Persuasion in an asymmetric information economy: a justification of Wald's maxmin preferences

Zhiwei Liu[1] Nicholas C. Yannelis[2]

[1]Capital University of Economics and Business

[2]The University of Iowa

Microeconomics Workshop, IES, Keio University

Motivation

In an asymmetric information exchange economy, de Castro and Yannelis (2018 JET) showed that

- when agents have the Wald's maxmin preferences, the conflict between efficiency and incentive compatibility no longer exists, i.e., all efficient allocations are also incentive compatible;
- furthermore, the Wald's maxmin are the only preferences under which all efficient allocations are incentive compatible.

Moreover, we know from the recent literature that if agents have the Wald's maxmin preferences, then core allocations, value allocations, Walrasian expectations equilibrium allocations and rational expectations equilibrium allocations are both incentive compatible and implementable. There results are false under the Bayesian preferences.

Motivation

However, the literature assumes that from the primitive, the agents are Wald's maxmin. Agents take the expected utility form as given.

Obviously, different expected utility functional forms provide different outcomes as the equilibrium allocations are computed adopting different expected utility operators.

Motivation: Preferences

Let Ω be a finite set of states of nature and $\omega \in \Omega$ a state of nature.

Each agent i has a partition $\mathcal{F}_i$ of Ω.

An element of the partition $\mathcal{F}_i$ is called an event, denoted by E_i.

In the interim, each agent observes an event in $\mathcal{F}_i$ that contains the realized state of nature.

Each agent i has a probability measure $\pi_i : \sigma(\mathcal{F}_i) \to [0, 1]$, where $\sigma(\mathcal{F}_i)$ is the algebra generated by agent i's partition.

Let Δ_i be the set of all probability measures over 2^Ω that agree with π_i. Formally,

$$\Delta_i = \left\{ \text{probability measure } \mu_i : 2^\Omega \to [0, 1] \mid \mu_i(A) = \pi_i(A), \forall A \in \sigma(\mathcal{F}_i) \right\}.$$

Motivation: Preferences

An allocation $x = (x_i)_{i \in I}$ is a mapping from Ω to $\mathbb{R}_+^{\ell \times N}$, where x_i is *agent* i*'s allocation*. Let L denote the set of allocations.

We postulate that each agent i's preferences on L are *maxmin* $\grave{a}$ la Gilboa and Schmeidler (1989 JME). Let P_i be agent i's multi-belief set which is a non-empty, closed and convex subset of Δ_i.

Given an allocation x, agent i's *ex ante expected utility* is

$$\min_{\mu_i \in P_i} \sum_{\omega \in \Omega} u_i \left(x_i \left(\omega \right), \omega \right) \mu_i \left(\omega \right). \tag{1}$$

Motivation: Preferences

This general multi-belief model includes both the Bayesian and the Wald's maxmin preferences in de Castro and Yannelis (2018 JET) as special cases.

Indeed, if agent i has a belief, i.e., $P_i = \{\mu_i\}$ is a singleton set, then (1) becomes (2). Clearly, the multi-belief preferences become the Bayesian preferences. Agent i's *ex ante expected utility* of an allocation x is

$$\sum_{\omega \in \Omega} u_i\left(x_i\left(\omega\right), \omega\right) \mu_i\left(\omega\right). \tag{2}$$

Motivation: Preferences

If $P_i = \Delta_i$, then the worst probability in the multi-belief set P_i should assign the whole weight to the worst state in each E_i.

In this case, the multi-belief preferences become the Wald's maxmin preferences in de Castro and Yannelis (2018 JET), where the following formulation (3) is equivalent to (1). Agent i's *ex ante expected utility* of an allocation x is

$$\sum_{E_i \in \mathcal{F}_i} \left(\min_{\omega \in E_i} u_i \left(x_i \left(\omega \right), \omega \right) \right) \mu_i \left(E_i \right). \tag{3}$$

Motivation

If agents start with the Gilboa and Schmeidler (1989 JME) preferences,

- can the Wald's maxmin preferences be justified by introducing a Designer who persuades agents to use the Wald's maxmin preferences so that agents can achieve superior outcomes?
- when is it always a good idea to persuade agents to use the Wald's maxmin preferences?

A couple of findings

With these questions in mind, we introduce persuasion devices (Kamenica and Gentzkow 2011 AER, Beauchene, Li and Li 2019 JET) in an exchange economy, where the agents have asymmetric information. To the best of our knowledge, our paper is the first one to do so.

- To reach an allocation that makes agents better off, the Designer may want to change the agents' preferences to Wald's maxmin by adopting a persuasion device.

A couple of findings

- In the face of a Designer who thinks that any belief can be the agents' priors, we show that it is a good idea for the Designer to persuade the agents to use the Wald's maxmin preferences. It is because that *the set of maxmin incentive compatible allocations* contains *the set of ex post incentive compatible allocations* and *the set of allocations that are Bayesian incentive compatible under all beliefs* as strict subsets.

Furthermore, these results hold

- even if the Designer can rule out impossible beliefs (i.e., beliefs that cannot be the agents' priors) based on the individually rationality conditions.
- even if we take into account that agents may randomize their choices.

Before using persuasion devices

Example 1: There are two agents, one good, and three possible states of nature $\Omega = \{a, b, c\}$. Each agent i has a partition of Ω, denoted by $\mathcal{F}_i$, where $i = 1, 2$:

$$\mathcal{F}_1 = \{\{a, b\}, \{c\}\}; \quad \mathcal{F}_2 = \{\{a, c\}, \{b\}\}.$$

The agents' random initial endowments are

$$(e_1(a), e_1(b), e_1(c)) = (5, 5, 2); \quad (e_2(a), e_2(b), e_2(c)) = (5, 2, 5).$$

The ex post utility function of each agent i is $u_i(c_i, \omega) = \sqrt{c_i}$ for all $\omega \in \Omega$, where c_i denotes agent i's consumption of the good.

Before using persuasion devices

Initially, each agent i has a prior $\mu_i \in \Delta_i$ with full support, i.e., $\mu_i(\omega) > 0$ for each $\omega \in \Omega$.

Suppose that the agents want to end up with an ex post efficient and individually rational allocation y, which provides insurance for them against low endowment realizations.

$$y = \begin{pmatrix} y_1(a) & y_1(b) & y_1(c) \\ y_2(a) & y_2(b) & y_2(c) \end{pmatrix} = \begin{pmatrix} 5 & 5-1.5 & 2+1.5 \\ 5 & 2+1.5 & 5-1.5 \end{pmatrix} = \begin{pmatrix} 5 & 3.5 & 3.5 \\ 5 & 3.5 & 3.5 \end{pmatrix}.$$

Recall that

$$\mathcal{F}_1 = \{\{a,b\},\{c\}\}; \quad \mathcal{F}_2 = \{\{a,c\},\{b\}\}.$$

Before using persuasion devices

We do not require the agents' priors to be common knowledge.

Given an allocation $y \neq e$ that the agents want to reach, the Designer's goal is to help the agents to reach it.

The Designer knows the agents' partitions $\mathcal{F}_i$, $i = 1, 2$, as the partitions are common knowledge. However, the Designer does not know the realized state of nature ω, the private information $E_i(\omega)$, $i = 1, 2$ of the agents, nor the agents' priors, except that the agents' priors have full support.

Before using persuasion devices

Since y is individually rational for both agents, the Designer knows that the agents' priors, μ_1 and μ_2, must satisfy

$$\mu_1(a)\sqrt{5} + \mu_1(b)\sqrt{3.5} + \mu_1(c)\sqrt{3.5} \geq \mu_1(a)\sqrt{5} + \mu_1(b)\sqrt{5} + \mu_1(c)\sqrt{2}, \quad (4)$$

for agent 1 and

$$\mu_2(a)\sqrt{5} + \mu_2(b)\sqrt{3.5} + \mu_2(c)\sqrt{3.5} \geq \mu_2(a)\sqrt{5} + \mu_2(b)\sqrt{2} + \mu_2(c)\sqrt{5} \quad (5)$$

for agent 2.

We show below that the allocation y is not Bayesian incentive compatible under all beliefs in $\Delta_{i,d}$ for each i, where
$\Delta_{1,d} = \left\{\mu \in \Delta^{full} : \mu \text{ satisfies } (4)\right\}$ and
$\Delta_{2,d} = \left\{\mu \in \Delta^{full} : \mu \text{ satisfies } (5)\right\}$.

Before using persuasion devices

It can be easily checked that every belief μ in Δ^{full} that has $\mu(b) = \mu(c)$ is in the set $\Delta_{i,d}$.

Let μ be a belief in $\Delta_{i,d}$. When agent 1 observes the event $\{a, b\}$, and reports the true event $\{a, b\}$, she gets $\sqrt{5}$ if the state is a, and she gets $\sqrt{3.5}$ if the state is b. Since she only knows that the realized state can be a or b, her interim Bayesian payoff is

$$\sqrt{5} \cdot \frac{\mu(a)}{\mu(\{a, b\})} + \sqrt{3.5} \cdot \frac{\mu(b)}{\mu(\{a, b\})}.$$

Before using persuasion devices

If agent 1 reports the lie $\{c\}$, then agent 2 believes her when the state is a. She gets $\sqrt{e_1(a) + y_1(c) - e_1(c)} = \sqrt{6.5}$. If the state is b, agent 2 knows the state and the agents' reports are incompatible.

Suppose that agent 1 is punished and ends up with a payoff of $\sqrt{5 - D_1} \leq \sqrt{5}$, where $D_1 \geq 0$, i.e., agent 1 gets less than her endowment $e_1(b) = 5$. Since she only knows that the realized state can be a or b, her interim Bayesian payoff is

$$\sqrt{6.5} \cdot \frac{\mu(a)}{\mu(\{a,b\})} + \sqrt{5 - D_1} \cdot \frac{\mu(b)}{\mu(\{a,b\})}.$$

Before using persuasion devices

For the allocation y to be Bayesian incentive compatible under all beliefs in $\Delta_{i,d}$ for each i, we need

$$\sqrt{5} \cdot \mu(a) + \sqrt{3.5} \cdot \mu(b) \geq \sqrt{6.5} \cdot \mu(a) + \sqrt{5 - D_1} \cdot \mu(b), \tag{6}$$

for agent 1 under all μ in $\Delta_{1,d}$. That is, when agent 1 observes $\{a, b\}$, she prefers to report $\{a, b\}$ instead of $\{c\}$.

Also, we need

$$\sqrt{5} \cdot \mu(a) + \sqrt{3.5} \cdot \mu(c) \geq \sqrt{6.5} \cdot \mu(a) + \sqrt{5 - D_2} \cdot \mu(c), \tag{7}$$

for agent 2 under all μ in $\Delta_{2,d}$, where $D_2 \geq 0$. That is, when agent 2 observes $\{a, c\}$, he prefers to report $\{a, c\}$ instead of $\{b\}$.

Before using persuasion devices

However, regardless of the values of $\sqrt{5 - D_1} \leq \sqrt{5}$ and $\sqrt{5 - D_2} \leq \sqrt{5}$, (6) and (7) cannot hold whenever $\mu(a)$ is large enough.

Thus, under the primitives of the economy, the allocation y is not Bayesian incentive compatible under all beliefs in $\Delta_{i,d}$ for each i.

Now, the Designer and the agents cannot be sure about the incentive compatibility of y. When an agent is not sure about the incentive compatibility of y, she is unwilling to transfer goods to the other agent according to $y - e$. The reason is that she knows that she may be cheated and lose her wealth without understanding it.

The agents end up consuming their initial endowment e which provides no insurance against low endowment realizations.

Bayesian persuasion

The Designer can alter the agents' beliefs by adopting a Bayesian persuasion device of Kamenica and Gentzkow (AER, 2011). That is,

- the Designer conducts an investigation about the unknown of the economy (i.e., the realized state of nature) and the Designer is required to truthfully report the outcome of the investigation.
- Upon observing the outcome, the agents update their beliefs from μ_i to a new belief in the set Δ, where Δ is the set of all probability measures over 2^Ω, i.e.,
  $\Delta = \{$ probability measure $\mu : 2^\Omega \to [0,1]\}$.

Bayesian persuasion

Formally, a Bayesian persuasion device $\left\langle M, \{q(\cdot \mid \omega)\}_{\omega\in\Omega} \text{ over } M \right\rangle$ consists of a finite set of outcomes M and a family of probability distributions $\{q(\cdot \mid \omega)\}_{\omega\in\Omega}$ over M. That is, for each ω, there is a probability distribution $q(\cdot \mid \omega)$ over M.

Upon observing an outcome $m \in M$, each agent i forms a posterior belief $\mu_{i,m}$, i.e., for each $\omega \in \Omega$,

$$\mu_{i,m}(\omega) = \frac{q(m \mid \omega)\,\mu_i(\omega)}{\sum_{\omega'\in\Omega} q(m \mid \omega')\,\mu_i(\omega')} \tag{8}$$

where μ_i is agent i's prior. That is, agent i's belief changes from μ_i to $\mu_{i,m}$.

Bayesian persuasion

We show in Example 2 below that if the Designer adopts a Bayesian persuasion device, then to make y incentive compatible, the Designer needs a Bayesian persuasion device that has God's type power.

That is, the Designer needs a Bayesian persuasion device that can remove private information by telling everyone the realized state of nature.

Bayesian persuasion

Example 2: The economy is the same as in Example 1. Agent 1 has no incentive to lie when she observes the event $\{c\}$, as this is the state that agent 1 needs insurance and must get something. Same with agent 2 when he observes $\{b\}$.

The Designer needs a Bayesian persuasion device that can change agent 1's belief so that when agent 1 observes $\{a, b\}$, she prefers to report the true event $\{a, b\}$ instead of $\{c\}$. At the same time, the Bayesian persuasion device should change agent 2's belief, so that when agent 2 observes $\{a, c\}$, he prefers to report the true event $\{a, c\}$ instead of $\{b\}$.

Bayesian persuasion

If the Designer adopts a Bayesian persuasion device that has an outcome m, such that $0 < q(m \mid \omega) \leq 1$ for all ω, then upon observing m, the Designer still does not know the value of $\mu_{i,m}(a)$ except $0 < \mu_{i,m}(a) < 1$.

Indeed, the Designer knows that μ_i is from the set Δ^{full} and μ_i satisfies the individually rational constraint (i.e., (4) for agent 1 and (5) for agent 2), but none of these conditions puts any restriction on $\mu_i(a)$, except that $0 < \mu_i(a) < 1$.
Then, it follows from

$$\mu_{i,m}(a) = \frac{q(m \mid a)\,\mu_i(a)}{q(m \mid a)\,\mu_i(a) + q(m \mid b)\,\mu_i(b) + q(m \mid c)\,\mu_i(c)} \tag{9}$$

that $\mu_{i,m}(a)$ can take any value between zero and one as well, i.e., $0 < \mu_{i,m}(a) < 1$.

Bayesian persuasion

Example 1 showed that the allocation y of Example 1 is not Bayesian incentive compatible, whenever the probability of state a is high enough. It follows that such a Bayesian persuasion device cannot make y Bayesian incentive compatible under all beliefs in $\Delta_{i,m,d}$ for each i, and thus y is not Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \cup_{m \in M} \Delta_{i,m,d}$ for each i.

Now, we look into Bayesian persuasion devices in which each outcome m can rule out some state. That is, for each $m \in M$, $q\left(m \mid \omega\right) = 0$ for some ω.

Suppose that the realized state is a, then agent 1 observes the event $\{a, b\}$ and agent 2 observes the event $\{a, c\}$. Also, suppose that the outcome of the persuasion device is m. Clearly, the Bayesian persuasion device must have $q\left(m \mid a\right) \neq 0$.

Bayesian persuasion

Then, there are three cases left:

Case 1: $q(m \mid b) = 0$, $q(m \mid c) = 0$. Now, for all $\mu \in \Delta^{full}$,

$$\mu_m(a) = \frac{q(m \mid a)\,\mu(a)}{q(m \mid a)\,\mu(a) + q(m \mid b)\,\mu(b) + q(m \mid c)\,\mu(c)} = \frac{q(m \mid a)\,\mu(a)}{q(m \mid a)\,\mu(a) + 0 + 0} = 1$$

and

$$\mu_m(\omega) = \frac{q(m \mid \omega)\,\mu(\omega)}{q(m \mid a)\,\mu(a) + q(m \mid b)\,\mu(b) + q(m \mid c)\,\mu(c)} = \frac{0}{q(m \mid a)\,\mu(a) + 0 + 0} = 0,$$

where $\omega = b, c$.

Clearly, by observing m, the Designer and the agents know that the realized state is a. Now, agents' information partitions changed. Every lie can be detected and punished. Thus, every agent prefers to report truthfully. According to the allocation y, the agents get $y(a)$, i.e., there is no trade.

Bayesian persuasion

Case 2: $q(m \mid b) = 0$, $q(m \mid c) \neq 0$. By

$$\mu_m(\omega) = \frac{q(m \mid \omega)\,\mu(\omega)}{\sum_{\omega' \in \Omega} q(m \mid \omega')\,\mu(\omega')},$$

we have that $\mu_m(a) \neq 0$, $\mu_m(b) = 0$, and $\mu_m(c) \neq 0$, for all $\mu \in \Delta^{full}$.

Now, agent 1 knows that the realized state cannot be b, and therefore the realized state must be a. However, the Designer and agent 2 only know that the realized state can be a or c.

Now, agent 1 has an incentive to misreport. Indeed, if she reports $\{c\}$, she gets $5 + 1.5$. If she reports the true event $\{a, b\}$, she gets only 5. Clearly, lying is better.

Case 3: $q\left(m \mid b\right) \neq 0$, $q\left(m \mid c\right) = 0$. By

$$\mu_m\left(\omega\right) = \frac{q\left(m \mid \omega\right)\mu\left(\omega\right)}{\sum_{\omega' \in \Omega} q\left(m \mid \omega'\right)\mu\left(\omega'\right)},$$

we have that $\mu_m\left(a\right) \neq 0$, $\mu_m\left(b\right) \neq 0$, and $\mu_m\left(c\right) = 0$, for all $\mu \in \Delta^{full}$.

Now, agent 2 knows that the realized state cannot be c, and therefore the realized state must be a. However, the Designer and agent 1 only know that the realized state can be a or b.

Now, agent 2 has an incentive to misreport. Indeed, if he reports $\{b\}$, he gets $5 + 1.5$. If he reports the true event $\{a, c\}$, he gets only 5. Clearly, lying is better.

Bayesian persuasion

Thus, the Designer needs a Bayesian persuasion device that has God's type power. That is, the device satisfies that for each outcome m that has $q(m \mid a) \neq 0$, the $q(m \mid b)$ and $q(m \mid c)$ must be zero. Otherwise, y may not be Bayesian incentive compatible.

Under this device, when state a is realized, the Designer and the agents have complete information. We can conclude that the only way for a Bayesian persuasion device to remove incentives to lie is to remove private information.

Ambiguous persuasion

However, we show in Example 3 below that if the Designer adopts an ambiguous persuasion device of Beauchene, Li and Li (JET, 2019) to persuade the agents to use the Wald's maxmin preferences, then the allocation y becomes incentive compatible and the device does not need to know the realized state of nature.

An ambiguous persuasion device consists of a finite set of outcomes M and a collection of families of (probability) distributions over M, in contrast with the Bayesian persuasion device which has only one family of probability distributions over M.

Ambiguous persuasion

Example 3: The economy is the same as in Example 1. An ambiguous persuasion device of Beauchene, Li and Li (JET, 2019) can change μ_i to a set of beliefs. The ambiguous persuasion device consists of a finite set of outcomes M and a collection of families of distributions over M, denoted by $co\,(Q^*)$. Suppose that $M = \{a, \text{ not } a\}$, furthermore the following two families of distributions over M, denoted by q and q', belong to the collection of families of distributions $co\,(Q^*)$:

$$q\,(m = \textit{ a } | \omega = a) = 1; \quad q\,(m = \textit{ not a } | \omega = a) = 0;$$

$$q\,(m = \textit{ a } | \omega = b) = 0; \quad q\,(m = \textit{ not a } | \omega = b) = 1;$$

$$q\,(m = \textit{ a } | \omega = c) = 0; \quad q\,(m = \textit{ not a } | \omega = c) = 1,$$

i.e., if q is used, the outcome m is accurate;

Ambiguous persuasion

$$q'\left(m = \textit{a} \mid \omega = a\right) = 0; \quad q'\left(m = \textit{not a} \mid \omega = a\right) = 1;$$

$$q'\left(m = \textit{a} \mid \omega = b\right) = 1; \quad q'\left(m = \textit{not a} \mid \omega = b\right) = 0;$$

$$q'\left(m = \textit{a} \mid \omega = c\right) = 1; \quad q'\left(m = \textit{not a} \mid \omega = c\right) = 0,$$

i.e., if q' is used, the outcome m is completely wrong.

The Designer does not know which family of probability distributions over M is accurate.

Ambiguous persuasion

The agents observe the outcome $m \in M$ of the persuasion device, and update their priors based on q and q' respectively. That is, if agent i, $i = 1, 2$, observes the outcome "$m = a$", then the set of posteriors is $\left\{\mu_{i,a}, \mu'_{i,a}\right\}$.

Agent i gets the posterior $\mu_{i,a}$ by updating μ_i based on q, i.e.,

$$\begin{aligned}
\mu_{i,a}(a) &= \frac{q(m=a \mid \omega=a)\,\mu_i(a)}{q(m=a \mid \omega=a)\,\mu_i(a) + q(m=a \mid \omega=b)\,\mu_i(b) + q(m=a \mid \omega=c)\,\mu_i(c)} \\
&= \frac{1 \times \mu_i(a)}{1 \times \mu_i(a) + 0 \times \mu_i(b) + 0 \times \mu_i(c)} = 1,
\end{aligned}$$

$$\begin{aligned}
\mu_{i,a}(b) &= \frac{q(m=a \mid \omega=b)\,\mu_i(b)}{q(m=a \mid \omega=a)\,\mu_i(a) + q(m=a \mid \omega=b)\,\mu_i(b) + q(m=a \mid \omega=c)\,\mu_i(c)} \\
&= \frac{0 \times \mu_i(b)}{1 \times \mu_i(a) + 0 \times \mu_i(b) + 0 \times \mu_i(c)} = 0,
\end{aligned}$$

$$\begin{aligned}
\mu_{i,a}(c) &= \frac{q(m=a \mid \omega=c)\,\mu_i(c)}{q(m=a \mid \omega=a)\,\mu_i(a) + q(m=a \mid \omega=b)\,\mu_i(b) + q(m=a \mid \omega=c)\,\mu_i(c)} \\
&= \frac{0 \times \mu_i(c)}{1 \times \mu_i(a) + 0 \times \mu_i(b) + 0 \times \mu_i(c)} = 0.
\end{aligned}$$

Ambiguous persuasion

Similarly, agent i gets the posterior $\mu'_{i,a}$ by updating μ_i based on q', where

$$\mu'_{i,a}(a) = 0, \qquad \mu'_{i,a}(b) = \frac{\mu_i(b)}{\mu_i(b) + \mu_i(c)}, \qquad \mu'_{i,a}(c) = \frac{\mu_i(c)}{\mu_i(b) + \mu_i(c)}.$$

If agent i observes "$m =$ not a", then the set of posteriors is $\left\{\mu_{i,\text{not a}}, \mu'_{i,\text{not a}}\right\}$, where

$$\mu_{i,\text{not a}}(a) = 0, \qquad \mu_{i,\text{not a}}(b) = \frac{\mu_i(b)}{\mu_i(b) + \mu_i(c)}, \qquad \mu_{i,\text{not a}}(c) = \frac{\mu_i(c)}{\mu_i(b) + \mu_i(c)};$$

$$\mu'_{i,\text{not a}}(a) = 1, \qquad \mu'_{i,\text{not a}}(b) = 0, \qquad \mu'_{i,\text{not a}}(c) = 0.$$

Ambiguous persuasion

The Designer does not know the values of $\mu_i(\omega)$, $\omega \in \Omega$. However, the Designer knows from $\mu_{i,a}$, $\mu'_{i,a}$, $\mu_{i,\text{not a}}$ and $\mu'_{i,\text{not a}}$ that regardless of the outcome of the ambiguous persuasion device, when agent 1 observes $\{a, b\}$, her multi-belief set contains all probability distributions over $\{a, b\}$.

The same holds for agent 2: regardless of the outcome of the ambiguous persuasion device, when agent 2 observes $\{a, c\}$, his multi-belief set contains all probability distributions over $\{a, c\}$.

That is, the Designer knows that both agents have the Wald's maxmin preferences.

Ambiguous persuasion

We show below that when both agents have the Wald's maxmin preferences, the allocation y becomes maxmin incentive compatible.

Suppose that agent 1 observes the event $\{a, b\}$. If agent 1 reports the true event $\{a, b\}$, she gets $\sqrt{5}$ if the state is a, and she gets $\sqrt{3.5}$ if the state is b. Since she only knows that the realized state could be a or b, her interim Wald's maxmin payoff is

$$\min\left\{\sqrt{5},\ \sqrt{3.5}\right\} = \sqrt{3.5}.$$

Ambiguous persuasion

If agent 1 reports the lie $\{c\}$, she gets $\sqrt{6.5}$ if the state is a, and she gets $\sqrt{3.5}$ if the state is b. Her interim Wald's maxmin payoff is

$$\min\left\{\sqrt{6.5},\ \sqrt{3.5}\right\} = \sqrt{3.5}.$$

It follows that agent 1 has no incentive to misreport the observed event, when she sees the event $\{a, b\}$.

When agent 1 observes the event $\{c\}$, she has no incentive to lie, as this is the state that she needs insurance and must get something: if she reports the true event $\{c\}$, she gets $\sqrt{3.5}$; if she reports the lie $\{a, b\}$, she gets nothing and her payoff is $\sqrt{2}$.

The same holds for agent 2. Thus, the allocation y is incentive compatible under the Wald's maxmin preferences (i.e., maxmin incentive compatible).

Randomization

Example 4 below shows that even if an ex post efficient allocation is maxmin incentive compatible, an agent may not want to report the true event with probability one, as reporting the true event with a positive probability that is less than one may give her a strictly higher interim Wald's maxmin payoff.

Example 4: There are two agents, 1 and 2, one good, and four states of nature $\Omega = \{a, b, c, d\}$. Each agent i has a partition of Ω, denoted by $\mathcal{F}_i$, where $i = 1, 2$:

$$\mathcal{F}_1 = \{\{a, b\}, \{c, d\}\}; \quad \mathcal{F}_2 = \{\{a, d\}, \{b, c\}\}.$$

The ex post utility function of each agent i is $u_i(c_i, \omega) = \sqrt{c_i}$ for all $\omega \in \Omega$, where c_i denotes agent i's consumption of the good.

The agents get 2.5 units of the good in each state, i.e., $e_i(\omega) = 2.5$, for each $\omega \in \Omega$ and for each i.

Randomization

The agents have the Wald's maxmin preferences.

Let x be a feasible allocation:

$$(x_1(a), x_1(b), x_1(c), x_1(d)) = (3, 2, 3, 2);$$

$$(x_2(a), x_2(b), x_2(c), x_2(d)) = (2, 3, 2, 3).$$

The allocation x is ex post efficient and maxmin incentive compatible.

However, when agent 1 observes the event $\{a, b\}$, she does not want to report the true event with probability one, as reporting the true event with a positive probability that is less than one gives her a strictly higher interim Wald's maxmin payoff.

Randomization

Indeed, let $\alpha \in [0,1]$ be the probability of reporting the event $\{a,b\}$ and $1-\alpha$ the probability of reporting the event $\{c,d\}$.

Then, agent 1 gets $\alpha\sqrt{3} + (1-\alpha)\sqrt{2}$ if the state is a, and she gets $\alpha\sqrt{2} + (1-\alpha)\sqrt{3}$ if the state is b. Since she only knows that the realized state could be a or b, her interim Wald's maxmin payoff is

$$\min\left\{\alpha\sqrt{3} + (1-\alpha)\sqrt{2},\ \alpha\sqrt{2} + (1-\alpha)\sqrt{3}\right\}.$$

- When $\alpha = 1$, i.e., she reports the true event $\{a,b\}$, her interim Wald's maxmin payoff is $\sqrt{2}$.
- When $\alpha = \frac{1}{2}$, i.e., she reports the true event $\{a,b\}$ with probability $\frac{1}{2}$, her interim Wald's maxmin payoff is $\frac{\sqrt{3}+\sqrt{2}}{2}$ which is strictly higher than $\sqrt{2}$.

Therefore, when agent 1 observes the event $\{a,b\}$, she has no incentive to report the true event with probability one.

Randomization

Example 5: We show below that the allocation y of Example 1 is mixed maxmin incentive compatible.

Indeed, suppose that agent 1 observes the event $\{a, b\}$. Now, let $\alpha \in [0, 1]$ be the probability of reporting the event $\{a, b\}$ and $1 - \alpha$ the probability of reporting the event $\{c\}$.

Then, agent 1 gets $\alpha\sqrt{5+0} + (1-\alpha)\sqrt{5+1.5}$ if the state is a, and she gets $\alpha\sqrt{5-1.5} + (1-\alpha)\sqrt{5-1.5}$ if the state is b. Since she only knows that the realized state could be a or b, her interim Wald's maxmin payoff is

$$\min\left\{\alpha\sqrt{5+0} + (1-\alpha)\sqrt{5+1.5},\ \alpha\sqrt{5-1.5} + (1-\alpha)\sqrt{5-1.5}\right\}$$
$$= \sqrt{5-1.5}.$$

That is, her interim Wald's maxmin payoff is $\sqrt{5-1.5}$ regardless of the value of α. It follows that agent 1 has no incentive to misreport the observed event, when she sees the event $\{a, b\}$.

Randomization

When Agent 1 observes the event $\{c\}$, she has no incentive to lie either.

Indeed, she gets $\alpha\sqrt{2+0} + (1-\alpha)\sqrt{2+1.5} < \sqrt{2+1.5}$ whenever α is not zero.

She gets $\sqrt{2+1.5}$, when she reports the true event $\{c\}$ (i.e., $\alpha = 0$).

Thus, reporting the true event is optimal.

The same argument holds for agent 2. We can conclude that the allocation y is mixed maxmin incentive compatible.

A better choice: Wald's maxmin preferences

Recall that the designer does not know the values of $\mu_i(\omega)$, $\omega \in \Omega$, except that $\mu_i(\omega) > 0$ for each $\omega \in \Omega$.

Suppose that the Designer knows that an allocation x is individually rational (under μ_i and Wald's maxmin). Based on this information, the Designer forms the set $\Delta_{i,d}$ by ruling out impossible beliefs from Δ^{full}. That is, the Designer thinks that any belief in $\Delta_{i,d}$ can be agent i's prior belief μ_i.

A better choice: Wald's maxmin preferences

In the face of a Designer who thinks that agent i's prior μ_i can be any belief in $\Delta_{i,d}$,

- One can focus on allocations that are incentive compatible under all beliefs in $\Delta_{i,d}$ for each i.
- Another choice is to focus on allocations that are ex post incentive compatible.
- We present a better choice, i.e., adopting an ambiguous persuasion device to persuade the agents to use the Wald's maxmin preferences. This change in preferences enlarges the set of incentive compatible allocations.

A better choice: Wald's maxmin preferences

- X_P denotes *the set of ex post efficient allocations that are ex post incentive compatible*.
- X_B denotes *the set of ex post efficient allocations that are incentive compatible under all beliefs in* $\Delta_{i,d}$.
- X_A denotes *the set of ex post efficient and maxmin incentive compatible allocations*
- X_{MA} denotes *the set of ex post efficient and mixed maxmin incentive compatible allocations*

A better choice: Wald's maxmin preferences

Main result 1: When $\Delta_{i,d} = \Delta^{full}$, we have that $X_P \subseteq X_B \subset X_{MA} \subset X_A$.

Main result 2: When $\Delta_{i,d} \subset \Delta^{full}$, we have that $X_B \subset X_{MA} \subset X_A$.

To sum up

- To reach an allocation that makes agents better off, the Designer may want to change the agents' preferences to Wald's maxmin by adopting a persuasion device.
- Furthermore, in the face of a Designer who thinks that an agent's prior can be any belief in Δ^{full}, it is always a good idea to persuade the agents to use the Wald's maxmin preferences, as the set of ex post efficient, individually rational and incentive compatible allocations becomes strictly larger. Moreover, this result remains true, even when we take into account that the agents may randomize over their choices.