} , then $\mathbf{x} \succ^i \mathbf{y}$.

We define a *social outcome* to be a device \mathbf{x} which assigns a real-valued “payoff” x^i to each individual i in \mathcal{I} .

(Formally, this means that \mathbf{x} is an \mathcal{I} -dimensional vector of real numbers.)

Let \mathcal{X}_{xp} be the set of feasible social outcomes (a subset of $\mathbb{R}^{\mathcal{I}}$).

Let \succeq_{xp} be an *ex post* social preference order on \mathcal{X}_{xp} . We will require:

Ex post Pareto: For all outcomes \mathbf{x} and \mathbf{y} in \mathcal{X}_{xp} , if $x^i \geq y^i$ for all i in \mathcal{I} , then $\mathbf{x} \succeq_{\text{xp}} \mathbf{y}$. If, also, $x^i > y^i$ for some i in \mathcal{I} , then $\mathbf{x} \succ_{\text{xp}} \mathbf{y}$.

Let \mathcal{S} be a finite set of states.

We now define an *individual prospect* to be a device \mathbf{x} which assigns a real-valued “payoff” x_s (e.g. an income or utility level) to each state $s \in \mathcal{S}$. (Formally, this means that \mathbf{x} is an \mathcal{S} -dimensional vector of real numbers.)

Let \mathcal{I} be a set of individuals. For all $i \in \mathcal{I}$, let \mathcal{X}^i be the set of individual prospects which are feasible for i (a subset of $\mathbb{R}^{\mathcal{S}}$).

Let \succeq^i be individual i 's *ex ante* preference order on \mathcal{X}^i .

Instead of the Savage SEU axioms, we will now only require:

Individual statewise dominance: For all \mathbf{x} and \mathbf{y} in \mathcal{X}^i , if $x_s \geq y_s$ for all s in \mathcal{S} , then $\mathbf{x} \succeq^i \mathbf{y}$. If, also, $x_s > y_s$ for some s in \mathcal{S} , then $\mathbf{x} \succ^i \mathbf{y}$.

We define a *social outcome* to be a device \mathbf{x} which assigns a real-valued “payoff” x^i to each individual i in \mathcal{I} .

(Formally, this means that \mathbf{x} is an \mathcal{I} -dimensional vector of real numbers.)

Let \mathcal{X}_{xp} be the set of feasible social outcomes (a subset of $\mathbb{R}^{\mathcal{I}}$).

Let \succeq_{xp} be an *ex post* social preference order on \mathcal{X}_{xp} . We will require:

Ex post Pareto: For all outcomes \mathbf{x} and \mathbf{y} in \mathcal{X}_{xp} , if $x^i \geq y^i$ for all i in \mathcal{I} , then $\mathbf{x} \succeq_{xp} \mathbf{y}$. If, also, $x^i > y^i$ for some i in \mathcal{I} , then $\mathbf{x} \succ_{xp} \mathbf{y}$.

Let \mathcal{S} be a finite set of states.

We now define an *individual prospect* to be a device \mathbf{x} which assigns a real-valued “payoff” x_s (e.g. an income or utility level) to each state $s \in \mathcal{S}$. (Formally, this means that \mathbf{x} is an \mathcal{S} -dimensional vector of real numbers.)

Let \mathcal{I} be a set of individuals. For all $i \in \mathcal{I}$, let \mathcal{X}^i be the set of individual prospects which are feasible for i (a subset of $\mathbb{R}^{\mathcal{S}}$).

Let \succeq^i be individual i 's *ex ante* preference order on \mathcal{X}^i .

Instead of the Savage SEU axioms, we will now only require:

Individual statewise dominance: For all \mathbf{x} and \mathbf{y} in \mathcal{X}^i , if $x_s \geq y_s$ for all s in \mathcal{S} , then $\mathbf{x} \succeq^i \mathbf{y}$. If, also, $x_s > y_s$ for some s in \mathcal{S} , then $\mathbf{x} \succ^i \mathbf{y}$.

We define a *social outcome* to be a device \mathbf{x} which assigns a real-valued “payoff” x^i to each individual i in \mathcal{I} .

(Formally, this means that \mathbf{x} is an \mathcal{I} -dimensional vector of real numbers.)

Let \mathcal{X}_{xp} be the set of feasible social outcomes (a subset of $\mathbb{R}^{\mathcal{I}}$).

Let \succeq_{xp} be an *ex post* social preference order on \mathcal{X}_{xp} . We will require:

Ex post Pareto: For all outcomes \mathbf{x} and \mathbf{y} in \mathcal{X}_{xp} , if $x^i \geq y^i$ for all i in \mathcal{I} , then $\mathbf{x} \succeq_{xp} \mathbf{y}$. If, also, $x^i > y^i$ for some i in \mathcal{I} , then $\mathbf{x} \succ_{xp} \mathbf{y}$.

Let \mathcal{S} be a finite set of states.

We now define an *individual prospect* to be a device \mathbf{x} which assigns a real-valued “payoff” x_s (e.g. an income or utility level) to each state $s \in \mathcal{S}$. (Formally, this means that \mathbf{x} is an \mathcal{S} -dimensional vector of real numbers.)

Let \mathcal{I} be a set of individuals. For all $i \in \mathcal{I}$, let \mathcal{X}^i be the set of individual prospects which are feasible for i (a subset of $\mathbb{R}^{\mathcal{S}}$).

Let \succeq^i be individual i 's *ex ante* preference order on \mathcal{X}^i .

Instead of the Savage SEU axioms, we will now only require:

Individual statewise dominance: For all \mathbf{x} and \mathbf{y} in \mathcal{X}^i , if $x_s \geq y_s$ for all s in \mathcal{S} , then $\mathbf{x} \succeq^i \mathbf{y}$. If, also, $x_s > y_s$ for some s in \mathcal{S} , then $\mathbf{x} \succ^i \mathbf{y}$.

We define a *social outcome* to be a device \mathbf{x} which assigns a real-valued “payoff” x^i to each individual i in \mathcal{I} .

(Formally, this means that \mathbf{x} is an \mathcal{I} -dimensional vector of real numbers.)

Let \mathcal{X}_{xp} be the set of feasible social outcomes (a subset of $\mathbb{R}^{\mathcal{I}}$).

Let \succeq_{xp} be an *ex post* social preference order on \mathcal{X}_{xp} . We will require:

Ex post Pareto: For all outcomes \mathbf{x} and \mathbf{y} in \mathcal{X}_{xp} , if $x^i \geq y^i$ for all i in \mathcal{I} , then $\mathbf{x} \succeq_{\text{xp}} \mathbf{y}$. If, also, $x^i > y^i$ for some i in \mathcal{I} , then $\mathbf{x} \succ_{\text{xp}} \mathbf{y}$.

Let \mathcal{S} be a finite set of states.

We now define an *individual prospect* to be a device \mathbf{x} which assigns a real-valued “payoff” x_s (e.g. an income or utility level) to each state $s \in \mathcal{S}$. (Formally, this means that \mathbf{x} is an \mathcal{S} -dimensional vector of real numbers.)

Let \mathcal{I} be a set of individuals. For all $i \in \mathcal{I}$, let \mathcal{X}^i be the set of individual prospects which are feasible for i (a subset of $\mathbb{R}^{\mathcal{S}}$).

Let \succeq^i be individual i 's *ex ante* preference order on \mathcal{X}^i .

Instead of the Savage SEU axioms, we will now only require:

Individual statewise dominance: For all \mathbf{x} and \mathbf{y} in \mathcal{X}^i , if $x_s \geq y_s$ for all s in \mathcal{S} , then $\mathbf{x} \succeq^i \mathbf{y}$. If, also, $x_s > y_s$ for some s in \mathcal{S} , then $\mathbf{x} \succ^i \mathbf{y}$.

We define a *social outcome* to be a device \mathbf{x} which assigns a real-valued “payoff” x^i to each individual i in \mathcal{I} .

(Formally, this means that \mathbf{x} is an \mathcal{I} -dimensional vector of real numbers.)

Let \mathcal{X}_{xp} be the set of feasible social outcomes (a subset of $\mathbb{R}^{\mathcal{I}}$).

Let \succeq_{xp} be an *ex post* social preference order on \mathcal{X}_{xp} . We will require:

Ex post Pareto: For all outcomes \mathbf{x} and \mathbf{y} in \mathcal{X}_{xp} , if $x^i \geq y^i$ for all i in \mathcal{I} , then $\mathbf{x} \succeq_{\text{xp}} \mathbf{y}$. If, also, $x^i > y^i$ for some i in \mathcal{I} , then $\mathbf{x} \succ_{\text{xp}} \mathbf{y}$.

Finally, we will define a *social prospect* to be a device \mathbf{X} which assigns a social outcome \mathbf{x}_s to each state $s \in \mathcal{S}$.

Equivalently, \mathbf{X} assigns an individual prospect \mathbf{x}^i to each person i in \mathcal{I} .

Equivalently, \mathbf{X} yields a real-valued “payoff” x_s^i to all i in \mathcal{I} , for all s in \mathcal{S} . (Formally, this means \mathbf{X} is an $\mathcal{I} \times \mathcal{S}$ matrix of real numbers).

Let \mathcal{X} be the set of feasible individual prospects (a subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$).

Let \succeq_{xa} be the *ex ante* social preference order on \mathcal{X} . We will require

For all social prospects \mathbf{X} and \mathbf{Y} in \mathcal{X} :

- ▶ **Ex ante Pareto:** If $\mathbf{x}^i \succeq \mathbf{y}^i$ for all i in \mathcal{I} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$.
If, furthermore, $\mathbf{x}^i \succ \mathbf{y}^i$ for some i in \mathcal{I} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Group Statewise Dominance:** If $\mathbf{x}_s \succeq_{\text{xp}} \mathbf{y}_s$ for all s in \mathcal{S} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$. If, also, $\mathbf{x}_s \succ_{\text{xp}} \mathbf{y}_s$ for some s in \mathcal{S} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Continuity:** The set $\{\mathbf{W} \in \mathcal{X}; \mathbf{W} \succeq_{\text{xa}} \mathbf{X}\}$ is closed.
The set $\{\mathbf{Z} \in \mathcal{X}; \mathbf{Z} \preceq_{\text{xa}} \mathbf{X}\}$ is also closed.

Finally, we will define a *social prospect* to be a device \mathbf{X} which assigns a social outcome \mathbf{x}_s to each state $s \in \mathcal{S}$.

Equivalently, \mathbf{X} assigns an individual prospect \mathbf{x}^i to each person i in \mathcal{I} .

Equivalently, \mathbf{X} yields a real-valued “payoff” x_s^i to all i in \mathcal{I} , for all s in \mathcal{S} . (Formally, this means \mathbf{X} is an $\mathcal{I} \times \mathcal{S}$ matrix of real numbers).

Let \mathcal{X} be the set of feasible individual prospects (a subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$).

Let \succeq_{xa} be the *ex ante* social preference order on \mathcal{X} . We will require

For all social prospects \mathbf{X} and \mathbf{Y} in \mathcal{X} :

- ▶ **Ex ante Pareto:** If $\mathbf{x}^i \succeq \mathbf{y}^i$ for all i in \mathcal{I} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$.
If, furthermore, $\mathbf{x}^i \succ \mathbf{y}^i$ for some i in \mathcal{I} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Group Statewise Dominance:** If $\mathbf{x}_s \succeq_{\text{xp}} \mathbf{y}_s$ for all s in \mathcal{S} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$. If, also, $\mathbf{x}_s \succ_{\text{xp}} \mathbf{y}_s$ for some s in \mathcal{S} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Continuity:** The set $\{\mathbf{W} \in \mathcal{X}; \mathbf{W} \succeq_{\text{xa}} \mathbf{X}\}$ is closed.
The set $\{\mathbf{Z} \in \mathcal{X}; \mathbf{Z} \preceq_{\text{xa}} \mathbf{X}\}$ is also closed.

Finally, we will define a *social prospect* to be a device \mathbf{X} which assigns a social outcome \mathbf{x}_s to each state $s \in \mathcal{S}$.

Equivalently, \mathbf{X} assigns an individual prospect \mathbf{x}^i to each person i in \mathcal{I} .

Equivalently, \mathbf{X} yields a real-valued “payoff” x_s^i to all i in \mathcal{I} , for all s in \mathcal{S} . (Formally, this means \mathbf{X} is an $\mathcal{I} \times \mathcal{S}$ matrix of real numbers).

Let \mathcal{X} be the set of feasible individual prospects (a subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$).

Let \succeq_{xa} be the *ex ante* social preference order on \mathcal{X} . We will require

For all social prospects \mathbf{X} and \mathbf{Y} in \mathcal{X} :

- ▶ **Ex ante Pareto:** If $\mathbf{x}^i \succeq \mathbf{y}^i$ for all i in \mathcal{I} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$.
If, furthermore, $\mathbf{x}^i \succ \mathbf{y}^i$ for some i in \mathcal{I} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Group Statewise Dominance:** If $\mathbf{x}_s \succeq_{\text{xp}} \mathbf{y}_s$ for all s in \mathcal{S} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$. If, also, $\mathbf{x}_s \succ_{\text{xp}} \mathbf{y}_s$ for some s in \mathcal{S} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Continuity:** The set $\{\mathbf{W} \in \mathcal{X}; \mathbf{W} \succeq_{\text{xa}} \mathbf{X}\}$ is closed.
The set $\{\mathbf{Z} \in \mathcal{X}; \mathbf{Z} \preceq_{\text{xa}} \mathbf{X}\}$ is also closed.

Finally, we will define a *social prospect* to be a device \mathbf{X} which assigns a social outcome \mathbf{x}_s to each state $s \in \mathcal{S}$.

Equivalently, \mathbf{X} assigns an individual prospect \mathbf{x}^i to each person i in \mathcal{I} .

Equivalently, \mathbf{X} yields a real-valued “payoff” x_s^i to all i in \mathcal{I} , for all s in \mathcal{S} . (Formally, this means \mathbf{X} is an $\mathcal{I} \times \mathcal{S}$ matrix of real numbers).

Let \mathcal{X} be the set of feasible individual prospects (a subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$).

Let \succeq_{xa} be the *ex ante* social preference order on \mathcal{X} . We will require

For all social prospects \mathbf{X} and \mathbf{Y} in \mathcal{X} :

- ▶ **Ex ante Pareto:** If $\mathbf{x}^i \succeq \mathbf{y}^i$ for all i in \mathcal{I} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$.
If, furthermore, $\mathbf{x}^i \succ \mathbf{y}^i$ for some i in \mathcal{I} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Group Statewise Dominance:** If $\mathbf{x}_s \succeq_{\text{xp}} \mathbf{y}_s$ for all s in \mathcal{S} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$. If, also, $\mathbf{x}_s \succ_{\text{xp}} \mathbf{y}_s$ for some s in \mathcal{S} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Continuity:** The set $\{\mathbf{W} \in \mathcal{X}; \mathbf{W} \succeq_{\text{xa}} \mathbf{X}\}$ is closed.
The set $\{\mathbf{Z} \in \mathcal{X}; \mathbf{Z} \preceq_{\text{xa}} \mathbf{X}\}$ is also closed.

Finally, we will define a *social prospect* to be a device \mathbf{X} which assigns a social outcome \mathbf{x}_s to each state $s \in \mathcal{S}$.

Equivalently, \mathbf{X} assigns an individual prospect \mathbf{x}^i to each person i in \mathcal{I} .

Equivalently, \mathbf{X} yields a real-valued “payoff” x_s^i to all i in \mathcal{I} , for all s in \mathcal{S} . (Formally, this means \mathbf{X} is an $\mathcal{I} \times \mathcal{S}$ matrix of real numbers).

Let \mathcal{X} be the set of feasible individual prospects (a subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$).

Let \succeq_{xa} be the *ex ante* social preference order on \mathcal{X} . We will require

For all social prospects \mathbf{X} and \mathbf{Y} in \mathcal{X} :

- ▶ **Ex ante Pareto:** If $x^i \succeq y^i$ for all i in \mathcal{I} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$.
If, furthermore, $x^i \succ y^i$ for some i in \mathcal{I} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Group Statewise Dominance:** If $\mathbf{x}_s \succeq_{\text{xp}} \mathbf{y}_s$ for all s in \mathcal{S} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$. If, also, $\mathbf{x}_s \succ_{\text{xp}} \mathbf{y}_s$ for some s in \mathcal{S} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Continuity:** The set $\{\mathbf{W} \in \mathcal{X}; \mathbf{W} \succeq_{\text{xa}} \mathbf{X}\}$ is closed.
The set $\{\mathbf{Z} \in \mathcal{X}; \mathbf{Z} \preceq_{\text{xa}} \mathbf{X}\}$ is also closed.

Finally, we will define a *social prospect* to be a device \mathbf{X} which assigns a social outcome \mathbf{x}_s to each state $s \in \mathcal{S}$.

Equivalently, \mathbf{X} assigns an individual prospect \mathbf{x}^i to each person i in \mathcal{I} .

Equivalently, \mathbf{X} yields a real-valued “payoff” x_s^i to all i in \mathcal{I} , for all s in \mathcal{S} . (Formally, this means \mathbf{X} is an $\mathcal{I} \times \mathcal{S}$ matrix of real numbers).

Let \mathcal{X} be the set of feasible individual prospects (a subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$).

Let \succeq_{xa} be the *ex ante* social preference order on \mathcal{X} . We will require

For all social prospects \mathbf{X} and \mathbf{Y} in \mathcal{X} :

- ▶ **Ex ante Pareto:** If $\mathbf{x}^i \succeq \mathbf{y}^i$ for all i in \mathcal{I} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$.
If, furthermore, $\mathbf{x}^i \succ \mathbf{y}^i$ for some i in \mathcal{I} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Group Statewise Dominance:** If $\mathbf{x}_s \succeq_{\text{xp}} \mathbf{y}_s$ for all s in \mathcal{S} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$. If, also, $\mathbf{x}_s \succ_{\text{xp}} \mathbf{y}_s$ for some s in \mathcal{S} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Continuity:** The set $\{\mathbf{W} \in \mathcal{X}; \mathbf{W} \succeq_{\text{xa}} \mathbf{X}\}$ is closed.
The set $\{\mathbf{Z} \in \mathcal{X}; \mathbf{Z} \preceq_{\text{xa}} \mathbf{X}\}$ is also closed.

Finally, we will define a *social prospect* to be a device \mathbf{X} which assigns a social outcome \mathbf{x}_s to each state $s \in \mathcal{S}$.

Equivalently, \mathbf{X} assigns an individual prospect \mathbf{x}^i to each person i in \mathcal{I} .

Equivalently, \mathbf{X} yields a real-valued “payoff” x_s^i to all i in \mathcal{I} , for all s in \mathcal{S} . (Formally, this means \mathbf{X} is an $\mathcal{I} \times \mathcal{S}$ matrix of real numbers).

Let \mathcal{X} be the set of feasible individual prospects (a subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$).

Let \succeq_{xa} be the *ex ante* social preference order on \mathcal{X} . We will require

For all social prospects \mathbf{X} and \mathbf{Y} in \mathcal{X} :

- ▶ **Ex ante Pareto:** If $\mathbf{x}^i \succeq \mathbf{y}^i$ for all i in \mathcal{I} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$.
If, furthermore, $\mathbf{x}^i \succ \mathbf{y}^i$ for some i in \mathcal{I} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Group Statewise Dominance:** If $\mathbf{x}_s \succeq_{\text{xp}} \mathbf{y}_s$ for all s in \mathcal{S} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$. If, also, $\mathbf{x}_s \succ_{\text{xp}} \mathbf{y}_s$ for some s in \mathcal{S} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Continuity:** The set $\{\mathbf{W} \in \mathcal{X}; \mathbf{W} \succeq_{\text{xa}} \mathbf{X}\}$ is closed.
The set $\{\mathbf{Z} \in \mathcal{X}; \mathbf{Z} \preceq_{\text{xa}} \mathbf{X}\}$ is also closed.

Finally, we will define a *social prospect* to be a device \mathbf{X} which assigns a social outcome \mathbf{x}_s to each state $s \in \mathcal{S}$.

Equivalently, \mathbf{X} assigns an individual prospect \mathbf{x}^i to each person i in \mathcal{I} .

Equivalently, \mathbf{X} yields a real-valued “payoff” x_s^i to all i in \mathcal{I} , for all s in \mathcal{S} . (Formally, this means \mathbf{X} is an $\mathcal{I} \times \mathcal{S}$ matrix of real numbers).

Let \mathcal{X} be the set of feasible individual prospects (a subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$).

Let \succeq_{xa} be the *ex ante* social preference order on \mathcal{X} . We will require

For all social prospects \mathbf{X} and \mathbf{Y} in \mathcal{X} :

- ▶ **Ex ante Pareto:** If $\mathbf{x}^i \succeq \mathbf{y}^i$ for all i in \mathcal{I} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$.
If, furthermore, $\mathbf{x}^i \succ \mathbf{y}^i$ for some i in \mathcal{I} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Group Statewise Dominance:** If $\mathbf{x}_s \succeq_{\text{xp}} \mathbf{y}_s$ for all s in \mathcal{S} , then $\mathbf{X} \succeq_{\text{xa}} \mathbf{Y}$. If, also, $\mathbf{x}_s \succ_{\text{xp}} \mathbf{y}_s$ for some s in \mathcal{S} , then $\mathbf{X} \succ_{\text{xa}} \mathbf{Y}$.
- ▶ **Continuity:** The set $\{\mathbf{W} \in \mathcal{X}; \mathbf{W} \succeq_{\text{xa}} \mathbf{X}\}$ is closed.
The set $\{\mathbf{Z} \in \mathcal{X}; \mathbf{Z} \preceq_{\text{xa}} \mathbf{X}\}$ is also closed.

Theorem. (M & P, 2013, extending Blackorby, Donaldson & Mongin, 2004) Let \mathcal{X} be an open, connected subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$ (also with two other connectivity conditions). For all i in \mathcal{I} , let \succeq^i be individual i 's ex ante preference order on \mathcal{X}^i . Let \succeq_{xp} be an ex post social preference order on \mathcal{X}_{xp} . Suppose

- ▶ either \succeq_{xp} satisfies **Ex post Pareto**;
- ▶ or \succeq^i satisfies **Individual Statewise Dominance**, for all i in \mathcal{I} .

Finally, let \succeq_{xa} be an ex ante social preference order on \mathcal{X} .

If \succeq_{xa} satisfies **Continuity**, **Group Statewise Dominance**, and **Ex ante Pareto**, then there is a (unique) probability P on S , and for all i in \mathcal{I} , there are (unique) increasing, continuous utility functions u^i , such that:

- (a) For all i in \mathcal{I} , the order \succeq^i maximizes the P -expected value of u^i .
- (b) \succeq_{xp} is represented by the utilitarian ex post social welfare function W_{xp} defined by $W_{\text{xp}}(\mathbf{x}) := \sum_{i \in \mathcal{I}} u^i(x^i)$, for all $\mathbf{x} \in \mathcal{X}_{\text{xp}}$.
- (c) The ex ante order \succeq_{xa} maximizes the P -expected value of W_{xp} .

Upshot: Even if we weaken the Savage axioms to Statewise Dominance (perhaps the weakest “rationality” axiom imaginable), we *still* get all the conclusions of the Bayesian Social Aggregation Theorem.

The ex post social welfare function W_{xp} is utilitarian, all agents are SEU maximizers and all agents must have the same beliefs.

Theorem. (M & P, 2013, extending Blackorby, Donaldson & Mongin, 2004) *Let \mathcal{X} be an open, connected subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$ (also with two other connectivity conditions). For all i in \mathcal{I} , let \succeq^i be individual i 's ex ante preference order on \mathcal{X}^i . Let \succeq_{xp} be an ex post social preference order on \mathcal{X}_{xp} . Suppose*

- ▶ *either \succeq_{xp} satisfies **Ex post Pareto**;*
- ▶ *or \succeq^i satisfies **Individual Statewise Dominance**, for all i in \mathcal{I} .*

Finally, let \succeq_{xa} be an ex ante social preference order on \mathcal{X} .

*If \succeq_{xa} satisfies **Continuity**, **Group Statewise Dominance**, and **Ex ante Pareto**, then there is a (unique) probability P on S , and for all i in \mathcal{I} , there are (unique) increasing, continuous utility functions u^i , such that:*

- (a) *For all i in \mathcal{I} , the order \succeq^i maximizes the P -expected value of u^i .*
- (b) *\succeq_{xp} is represented by the utilitarian ex post social welfare function W_{xp} defined by $W_{\text{xp}}(\mathbf{x}) := \sum_{i \in \mathcal{I}} u^i(x^i)$, for all $\mathbf{x} \in \mathcal{X}_{\text{xp}}$.*
- (c) *The ex ante order \succeq_{xa} maximizes the P -expected value of W_{xp} .*

Upshot: Even if we weaken the Savage axioms to Statewise Dominance (perhaps the weakest “rationality” axiom imaginable), we *still* get all the conclusions of the Bayesian Social Aggregation Theorem.

The ex post social welfare function W_{xp} is utilitarian, all agents are SEU maximizers and *all agents must have the same beliefs*.

Theorem. (M & P, 2013, extending Blackorby, Donaldson & Mongin, 2004) *Let \mathcal{X} be an open, connected subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$ (also with two other connectivity conditions). For all i in \mathcal{I} , let \succeq^i be individual i 's ex ante preference order on \mathcal{X}^i . Let \succeq_{xp} be an ex post social preference order on \mathcal{X}_{xp} . Suppose*

- ▶ *either \succeq_{xp} satisfies **Ex post Pareto**;*
- ▶ *or \succeq^i satisfies **Individual Statewise Dominance**, for all i in \mathcal{I} .*

Finally, let \succeq_{xa} be an ex ante social preference order on \mathcal{X} .

*If \succeq_{xa} satisfies **Continuity**, **Group Statewise Dominance**, and **Ex ante Pareto**, then there is a (unique) probability P on S , and for all i in \mathcal{I} , there are (unique) increasing, continuous utility functions u^i , such that:*

- (a) *For all i in \mathcal{I} , the order \succeq^i maximizes the P -expected value of u^i .*
- (b) *\succeq_{xp} is represented by the utilitarian ex post social welfare function W_{xp} defined by $W_{xp}(\mathbf{x}) := \sum_{i \in \mathcal{I}} u^i(x^i)$, for all $\mathbf{x} \in \mathcal{X}_{xp}$.*
- (c) *The ex ante order \succeq_{xa} maximizes the P -expected value of W_{xp} .*

Upshot: Even if we weaken the Savage axioms to Statewise Dominance (perhaps the weakest “rationality” axiom imaginable), we *still* get all the conclusions of the Bayesian Social Aggregation Theorem.

The ex post social welfare function W_{xp} is utilitarian, all agents are SEU maximizers and *all agents must have the same beliefs*.

Theorem. (M & P, 2013, extending Blackorby, Donaldson & Mongin, 2004) *Let \mathcal{X} be an open, connected subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$ (also with two other connectivity conditions). For all i in \mathcal{I} , let \succeq^i be individual i 's ex ante preference order on \mathcal{X}^i . Let \succeq_{xp} be an ex post social preference order on \mathcal{X}_{xp} . Suppose*

- ▶ *either \succeq_{xp} satisfies **Ex post Pareto**;*
- ▶ *or \succeq^i satisfies **Individual Statewise Dominance**, for all i in \mathcal{I} .*

Finally, let \succeq_{xa} be an ex ante social preference order on \mathcal{X} .

*If \succeq_{xa} satisfies **Continuity**, **Group Statewise Dominance**, and **Ex ante Pareto**, then there is a (unique) probability P on S , and for all i in \mathcal{I} , there are (unique) increasing, continuous utility functions u^i , such that:*

- (a) For all i in \mathcal{I} , the order \succeq^i maximizes the P -expected value of u^i .*
- (b) \succeq_{xp} is represented by the utilitarian ex post social welfare function W_{xp} defined by $W_{\text{xp}}(\mathbf{x}) := \sum_{i \in \mathcal{I}} u^i(x^i)$, for all $\mathbf{x} \in \mathcal{X}_{\text{xp}}$.*
- (c) The ex ante order \succeq_{xa} maximizes the P -expected value of W_{xp} .*

Upshot: Even if we weaken the Savage axioms to Statewise Dominance (perhaps the weakest “rationality” axiom imaginable), we *still* get all the conclusions of the Bayesian Social Aggregation Theorem.

The ex post social welfare function W_{xp} is utilitarian, all agents are SEU maximizers and *all agents must have the same beliefs*.

Theorem. (M & P, 2013, extending Blackorby, Donaldson & Mongin, 2004) *Let \mathcal{X} be an open, connected subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$ (also with two other connectivity conditions). For all i in \mathcal{I} , let \succeq^i be individual i 's ex ante preference order on \mathcal{X}^i . Let \succeq_{xp} be an ex post social preference order on \mathcal{X}_{xp} . Suppose*

- ▶ *either \succeq_{xp} satisfies **Ex post Pareto**;*
- ▶ *or \succeq^i satisfies **Individual Statewise Dominance**, for all i in \mathcal{I} .*

Finally, let \succeq_{xa} be an ex ante social preference order on \mathcal{X} .

*If \succeq_{xa} satisfies **Continuity**, **Group Statewise Dominance**, and **Ex ante Pareto**, then there is a (unique) probability P on S , and for all i in \mathcal{I} , there are (unique) increasing, continuous utility functions u^i , such that:*

- (a) *For all i in \mathcal{I} , the order \succeq^i maximizes the P -expected value of u^i .*
- (b) *\succeq_{xp} is represented by the utilitarian ex post social welfare function W_{xp} defined by $W_{xp}(\mathbf{x}) := \sum_{i \in \mathcal{I}} u^i(x^i)$, for all $\mathbf{x} \in \mathcal{X}_{xp}$.*
- (c) *The ex ante order \succeq_{xa} maximizes the P -expected value of W_{xp} .*

Upshot: Even if we weaken the Savage axioms to Statewise Dominance (perhaps the weakest “rationality” axiom imaginable), we *still* get all the conclusions of the Bayesian Social Aggregation Theorem.

The ex post social welfare function W_{xp} is utilitarian, all agents are SEU maximizers and *all agents must have the same beliefs*.

Theorem. (M & P, 2013, extending Blackorby, Donaldson & Mongin, 2004) *Let \mathcal{X} be an open, connected subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$ (also with two other connectivity conditions). For all i in \mathcal{I} , let \succeq^i be individual i 's ex ante preference order on \mathcal{X}^i . Let \succeq_{xp} be an ex post social preference order on \mathcal{X}_{xp} . Suppose*

- ▶ *either \succeq_{xp} satisfies **Ex post Pareto**;*
- ▶ *or \succeq^i satisfies **Individual Statewise Dominance**, for all i in \mathcal{I} .*

Finally, let \succeq_{xa} be an ex ante social preference order on \mathcal{X} .

*If \succeq_{xa} satisfies **Continuity**, **Group Statewise Dominance**, and **Ex ante Pareto**, then there is a (unique) probability P on \mathcal{S} , and for all i in \mathcal{I} , there are (unique) increasing, continuous utility functions u^i , such that:*

- (a)** *For all i in \mathcal{I} , the order \succeq^i maximizes the P -expected value of u^i .*
- (b)** *\succeq_{xp} is represented by the utilitarian ex post social welfare function W_{xp} defined by $W_{\text{xp}}(\mathbf{x}) := \sum_{i \in \mathcal{I}} u^i(x^i)$, for all $\mathbf{x} \in \mathcal{X}_{\text{xp}}$.*
- (c)** *The ex ante order \succeq_{xa} maximizes the P -expected value of W_{xp} .*

Upshot: Even if we weaken the Savage axioms to Statewise Dominance (perhaps the weakest “rationality” axiom imaginable), we *still* get all the conclusions of the Bayesian Social Aggregation Theorem.

The ex post social welfare function W_{xp} is utilitarian, all agents are SEU maximizers and *all agents must have the same beliefs*.

Theorem. (M & P, 2013, extending Blackorby, Donaldson & Mongin, 2004) *Let \mathcal{X} be an open, connected subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$ (also with two other connectivity conditions). For all i in \mathcal{I} , let \succeq^i be individual i 's ex ante preference order on \mathcal{X}^i . Let \succeq_{xp} be an ex post social preference order on \mathcal{X}_{xp} . Suppose*

- ▶ *either \succeq_{xp} satisfies **Ex post Pareto**;*
- ▶ *or \succeq^i satisfies **Individual Statewise Dominance**, for all i in \mathcal{I} .*

Finally, let \succeq_{xa} be an ex ante social preference order on \mathcal{X} .

*If \succeq_{xa} satisfies **Continuity**, **Group Statewise Dominance**, and **Ex ante Pareto**, then there is a (unique) probability P on \mathcal{S} , and for all i in \mathcal{I} , there are (unique) increasing, continuous utility functions u^i , such that:*

(a) *For all i in \mathcal{I} , the order \succeq^i maximizes the P -expected value of u^i .*

(b) *\succeq_{xp} is represented by the utilitarian ex post social welfare function W_{xp} defined by $W_{xp}(\mathbf{x}) := \sum_{i \in \mathcal{I}} u^i(x^i)$, for all $\mathbf{x} \in \mathcal{X}_{xp}$.*

(c) *The ex ante order \succeq_{xa} maximizes the P -expected value of W_{xp} .*

Upshot: Even if we weaken the Savage axioms to Statewise Dominance (perhaps the weakest “rationality” axiom imaginable), we *still* get all the conclusions of the Bayesian Social Aggregation Theorem.

The ex post social welfare function W_{xp} is utilitarian, all agents are SEU maximizers and *all agents must have the same beliefs*.

Theorem. (M & P, 2013, extending Blackorby, Donaldson & Mongin, 2004) *Let \mathcal{X} be an open, connected subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$ (also with two other connectivity conditions). For all i in \mathcal{I} , let \succeq^i be individual i 's ex ante preference order on \mathcal{X}^i . Let \succeq_{xp} be an ex post social preference order on \mathcal{X}_{xp} . Suppose*

- ▶ *either \succeq_{xp} satisfies **Ex post Pareto**;*
- ▶ *or \succeq^i satisfies **Individual Statewise Dominance**, for all i in \mathcal{I} .*

Finally, let \succeq_{xa} be an ex ante social preference order on \mathcal{X} .

*If \succeq_{xa} satisfies **Continuity**, **Group Statewise Dominance**, and **Ex ante Pareto**, then there is a (unique) probability P on \mathcal{S} , and for all i in \mathcal{I} , there are (unique) increasing, continuous utility functions u^i , such that:*

- (a)** *For all i in \mathcal{I} , the order \succeq^i maximizes the P -expected value of u^i .*
- (b)** *\succeq_{xp} is represented by the utilitarian ex post social welfare function W_{xp} defined by $W_{\text{xp}}(\mathbf{x}) := \sum_{i \in \mathcal{I}} u^i(x^i)$, for all $\mathbf{x} \in \mathcal{X}_{\text{xp}}$.*
- (c)** *The ex ante order \succeq_{xa} maximizes the P -expected value of W_{xp} .*

Upshot: Even if we weaken the Savage axioms to Statewise Dominance (perhaps the weakest “rationality” axiom imaginable), we *still* get all the conclusions of the Bayesian Social Aggregation Theorem.

The ex post social welfare function W_{xp} is utilitarian, all agents are SEU maximizers and *all agents must have the same beliefs*.

Theorem. (M & P, 2013, extending Blackorby, Donaldson & Mongin, 2004) *Let \mathcal{X} be an open, connected subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$ (also with two other connectivity conditions). For all i in \mathcal{I} , let \succeq^i be individual i 's ex ante preference order on \mathcal{X}^i . Let \succeq_{xp} be an ex post social preference order on \mathcal{X}_{xp} . Suppose*

- ▶ *either \succeq_{xp} satisfies **Ex post Pareto**;*
- ▶ *or \succeq^i satisfies **Individual Statewise Dominance**, for all i in \mathcal{I} .*

Finally, let \succeq_{xa} be an ex ante social preference order on \mathcal{X} .

*If \succeq_{xa} satisfies **Continuity**, **Group Statewise Dominance**, and **Ex ante Pareto**, then there is a (unique) probability P on S , and for all i in \mathcal{I} , there are (unique) increasing, continuous utility functions u^i , such that:*

- (a) For all i in \mathcal{I} , the order \succeq^i maximizes the P -expected value of u^i .*
- (b) \succeq_{xp} is represented by the utilitarian ex post social welfare function W_{xp} defined by $W_{\text{xp}}(\mathbf{x}) := \sum_{i \in \mathcal{I}} u^i(x^i)$, for all $\mathbf{x} \in \mathcal{X}_{\text{xp}}$.*
- (c) The ex ante order \succeq_{xa} maximizes the P -expected value of W_{xp} .*

Upshot: Even if we weaken the Savage axioms to Statewise Dominance (perhaps the weakest “rationality” axiom imaginable), we *still* get all the conclusions of the Bayesian Social Aggregation Theorem.

The ex post social welfare function W_{xp} is utilitarian, all agents are SEU maximizers and all agents must have the same beliefs.

Theorem. (M & P, 2013, extending Blackorby, Donaldson & Mongin, 2004) *Let \mathcal{X} be an open, connected subset of $\mathbb{R}^{\mathcal{I} \times \mathcal{S}}$ (also with two other connectivity conditions). For all i in \mathcal{I} , let \succeq^i be individual i 's ex ante preference order on \mathcal{X}^i . Let \succeq_{xp} be an ex post social preference order on \mathcal{X}_{xp} . Suppose*

- ▶ *either \succeq_{xp} satisfies **Ex post Pareto**;*
- ▶ *or \succeq^i satisfies **Individual Statewise Dominance**, for all i in \mathcal{I} .*

Finally, let \succeq_{xa} be an ex ante social preference order on \mathcal{X} .

*If \succeq_{xa} satisfies **Continuity**, **Group Statewise Dominance**, and **Ex ante Pareto**, then there is a (unique) probability P on \mathcal{S} , and for all i in \mathcal{I} , there are (unique) increasing, continuous utility functions u^i , such that:*

- (a) For all i in \mathcal{I} , the order \succeq^i maximizes the P -expected value of u^i .*
- (b) \succeq_{xp} is represented by the utilitarian ex post social welfare function W_{xp} defined by $W_{\text{xp}}(\mathbf{x}) := \sum_{i \in \mathcal{I}} u^i(x^i)$, for all $\mathbf{x} \in \mathcal{X}_{\text{xp}}$.*
- (c) The ex ante order \succeq_{xa} maximizes the P -expected value of W_{xp} .*

Upshot: Even if we weaken the Savage axioms to Statewise Dominance (perhaps the weakest “rationality” axiom imaginable), we *still* get all the conclusions of the Bayesian Social Aggregation Theorem.

The ex post social welfare function W_{xp} is utilitarian, all agents are SEU maximizers, and *all agents must have the same beliefs*.

Appendix C

Complete specification

in the

Diamond Paradox

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story—or stipulate that these effects are *already accounted for* in the payoff tables.

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story—or stipulate that these effects are *already accounted for* in the payoff tables.

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story—or stipulate that these effects are *already accounted for* in the payoff tables.

To see this, suppose that, in addition to receiving the stated payoffs from the candy at time 1, each person gets an *anticipation* utility of $p\alpha$ at time 0, where p is the probability they will receive the candy at time 1.

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	$1 + \alpha/2$	$\alpha/2$
Bob	$\alpha/2$	$1 + \alpha/2$

G	h	t
Ann	$1 + \alpha$	$1 + \alpha$
Bob	0	0

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story—or stipulate that these effects are *already accounted for* in the payoff tables.

To see this, suppose that, in addition to receiving the stated payoffs from the candy at time 1, each person gets an *anticipation* utility of $p\alpha$ at time 0, where p is the probability they will receive the candy at time 1. Then their **total** payoffs (time 0 + time 1) are as shown in the new table.

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	$1 + \alpha/2$	$\alpha/2$
Bob	$\alpha/2$	$1 + \alpha/2$

G	h	t
Ann	$1 + \alpha$	$1 + \alpha$
Bob	0	0

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story—or stipulate that these effects are *already accounted for* in the payoff tables.

To see this, suppose that, in addition to receiving the stated payoffs from the candy at time 1, each person gets an *anticipation* utility of $p\alpha$ at time 0, where p is the probability they will receive the candy at time 1.

Then their *total* payoffs (time 0 + time 1) are as shown in the new table. Now F yields a more egalitarian outcome than G in both states of nature.

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	$1 + \alpha/2$	$\alpha/2$
Bob	$\alpha/2$	$1 + \alpha/2$

G	h	t
Ann	$1 + \alpha$	$1 + \alpha$
Bob	0	0

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story—or stipulate that these effects are *already accounted for* in the payoff tables.

To see this, suppose that, in addition to receiving the stated payoffs from the candy at time 1, each person gets an *anticipation* utility of $p\alpha$ at time 0, where p is the probability they will receive the candy at time 1.

Then their *total* payoffs (time 0 + time 1) are as shown in the new table.

Now F yields a more egalitarian outcome than G in both states of nature.

So if \succ_{xp} is any egalitarian *ex post* SWO, then $F(h) \succ_{xp} G(h)$ and

$F(t) \succ_{xp} G(t)$. Hence $F(h) \succ_{xa} G(h)$ by Statewise Dominance.

F	h	t
Ann	$1 + \alpha/2$	$\alpha/2$
Bob	$\alpha/2$	$1 + \alpha/2$

G	h	t
Ann	$1 + \alpha$	$1 + \alpha$
Bob	0	0

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story—or stipulate that these effects are *already accounted for* in the payoff tables.

To see this, suppose that, in addition to receiving the stated payoffs from the candy at time 1, each person gets an *anticipation* utility of $p\alpha$ at time 0, where p is the probability they will receive the candy at time 1.

Then their *total* payoffs (time 0 + time 1) are as shown in the new table.

Now F yields a more egalitarian outcome than G in both states of nature.

So if \succ_{xp} is any egalitarian *ex post* SWO, then $F(h) \succ_{xp} G(h)$ and

$F(t) \succ_{xp} G(t)$. Hence $F(h) \succ_{xa} G(h)$ by Statewise Dominance.

So the Diamond Paradox only works if we exclude “anticipation” effects

(Deschamps and Gevers, 1979)

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	0	0

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story —or stipulate that these effects are *already accounted for* in the payoff tables.

On the other hand, suppose that Bob felt some “bitterness” $-\beta$ due to his unfair treatment in prospect G .

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	$-\beta$	$-\beta$

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story —or stipulate that these effects are *already accounted for* in the payoff tables.

On the other hand, suppose that Bob felt some “bitterness” $-\beta$ due to his unfair treatment in prospect G .

Then his **actual** payoffs would be as shown in the revised table.

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	$-\beta$	$-\beta$

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story—or stipulate that these effects are *already accounted for* in the payoff tables.

On the other hand, suppose that Bob felt some “bitterness” $-\beta$ due to his unfair treatment in prospect G .

Then his *actual* payoffs would be as shown in the revised table.

So if \succ_{xp} is any impartial, Pareto *ex post* SWO, then $F(h) \succ_{\text{xp}} G(h)$ and $F(t) \succ_{\text{xp}} G(t)$. Hence $F(h) \succ_{\text{xa}} G(h)$ by Statewise Dominance.

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	1	0
Bob	0	1

G	h	t
Ann	1	1
Bob	$-\beta$	$-\beta$

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story—or stipulate that these effects are *already accounted for* in the payoff tables.

On the other hand, suppose that Bob felt some “bitterness” $-\beta$ due to his unfair treatment in prospect G .

Then his *actual* payoffs would be as shown in the revised table.

So if \succ_{xp} is any impartial, Pareto *ex post* SWO, then $F(h) \succ_{\text{xp}} G(h)$ and $F(t) \succ_{\text{xp}} G(t)$. Hence $F(h) \succ_{\text{xa}} G(h)$ by Statewise Dominance.

So the Diamond Paradox only works if we exclude Bob’s bitterness.

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	1	0
Bob	0	1
Carol		

G	h	t
Ann	1	1
Bob	0	0
Carol	$-\gamma$	$-\gamma$

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story—or stipulate that these effects are *already accounted for* in the payoff tables.

Or, suppose we include the decision-maker, Carol, as a third agent.

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	1	0
Bob	0	1
Carol		

G	h	t
Ann	1	1
Bob	0	0
Carol	$-\gamma$	$-\gamma$

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story—or stipulate that these effects are *already accounted for* in the payoff tables.

Or, suppose we include the decision-maker, Carol, as a third agent.

Carol may feel some *guilt* $-\gamma$ if she treats Bob unfairly.

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	1	0
Bob	0	1
Carol	0	0

G	h	t
Ann	1	1
Bob	0	0
Carol	$-\gamma$	$-\gamma$

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story—or stipulate that these effects are *already accounted for* in the payoff tables.

Or, suppose we include the decision-maker, Carol, as a third agent.

Carol may feel some *guilt* $-\gamma$ if she treats Bob unfairly.

Then her payoffs would be as shown in the revised table.

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	1	0
Bob	0	1
Carol	0	0

G	h	t
Ann	1	1
Bob	0	0
Carol	$-\gamma$	$-\gamma$

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story—or stipulate that these effects are *already accounted for* in the payoff tables.

Or, suppose we include the decision-maker, Carol, as a third agent.

Carol may feel some *guilt* $-\gamma$ if she treats Bob unfairly.

Then her payoffs would be as shown in the revised table.

So if \succ_{xp} is any impartial, Pareto *ex post* SWO, then $F(h) \succ_{xp} G(h)$ and $F(t) \succ_{xp} G(t)$. Hence $F(h) \succ_{xa} G(h)$ by Statewise Dominance.

Complete specification in the Diamond Paradox

(43/47)

F	h	t
Ann	1	0
Bob	0	1
Carol	0	0

G	h	t
Ann	1	1
Bob	0	0
Carol	$-\gamma$	$-\gamma$

It is important to stipulate that these payoff tables provide a **complete description** of the problem.

We must exclude any *anticipation*, *bitterness*, or *guilt* from the story—or stipulate that these effects are *already accounted for* in the payoff tables.

Or, suppose we include the decision-maker, Carol, as a third agent.

Carol may feel some *guilt* $-\gamma$ if she treats Bob unfairly.

Then her payoffs would be as shown in the revised table.

So if \succ_{xp} is any impartial, Pareto *ex post* SWO, then $F(h) \succ_{\text{xp}} G(h)$ and $F(t) \succ_{\text{xp}} G(t)$. Hence $F(h) \succ_{\text{xa}} G(h)$ by Statewise Dominance.

So the Diamond Paradox only works if we exclude Carol's guilt from the model (Sen,1985).

Appendix D

Hybrids of *ex ante* egalitarianism
and
ex post egalitarianism

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

<i>F</i>	<i>h</i>	<i>t</i>
Ann	1	0
Bob	0	1

MoM: 0.4

<i>G</i>	<i>h</i>	<i>t</i>
Ann	1	1
Bob	0	0

MoM: 0.3

<i>E</i>	<i>h</i>	<i>t</i>
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

The \mathbf{A} -weighted mean of F is $0.3 \cdot 1 + 0.4 \cdot 0 + 0.2 \cdot 0 + 0.1 \cdot 1 = 0.4$.

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

The **A-weighted mean** of F is $0.3 \cdot 1 + 0.4 \cdot 0 + 0.2 \cdot 0 + 0.1 \cdot 1 = 0.4$.

The **B-weighted mean** of F is $0.4 \cdot 1 + 0.3 \cdot 0 + 0.1 \cdot 0 + 0.2 \cdot 1 = 0.6$.

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

The **A**-weighted mean of F is $0.3 \cdot 1 + 0.4 \cdot 0 + 0.2 \cdot 0 + 0.1 \cdot 1 = 0.4$.

The **B**-weighted mean of F is $0.4 \cdot 1 + 0.3 \cdot 0 + 0.1 \cdot 0 + 0.2 \cdot 1 = 0.6$.

The **C-weighted mean** of F is $0.1 \cdot 1 + 0.2 \cdot 0 + 0.4 \cdot 0 + 0.3 \cdot 1 = 0.4$.

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

The \mathbf{A} -weighted mean of F is $0.3 \cdot 1 + 0.4 \cdot 0 + 0.2 \cdot 0 + 0.1 \cdot 1 = 0.4$.

The \mathbf{B} -weighted mean of F is $0.4 \cdot 1 + 0.3 \cdot 0 + 0.1 \cdot 0 + 0.2 \cdot 1 = 0.6$.

The \mathbf{C} -weighted mean of F is $0.1 \cdot 1 + 0.2 \cdot 0 + 0.4 \cdot 0 + 0.3 \cdot 1 = 0.4$.

The \mathbf{D} -weighted mean of F is $0.2 \cdot 1 + 0.1 \cdot 0 + 0.3 \cdot 0 + 0.4 \cdot 1 = 0.6$.

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

The \mathbf{A} -weighted mean of F is $0.3 \cdot 1 + 0.4 \cdot 0 + 0.2 \cdot 0 + 0.1 \cdot 1 = 0.4$.
 The \mathbf{B} -weighted mean of F is $0.4 \cdot 1 + 0.3 \cdot 0 + 0.1 \cdot 0 + 0.2 \cdot 1 = 0.6$.
 The \mathbf{C} -weighted mean of F is $0.1 \cdot 1 + 0.2 \cdot 0 + 0.4 \cdot 0 + 0.3 \cdot 1 = 0.4$.
 The \mathbf{D} -weighted mean of F is $0.2 \cdot 1 + 0.1 \cdot 0 + 0.3 \cdot 0 + 0.4 \cdot 1 = 0.6$.
 Thus, the **min-of-means** for F is $\min\{0.4, 0.6\} = 0.4$.

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

The **A-weighted mean** of G is $0.3 \cdot 1 + 0.4 \cdot 1 + 0.2 \cdot 0 + 0.1 \cdot 0 = 0.7$.

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

The \mathbf{A} -weighted mean of G is $0.3 \cdot 1 + 0.4 \cdot 1 + 0.2 \cdot 0 + 0.1 \cdot 0 = 0.7$.

The \mathbf{B} -weighted mean of G is $0.4 \cdot 1 + 0.3 \cdot 1 + 0.1 \cdot 0 + 0.2 \cdot 0 = 0.7$.

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

The **A**-weighted mean of G is $0.3 \cdot 1 + 0.4 \cdot 1 + 0.2 \cdot 0 + 0.1 \cdot 0 = 0.7$.

The **B-weighted mean** of G is $0.4 \cdot 1 + 0.3 \cdot 1 + 0.1 \cdot 0 + 0.2 \cdot 0 = 0.7$.

The **C**-weighted mean of G is $0.1 \cdot 1 + 0.2 \cdot 1 + 0.4 \cdot 0 + 0.3 \cdot 0 = 0.3$.

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

The **A**-weighted mean of G is $0.3 \cdot 1 + 0.4 \cdot 1 + 0.2 \cdot 0 + 0.1 \cdot 0 = 0.7$.

The **B**-weighted mean of G is $0.4 \cdot 1 + 0.3 \cdot 1 + 0.1 \cdot 0 + 0.2 \cdot 0 = 0.7$.

The **C-weighted mean** of G is $0.1 \cdot 1 + 0.2 \cdot 1 + 0.4 \cdot 0 + 0.3 \cdot 0 = 0.3$.

The **D**-weighted mean of G is $0.2 \cdot 1 + 0.1 \cdot 1 + 0.3 \cdot 0 + 0.4 \cdot 0 = 0.3$.

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

The **A**-weighted mean of G is $0.3 \cdot 1 + 0.4 \cdot 1 + 0.2 \cdot 0 + 0.1 \cdot 0 = 0.7$.

The **B**-weighted mean of G is $0.4 \cdot 1 + 0.3 \cdot 1 + 0.1 \cdot 0 + 0.2 \cdot 0 = 0.7$.

The **C**-weighted mean of G is $0.1 \cdot 1 + 0.2 \cdot 1 + 0.4 \cdot 0 + 0.3 \cdot 0 = 0.3$.

The **D-weighted mean** of G is $0.2 \cdot 1 + 0.1 \cdot 1 + 0.3 \cdot 0 + 0.4 \cdot 0 = 0.3$.

Thus, the min-of-means for G is $\min\{0.7, 0.3\} = 0.3$.

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

Thus, $\text{MoM}(F) > \text{MoM}(G)$, so F is ranked above G , in accord with our *ex ante* egalitarian intuitions.

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$

MoM: 0.5

Thus, $\text{MoM}(F) > \text{MoM}(G)$, so F is ranked above G , in accord with our *ex ante* egalitarian intuitions.

Finally, the **min-of-means** for E is $\frac{1}{2}(0.1 + 0.2 + 0.3 + 0.4) = 0.5$.

Ben-Porath, Gilboa & Schmeidler (1997) axiomatically characterized the so-called *min of means* method of evaluating social prospects.

To illustrate this, consider the following four matrices of *weightings*:

$$\mathbf{A} = \begin{bmatrix} 0.3 & 0.4 \\ 0.2 & 0.1 \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0.4 & 0.3 \\ 0.1 & 0.2 \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0.1 & 0.2 \\ 0.4 & 0.3 \end{bmatrix}, \mathbf{D} = \begin{bmatrix} 0.2 & 0.1 \\ 0.3 & 0.4 \end{bmatrix}.$$

Consider the following three social prospects:

F	h	t
Ann	1	0
Bob	0	1

MoM: 0.4

G	h	t
Ann	1	1
Bob	0	0

MoM: 0.3

E	h	t
Ann	$1/2$	$1/2$
Bob	$1/2$	$1/2$


MoM: 0.5

Thus, $\text{MoM}(F) > \text{MoM}(G)$, so F is ranked above G , in accord with our *ex ante* egalitarian intuitions.

Finally, the min-of-means for E is $\frac{1}{2}(0.1 + 0.2 + 0.3 + 0.4) = 0.5$.

Thus, $\text{MoM}(E) > \text{MoM}(F)$, so E is ranked above F , in accord with our *ex post* egalitarian intuitions.

Gajdos & Maurin: weighted cross-iterative evaluations (46/47)



Let F be a social prospect (indexed by states and people).

Let ω be an *ex post* social welfare function, which acts on *ex post* social outcomes expressed as utility vectors.

Apply ω to each outcome of F , to get a vector $\omega(F)$ *ex post* social welfare levels, indexed by states.


Let ϵ be a sort of “generalized expectation” operator, which acts on *ex ante* prospects, expressed as state-contingent payoff vectors.

Apply ϵ to each individual prospect contained in F , to get a vector $\epsilon(F)$ of “expected utilities”, indexed by individuals.

We then define $\omega * \epsilon(F) := \omega[\epsilon(F)]$ (a sort of generalized XAE evaluation). Likewise, we define $\epsilon * \omega(F) := \epsilon[\omega(F)]$ (a generalized XPE evaluation).

Finally, define $\Phi(F) := \lambda \epsilon * \omega(F) + (1 - \lambda) \omega * \epsilon(F)$, for some $0 \leq \lambda \leq 1$.

Gajdos & Maurin: weighted cross-iterative evaluations (46/47)



Let F be a social prospect (indexed by states and people).

Let ω be an *ex post* social welfare function, which acts on *ex post* social outcomes expressed as utility vectors.

Apply ω to each outcome of F , to get a vector $\omega(F)$ *ex post* social welfare levels, indexed by states.

Let ϵ be a sort of “generalized expectation” operator, which acts on *ex ante* prospects, expressed as state-contingent payoff vectors.


Apply ϵ to each individual prospect contained in F , to get a vector $\epsilon(F)$ of “expected utilities”, indexed by individuals.

We then define $\omega * \epsilon(F) := \omega[\epsilon(F)]$ (a sort of generalized XAE evaluation).

Likewise, we define $\epsilon * \omega(F) := \epsilon[\omega(F)]$ (a generalized XPE evaluation).

Finally, define $\Phi(F) := \lambda \epsilon * \omega(F) + (1 - \lambda) \omega * \epsilon(F)$, for some $0 \leq \lambda \leq 1$.

Gajdos & Maurin: weighted cross-iterative evaluations (46/47)



Let F be a social prospect (indexed by states and people).

Let ω be an *ex post* social welfare function, which acts on *ex post* social outcomes expressed as utility vectors.

Apply ω to each outcome of F , to get a vector $\omega(F)$ *ex post* social welfare levels, indexed by states.

Let ϵ be a sort of “generalized expectation” operator, which acts on *ex ante* prospects, expressed as state-contingent payoff vectors.


Apply ϵ to each individual prospect contained in F , to get a vector $\epsilon(F)$ of “expected utilities”, indexed by individuals.

We then define $\omega * \epsilon(F) := \omega[\epsilon(F)]$ (a sort of generalized XAE evaluation).

Likewise, we define $\epsilon * \omega(F) := \epsilon[\omega(F)]$ (a generalized XPE evaluation).

Finally, define $\Phi(F) := \lambda \epsilon * \omega(F) + (1 - \lambda) \omega * \epsilon(F)$, for some $0 \leq \lambda \leq 1$.

Gajdos & Maurin: weighted cross-iterative evaluations (46/47)



Let F be a social prospect (indexed by states and people).

Let ω be an *ex post* social welfare function, which acts on *ex post* social outcomes expressed as utility vectors.

Apply ω to each outcome of F , to get a vector $\omega(F)$ *ex post* social welfare levels, indexed by states.

Let ϵ be a sort of “generalized expectation” operator, which acts on *ex ante* prospects, expressed as state-contingent payoff vectors.


Apply ϵ to each individual prospect contained in F , to get a vector $\epsilon(F)$ of “expected utilities”, indexed by individuals.

We then define $\omega * \epsilon(F) := \omega[\epsilon(F)]$ (a sort of generalized XAE evaluation).

Likewise, we define $\epsilon * \omega(F) := \epsilon[\omega(F)]$ (a generalized XPE evaluation).

Finally, define $\Phi(F) := \lambda \epsilon * \omega(F) + (1 - \lambda) \omega * \epsilon(F)$, for some $0 \leq \lambda \leq 1$.

Gajdos & Maurin: weighted cross-iterative evaluations (46/47)



Let F be a social prospect (indexed by states and people).

Let ω be an *ex post* social welfare function, which acts on *ex post* social outcomes expressed as utility vectors.

Apply ω to each outcome of F , to get a vector $\omega(F)$ *ex post* social welfare levels, indexed by states.

Let ϵ be a sort of “generalized expectation” operator, which acts on *ex ante* prospects, expressed as state-contingent payoff vectors.


Apply ϵ to each individual prospect contained in F , to get a vector $\epsilon(F)$ of “expected utilities”, indexed by individuals.

We then define $\omega * \epsilon(F) := \omega[\epsilon(F)]$ (a sort of generalized XAE evaluation).

Likewise, we define $\epsilon * \omega(F) := \epsilon[\omega(F)]$ (a generalized XPE evaluation).

Finally, define $\Phi(F) := \lambda \epsilon * \omega(F) + (1 - \lambda) \omega * \epsilon(F)$, for some $0 \leq \lambda \leq 1$.

Gajdos & Maurin: weighted cross-iterative evaluations (46/47)



Let F be a social prospect (indexed by states and people).

Let ω be an *ex post* social welfare function, which acts on *ex post* social outcomes expressed as utility vectors.

Apply ω to each outcome of F , to get a vector $\omega(F)$ *ex post* social welfare levels, indexed by states.

Let ϵ be a sort of “generalized expectation” operator, which acts on *ex ante* prospects, expressed as state-contingent payoff vectors.


Apply ϵ to each individual prospect contained in F , to get a vector $\epsilon(F)$ of “expected utilities”, indexed by individuals.

We then define $\omega * \epsilon(F) := \omega[\epsilon(F)]$ (a sort of generalized XAE evaluation).

Likewise, we define $\epsilon * \omega(F) := \epsilon[\omega(F)]$ (a generalized XPE evaluation).

Finally, define $\Phi(F) := \lambda \epsilon * \omega(F) + (1 - \lambda) \omega * \epsilon(F)$, for some $0 \leq \lambda \leq 1$.

Gajdos & Maurin: weighted cross-iterative evaluations (46/47)



Let F be a social prospect (indexed by states and people).

Let ω be an *ex post* social welfare function, which acts on *ex post* social outcomes expressed as utility vectors.

Apply ω to each outcome of F , to get a vector $\omega(F)$ *ex post* social welfare levels, indexed by states.

Let ϵ be a sort of “generalized expectation” operator, which acts on *ex ante* prospects, expressed as state-contingent payoff vectors.


Apply ϵ to each individual prospect contained in F , to get a vector $\epsilon(F)$ of “expected utilities”, indexed by individuals.

We then define $\omega * \epsilon(F) := \omega[\epsilon(F)]$ (a sort of generalized XAE evaluation).

Likewise, we define $\epsilon * \omega(F) := \epsilon(\omega[F])$ (a generalized XPE evaluation).

Finally, define $\phi(F) := \lambda \epsilon * \omega(F) + (1 - \lambda) \omega * \epsilon(F)$, for some $0 \leq \lambda \leq 1$.

Gajdos & Maurin: weighted cross-iterative evaluations (47/47)



BG&S ('97) showed that, if ω and ϵ are M.o.M's, then ϕ is also an M.o.M.

Gajdos and Maurin (2004) investigated other *ex ante* social evaluations with a similar structure (where ϵ and ω are *not* necessarily M.o.M's).

They allowed λ to depend on F (subject to certain restrictions), and referred to ϕ as a *weighted cross-iterative* (WCI) evaluation.

In effect, ϕ is a sort of “compromise” between XAE and XPE.

G&M axiomatically characterized various other families of WCI evaluations.

For example, certain axioms single out social evaluations of the form

$$\phi(F) := \lambda \cdot \underline{M}(F) + (1 - \lambda) \overline{M}(F), \quad \text{where}$$

$$\underline{M}(F) := \min \{ \epsilon * \omega(F), \omega * \epsilon(F) \} \text{ and } \overline{M}(F) := \max \{ \epsilon * \omega(F), \omega * \epsilon(F) \}$$

Gajdos & Maurin: weighted cross-iterative evaluations (47/47)

$$\begin{array}{lcl}
 \text{Social prospect: } F & = & \begin{array}{c} \text{A} \\ \hline \begin{array}{|c|c|c|c|} \hline & & & \\ \hline & & & \\ \hline \end{array} \\ \hline \text{B} \end{array} \xrightarrow{\epsilon} \epsilon(F) = \begin{array}{|c|} \hline \\ \hline \end{array} \xrightarrow{\omega} \omega * \epsilon(F) \\
 \\
 \text{ } & & \downarrow \\
 \text{ } & & \downarrow \\
 \text{ } & & \downarrow \\
 \text{ } & & \epsilon * \omega(F) \longrightarrow \phi(F) = \lambda \epsilon * \omega(F) + (1 - \lambda) \omega * \epsilon(F)
 \end{array}$$

$\omega(F) = \begin{array}{|c|c|c|c|} \hline & & & \\ \hline \end{array}$

BG&S ('97) showed that, if ω and ϵ are M.o.M's, then ϕ is also an M.o.M. Gajdos and Maurin (2004) investigated other *ex ante* social evaluations with a similar structure (where ϵ and ω are *not* necessarily M.o.M's).

They allowed λ to depend on F (subject to certain restrictions), and referred to ϕ as a *weighted cross-iterative* (WCI) evaluation.

In effect, ϕ is a sort of “compromise” between XAE and XPE.


G&M axiomatically characterized various other families of WCI evaluations.

For example, certain axioms single out social evaluations of the form

$$\phi(F) := \lambda \cdot \underline{M}(F) + (1 - \lambda) \overline{M}(F), \quad \text{where}$$

$$\underline{M}(F) := \min \{ \epsilon * \omega(F), \omega * \epsilon(F) \} \text{ and } \overline{M}(F) := \max \{ \epsilon * \omega(F), \omega * \epsilon(F) \}$$

Gajdos & Maurin: weighted cross-iterative evaluations (47/47)



BG&S ('97) showed that, if ω and ϵ are M.o.M's, then ϕ is also an M.o.M. Gajdos and Maurin (2004) investigated other *ex ante* social evaluations with a similar structure (where ϵ and ω are *not* necessarily M.o.M's).

They allowed λ to depend on F (subject to certain restrictions), and referred to ϕ as a **weighted cross-iterative** (WCI) evaluation.

In effect, ϕ is a sort of "compromise" between XAE and XPE.


G&M axiomatically characterized various other families of WCI evaluations.

For example, certain axioms single out social evaluations of the form

$$\phi(F) := \lambda \cdot \underline{M}(F) + (1 - \lambda) \overline{M}(F), \quad \text{where}$$

$$\underline{M}(F) := \min \{ \epsilon * \omega(F), \omega * \epsilon(F) \} \text{ and } \overline{M}(F) := \max \{ \epsilon * \omega(F), \omega * \epsilon(F) \}$$

Gajdos & Maurin: weighted cross-iterative evaluations (47/47)



BG&S ('97) showed that, if ω and ϵ are M.o.M's, then ϕ is also an M.o.M. Gajdos and Maurin (2004) investigated other *ex ante* social evaluations with a similar structure (where ϵ and ω are *not* necessarily M.o.M's).

They allowed λ to depend on F (subject to certain restrictions), and referred to ϕ as a *weighted cross-iterative* (WCI) evaluation.

In effect, ϕ is a sort of “compromise” between XAE and XPE.


G&M axiomatically characterized various other families of WCI evaluations.

For example, certain axioms single out social evaluations of the form

$$\phi(F) := \lambda \cdot \underline{M}(F) + (1 - \lambda) \overline{M}(F), \quad \text{where}$$

$$\underline{M}(F) := \min \{ \epsilon * \omega(F), \omega * \epsilon(F) \} \text{ and } \overline{M}(F) := \max \{ \epsilon * \omega(F), \omega * \epsilon(F) \}$$

Gajdos & Maurin: weighted cross-iterative evaluations (47/47)



BG&S ('97) showed that, if ω and ϵ are M.o.M's, then ϕ is also an M.o.M. Gajdos and Maurin (2004) investigated other *ex ante* social evaluations with a similar structure (where ϵ and ω are *not* necessarily M.o.M's).

They allowed λ to depend on F (subject to certain restrictions), and referred to ϕ as a *weighted cross-iterative* (WCI) evaluation.

In effect, ϕ is a sort of “compromise” between XAE and XPE.


G&M axiomatically characterized various other families of WCI evaluations.

For example, certain axioms single out social evaluations of the form

$$\phi(F) := \lambda \cdot \underline{M}(F) + (1 - \lambda) \overline{M}(F), \quad \text{where}$$

$$\underline{M}(F) := \min \{ \epsilon * \omega(F), \omega * \epsilon(F) \} \text{ and } \overline{M}(F) := \max \{ \epsilon * \omega(F), \omega * \epsilon(F) \}$$

Gajdos & Maurin: weighted cross-iterative evaluations (47/47)



BG&S ('97) showed that, if ω and ϵ are M.o.M's, then ϕ is also an M.o.M. Gajdos and Maurin (2004) investigated other *ex ante* social evaluations with a similar structure (where ϵ and ω are *not* necessarily M.o.M's).

They allowed λ to depend on F (subject to certain restrictions), and referred to ϕ as a *weighted cross-iterative* (WCI) evaluation.

In effect, ϕ is a sort of “compromise” between XAE and XPE.

G&M axiomatically characterized various other families of WCI evaluations.

For example, certain axioms single out social evaluations of the form

$$\phi(F) := \lambda \cdot \underline{M}(F) + (1 - \lambda) \overline{M}(F), \quad \text{where}$$

$$\underline{M}(F) := \min \{ \epsilon * \omega(F), \omega * \epsilon(F) \} \text{ and } \overline{M}(F) := \max \{ \epsilon * \omega(F), \omega * \epsilon(F) \},$$