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

Zhiwei Liu† Nicholas C. Yannelis ‡

Abstract

To justify the use of Wald's maxmin preferences in an asymmetric information economy, we introduce a mechanism designer who can convince/persuade agents to adopt Wald's maxmin preferences. We show that more efficient and individually rational allocations become incentive compatible if the mechanism designer persuades agents to use Wald's maxmin preferences instead of Bayesian preferences. Thus, we justify the Wald's maxmin preferences by showing that agents can be persuaded to use them in order to enlarge the set of efficient, individually rational and incentive compatible allocations.

Keywords: Persuasion, Efficient, Individually rational, Incentive compatibility, Wald's maxmin preferences.

JEL Classification Number: D51, D81, D82, D83

*We thank Andy Zapechelnyuk, Joel Sobel, Bart Taub, Sayre P. Satterwhite and two referees and the participants at "Seminar Series on Communication and Persuasion" for helpful comments, discussions and suggestions. Zhiwei Liu gratefully acknowledges the financial support from the National Natural Science Foundation of China [grant No. 71703110].

†International School of Economics and Management, Capital University of Economics and Business, Beijing, China

Email: liuzhiwei@cueb.edu.cn

‡Department of Economics, Henry B. Tippie College of Business, The University of Iowa, 108 John Pappajohn Business Building, Iowa City, IA 52242-1994.

Email: nicholasyannelis@gmail.com

1 Introduction

When agents have the Wald's maxmin preferences, Castro and Yannelis (2018) showed that the conflict between efficiency and incentive compatibility no longer exists, i.e., all efficient allocations are also incentive compatible.[1] However, the literature assumes that from the primitive, the agents are Wald's maxmin, as Wald's maxmin are the only preferences under which all efficient allocations are incentive compatible. In other words, in the literature we see that one decides a priori the functional form of the expected utility operator for each agent and no choice is given to them. Obviously, different expected utility functional forms provide different outcomes as the equilibrium allocations are computed adopting different expected utility operators. However, an agent cannot make a choice as the expected utility form is taken as given. If the agents start with the Gilboa and Schmeidler (1989) preferences, can a mechanism designer (thereafter, Designer) give a choice to the agents to adopt a specific expected utility form, i.e., the Wald's maxmin expected utility by persuading them to do so? In particular, can the Wald's maxmin preferences be justified by introducing a Designer who convinces/persuades agents to use the Wald's maxmin preferences to provide superior outcomes? When is it always a good idea to persuade agents to use the Wald's maxmin preferences? With these questions in mind, we introduce persuasion devices in an exchange economy, where the agents have asymmetric information. Our paper is the first one that introduces persuasion in an asymmetric information exchange economy. Furthermore, we compare Bayesian persuasion devices of Kamenica and Gentzkow (2011) with ambiguous persuasion devices of Beauchêne, Li, and Li (2019). We list our main contributions below.

First, we show that changing to Wald's maxmin preferences can make agents better off. More efficient and individually rational allocations become incentive compatible if the Designer persuades the agents to use the Wald's maxmin preferences instead of the Bayesian preferences.[2] Interestingly, there are allocations that are strictly better than the random initial endowment. However, these allocations are not Bayesian incentive compatible. Thus, if the agents have the Bayesian preferences, they cannot reach such allocations. After the agents change to Wald's maxmin preferences, these allocations become maxmin incentive compatible, and thus they can be reached (Liu and Yannelis (2021)).[3] In other words, now the agents can reach these

[1] When coalitional manipulations (Guo and Yannelis (2020)) and mixed strategy deviations (Liu, Song, and Yannelis (2020)) are considered, the results of Castro and Yannelis (2018) still hold.

[2] We do not assume that the Designer has a Bayesian persuasion device that has God's type power (i.e., the device can remove private information by telling everyone the realized state of nature), as it can be prohibitively costly to have such a persuasion device. Furthermore, we refer interested readers to Degan and Li (2021) on persuasion with costly precision.

[3] Liu and Yannelis (2021) showed that every maxmin incentive compatible allocation is implementable as a maxmin equilibrium. In a maxmin equilibrium, agents use maxmin strategies as in Decerf and Riedel (2020). That is, every agent maximizes her payoff that takes into account the worst state that can occur and also the worst strategy of all the other agents against her.

allocations that are strictly better than their random initial endowment, and cannot be reached under the Bayesian preferences.

Second, in the face of a Designer who thinks that any belief can be the agents' priors, we show in Theorem 1 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, the result of Theorem 1 holds, even if the Designer can rule out impossible beliefs (i.e., beliefs that cannot be the agents' priors) based on the individually rationality conditions. It is exactly for these reasons we remarked that Wald's maxmin preferences provide superior outcomes for all agents.

Third, we introduce randomization and indicate in Example 6 that randomizing over choices improves welfare. In particular, we show that under the Wald's maxmin preferences, even if an 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. Furthermore, such a profitable unilateral deviation brings Pareto improvements to the agents. Since randomization may increase agents' interim Wald's maxmin payoffs and bring Pareto improvements, we take into account that the agents may randomize over their choices. It follows that the set of *mixed maxmin incentive compatible* allocations is a strict subset of the set of maxmin incentive compatible allocations. Nevertheless, we show that Theorem 1 above remains true, i.e., Wald's maxmin preferences provide superior outcomes for all agents. It is because that the set of mixed 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.

The paper is organized as follows. Section 2 and Section 3 define an asymmetric information exchange economy and incentive compatible notions. In Section 4, we introduce persuasion devices in the asymmetric information exchange economy and show that convincing/persuading agents to use the Wald's maxmin preferences leads to superior outcomes. In Section 5, we look at when it is always a good idea to use the Wald's maxmin preferences. Finally, we conclude in Section 6. The proofs of our results are collected in the Appendix.

2 Asymmetric information exchange economy

Let $\mathbb{R}_+^\ell$ denote the ℓ-goods commodity space and I the set of N agents, i.e., $I = \{1, \cdots, N\}$. Let Ω be a finite set of states of nature and $\omega \in \Omega$ a state of nature. Agent i's random initial endowment is a mapping from the set of states of nature to the commodity space, i.e., $e_i : \Omega \to \mathbb{R}_+^\ell$. *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. Let $u_i : \mathbb{R}_+^\ell \times \Omega \to \mathbb{R}$ denote agent i's ex post utility function, taking the form of $u_i(c_i, \omega)$ where c_i denotes agent i's consumption. For each ω, the function $u_i(\cdot, \omega)$ is continuous and bounded.

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. Each event is a maximal set of states that agent i cannot distinguish. In the interim, each agent observes an event in $\mathcal{F}_i$ that contains the realized state of nature. That is, if state ω occurs, agent i only knows that the event $E_i(\omega)$ has occurred, where $E_i(\omega)$ denotes the element of $\mathcal{F}_i$ that contains the state ω. The event $E_i(\omega)$ is agent i's private information. We impose the standard no redundant state assumption. That is, when a state occurs and all agents truthfully report their private information, they will know the realized state:

Assumption 1: *For each* ω, $\bigcap_{j \in I} E_j(\omega) = \{\omega\}$.

Since the agents observe events in the interim, it is natural to assume that at ex ante each agent is able to form a probability assessment over her partition. That is, 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. Each π_i is a well defined probability measure, but it is not defined on every state of nature. Indeed, if $E_i = \{\omega, \omega'\}$ with $\omega \neq \omega'$, then the probability of the event E_i is well defined, but not the probability of the event $\{\omega\}$ or the event $\{\omega'\}$.

Assumption 2: *For each* i *and for each event* $E_i \in \mathcal{F}_i$, $\pi_i(E_i) > 0$.

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

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

We postulate that each agent i's preferences on L are *maxmin à la* Gilboa and Schmeidler (1989). Let P_i be agent i's multi-belief set which is a non-empty, closed and convex subset of Δ_i.[4] We focus on two

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

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

special cases of the maxmin preferences: $P_i = \{\mu_i\}$ and $P_i = \Delta_i$. When $P_i = \{\mu_i\}$, agent i has the Bayesian preferences. Agent i's *ex ante expected utility* of an allocation x is

$$\sum_{\omega \in \Omega} u_i (x_i (\omega), \omega) \mu_i (\omega) . \tag{3}$$

When $P_i = \Delta_i$, agent i has the Wald's maxmin preferences in Castro and Yannelis (2018). Agent i's *ex ante expected utility* of an allocation x is

$$\sum_{E_i \in \mathcal{F}_i} \left(\min_{\omega \in E_i} u_i (x_i (\omega), \omega) \right) \mu_i (E_i) . \tag{4}$$

Initially, each agent i has a prior $\mu_i \in \Delta_i$ with full support, i.e., $\mu_i (\omega) > 0$ for each $\omega \in \Omega$. We do not require the agents' priors to be common knowledge. Given an allocation $x \neq e$ that the agents want to reach, the Designer's goal is to help the agents to reach it. The Designer knows the agents' partitions $\mathcal{F}_i$, $i \in I$, as the partitions are common knowledge. However, the Designer does not know the realized state of nature ω, the private information $E_1 (\omega), \cdots, E_N (\omega)$ of the agents, nor the agents' priors, except that the agents' priors have full support.[5]

The Designer can adopt a Bayesian persuasion device of Kamenica and Gentzkow (2011) or an ambiguous persuasion device of Beauchêne, Li, and Li (2019) to change μ_i. A Bayesian persuasion device can change μ_i to a new belief in the set Δ, where Δ is the set of all probability measures over 2^{Ω}, i.e., $\Delta = \big\{ \text{probability measure } \mu : 2^{\Omega} \to [0, 1] \big\}$.[6] Clearly, a change in an agent's belief leads to a change in her preferences, but she is still Bayesian. An ambiguous persuasion device can change μ_i to a set of beliefs. Now, agent i uses (2) with a non-singleton P_i instead of (3). Obviously, the agent's preferences change from Bayesian to non-Bayesian. In other words, the agents can be persuaded to change preferences. We show below that all agents become better off if the agents are persuaded to use the Wald's maxmin preferences. More generally, we show below that the set of Pareto optimal, individually rational and incentive compatible allocations become larger under the Wald's maxmin preferences.

An asymmetric information exchange economy $\mathcal{E}$ is the set $\mathcal{E} = \{\Omega, (\mu_i, \mathcal{F}_i, e_i, u_i) : i \in I\}$. In this econ-

Let x and y be two allocations from L. If the ex ante expected utility of x_i is larger than that of y_i, then agent i prefers x_i to y_i, $x_i \succeq_i y_i$. By "Larger than", we mean "greater than or equal to". Moreover, she strictly prefers x_i to y_i, $x_i \succ_i y_i$, if she prefers x_i to y_i but not the reverse, i.e., $x_i \succeq_i y_i$ but $y_i \not\succeq_i x_i$. This general multi-belief model includes both the Bayesian and the Wald's maxmin preferences in Castro and Yannelis (2018) as special cases. Indeed, if agent i has a belief, i.e., $P_i = \{\mu_i\}$ is a singleton set, then (2) becomes (3). Clearly, the multi-belief preferences become the Bayesian 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 Castro and Yannelis (2018), where the following formulation (4) is equivalent to (2).

[5] Hu and Weng (2020) has a similar assumption.

[6] The Designer does not know the values of $\mu_i (\omega)$, $\omega \in \Omega$, thus he may not know the values of the new belief either.

omy, an allocation x is *feasible* if at every state $\omega \in \Omega$, the sum of consumptions is the same as the sum of the endowments, i.e., $\sum_{i=1}^{N} x_i(\omega) = \sum_{i=1}^{N} e_i(\omega)$. A feasible allocation is *ex post efficient* if there is no Pareto improvement at ex post, that is, there does not exist a state of nature at which some agent can be made strictly better off without hurting other agents.

Definition 1: *A feasible allocation $x = (x_i)_{i\in I}$ is ex post efficient, if there does not exist another feasible allocation $y = (y_i)_{i\in I}$ and a state ω, such that $u_i(y_i(\omega), \omega) \geq u_i(x_i(\omega), \omega)$ for all i, and $u_i(y_i(\omega), \omega) > u_i(x_i(\omega), \omega)$ for at least one i.*

Furthermore, a feasible allocation x is individually rational if every agent prefers x to their random initial endowment e. Formally,

Definition 2: *A feasible allocation x is said to be individually rational if*

$$\min_{\mu_i \in P_i} \sum_{\omega\in\Omega} u_i(x_i(\omega), \omega)\,\mu_i(\omega) \geq \min_{\mu_i \in P_i} \sum_{\omega\in\Omega} u_i(e_i(\omega), \omega)\,\mu_i(\omega)$$

for all i. When $P_i = \{\mu_i\}$, the allocation x is Bayesian individually rational. When $P_i = \Delta_i$, the allocation x is maxmin individually rational.

Remark 1: *We focus on ex post efficient and individually rational allocations, as a lot of equilibrium notions under asymmetric information are ex post efficient and individually rational, for example, core allocations, value allocations, Walrasian expectations equilibrium allocations and rational expectations equilibrium allocations. These notions are neither incentive compatible nor implementable under the Bayesian preferences, but they are both incentive compatible and implementable under the Wald's maxmin preferences, as were shown by Castro, Pesce, and Yannelis (2011), Glycopantis and Yannelis (2018), Pram (2020), Lombardi and Yoshihara (2020), Guo and Yannelis (2021), Angelopoulos and Koutsougeras (2015), He and Yannelis (2015), Castro, Pesce, and Yannelis (2020), Qin and Yang (2020)* and *Liu (2016).*

3 Incentive compatibility

Given the random initial endowment e, if agents want to end up with a feasible allocation $x \neq e$, transfers need to take place. Since we allow both e and x to depend on the state of nature ω, the transfers may depend on ω as well. In the interim, each agent i privately observes an event $E_i(\omega)$ that contains the realized state of nature ω. Thus, to end up with the correct transfer, it is necessary to pool their private information.

Therefore, we assume that each agent i decides which event $\hat{E}_i \in \mathcal{F}_i$ to report after learning $E_i(\omega)$, in order that they may end up with the correct transfer. By doing so, agents may misreport their true events.

Formally, let x be an allocation. Let $t\left(\hat{E}_1, \cdots, \hat{E}_N\right)$ denote the transfer among the agents, when every agent i reports $\hat{E}_i \in \mathcal{F}_i$. By Assumption 1, $\cap_{i\in I}\hat{E}_i$ is either an empty set or a singleton set. When $\cap_{i\in I}\hat{E}_i = \{\hat{\omega}\}$ for some $\hat{\omega}$ in Ω, let $t\left(\hat{E}_1, \cdots, \hat{E}_N\right) = x(\hat{\omega}) - e(\hat{\omega})$. Clearly, if the state is ω, and the agents report truthfully, i.e., $\hat{E}_i = E_i(\omega)$ for each i, then the agents end up with the transfer $t\left(\hat{E}_1, \cdots, \hat{E}_N\right) = x(\omega) - e(\omega)$. This transfer is correct, since the agents reach $x(\omega)$ after this transfer, i.e., $e(\omega) + t\left(\hat{E}_1, \cdots, \hat{E}_N\right) = e(\omega) + x(\omega) - e(\omega) = x(\omega)$. When $\cap_{i\in I}\hat{E}_i = \emptyset$, we say that the agents' reports are incompatible. There are many ways to define the transfer $t\left(\hat{E}_1, \cdots, \hat{E}_N\right)$ when $\cap_{i\in I}\hat{E}_i = \emptyset$. Some choices are no transfer (i.e., every agent keeps her endowment), imposing the worst possible transfer, or randomly assigning a transfer, just to name a few (see for example, Glycopantis, Muir, and Yannelis (2001), Liu (2016), Castro, Pesce, and Yannelis (2020) and Castro, Liu, and Yannelis (2017b)). Since we allow the agents' endowments to vary with the state of nature, we impose the following feasibility condition as in Castro, Liu, and Yannelis (2017a,b), Moreno-García and Torres-Martínez (2020): every transfer under x is feasible, i.e., $e_i(\omega) + x_i(\hat{\omega}) - e_i(\hat{\omega}) \in \mathbb{R}^\ell_+$, for each i, ω, $\hat{\omega}$. Clearly, if each e_i is constant, then the feasibility condition above is automatically satisfied. Let $E_{-i}(\omega)$ denote $(E_1(\omega), \cdots, E_{i-1}(\omega), E_{i+1}(\omega), \cdots, E_N(\omega))$, also let $\hat{E}_i \cap E_{-i}(\omega)$ denote $\hat{E}_i \cap E_1(\omega) \cap \cdots \cap E_{i-1}(\omega) \cap E_{i+1}(\omega) \cap \cdots \cap E_N(\omega)$. For any allocation x and for any $\hat{E}_i$ in $\mathcal{F}_i$, let

$$x_i^{\hat{E}_i}(\omega) = e_i(\omega) + t_i\left(\hat{E}_i, E_{-i}(\omega)\right) = \begin{cases} e_i(\omega) + x_i(\hat{\omega}) - e_i(\hat{\omega}) & \text{if } \hat{E}_i \cap E_{-i}(\omega) = \{\hat{\omega}\} \\ e_i(\omega) + D_i & \text{if } \hat{E}_i \cap E_{-i}(\omega) = \emptyset, \end{cases}$$

where $D_i \in \mathbb{R}^\ell$ is agent i's transfer when the agents' reports are incompatible.

If agents have the Bayesian preferences, then the incentive compatibility notion is standard: an allocation x is incentive compatible if no agent can improve her interim Bayesian payoff by falsely reporting her private information. Formally,

Definition 3: *An allocation x is Bayesian incentive compatible, if for each i, and for each $E_i \in \mathcal{F}_i$,*

$$\sum_{\omega\in\Omega} u_i(x_i(\omega), \omega)\, \mu_i(\omega \mid E_i) \geq \sum_{\omega\in\Omega} u_i\left(x_i^{\hat{E}_i}(\omega), \omega\right) \mu_i(\omega \mid E_i), \tag{5}$$

for all $\hat{E}_i \in \mathcal{F}_i$, where $\mu_i(\omega \mid E_i)$ denotes agent i's conditional probability for the state of nature being ω, *given that she observes E_i, i.e.,*

$$\mu_i(\omega \mid E_i) = \begin{cases} \frac{\mu_i(\omega')}{\mu_i(E_i)} & \text{if } \omega' \in E_i \\ 0 & \text{if } \omega' \notin E_i. \end{cases}$$

Recall that the designer does not know the values of $\mu_i(\omega)$, $\omega \in \Omega$, except that $\mu_i(\omega) > 0$ for each $\omega \in \Omega$. Let Δ^{full} be the set of beliefs that have full support: $\mu \in \Delta$ is in the set Δ^{full} if and only if $\mu(\omega) > 0$ for all $\omega \in \Omega$. Furthermore, let $\Delta_{i,d} \subseteq \Delta$ denote the set of beliefs that the Designer cannot distinguish. That is, the Designer thinks that any belief $\mu \in \Delta_{i,d}$ can be agent i's belief. In particular, if no persuasion device is used, we use $\Delta_{i,d}$ to denote the set of beliefs that the Designer thinks that any belief $\mu \in \Delta_{i,d} \subseteq \Delta^{full}$ can be agent i's prior μ_i. If a Bayesian persuasion device of Kamenica and Gentzkow (2011) is used, we use $\Delta_{i,d}$ to denote the set of beliefs that the Designer thinks that any belief $\mu \in \Delta_{i,d} \subseteq \Delta$ can be agent i's belief after the persuasion. When $\Delta_{i,d}$ contains only one belief μ, the Designer knows that an allocation x is Bayesian incentive compatible if and only if x satisfies Definition 3. When $\Delta_{i,d}$ contains more than one belief and the agents have Bayesian preferences, the Designer knows that an allocation x is Bayesian incentive compatible if and only if x satisfies Definition 3 under all beliefs in $\Delta_{i,d}$ for each i.[7]

If the Designer adopts an ambiguous persuasion device of Beauchêne, Li, and Li (2019) to convince/persuade the agents to use the Wald's maxmin preferences, then the agents maximize their interim Wald's maxmin payoffs. An allocation is incentive compatible, if no agent can improve her interim Wald's maxmin payoff by falsely reporting her private information.[8]

Definition 4: *Suppose that agents have the Wald's maxmin preferences. An allocation x is maxmin incentive compatible, if for each i, and for each $E_i \in \mathcal{F}_i$,*

$$\min_{\omega \in E_i} u_i(x_i(\omega), \omega) \geq \min_{\omega \in E_i} u_i\left(x_i^{\hat{E}_i}(\omega), \omega\right), \tag{6}$$

for all $\hat{E}_i \in \mathcal{F}_i$.

[7]In face of such uncertainty, i.e., $\Delta_{i,d}$, the Designer (sender) can use the worst belief in $\Delta_{i,d}$, as in Carrasco, Luz, Monteiro, and Moreira (2019) and Hu and Weng (2020). Another choice is that he can apply the principle of maximum entropy to resolve his ignorance (Zapechelnyuk and Kolotilin (2021)). However, in our paper, the Designer's goal is to remove the agents' incentives to lie. Thus, we assume that the Designer focuses on allocations that are Bayesian incentive compatible under all beliefs in $\Delta_{i,d}$ for each i.

[8]As in the standard Bayesian incentive compatibility notion (Definition 3), we use $\geq$ in the definition of maxmin incentive compatibility (Definition 4). The two incentive compatibility notions only differ in the agents' preferences. This formulation allows us to compare what happens to an allocation's incentive compatibility when the agents' preferences change (Sections 4 and 5 below).

4 Persuasion in an asymmetric information exchange economy

We show in Example 1 below that under the primitives of the economy, the agents want to end up with an individually rational allocation y which provides insurance for them against low endowment realizations. However, such an allocation is not Bayesian incentive compatible under all beliefs in $\Delta_{i,d}$ for each i. Can the Designer make the allocation y incentive compatible by adopting a persuasion device? We use persuasion devices as in the literature and compare the use of different persuasion devices in Examples 2 and 3. If the Designer adopts a Bayesian persuasion device of Kamenica and Gentzkow (2011), then to make y incentive compatible, the Designer needs a device that has God's type power: the Bayesian persuasion device needs to know the realized state of nature. However, as we will show below, this is not the case if the Designer adopts an ambiguous persuasion device of Beauchêne, Li, and Li (2019) to convince/persuade the agents to use the Wald's maxmin preferences. In particular, the ambiguous persuasion device does not need to know the realized state of nature, under which the allocation y is incentive compatible.

4.1 Before using persuasion devices

We show in Example 1 below that the agents want to end up with an efficient and individually rational allocation $y \neq e$ which provides insurance for them against low endowment realizations. However, the allocation y is not Bayesian incentive compatible under all beliefs in $\Delta_{i,d}$ for each i. If y is not Bayesian incentive compatible, then an agent may think that she is getting a lot according to y, but she may be cheated and in fact she is getting very little. Thus, when agents are not sure about the incentive compatibility of y, they are unwilling to transfer goods to each other according to $y - e$. That is, they end up consuming their random initial endowment e which may give them low endowment realizations.

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 $i = 1, 2$:*

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

For example, if state a occurs, agent 1 observes the event $\{a, b\}$ which is her private information in the interim. At the same time, agent 2 observes the event $\{a, c\}$ which is his private information in the interim. The agents' random initial endowment is

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

That is, in state a, agents 1 and 2 receive 5 units of the good each, etc. 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.*

Clearly, agent 1's endowment is low in state c, and agent 2's endowment is low in state b. The agents would want to sign a contract at ex ante to insure each other against low endowment realizations. Such a contract specifies the agents' transfer at each state of nature, and results in an allocation that provides risk sharing. Suppose that the agents want to end up with an ex post efficient and individually rational allocation y: [9]

$$y = \begin{pmatrix} y_1(a) & y_1(b) & y_1(c) \\ y_2(a) & y_2(b) & y_2(c) \end{pmatrix} = \begin{pmatrix} 5 & 5-1.5 & 2+1.5 \\ 5 & 2+1.5 & 5-1.5 \end{pmatrix} = \begin{pmatrix} 5 & 3.5 & 3.5 \\ 5 & 3.5 & 3.5 \end{pmatrix}. \tag{7}$$

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

$$\mu_1(a)\sqrt{5} + \mu_1(b)\sqrt{3.5} + \mu_1(c)\sqrt{3.5} \geq \mu_1(a)\sqrt{5} + \mu_1(b)\sqrt{5} + \mu_1(c)\sqrt{2}, \tag{8}$$

for agent 1 and

$$\mu_2(a)\sqrt{5} + \mu_2(b)\sqrt{3.5} + \mu_2(c)\sqrt{3.5} \geq \mu_2(a)\sqrt{5} + \mu_2(b)\sqrt{2} + \mu_2(c)\sqrt{5} \tag{9}$$

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* $\Delta_{1,d} = \{\mu \in \Delta^{full} : \mu \text{ satisfies } (8)\}$ *and* $\Delta_{2,d} = \{\mu \in \Delta^{full} : \mu \text{ satisfies } (9)\}$.

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*

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

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*

[9] It is easy to check that if the agents' priors are the same $\mu_1 = \mu_2$, then y is ex ante efficient regardless of the values of μ_1 and μ_2.

payoff is

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

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

$$\sqrt{5} \cdot \mu(a) + \sqrt{3.5} \cdot \mu(b) \geq \sqrt{6.5} \cdot \mu(a) + \sqrt{5 - D_1} \cdot \mu(b), \tag{10}$$

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

$$\sqrt{5} \cdot \mu(a) + \sqrt{3.5} \cdot \mu(c) \geq \sqrt{6.5} \cdot \mu(a) + \sqrt{5 - D_2} \cdot \mu(c), \tag{11}$$

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 $\{b\}$. However, regardless of the values of $\sqrt{5 - D_1} \leq \sqrt{5}$ and $\sqrt{5 - D_2} \leq \sqrt{5}$, (10) and (11) 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.

*As a special case, if we let $-D_i$ equal to the worst possible transfer of agent i, i.e., $-D_i = \min_{\omega \in \Omega} y_i(\omega) - e_i(\omega) = -1.5$, $i = 1, 2$, then it is clear from (10) and (11) that the allocation y is not Bayesian incentive compatible under **any** belief in Δ^{full}.*

Since y is not Bayesian incentive compatible under all beliefs in $\Delta_{i,d}$ for each i, the Designer and the agents cannot be sure about the incentive compatibility of y. That is, an agent may be cheated and lose her wealth without understanding it. For example, when state a occurs, agent 2 tells agent 1 that the state is b. Now, agent 1 is in the event $\{a,b\}$, and she loses 1.5 units of the good in state a without understanding it. Thus, 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 agents end up consuming their initial endowment e which provides no insurance against low endowment realizations.

4.2 Bayesian persuasion

An agent's belief plays an important role in the Bayesian incentive compatibility notion. An allocation that is not Bayesian incentive compatible may become Bayesian incentive compatible, if the agents' beliefs change. The Designer can alter the agents' beliefs by adopting a Bayesian persuasion device of Kamenica and Gentzkow (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. It follows that the outcome of the investigation (i.e., the outcome of a persuasion device) can change the beliefs of the agents. 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 $\omega \in \Omega$,

$$\mu_{i,m}(\omega) = \frac{q(m \mid \omega)\,\mu_i(\omega)}{\sum_{\omega'\in\Omega} q(m \mid \omega')\,\mu_i(\omega')} \tag{12}$$

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

Can the Designer make the allocation y of Example 1 incentive compatible by adopting a Bayesian persuasion device? Since the Designer does not know the values of $\mu_i(\omega)$, $\omega \in \Omega$, he may not know the values of $\mu_{i,m}(\omega)$, $\omega \in \Omega$ either. Given an outcome m and the set Δ^{full}, let Δ^{full}_m be the set of posteriors

$$\Delta^{full}_m = \left\{\mu_m \mid \mu \in \Delta^{full}\right\} \subseteq \Delta,$$

where for each $\omega \in \Omega$,

$$\mu_m(\omega) = \frac{q(m \mid \omega)\,\mu(\omega)}{\sum_{\omega'\in\Omega} q(m \mid \omega')\,\mu(\omega')}. \tag{13}$$

Then, upon observing m, the set of beliefs that the Designer cannot distinguish $\Delta_{i,m,d}$ is a subset of Δ^{full}_m. That is, the Designer thinks that any belief in $\Delta_{i,m,d}$ can be $\mu_{i,m}$. An allocation y is incentive compatible under a Bayesian persuasion device if and only if the allocation is Bayesian incentive compatible under all beliefs in $\Delta_{i,m,d}$, for each possible $m \in M$ and for each i, i.e., y is Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \cup_{m\in M}\Delta_{i,m,d}$, for each i. We show in Example 2 below that the only way for the Designer to make y incentive compatible by adopting a Bayesian persuasion device is to adopt a 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.

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 $\{b\}$.*

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 $\mu_{i,m}(a)$. That is, the Designer only knows that $\mu_{i,m}(a)$ is between zero and one, i.e., $0 < \mu_{i,m}(a) < 1$. Indeed, the Designer knows that μ_i is from the set Δ^{full} and μ_i satisfies the individually rational constraint (i.e., (8) for agent 1 and (9) 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

$$\mu_{i,m}(a) = \frac{q(m \mid a)\,\mu_i(a)}{q(m \mid a)\,\mu_i(a) + q(m \mid b)\,\mu_i(b) + q(m \mid c)\,\mu_i(c)} \tag{14}$$

that $\mu_{i,m}(a)$ can take any value between zero and one as well, i.e., $0 < \mu_{i,m}(a) < 1$. 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 $m \in M$, $q(m \mid \omega) = 0$ 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 $q(m \mid a) \neq 0$. 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 $\omega = b, c$. That is, we have that $\mu_m(a) = 1$, $\mu_m(b) = 0$, and $\mu_m(c) = 0$. 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.

Case 2: $q(m \mid b) = 0$, $q(m \mid c) \neq 0$. By (13), 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: $q(m \mid b) \neq 0$, $q(m \mid c) = 0$. By (13), we have that $\mu_m(a) \neq 0$, $\mu_m(b) \neq 0$, and $\mu_m(c) = 0$, 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.

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.

Remark 2: *In Example 2, the Designer adopts a Bayesian persuasion device of Kamenica and Gentzkow (2011). The message/outcome of the persuasion device m is observed by all the agents. Alternatively, the Designer can adopt a persuasion device that sends separate messages to each agent. Now, there is a finite set of outcomes M_i for each agent i, and a family of probability distributions $\{q(\cdot \mid \omega)\}_{\omega \in \Omega}$ over $\times_{i=1}^{N} M_i$, i.e., for each ω, there is a probability distribution $q(\cdot \mid \omega)$ over $\times_{i=1}^{N} M_i$. When a state ω is realized, a profile of messages/outcomes $(m_1, \cdots, m_N)$ is generated according to $q(\cdot \mid \omega)$. Then, each agent i privately observes message/outcome m_i. Upon observing m_i, agent i forms a posterior belief $\mu_{i,m_i}(\omega)$ for each ω based on (12), except now we use $q(m_i \mid \omega')$ instead of $q(m \mid \omega')$, and we get $q(m_i \mid \omega')$ from the probability distribution $q(m_1, \cdots, m_N \mid \omega')$. Now, to insure truth telling of all the agents in Example 2, the Designer needs the persuasion device to have God's type power in state a. That is, when state a is realized, the persuasion device needs to let every agent know that the state is a, otherwise some agent may have an incentive to lie. Indeed, suppose that the realized state is a. Also, suppose that agent 1 privately observes m_1 and agent 2 privately observes m_2. Clearly, $q(m_1, m_2 \mid a)$ is not zero. It follows from (12) that $\mu_{i,m_i}(a)$ is not zero for both agents. If $\mu_{i,m_i}(a)$ is not one either, then the Designer does not know the value of $\mu_{i,m_i}(a)$, except that $0 < \mu_{i,m_i}(a) < 1$. Indeed, since $\mu_i(a)$ can take any value between zero and one $0 < \mu_i(a) < 1$, then according to (12), $\mu_{i,m_i}(a)$ can take any value between zero and one $0 < \mu_{i,m_i}(a) < 1$. We know from Example 1 that agent i lies whenever the probability of state a is high enough. Thus, to insure truth telling, it is necessary to have $\mu_{i,m_i}(a) = 1$ for all i. In other words, when state a is realized, to insure truth* *telling of all the agents, the persuasion device needs to let every agent know that state a is realized. This information requirement may not be realistic.*

4.3 Ambiguous persuasion

We show in Example 3 below that if the Designer adopts an ambiguous persuasion device to convince/persuade the agents to use the Wald's maxmin preferences, then the allocation y of Example 1 becomes maxmin incentive compatible.

An ambiguous persuasion device of Beauchêne, Li, and Li (2019) 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. Let q denote a family of probability distributions $\{q(\cdot \mid \omega)\}_{\omega \in \Omega}$ over M. Following Beauchêne, Li, and Li (2019), let Q denote a finite set of families of distributions over M, i.e., $Q = \{q^1, \cdots, q^K\}$, and let $co(Q)$ denote the convex hull of Q.[10] Formally, an ambiguous persuasion device is $\langle M, co(Q) \rangle$. Now, upon observing an outcome $m \in M$, agent i forms a posterior belief over Ω according to (12), based on each family of probability distributions q^k in Q, where $k = 1, \cdots, K$. Let $\mu^k_{i,m}$ denote the posterior belief over Ω associated with the outcome m that is induced by q^k in Q. That is, for each $\omega \in \Omega$,

$$\mu^k_{i,m}(\omega) = \frac{q^k(m \mid \omega)\, \mu_i(\omega)}{\sum_{\omega' \in \Omega} q^k(m \mid \omega')\, \mu_i(\omega')}.$$

Let $\mathbb{M}_{i,m}$ denote the set of posterior beliefs, when the outcome is m, i.e.,

$$\mathbb{M}_{i,m} = \left\{\mu^k_{i,m} : q^k \in Q\right\}.$$

Then, the agents' multi-belief set is the convex hull of $\mathbb{M}_{i,m}$, denoted by $co(\mathbb{M}_{i,m})$.[11] If Q is chosen such that, regardless of m, the set $\{\mu_{i,m}(\cdot \mid E_i) : \mu_{i,m} \in co(\mathbb{M}_{i,m})\}$ contains all probability distributions over E_i whenever E_i is a non-singleton set, then agent i has the Wald's maxmin preferences.

Example 3: *The economy is the same as in Example 1. We show below that if the Designer adopts an ambiguous persuasion device to convince/persuade the agents to use the Wald's maxmin preferences, then the allocation y (i.e., (7) above) of Example 1 is maxmin incentive compatible. The ambiguous persuasion*

[10] Beauchêne, Li, and Li (2019) pointed out that considering convex hull of the probabilities is by convention of the maxmin model, as only the convex hull of beliefs can be identified.

[11] "Maximizing her expected utility by taking into account the worst probability in $\mathbb{M}_{i,m}$" is equivalent to "maximizing her expected utility by taking into account the worst probability in $co(\mathbb{M}_{i,m})$".

device consists of a finite set of outcomes M and a collection of families of distributions over M, denoted by $co(Q^*)$. Suppose that $M = \{a,\ not\ a\}$, furthermore the following two families of distributions over M, denoted by q and q', belong to the collection of families of distributions $co(Q^*)$:

$$q(m = \ a \ | \omega = a) = 1; \quad q(m = \ not\ a \ | \omega = b) = 1; \quad q(m = \ not\ a \ | \omega = c) = 1,$$

$$q(m = \ not\ a \ | \omega = a) = 0; \quad q(m = \ a \ | \omega = b) = 0; \quad q(m = \ a \ | \omega = c) = 0,$$

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

$$q'(m = not\ a \ | \omega = a) = 1; \quad q'(m = \ a \ | \omega = b) = 1; \quad q'(m = \ a \ | \omega = c) = 1,$$

$$q'(m = \ a \ | \omega = a) = 0; \quad q'(m = not\ a \ | \omega = b) = 0; \quad q'(m = not\ a \ | \omega = c) = 0,$$

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. The agents observe the outcome $m \in M$ of the persuasion device. Then, each agent i updates μ_i 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 $\{\mu_{i,a}, \mu'_{i,a}\}$. Agent i gets the posterior $\mu_{i,a}$ by updating μ_i based on q, i.e.,

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

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

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

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

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

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

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 when agent 1 observes* $\{a, b\}$, *her multi-belief set contains all probability distributions over* $\{a, b\}$. *Indeed, if agent 1 observes the outcome "*$m = a$*", her multi-belief set is the convex hull of* $\left\{\mu_{1,a}, \mu'_{1,a}\right\}$. *Take* μ_a *from the convex hull of* $\left\{\mu_{1,a}, \mu'_{1,a}\right\}$, *i.e.,* $\mu_a = \alpha\mu_{1,a} + (1-\alpha)\mu'_{1,a}$, *where* $0 \le \alpha \le 1$. *Then her conditional probability* $\mu_a(\cdot \mid \{a, b\})$ *is*

$$\mu_a(\omega \mid \{a,b\}) = \begin{cases} \frac{\mu_a(a)}{\mu_a(a)+\mu_a(b)} = \frac{\alpha}{\alpha + (1-\alpha)\frac{\mu_1(b)}{\mu_1(b)+\mu_1(c)}} & \text{if } \omega = a \\ \frac{\mu_a(b)}{\mu_a(a)+\mu_a(b)} = \frac{(1-\alpha)\frac{\mu_1(b)}{\mu_1(b)+\mu_1(c)}}{\alpha + (1-\alpha)\frac{\mu_1(b)}{\mu_1(b)+\mu_1(c)}} & \text{if } \omega = b \\ 0 & \text{if } \omega \notin \{a,b\}, \end{cases} \tag{15}$$

where $0 \le \alpha \le 1$. *Clearly, when* $\alpha = 0$, *we have* $\mu_a(a \mid \{a,b\}) = 0$; *when* $\alpha = 1$, *we have* $\mu_a(a \mid \{a,b\}) = 1$; *and when* $0 < \alpha < 1$, *we have* $0 < \mu_a(a \mid \{a,b\}) < 1$. *That is,* $\mu_a(a \mid \{a,b\})$ *can take any value between and include zero and one. We can conclude that when agent 1 observes "*$m = a$*" and* $\{a, b\}$, *her multi-belief set contains all probability distributions over* $\{a, b\}$. *Similarly, if agent 1 observes the message "*$m = \text{not } a$*" and the event* $\{a, b\}$, *her multi-belief set contains all probability distributions over* $\{a, b\}$. *Thus, 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. That is, 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.*

We show below that when both agents have the Wald's maxmin preferences, the allocation y *becomes maxmin incentive compatible.*[12] *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*

[12] Given the allocation y, for every $\hat{E}_i \in \mathcal{F}_i$, let

$$y_i^{\hat{E}_i}(\omega) = e_i(\omega) + t_i\left(\hat{E}_i, E_{-i}(\omega)\right) = \begin{cases} e_i(\omega) + x_i(\hat{\omega}) - e_i(\hat{\omega}) & \text{if } \hat{E}_i \cap E_{-i}(\omega) = \{\hat{\omega}\} \\ e_i(\omega) + [\min_{\omega \in \Omega} y_i(\omega) - e_i(\omega)] & \text{if } \hat{E}_i \cap E_{-i}(\omega) = \emptyset. \end{cases}$$

That is, when the agents' reports are incompatible, we let the transfer of agent i equal to the worst possible transfer of agent i, i.e., $-D_i = \min_{\omega \in \Omega} y_i(\omega) - e_i(\omega) = -1.5$.

realized state could be a or b, her interim Wald's maxmin payoff is

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

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. Indeed, 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. Clearly, the allocation y in equation (7) above is incentive compatible under the Wald's maxmin preferences (i.e., maxmin incentive compatible), contrary to the Bayesian case in Sections 4.1 and 4.2.

5 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 Δ^{full}, one can focus on allocations that are Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i. Indeed, the Designer is certain that an allocation x is incentive compatible, if and only if x is Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i. Another choice is to focus on allocations that are ex post incentive compatible.[13] Ex post incentive compatibility was discussed as "uniform incentive compatibility" by Holmström and Myerson (1983). It turns out that if an allocation is ex post incentive compatible, then it is Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i. That is, the set of ex post incentive compatible allocations is smaller than the set of allocations that are Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i.[14] In Subsection 5.1, we present a better choice, i.e., adopting an ambiguous persuasion device to convince/persuade the agents to use the Wald's maxmin

[13] An allocation x is ex post incentive compatible, if for each i, and for each ω,

$$u_i\left(x_i\left(\omega\right), \omega\right) \geq u_i\left(x_i^{\hat{E}_i}\left(\omega\right), \omega\right),$$

for all $\hat{E}_i \in \mathcal{F}_i$. That is, it requires each agent i to prefer truth telling at each state ω if all the other agents also report truthfully.

[14] Focusing on dominant strategy incentive compatible is also a good choice. Compared to Bayesian incentive compatible, the dominant strategy incentive compatible notion is arguably a rather strong notion. However, a recent work by Kushnir and Liu (2019) extends the equivalence between Bayesian incentive compatible and dominant strategy incentive compatible to environments with nonlinear utilities satisfying a property of increasing differences over distributions and a convex-valued assumption.

preferences. Even though the Designer thinks that agent i's prior μ_i can be any belief in Δ^{full}, he can adopt an ambiguous persuasion device to change the agents' preferences. This change in preferences enlarges the set of incentive compatible allocations. In particular, the set of maxmin incentive compatible allocations contains the set of allocations that are Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i as a strict subset. Consequently, the set of maxmin incentive compatible allocations contains the set of ex post incentive compatible allocations as a strict subset. Moreover, more ex post efficient and individually rational allocations are incentive compatible under the Wald's maxmin preferences.

We show in Subsection 5.2 that the results of Subsection 5.1 hold, even if the Designer can rule out impossible beliefs based on the maxmin individually rationality of an allocation, i.e., $\Delta_{i,d}$ is a strict subset of Δ^{full}. Furthermore, we show in Subsection 5.3 that the results of Subsection 5.1 hold, when we take into account randomized action choices.

5.1 Incentive compatibility when $\Delta_{i,d} = \Delta^{full}$

We show in Example 4 below that an ex post efficient allocation x can be Bayesian individually rational under all beliefs in Δ^{full}. Thus, the Designer thinks that any belief in $\Delta_{i,d} = \Delta^{full}$ can be agent i's prior belief, for each i. Furthermore, the allocation is not Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i. However, the same allocation is both individually rational and incentive compatible under the Wald's maxmin preferences, i.e., x is both maxmin individually rational and maxmin incentive compatible.

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

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

For example, if state a occurs, agent 1 observes the event $\{a, b\}$ which is her private information in the interim. At the same time, agent 2 observes the event $\{a, c\}$ which is his private information in the interim. Agent 1's random initial endowment is $e_1(\omega) = (8, 0)$ for all $\omega \in \Omega$ and agent 2's random initial endowment is $e_2(\omega) = (0, 8)$ for all $\omega \in \Omega$. That is, agent 1 is endowed with 8 units of good 1 in each state and agent 2 is endowed with 8 units of good 2 in each state. Let c_i^1 and c_i^2 denote agent i's consumption of good 1 and 2 respectively. The ex post utility function of agent 1 is $u_1\left(c_1^1, c_1^2, \omega\right) = \sqrt{2.1}\sqrt{c_1^1} + \sqrt{2.1}\sqrt{c_1^2}$, $\omega = a, b$, and $u_1\left(c_1^1, c_1^2, \omega\right) = \sqrt{c_1^1} + \sqrt{3}\sqrt{c_1^2}$, $\omega = c, d$. The ex post utility function of agent 2 is $u_2\left(c_2^1, c_2^2, \omega\right) = c_2^1 + c_2^2$, $\omega = a, c$, and $u_2\left(c_2^1, c_2^2, \omega\right) = \sqrt{1.6}\sqrt{c_2^1} + \sqrt{2}\sqrt{c_2^2}$, $\omega = b, d$.

Let allocation x be

$$\left(x_1(a), x_1(b), x_1(c), x_1(d)\right) = \left((4,4), (4.54054, 4.09756), (2,6), (3.07692, 4.8)\right);$$

$$\left(x_2(a), x_2(b), x_2(c), x_2(d)\right) = \left((4,4), (3.45946, 3.90244), (6,2), (4.92308, 3.2)\right).$$

That is, in state a, agent 1 consumes 4 units of each good, etc.

The allocation x is feasible and ex post efficient. Comparing with the random initial endowment e, the allocation x makes every agent better off in each state. Indeed, take agent 1 and state a as an example, the random initial endowment gives her a payoff of $\sqrt{2.1}\sqrt{8} = 4.0988$, whereas the allocation x gives her a strictly higher payoff of $2\sqrt{2.1}\sqrt{4} = 5.7966$. It follows that x is Bayesian individually rational under all beliefs in Δ^{full}. Furthermore, x is maxmin individually rational.

However, x is not Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i. Let agent 1's prior belief be

$$\mu_1(a) = \mu_1(b) = 0.25; \quad \mu_1(c) = 0.1; \quad \mu_1(d) = 0.4.$$

The allocation x is not Bayesian incentive compatible. Indeed, for agent 1 and event $\{c,d\}$, she gets

$$\left(\sqrt{2} + \sqrt{3}\sqrt{6}\right)\frac{\mu_1(c)}{\mu_1(c)+\mu_1(d)} + \left(\sqrt{3.07692} + \sqrt{3}\sqrt{4.8}\right)\frac{\mu_1(d)}{\mu_1(c)+\mu_1(d)}$$

$$= \left(\sqrt{2} + \sqrt{3}\sqrt{6}\right)\frac{0.1}{0.5} + \left(\sqrt{3.07692} + \sqrt{3}\sqrt{4.8}\right)\frac{0.4}{0.5} = 5.5705$$

by reporting the true event $\{c,d\}$. However, she gets a strictly higher payoff

$$\left(\sqrt{4} + \sqrt{3}\sqrt{4}\right)\frac{\mu_1(c)}{\mu_1(c)+\mu_1(d)} + \left(\sqrt{4.54054} + \sqrt{3}\sqrt{4.09756}\right)\frac{\mu_1(d)}{\mu_1(c)+\mu_1(d)}$$

$$= \left(\sqrt{4} + \sqrt{3}\sqrt{4}\right)\frac{0.1}{0.5} + \left(\sqrt{4.54054} + \sqrt{3}\sqrt{4.09756}\right)\frac{0.4}{0.5} = 5.6024$$

by reporting the lie $\{a,b\}$. That is, x is not Bayesian incentive compatible.

However, the allocation x is maxmin incentive compatible. Take the same agent and the same event as above. Upon observing the event $\{c,d\}$, agent 1 gets

$$\min\left\{\sqrt{2} + \sqrt{3}\sqrt{6}, \sqrt{3.07692} + \sqrt{3}\sqrt{4.8}\right\} = 5.5488$$

by reporting the true event $\{c,d\}$. She gets

$$\min\left\{\sqrt{4}+\sqrt{3}\sqrt{4},\sqrt{4.54054}+\sqrt{3}\sqrt{4.09756}\right\}=5.4641$$

by reporting the lie $\{a,b\}$. Clearly, reporting the true event $\{c,d\}$ is strictly better. It can be checked that both agents strictly prefer to report the true events under the Wald's maxmin preferences. Thus, x is maxmin incentive compatible.

In fact, 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 in $\Delta_{i,d}=\Delta^{full}$ for each i as strict subsets (Theorem 1 below).

Denote by X_P *the set of ex post efficient allocations that are ex post incentive compatible.* Also, denote by X_B *the set of ex post efficient allocations that are Bayesian incentive compatible under all beliefs in* $\Delta_{i,d}=\Delta^{full}$ *for each* i. That is, every allocation x in X_B satisfies Definition 3 under all beliefs in $\Delta_{i,d}=\Delta^{full}$ for each i and Definition 1. Clearly, we have that X_P is a subset of X_B, i.e., $X_P\subseteq X_B$.

If the Designer adopts an ambiguous persuasion device $\langle M, co\,(Q)\rangle$ to convince/persuade the agents to use the Wald's maxmin preferences, then each agent i maximizes her interim Wald's maxmin payoff after observing E_i. An allocation is incentive compatible if it is maxmin incentive compatible. Denote by X_A *the set of ex post efficient and maxmin incentive compatible allocations.* That is, every allocation x in X_A satisfies Definition 4 and Definition 1.

Theorem 1 below shows that every ex post efficient allocation that is Bayesian incentive compatible under all beliefs in $\Delta_{i,d}=\Delta^{full}$ for each i is maxmin incentive compatible, but the reverse is not true. That is, we have that X_B is a strict subset of X_A, denoted by $X_B\subset X_A$. Theorem 1 holds regardless of the definition of the transfer $t\left(\hat{E}_1,\cdots,\hat{E}_N\right)$ where $\cap_{i\in I}\hat{E}_i=\emptyset$. Thus, we have $X_P\subseteq X_B\subset X_A$.

Theorem 1: *If an ex post efficient allocation x is Bayesian incentive compatible under all beliefs in $\Delta_{i,d}=\Delta^{full}$ for each i, then x is maxmin incentive compatible. The reverse is not true. That is, X_B is a strict subset of X_A, $X_B\subset X_A$.*

From Examples 1, 3 and 4 above, we know that an ex post efficient and maxmin incentive compatible allocation may not be Bayesian incentive compatible under all beliefs in $\Delta_{i,d}=\Delta^{full}$ for each i. Thus, X_A is not a subset of X_B. Suppose that an ex post efficient allocation x is not maxmin incentive compatible. Then, there exists an agent, a state ω and a lie, such that reporting the lie gives the agent a strictly higher ex post payoff in the state ω than reporting the true event. In our economy, it follows that x cannot be Bayesian incentive compatible when the probability of ω is high enough. Thus, if an ex post efficient allocation is not maxmin incentive compatible, then it cannot be Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i. In other words, X_B is a subset of X_A. We can conclude that X_B is a strict subset of X_A. The formal proof of Theorem 1 is in the Appendix.

Remark 3: *From Example 4 and Theorem 1, we can conclude that more ex post efficient and individually rational allocations are incentive compatible under the Wald's maxmin preferences. Indeed, let x be an ex post efficient allocation that is Bayesian individually rational under all beliefs in Δ^{full}. Also, let x be maxmin individually rational. From Theorem 1, we know that if x is Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i, then x is maxmin incentive compatible. Furthermore, from Example 4, we know that such an allocation can be maxmin incentive compatible, without being Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i. Thus, Wald's maxmin preferences make more ex post efficient and individually rational allocations incentive compatible.*

Remark 4: *In general, a Bayesian persuasion device of Kamenica and Gentzkow (2011) has an outcome $m \in M$, such that $0 < q(m \mid \omega) \leq 1$ for all ω. Under such a Bayesian persuasion device, if the outcome is m, then the set of posteriors $\Delta_{i,m,d} = \Delta^{full}$ for each i. For an allocation to be incentive compatible under this Bayesian persuasion device, the allocation needs to be Bayesian incentive compatible under all beliefs in Δ^{full} for each i. Thus, by Theorem 1, we know that more ex post efficient allocations are incentive compatible under the ambiguous persuasion device $\langle M, co(Q)\rangle$ than under any Bayesian persuasion device that has an outcome m, such that $0 < q(m \mid \omega) \leq 1$ for all ω.*

5.2 Incentive compatibility when $\Delta_{i,d} \subset \Delta^{full}$

Recall that the designer does not know the values of $\mu_i(\omega)$, $\omega \in \Omega$, except that $\mu_i(\omega) > 0$ for each $\omega \in \Omega$. Hence, the Designer does not know Δ_i (i.e., (1)) either. Suppose that the agents want to end up with an ex post efficient allocation x. Then, the Designer knows that x must be Bayesian individually rational. Furthermore, suppose that the Designer learns from the agents that the allocation x is maxmin individually rational (Definition 2) too. Based on this information, the Designer forms the set $\Delta_{i,d}$ by ruling out impossible beliefs from Δ^{full}. That is, a belief $\mu \in \Delta^{full}$ is in the set $\Delta_{i,d}$ if and only if for agent i the allocation x is Bayesian individually rational under μ, and maxmin individually rational under $\mu(E_i)$, $E_i \in \mathcal{F}_i$. In other words, the Designer thinks that any belief in $\Delta_{i,d}$ can be agent i's prior belief μ_i. If

$\Delta_{i,d} = \Delta^{full}$, then we have Theorem 1. In this subsection, we focus on the case in which $\Delta_{i,d}$ is a strict subset of Δ^{full}. Clearly, when the set $\Delta_{i,d}$ becomes smaller, the set of allocations that are Bayesian incentive compatible under all beliefs in $\Delta_{i,d}$ becomes larger.

We assume that each agent knows her endowment and utility function in the interim. Moreover, e_i and u_i do not reveal more information than E_i. That is, we assume that both e_i and u_i are $\mathcal{F}_i$-measurable. Then, we have that $e_i(\cdot)$ is constant on each element of $\mathcal{F}_i$: $e_i(\omega) = e_i(\omega')$ whenever ω and ω' are in the same event, i.e., $E_i(\omega) = E_i(\omega')$. Also, given any $c_i \in \mathbb{R}^{\ell}_{+}$, whenever $E_i(\omega) = E_i(\omega')$, we have $u_i(c_i, \omega) = u_i(c_i, \omega')$. Assuming e_i and u_i to be $\mathcal{F}_i$-measurable is more general than being constant.

Lemma 1: *Suppose that e_i and u_i are $\mathcal{F}_i$-measurable for each i. If an ex post efficient allocation x is maxmin individually rational, then x is Bayesian individually rational under each belief in Δ_i (i.e., (1)) for each i.*

The intuition behind Lemma 1 is simple. When agent i's multi-belief set P_i equals to Δ_i, agent i has the Wald's maxmin preferences. The worst belief in P_i puts the whole weight to the worst state in each event E_i. Since e_i and u_i are $\mathcal{F}_i$-measurable for each i, the maxmin payoff of e_i is the same as the Bayesian payoff of e_i for all beliefs in Δ_i. Hence, if an allocation is maxmin individually rational, then this allocation is Bayesian individually rational under the worst belief in Δ_i. It follows that this allocation is Bayesian individually rational under all beliefs in Δ_i. The formal proof of Lemma 1 is in the Appendix.

Remark 5: *Recall that the Designer does not know Δ_i, but the Designer learns that the allocation x is maxmin individually rational. Then, by Lemma 1, the Designer knows that the set $\Delta_i \cap \Delta^{full}$ is contained in $\Delta_{i,d}$.*

As $\Delta_i \cap \Delta^{full}$ is a subset of $\Delta_{i,d}$, it is clear that if an allocation is not Bayesian incentive compatible under all beliefs in $\Delta_i \cap \Delta^{full}$ for each i, then it cannot be Bayesian incentive compatible under all beliefs in $\Delta_{i,d}$ for each i. Now, we compare "Bayesian incentive compatible under all beliefs in $\Delta_i \cap \Delta^{full}$ for each i" with "maxmin incentive compatibility". With the help of Example 5 below, we show that an ex post efficient allocation x may not be Bayesian incentive compatible under all beliefs in $\Delta_i \cap \Delta^{full}$ for each i, but it is maxmin incentive compatible.

Example 5: *There are two agents, two goods, and four possible states of nature $\Omega = \{a, b, c, d\}$. Each agent i has a partition of Ω, denoted by $\mathcal{F}_i$, where $i = 1, 2$,*

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

For example, if state a occurs, agent 1 observes the event $\{a, b\}$ which is her private information in the interim. At the same time, agent 2 observes the event $\{a, c\}$ which is his private information in the interim. Let the agents' random initial endowment e be

$$(e_1(a), e_1(b), e_1(c), e_1(d)) = ((8, 10), (8, 10), (6, 12), (6, 12));$$

$$(e_2(a), e_2(b), e_2(c), e_2(d)) = ((8, 10), (12, 6), (8, 10), (12, 6)).$$

That is, in state a, agent 1 receives 8 units of good 1 and 10 units of good 2, etc. Let c_i^1 and c_i^2 denote agent i's consumption of good 1 and 2 respectively. The ex post utility function of agent i is $u_i\left(c_i^1, c_i^2, \omega\right) = \sqrt{c_i^1} + \sqrt{c_i^2}$, for each $\omega \in \Omega$.

Let allocation x be

$$(x_1(a), x_1(b), x_1(c), x_1(d)) = ((8, 10), (10, 8), (7, 11), (9, 9));$$

$$(x_2(a), x_2(b), x_2(c), x_2(d)) = ((8, 10), (10, 8), (7, 11), (9, 9)).$$

That is, in state a, agent 1 consumes 8 units of good 1 and 10 units of good 2, etc. The allocation x is feasible and ex post efficient. Suppose that the agents want to end up with the allocation x. Then, the Designer knows that x is Bayesian individually rational. Also, suppose that the Designer learns from the agents that the allocation x is maxmin individually rational. Then, the set $\Delta_{i,d}$ consists of beliefs $\mu \in \Delta^{full}$ such that for agent i the allocation x is Bayesian individually rational under μ, and maxmin individually rational under $\mu(E_i)$, $E_i \in \mathcal{F}_i$.[15]

Suppose that each agent i's prior belief is

$$\mu_i(a) = 0.4; \quad \mu_i(b) = 0.1; \quad \mu_i(c) = 0.25; \quad \mu_i(d) = 0.25.$$

It can be checked that the allocation x is Bayesian individually rational under μ_i, $i = 1, 2$. Also, x is maxmin individually rational. Indeed, for agent 1, we have that

$$\min\{u_1(e_1(a), a), u_1(e_1(b), b)\}[\mu_1(a) + \mu_1(b)] + \min\{u_1(e_1(c), c), u_1(e_1(d), d)\}[\mu_1(c) + \mu_1(d)] =$$

$$\min\{5.9907, 5.9907\}\, 0.5 + \min\{5.9136, 5.9136\}\, 0.5 = 5.9521 <$$

[15] It can be checked that for agent 2, $\Delta_{2,d}$ is a strict subset of Δ^{full}.

$$\min\{u_1(x_1(a),a),u_1(x_1(b),b)\}[\mu_1(a)+\mu_1(b)]+\min\{u_1(x_1(c),c),u_1(x_1(d),d)\}[\mu_1(c)+\mu_1(d)]=$$

$$\min\{5.9907,5.9907\}\,0.5+\min\{5.9624,6\}\,0.5=5.9765.$$

The same holds for agent 2.[16] *Hence, x is maxmin individually rational. Clearly, μ_i is in $\Delta_{i,d}$, as the Designer cannot rule out μ_i based on what he knows.*

However, x is not Bayesian incentive compatible under μ_i. Indeed, for agent 1 and event $\{a,b\}$, she gets

$$\left(\sqrt{8}+\sqrt{10}\right)\frac{\mu_1(a)}{\mu_1(a)+\mu_1(b)}+\left(\sqrt{10}+\sqrt{8}\right)\frac{\mu_1(b)}{\mu_1(a)+\mu_1(b)}=\sqrt{8}+\sqrt{10}=5.9907$$

by reporting the true event $\{a,b\}$. However, she gets a strictly higher payoff

$$\left(\sqrt{9}+\sqrt{9}\right)\frac{\mu_1(a)}{\mu_1(a)+\mu_1(b)}+\left(\sqrt{11}+\sqrt{7}\right)\frac{\mu_1(b)}{\mu_1(a)+\mu_1(b)}=\left(\sqrt{9}+\sqrt{9}\right)\frac{0.4}{0.5}+\left(\sqrt{11}+\sqrt{7}\right)\frac{0.1}{0.5}=5.9925$$

by reporting the lie $\{c,d\}$. That is, x is not Bayesian incentive compatible.

The allocation x is maxmin incentive compatible. Take the same agent and the same event as above. Upon observing the event $\{a,b\}$, agent 1 gets

$$\min\left\{\sqrt{8}+\sqrt{10},\sqrt{10}+\sqrt{8}\right\}=\sqrt{8}+\sqrt{10}=5.9907$$

by reporting the true event $\{a,b\}$. She gets

$$\min\left\{\sqrt{9}+\sqrt{9},\sqrt{11}+\sqrt{7}\right\}=5.9624$$

by reporting the lie $\{c,d\}$. Clearly, reporting the true event $\{a,b\}$ is strictly better. It can be checked that both agents strictly prefer to report the true events under the Wald's maxmin preferences. That is, x is maxmin incentive compatible.

By the definition of Δ_i, we know that agent i's prior μ_i is in the set $\Delta_i\cap\Delta^{full}$. Hence, we showed that an ex post efficient allocation x may not be Bayesian incentive compatible under all beliefs in $\Delta_i\cap\Delta^{full}$ for

[16] For agent 2, we have that

$$\min\{u_2(e_2(a),a),u_2(e_2(c),c)\}[\mu_2(a)+\mu_2(c)]+\min\{u_2(e_2(b),b),u_2(e_2(d),d)\}[\mu_2(b)+\mu_2(d)]=$$
$$\min\{5.9907,5.9907\}\,0.65+\min\{5.9136,5.9136\}\,0.35=5.9637<$$
$$\min\{u_2(x_2(a),a),u_2(x_2(c),c)\}[\mu_2(a)+\mu_2(c)]+\min\{u_2(x_2(b),b),u_2(x_2(d),d)\}[\mu_2(b)+\mu_2(d)]=$$
$$\min\{5.9907,5.9624\}\,0.65+\min\{5.9907,6\}\,0.35=5.9723.$$

each i, but it is maxmin incentive compatible.

Moreover, if the Designer learns that $\Delta_{i,d}$ contains $\Delta_i \cap \Delta^{full}$ as a subset, then persuading the agents to use the Wald's maxmin preferences can enlarge the set of incentive compatible allocations. Formally,

Corollary 1: *If an ex post efficient allocation x is Bayesian incentive compatible under all beliefs in $\Delta_i \cap \Delta^{full}$ for each i, then x is maxmin incentive compatible. The reverse is not true.*[17]

In other words, Corollary 1 says that the set of maxmin incentive compatible allocations contains the set of allocations that are Bayesian incentive compatible under all beliefs in $\Delta_i \cap \Delta^{full}$ for each i as a strict subset.

Remark 6: *Since $\Delta_{i,d}$ contains $\Delta_i \cap \Delta^{full}$ as a subset by Lemma 1, then the set of allocations that are Bayesian incentive compatible under all beliefs in $\Delta_{i,d}$ for each i is a subset of the set of allocations that are Bayesian incentive compatible under all beliefs in $\Delta_i \cap \Delta^{full}$ for each i. Now, it follows from Corollary 1 that the set of maxmin incentive compatible allocations contains the set of allocations that are Bayesian incentive compatible under all beliefs in $\Delta_{i,d}$ for each i as a strict subset. Thus, changing to Wald's maxmin preferences makes the set of incentive compatible allocations larger.*

Remark 7: *In this subsection, we focus on the case in which $\Delta_{i,d}$ is a strict subset of Δ^{full}. From Example 5 and Corollary 1, we know that more ex post efficient and individually rational allocations are incentive compatible under the Wald's maxmin preferences.*

5.3 Randomization

As pointed out by Raiffa (1961), and rigorously showed by Saito (2015), Ke and Zhang (2020) and Liu, Song, and Yannelis (2020), an agent with the Wald's maxmin preferences may strictly prefer randomizing over her choices to not randomizing. This is because randomization may smooth out her payoff at every state of nature and strictly increase her interim Wald's maxmin payoff. Example 6 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. Furthermore, such a profitable unilateral deviation brings Pareto improvements to the agents.

[17] The first part of proof is the same as the proof of Theorem 1, except now many beliefs in $\Delta_i \cap \Delta^{full}$ coincide with the constructed belief $\mu^{k=K}$, where $\mu^{k=K}$ is defined in the proof of Theorem 1. Furthermore, Example 5 proves the second part of Corollary 1.

Example 6: *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* $i = 1, 2$:

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

For example, if state a *occurs, agent 1 observes the event* $\{a, b\}$ *which is her private information in the interim. At the same time, agent 2 observes the event* $\{a, d\}$ *which is his private information in the interim. 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. *Furthermore, the agents have the Wald's maxmin preferences. Let* x *be a feasible allocation:*

$$(x_1(a), x_1(b), x_1(c), x_1(d)) = (3, 2, 3, 2); \quad (x_2(a), x_2(b), x_2(c), x_2(d)) = (2, 3, 2, 3).$$

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. Indeed, let* $\alpha \in [0, 1]$ *be the probability of reporting the event* $\{a, b\}$ *and* $1 - \alpha$ *the probability of reporting the event* $\{c, d\}$. *Clearly,* $\alpha = 1$ *means that agent 1 reports the true event, and* $\alpha = 0$ *means that she reports a lie. 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*

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

Clearly, 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.*

Furthermore, we show that agent 1's profitable unilateral deviation brings Pareto improvements to the agents. Suppose that agent 1 reports the event $\{a, b\}$ *with probability* $\alpha = \frac{1}{2}$ *when she observes the event* $\{a, b\}$, *and she reports the event* $\{c, d\}$ *with probability one when she observes the event* $\{c, d\}$, *while agent 2 always reports the true events. We calculate the agents' interim Wald's maxmin payoffs below. We know from above that when agent 1 observes the event* $\{a, b\}$, *her interim Wald's maxmin payoff of using* $\alpha = \frac{1}{2}$ *is* $\frac{\sqrt{3}+\sqrt{2}}{2}$ *which is strictly higher than her interim Wald's maxmin payoff of* x, *i.e.,* $\sqrt{2}$. *When agent 1* *observes the event $\{c, d\}$, her interim Wald's maxmin payoff is the same as the interim Wald's maxmin payoff of x. For agent 2, when he observes the event $\{a, d\}$, his interim Wald's maxmin payoff is*

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

which is strictly higher than his interim Wald's maxmin payoff of x,

$$\min\left\{\sqrt{2}, \sqrt{3}\right\} = \sqrt{2}.$$

When agent 2 observes the event $\{b, c\}$, his interim Wald's maxmin payoff is

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

which is the same as his interim Wald's maxmin payoff of x,

$$\min\left\{\sqrt{3}, \sqrt{2}\right\} = \sqrt{2}.$$

We can conclude that agent 1's profitable unilateral deviation brings Pareto improvements to the agents.

Since randomization may increase agents' interim Wald's maxmin payoffs and bring Pareto improvements, we take into account that the agents may randomize over their choices. In particular, we adopt the mixed maxmin incentive compatible notion of Liu, Song, and Yannelis (2020) which is a stronger notion than maxmin incentive compatible (Definition 4).[18] An allocation is mixed maxmin incentive compatible if no agent can improve her interim Wald's maxmin payoff by lying about her observed event with a strictly positive probability. Nevertheless, the results of Section 4 and Subsections 5.1 and 5.2 hold under randomization. In particular, we show that the set of ex post efficient and Bayesian incentive compatible (under all beliefs in $\Delta_{i,d}$ for each i) allocations is a subset of the set of ex post efficient and mixed maxmin incentive compatible allocations.

Let α_i be a probability distribution over $\mathcal{F}_i$.

Definition 5: *An allocation x is mixed maxmin incentive compatible, if for each agent i, and for each*

[18] Liu, Song, and Yannelis (2020) showed that every mixed maxmin incentive compatible allocation is maxmin incentive compatible, but a maxmin incentive compatible allocation may not be mixed maxmin incentive compatible.

$E_i \in \mathcal{F}_i$,

$$\min_{\omega \in E_i} u_i\left(x_i\left(\omega\right), \omega\right) \geq \min_{\omega \in E_i} \sum_{\hat{E}_i \in \mathcal{F}_i} u_i\left(x_i^{\hat{E}_i}\left(\omega\right), \omega\right) \alpha_i\left(\hat{E}_i\right), \tag{16}$$

for all α_i.

Denote by X_{MA} *the set of ex post efficient and mixed maxmin incentive compatible allocations.* That is, every allocation x in X_{MA} satisfies Definition 5 and Definition 1.

Remark 8: *Clearly, every mixed maxmin incentive compatible allocation is maxmin incentive compatible. We know from Example 6 that an ex post efficient and maxmin incentive compatible allocation may not be mixed maxmin incentive compatible. Thus, we know that the set of ex post efficient and mixed maxmin incentive compatible allocations is a strict subset of the set of ex post efficient and maxmin incentive compatible allocations, i.e.,* X_{MA} *is a strict subset of* X_A*, denoted by* $X_{MA} \subset X_A$.

We show in Example 7 below that the allocation y of Example 1 (equation (7)) is not only maxmin incentive compatible (see Example 3 above), but also mixed maxmin incentive compatible. Thus, if the Designer adopts an ambiguous persuasion device to convince/persuade the agents to use the Wald's maxmin preferences as in Example 3, then given y, no agent has an incentive to lie.

Example 7: *Recall that the allocation* y *of Example 1 is not Bayesian incentive compatible under all beliefs in* $\Delta_{i,d} = \Delta^{full}$ *for each* i*. Continue with Example 3 above. 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* $\{a, b\}$*. 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.*

Remark 9: *An allocation x is mixed Bayesian incentive compatible, if for each agent i, and for each $E_i \in \mathcal{F}_i$,*

$$\sum_{\omega \in \Omega} u_i\left(x_i(\omega), \omega\right) \mu_i\left(\omega \mid E_i\right) \geq \sum_{\omega \in \Omega} \sum_{\hat{E}_i \in \mathcal{F}_i} u_i\left(x_i^{\hat{E}_i}(\omega), \omega\right) \alpha_i\left(\hat{E}_i\right) \mu_i\left(\omega \mid E_i\right), \tag{17}$$

for all α_i. Comparing with the Definition 3, we know that under the Bayesian preferences, an allocation x is Bayesian incentive compatible, if and only if it is mixed Bayesian incentive compatible.

We show below that the set X_B is a strict subset of the set X_{MA}.

Theorem 2: *If an ex post efficient allocation x is Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i, then x is mixed maxmin incentive compatible. The reverse is not true. That is, X_B is a strict subset of X_{MA}, $X_B \subset X_{MA}$.*

From Example 7 above, we know that the ex post efficient allocation y is mixed maxmin incentive compatible, but it is not Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i. Thus, X_{MA} is not a subset of X_B. Suppose that an ex post efficient allocation x is not mixed maxmin incentive compatible. Then, there exists an agent i, a state ω and a lie α_i, such that reporting α_i is strictly better than reporting the true event in the state ω. Since α_i is a probability distribution over $\mathcal{F}_i$, there must exist an event, such that reporting this event is strictly better than reporting the true event in the state ω. Then, it follows from the proof of Theorem 1 that X_B is a subset of X_{MA}. We can conclude that X_B is a strict subset of X_{MA}. The formal proof of Theorem 2 is in the Appendix.

Thus, even if the set of ex post efficient and mixed maxmin incentive compatible allocations (i.e., X_{MA}) is a strict subset of the set of ex post efficient and maxmin incentive compatible allocations (i.e., X_A), the set X_{MA} is still larger than the set of ex post efficient allocations that are Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i. (i.e., X_B). We can conclude that $X_P \subseteq X_B \subset X_{MA} \subset X_A$.

Furthermore, by Example 5 and the proofs of Theorem 2 and Corollary 1, we have the following corollary.[19]

Corollary 2: *If an ex post efficient allocation x is Bayesian incentive compatible under all beliefs in $\Delta_i \cap \Delta^{full}$ for each i, then x is mixed maxmin incentive compatible. However, the reverse is not true.*

[19] It can be checked that the allocation x of Example 5 is mixed maxmin incentive compatible.

6 Concluding remarks

We show that more efficient and individually rational allocations become incentive compatible if the Designer persuades the agents to use the Wald's maxmin preferences instead of the Bayesian preferences. In other words, the agents become better off under the Wald's maxmin preferences, as they can reach efficient and individually rational allocations that cannot be reached under the Bayesian preferences. Thus, the Wald's maxmin preferences provide superior outcomes for all agents. 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 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. Thus, we justify the use of the Wald's maxmin preferences by showing that the agents can be persuaded to use the Wald's maxmin preferences in order to enlarge the set of incentive compatible allocations.

7 Appendix

7.1 Proof of Theorem 1

Proof. We show that every allocation in X_B is an allocation in X_A. Therefore, X_B is a subset of X_A, i.e., $X_B \subseteq X_A$. Let x be an allocation from X_B. Then, x is ex post efficient. If x is not maxmin incentive compatible, then there must exist an agent i, an event E_i and a lie $\hat{E}_i \neq E_i$, such that

$$\min_{\omega \in E_i} u_i\left(x_i\left(\omega\right), \omega\right) < \min_{\omega \in E_i} u_i\left(e_i\left(\omega\right) + t_i\left(\hat{E}_i, E_{-i}\left(\omega\right)\right), \omega\right). \tag{18}$$

Now, let ω^* and ω^{**} in E_i be such that

$$u_i\left(x_i\left(\omega^*\right), \omega^*\right) = \min_{\omega \in E_i} u_i\left(x_i\left(\omega\right), \omega\right)$$

and

$$u_i\left(e_i\left(\omega^{**}\right) + t_i\left(\hat{E}_i, E_{-i}\left(\omega^{**}\right)\right), \omega^{**}\right) = \min_{\omega \in E_i} u_i\left(e_i\left(\omega\right) + t_i\left(\hat{E}_i, E_{-i}\left(\omega\right)\right), \omega\right).$$

Case one: if E_i is a singleton event, then $\omega^* = \omega^{**}$. By (18), we have

$$u_i\left(x_i\left(\omega^*\right), \omega^*\right) < u_i\left(e_i\left(\omega^*\right) + t_i\left(\hat{E}_i, E_{-i}\left(\omega^*\right)\right), \omega^*\right). \tag{19}$$

Then, x is not Bayesian incentive compatible regardless of the belief μ in Δ^{full}. This contradicts with the fact that x belongs to the set X_B.

Case two: if E_i is not a singleton event and $\omega^* = \omega^{**}$, then by (18) we have

$$u_i\left(x_i\left(\omega^*\right), \omega^*\right) < u_i\left(e_i\left(\omega^*\right) + t_i\left(\hat{E}_i, E_{-i}\left(\omega^*\right)\right), \omega^*\right). \tag{20}$$

We show below that the interim Bayesian expected utility is continuous in probability. Let $\Delta^{E_i} = \{\mu\left(\cdot \mid E_i\right) : \mu \in \Delta\}$ be the set of all conditional probabilities of agent i, when she observes E_i. Given a conditional probability $\mu^k\left(\cdot \mid E_i\right) \in \Delta^{E_i}$, the interim Bayesian expected utility is

$$EU^k\left(x_i \mid E_i\right) = \sum_{\omega \in \Omega} u_i\left(x_i\left(\omega\right), \omega\right) \mu^k\left(\omega \mid E_i\right). \tag{21}$$

Since the set of states of nature Ω is finite, for ease of notation, we also use $\mu\left(\cdot \mid E_i\right)$ to denote the vector

$(\mu(\omega \mid E_i))_{\omega\in\Omega}$, which is in $\mathbb{R}^{|\Omega|}$ and $|\Omega|$ is the cardinality of Ω. For every $\left\{\mu^k(\cdot \mid E_i)\right\}_{k=1}^{\infty} \to \mu(\cdot \mid E_i)$, it must be $\left\{\mu^k(\omega \mid E_i)\right\}_{k=1}^{\infty} \to \mu(\omega \mid E_i)$ for each ω in Ω. It follows that we have $\left\{EU^k(x_i \mid E_i)\right\}_{k=1}^{\infty} \to EU(x_i \mid E_i)$ by the additivity of the limit operation. Therefore, the interim Bayesian expected utility is continuous at $\mu(\cdot \mid E_i)$. Since $\mu(\cdot \mid E_i)$ is an arbitrary conditional probability in Δ^{E_i}, we have that EU is continuous on Δ^{E_i}.

Now, let a $\mu^k(\cdot \mid E_i)$ be such that $\mu^k(\omega^* \mid E_i) = \frac{k}{k+1}$, $\mu^k(\omega' \mid E_i) = \frac{1}{(k+1)(|E_i|-1)}$ for all $\omega' \in E_i$, $\omega' \neq \omega^*$, and $\mu^k(\omega' \mid E_i) = 0$ for all $\omega' \notin E_i$. Clearly, when k goes to infinity, then $\mu^k(\cdot \mid E_i)$ converges to the $\mu(\cdot \mid E_i)$ which has $\mu(\omega^* \mid E_i) = 1$. Since the interim Bayesian expected utility is continuous in probability, we have that as k goes to infinity,

$$\sum_{\omega\in\Omega} u_i\left(e_i(\omega) + t_i\left(\hat{E}_i, E_{-i}(\omega)\right), \omega\right)\mu^k(\omega \mid E_i)$$

converges to

$$u_i\left(e_i(\omega^*) + t_i\left(\hat{E}_i, E_{-i}(\omega^*)\right), \omega^*\right).$$

Also,

$$\sum_{\omega\in\Omega} u_i(x_i(\omega), \omega)\mu^k(\omega \mid E_i)$$

converges to

$$u_i(x_i(\omega^*), \omega^*).$$

Now, in view of (20), there exists an integer K, such that $\mu^{k=K}(\cdot \mid E_i)$ assigns sufficient weight to ω^* and

$$\sum_{\omega\in\Omega} u_i\left(e_i(\omega) + t_i\left(\hat{E}_i, E_{-i}(\omega)\right), \omega\right)\mu^{k=K}(\omega \mid E_i) > \sum_{\omega\in\Omega} u_i(x_i(\omega), \omega)\mu^{k=K}(\omega \mid E_i). \tag{22}$$

That is, there exists a μ in Δ^{full} under which x is not Bayesian incentive compatible. This contradicts with the fact that x belongs to the set X_B. Thus, we can conclude that x must be maxmin incentive compatible.

<u>Case three:</u> if $\omega^* \neq \omega^{**}$, then it must be

$$u_i(x_i(\omega^*), \omega^*) < u_i\left(e_i(\omega^{**}) + t_i\left(\hat{E}_i, E_{-i}(\omega^{**})\right), \omega^{**}\right)$$

$$\leq u_i\left(e_i(\omega^*) + t_i\left(\hat{E}_i, E_{-i}(\omega^*)\right), \omega^*\right).$$

One can proceed as in Case two to obtain a contradiction.

We can conclude that if an allocation x is Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$

for each i, then x is maxmin incentive compatible. That is, X_B is a subset of X_A, i.e., $X_B \subseteq X_A$.

Furthermore, the allocation y of Example 1 (equation (7)) is ex post efficient. Examples 1 and 3 above showed that y is maxmin incentive compatible, but it is not Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i. Thus, the set of ex post efficient and maxmin incentive compatible allocations contains the set of ex post efficient allocations that are Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i as a strict subset, i.e., $X_B \subset X_A$. $\square$

7.2 Proof of Lemma 1

Proof. Suppose that x is not Bayesian individually rational under each belief in Δ_i for each i. Then, there exists an agent i and a belief μ in Δ_i, such that

$$\sum_{\omega\in\Omega} u_i\left(e_i\left(\omega\right),\omega\right)\mu\left(\omega\right) > \sum_{\omega\in\Omega} u_i\left(x_i\left(\omega\right),\omega\right)\mu\left(\omega\right). \tag{23}$$

From (1), we know that $\pi_i\left(E_i\right) = \mu\left(E_i\right)$ for each $E_i \in \mathcal{F}_i$. Thus, we have

$$\sum_{\omega\in\Omega} u_i\left(e_i\left(\omega\right),\omega\right)\mu\left(\omega\right) = \sum_{E_i\in\mathcal{F}_i}\left(\sum_{\omega\in E_i} u_i\left(e_i\left(\omega\right),\omega\right)\frac{\mu\left(\omega\right)}{\mu\left(E_i\right)}\right)\pi_i\left(E_i\right)$$

and

$$\sum_{\omega\in\Omega} u_i\left(x_i\left(\omega\right),\omega\right)\mu\left(\omega\right) = \sum_{E_i\in\mathcal{F}_i}\left(\sum_{\omega\in E_i} u_i\left(x_i\left(\omega\right),\omega\right)\frac{\mu\left(\omega\right)}{\mu\left(E_i\right)}\right)\pi_i\left(E_i\right).$$

Then, by (23) we have that

$$\sum_{E_i\in\mathcal{F}_i}\left(\sum_{\omega\in E_i} u_i\left(e_i\left(\omega\right),\omega\right)\frac{\mu\left(\omega\right)}{\mu\left(E_i\right)}\right)\pi_i\left(E_i\right) > \sum_{E_i\in\mathcal{F}_i}\left(\sum_{\omega\in E_i} u_i\left(x_i\left(\omega\right),\omega\right)\frac{\mu\left(\omega\right)}{\mu\left(E_i\right)}\right)\pi_i\left(E_i\right). \tag{24}$$

By the definition of a minimum, we have that

$$\sum_{E_i\in\mathcal{F}_i}\left(\sum_{\omega\in E_i} u_i\left(x_i\left(\omega\right),\omega\right)\frac{\mu\left(\omega\right)}{\mu\left(E_i\right)}\right)\pi_i\left(E_i\right) \geq \sum_{E_i\in\mathcal{F}_i}\left\{\min_{\omega\in E_i} u_i\left(x_i\left(\omega\right),\omega\right)\right\}\pi_i\left(E_i\right). \tag{25}$$

Furthermore, since e_i and u_i are $\mathcal{F}_i$-measurable for each i, we have that

$$\sum_{E_i\in\mathcal{F}_i}\left(\sum_{\omega\in E_i} u_i\left(e_i\left(\omega\right),\omega\right)\frac{\mu\left(\omega\right)}{\mu\left(E_i\right)}\right)\pi_i\left(E_i\right) = \sum_{E_i\in\mathcal{F}_i}\left\{\min_{\omega\in E_i} u_i\left(e_i\left(\omega\right),\omega\right)\right\}\pi_i\left(E_i\right). \tag{26}$$

Now, combining (24), (25) and (26), we have that

$$\sum_{E_i \in \mathcal{F}_i} \left\{ \min_{\omega \in E_i} u_i \left(e_i \left(\omega \right), \omega \right) \right\} \pi_i \left(E_i \right) > \sum_{E_i \in \mathcal{F}_i} \left\{ \min_{\omega \in E_i} u_i \left(x_i \left(\omega \right), \omega \right) \right\} \pi_i \left(E_i \right).$$

That is, the allocation x is not maxmin individually rational. This allows us to conclude that if an allocation x is maxmin individually rational under $P_i = \Delta_i$ for each i, then x is Bayesian individually rational under each belief in Δ_i for each i. □

7.3 Proof of Theorem 2

Proof. We show that every allocation in X_B is an allocation in X_{MA}. Therefore, X_B is a subset of X_{MA}, i.e., $X_B \subseteq X_{MA}$. Let x be an allocation from X_B. Then, x is ex post efficient. If x is not mixed maxmin incentive compatible, then there must exist an agent i, an event E_i and an α_i, such that

$$\min_{\omega \in E_i} \sum_{\hat{E}_i \in \mathcal{F}_i} u_i \left(e_i \left(\omega \right) + t_i \left(\hat{E}_i, E_{-i} \left(\omega \right) \right), \omega \right) \alpha_i \left(\hat{E}_i \right) > \min_{\omega \in E_i} u_i \left(x_i \left(\omega \right), \omega \right). \tag{27}$$

Now, let ω^* and ω^{**} in E_i be such that

$$u_i \left(x_i \left(\omega^* \right), \omega^* \right) = \min_{\omega \in E_i} u_i \left(x_i \left(\omega \right), \omega \right)$$

and

$$\sum_{\hat{E}_i \in \mathcal{F}_i} u_i \left(e_i \left(\omega^{**} \right) + t_i \left(\hat{E}_i, E_{-i} \left(\omega^{**} \right) \right), \omega^{**} \right) \alpha_i \left(\hat{E}_i \right) = \min_{\omega \in E_i} \sum_{\hat{E}_i \in \mathcal{F}_i} u_i \left(e_i \left(\omega \right) + t_i \left(\hat{E}_i, E_{-i} \left(\omega \right) \right), \omega \right) \alpha_i \left(\hat{E}_i \right).$$

Case one: if $\omega^* = \omega^{**}$, then by (27) we have

$$u_i \left(x_i \left(\omega^* \right), \omega^* \right) < \sum_{\hat{E}_i \in \mathcal{F}_i} u_i \left(e_i \left(\omega^* \right) + t_i \left(\hat{E}_i, E_{-i} \left(\omega^* \right) \right), \omega^* \right) \alpha_i \left(\hat{E}_i \right). \tag{28}$$

By definition, we have $0 \leq \alpha_i \left(\hat{E}_i \right) \leq 1$ and $\sum_{\hat{E}_i \in \mathcal{F}_i} \alpha_i \left(\hat{E}_i \right) = 1$. Then (28) implies that there exists an $\hat{E}_i$ in $\mathcal{F}_i$, such that

$$u_i \left(x_i \left(\omega^* \right), \omega^* \right) < u_i \left(e_i \left(\omega^* \right) + t_i \left(\hat{E}_i, E_{-i} \left(\omega^* \right) \right), \omega^* \right). \tag{29}$$

Now, it is the same as the Case one and Case two of Theorem 1: (29) contradicts with the fact that x belongs to the set X_B. Thus, we can conclude that x must be mixed maxmin incentive compatible.

Case two: if $\omega^* \neq \omega^{**}$, then it must be

$$u_i\left(x_i\left(\omega^*\right), \omega^*\right) < \sum_{\hat{E}_i \in \mathcal{F}_i} u_i\left(e_i\left(\omega^{**}\right) + t_i\left(\hat{E}_i, E_{-i}\left(\omega^{**}\right)\right), \omega^{**}\right) \alpha_i\left(\hat{E}_i\right)$$

$$\leq \sum_{\hat{E}_i \in \mathcal{F}_i} u_i\left(e_i\left(\omega^*\right) + t_i\left(\hat{E}_i, E_{-i}\left(\omega^*\right)\right), \omega^*\right) \alpha_i\left(\hat{E}_i\right).$$

Now, one can proceed as in Case one above to obtain a contradiction.

We can conclude that if an allocation x is Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i, then x is mixed maxmin incentive compatible. That is, X_B is a subset of X_{MA}, $X_B \subseteq X_{MA}$.

Furthermore, from Example 7, we know that the ex post efficient allocation y is not Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i, but it is mixed maxmin incentive compatible. Thus, the set of ex post efficient and mixed maxmin incentive compatible allocations contains the set of ex post efficient allocations that are Bayesian incentive compatible under all beliefs in $\Delta_{i,d} = \Delta^{full}$ for each i as a strict subset, i.e., $X_B \subset MX_A$. $\square$

References

A. Angelopoulos and L. C. Koutsougeras. Value allocation under ambiguity. *Economic Theory*, 59:147–167, 2015.

D. Beauchêne, J. Li, and M. Li. Ambiguous persuasion. *Journal of Economic Theory*, 179:312–365, 2019.

V. Carrasco, V. F. Luz, P. K. Monteiro, and H. Moreira. Robust mechanisms: the curvature case. *Economic Theory*, 68:203–222, 2019.

L. I. De Castro and N. C. Yannelis. Uncertainty, efficiency and incentive compatibility: Ambiguity solves the conflict between efficiency and incentive compatibility. *Journal of Economic Theory*, 177:678–707, 2018.

L. I. De Castro, M. Pesce, and N. C. Yannelis. Core and equilibria under ambiguity. *Economic Theory*, 48: 519–548, 2011.

L. I. De Castro, Z. Liu, and N. C. Yannelis. Implementation under ambiguity. *Games and Economic Behavior*, 101:20–33, 2017a.

L. I. De Castro, Z. Liu, and N. C. Yannelis. Ambiguous implementation: the partition model. *Economic Theory*, 63:233–261, 2017b.

L. I. De Castro, M. Pesce, and N. C. Yannelis. A new approach to the rational expectations equilibrium: existence, optimality and incentive compatibility. *Annals of Finance*, 16:1–61, 2020.

B. Decerf and F. Riedel. Purification and disambiguation of Ellsberg equilibria. *Economic Theory*, 69: 595–636, 2020.

A. Degan and M. Li. Persuasion with costly precision. *Economic Theory*, https://doi.org/10.1007/s00199-021-01346-9, 2021.

I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989.

D. Glycopantis and N. C. Yannelis. The maximin equilibrium and the PBE under ambiguity. *Economic Theory Bulletin*, 6:183–199, 2018.

D. Glycopantis, A. Muir, and N. C. Yannelis. An extensive form interpretation of the private core. *Economic Theory*, 18:293–319, 2001.

H. Guo and N. C. Yannelis. Incentive compatibility under ambiguity. *Economic Theory*, https://doi.org/10.1007/s00199-020-01304-x, 2020.

H. Guo and N. C. Yannelis. Full implementation under ambiguity. *American Economic Journal: Microeconomics*, 13:148–178, 2021.

W. He and N. C. Yannelis. Equilibrium theory under ambiguity. *Journal on Mathematical Economics*, 61: 86–95, 2015.

B. Holmström and R. B. Myerson. Efficient and durable decision rules with incomplete information. *Econometrica*, 51:1799–1819, 1983.

J. Hu and X. Weng. Robust persuasion of a privately informed receiver. *Economic Theory*, https://doi.org/10.1007/s00199-020-01299-5, 2020.

E. Kamenica and M. Gentzkow. Bayesian persuasion. *American Economic Review*, 101:2590–2615, 2011.

S. Ke and Q. Zhang. Randomization and ambiguity aversion. *Econometrica*, 88:1159–1195, 2020.

A. Kushnir and S. Liu. On the equivalence of Bayesian and dominant strategy implementation for environments with nonlinear utilities. *Economic Theory*, 67:617–644, 2019.

Z. Liu. Implementation of maximin rational expectations equilibrium. *Economic Theory*, 62:813–837, 2016.

Z. Liu and N. C. Yannelis. Implementing maxmin incentive compatible allocations. *working paper*, 2021.

Z. Liu, X. Song, and N. C. Yannelis. Randomization under ambiguity: Efficiency and incentive compatibility. *Journal of Mathematical Economics*, 90:1–11, 2020.

M. Lombardi and N. Yoshihara. Partially-honest Nash implementation: a full characterization. *Economic Theory*, 70:871–904, 2020.

E. Moreno-García and J. P. Torres-Martínez. Information within coalitions: risk and ambiguity. *Economic Theory*, 69:125–147, 2020.

K. Pram. Weak implementation. *Economic Theory*, 69:569–594, 2020.

CZ. Qin and X. Yang. On the equivalence of rational expectations equilibrium with perfect Bayesian equilibrium. *Economic Theory*, 69:1127–1146, 2020.

H. Raiffa. Risk, ambiguity, and the Savage axioms: Comment. *Quarterly Journal of Economics*, 75:690–694, 1961.

K. Saito. Preferences for flexibility and randomization under uncertainty. *American Economic Review*, 105: 1246–1271, 2015.

A. Zapechelnyuk and A. Kolotilin. Persuasion under insufficient reason. *working paper*, 2021.