Economic Theory Bulletin
https://doi.org/10.1007/s40505-021-00206-y

RESEARCH ARTICLE

Mixed strategy implementation under ambiguity

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

Received: 3 June 2021 / Accepted: 27 July 2021

Abstract
We extend the previous work of De Castro et al. (2017a, b) into mixed strategies.

Keywords Wald's maxmin preferences · Mixed strategy deviations · Implementation

JEL classification D51 · D81 · D82

1 Introduction

Under the Wald's maxmin preferences, Liu et al. (2020) showed that an incentive compatible allocation may not be mixed incentive compatible.[1] It follows that the set of incentive compatible allocations contains the set of mixed incentive compatible allocations as a strict subset. Furthermore, we know from Liu et al. (2020) that *efficient allocations* of De Castro and Yannelis (2018); De Castro et al. (2017a, b) and *interim maxmin value allocations* of Angelopoulos and Koutsougeras (2015) are not only incentive compatible, but also mixed incentive compatible under the Wald's maxmin preferences.[2] However, when we take into account that agents may randomize

We thank a referee for helpful questions, comments and suggestions. Zhiwei Liu gratefully acknowledges the financial support from the National Natural Science Foundation of China [Grant No. 71703110].

✉ Zhiwei Liu
liuzhiwei@cueb.edu.cn

Nicholas C. Yannelis
nicholasyannelis@gmail.com

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

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

[1] Mixed incentive compatibility means that no agent wants to unilaterally deviate from reporting the true event to a lottery. Hence, mixed incentive compatibility is stronger than incentive compatibility as the latter means that no agent wants to unilaterally deviate from reporting the true event to another event.

[2] Moreover, mixed incentive compatibility of a maxmin rational expectations equilibrium allocation follows from De Castro et al. (2020), Liu (2014) and Liu et al. (2020).

Published online: 06 September 2021

their choices, are these allocations implementable? What is the relationship between incentive compatibility and implementation?

When a Designer asks agents to report their privately observed events in a direct revelation mechanism, no agent knows the reports of other agents. Facing these uncertainties, agents may adopt a criterion *à la* Wald (1950). That is, each agent maximizes the worst case payoff, i.e., the payoff that takes into account not only the worst state that can occur, but also the worst strategy of all the other agents against her.

It turns out that if an allocation is incentive compatible and agents maximize their worst case payoffs, then agents may strictly prefer to misreport true events with a strictly positive probability (Example 1 below). Then, does mixed incentive compatibility imply that truth telling gives every agent the best worst payoff? In other words, if agents maximize their worst case payoffs, can they reach mixed incentive compatible allocations? We show that the answers are positive. In particular, we show that every mixed incentive compatible allocation is implementable as a maxmin equilibrium if and only if punishments *induce truth telling with respect to mixed strategy deviations*. That is, punishments must not make "the worst case payoff of truth telling" strictly less than "the worst case payoff of not being truthful with probability one". If an economy can be represented by the standard type model of the implementation literature, then we know that agents' reports are always compatible.[3] Thus, the Designer does not need to specify punishments. It follows that in such an economy, every mixed incentive compatible allocation is implementable as a maxmin equilibrium.

The paper is organized as follows. Sections 2 and 3 define an asymmetric information exchange economy, incentive compatibility and implementation. In Sect. 4, we show that each mixed incentive compatible allocation is implementable as a maxmin equilibrium. Finally, we conclude in Sect. 5. The proof of our result is collected in the Appendix.

2 Asymmetric information exchange economy

We consider an asymmetric information exchange economy in which agents have the Wald's maxmin preferences as in De Castro and Yannelis (2018). Let $\mathbb{R}_+^\ell$ denote the ℓ-goods commodity space and $I = \{1, \cdots, N\}$ the set of N agents. Let Ω be a finite set of states of nature and $\omega \in \Omega$ is a state of nature. Let $u_i : \mathbb{R}_+^\ell \times \Omega \to \mathbb{R}$ be agent i's ex post utility function, taking the form of $u_i(c_i, \omega)$ where $c_i \in \mathbb{R}_+^\ell$ is agent i's consumption. Agent i's random initial endowment is a mapping from the set of states of nature to the commodity space, i.e., $e_i : \Omega \to \mathbb{R}_+^\ell$. Agent i's allocation is a mapping from Ω to $\mathbb{R}_+^\ell$. Then, *an allocation* $x = (x_i)_{i \in I}$ is a mapping from Ω to $\mathbb{R}_+^{\ell \times N}$. We use L to denote the set of allocations.

Agents have asymmetric information with respect to the realized state of nature. That is, when a state of nature is realized, each agent observes an event that contains the realized state of nature. Formally, each agent i has a partition $\mathcal{F}_i$ of Ω. Each event E_i in the partition $\mathcal{F}_i$ represents a maximal set of states that agent i cannot distinguish. If ω is the realized state of nature, agent i observes the event $E_i(\omega)$, where $E_i(\omega)$

[3] De Castro et al. (2017a) adopts the type model.

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, $\bigcap_{j\in I} E_j(\omega) = \{\omega\}$.

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. However, if ω and ω' are in the same event $E_i \in \mathcal{F}_i$, then neither $\{\omega\}$ nor $\left\{\omega'\right\}$ is in the algebra $\sigma(\mathcal{F}_i)$. Hence, the probability $\pi_i(\{\omega\})$ or $\pi_i\left(\left\{\omega'\right\}\right)$ is not defined. Let Δ_i be the set of all probability measures over 2^{Ω} that agree with π_i. Formally,

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

We postulate that each agent i's preferences on L are *maxmin à la* Gilboa and Schmeidler (1989). Given an allocation x, agent i's ex ante expected utility is

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

where P_i is a non-empty, closed and convex subset of Δ_i, i.e., P_i is agent i's multi-belief set. 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, $x_i \succeq_i y_i$.[4] 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 the Wald's maxmin preferences in De Castro and Yannelis (2018) as special cases. Indeed, if $P_i = \Delta_i$, then the worst probability in the multi-belief set P_i should assign the whole weight to the worst state in each E_i. In this case, the multi-belief preferences become the Wald's maxmin preferences in De Castro and Yannelis (2018), where the following formulation is equivalent to (1),

$$\sum_{E_i\in\mathcal{F}_i} \left(\min_{\omega\in E_i} u_i(x_i(\omega), \omega)\right) \mu_i(E_i). \tag{2}$$

Then, when agent i observes an event E_i, she prefers x_i to y_i, if

$$\min_{\omega\in E_i} u_i(x_i(\omega), \omega) \geq \min_{\omega\in E_i} u_i(y_i(\omega), \omega). \tag{3}$$

In this paper, we focus on the Wald's maxmin preferences. The interest of the Wald's maxmin preferences comes from that only under these preferences any efficient allocation is incentive compatible (De Castro and Yannelis 2018). Furthermore, (Liu and Yannelis 2021) showed that even if agents start with (Gilboa and Schmeidler

[4] "Larger than" means "greater than or equal to".

1989) preferences, the agents can be persuaded to use the Wald's maxmin preferences in order to enlarge the set of efficient, individually rational and incentive compatible allocations.

An asymmetric information exchange economy $\mathcal{E}$ is the set $\mathcal{E} = \{\Omega, (u_i, e_i, \mathcal{F}_i, P_i) : i \in I\}$. In this economy, 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., $\sum_{i=1}^{N} x_i(\omega) = \sum_{i=1}^{N} e_i(\omega)$. Furthermore, 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.

3 Mixed maxmin incentive compatibility and implementation

Suppose that agents want to end up with a feasible allocation x that is different from their random initial endowment e. Then, transfers need to take place. Since both e and x depend on the state of nature ω, the transfers may depend on ω as well. Recall that the agents have asymmetric information with respect to the realized state of nature. Thus, it is necessary to pool their private information so that they may know the realized state of nature and end up with the correct transfers. We assume that each agent i decides what to report to the Designer after observing $E_i(\omega)$.

It turns out that a Wald's maxmin agent can get an averaged payoff in each state of nature by randomizing her choices. Thus, randomization may help her to avoid extremely low payoffs ((Raiffa 1961), (Saito 2015), (Ke and Zhang 2020)). Indeed, suppose that reporting the truth gives an agent a payoff of one in state a and zero in state b. Moreover, reporting a lie gives the agent a payoff of zero in state a and one in state b. Then, mixing between the truth and the lie in state a gives the agent an averaged payoff which is strictly higher than zero. The same holds for state b. Thus, a Wald's maxmin agent may strictly prefer to randomize her choices. Liu et al. (2020) showed that even if a Wald's maxmin agent has no incentive to unilaterally deviate from reporting the truth $E_i(\omega)$ to reporting a lie $\hat{E}_i \neq E_i(\omega)$, she may strictly prefer to unilaterally deviate from truth telling to a lottery over $\mathcal{F}_i$. The outcome/realization of the lottery is an event $\hat{E}_i$ in $\mathcal{F}_i$, thus it can be a lie. Furthermore, Liu and Yannelis (2021) showed that such a profitable unilateral deviation brings Pareto improvements to the agents. Thus, we take into account that a Wald's maxmin agent may flip a coin, roll dice or randomize in her mind to decide which event to report. In particular, upon observing $E_i(\omega)$, agent i may choose to use a lottery over $\mathcal{F}_i$ and report the outcome/realization of the lottery.

The Designer does not know the realized state of nature ω, nor the agents' privately observed events $E_i(\omega)$, $i \in I$. However, the Designer observes the agents' reports. Thus, the Designer sets the transfers based on the agents' reports.

Formally, let $\left(\hat{E}_1, \cdots, \hat{E}_N\right)$ denote a report profile, where $\hat{E}_i \in \mathcal{F}_i$ is agent i's report, $i \in \{1, \cdots, N\}$. By Assumption 1, we know that the agents' reports are either compatible (i.e., $\bigcap_{j\in I} \hat{E}_j$ is a singleton set) or incompatible (i.e., $\bigcap_{j\in I} \hat{E}_j$ is an empty set). Let $t\left(\hat{E}_1, \cdots, \hat{E}_N\right)$ denote the transfers among the agents, when every agent i reports $\hat{E}_i \in \mathcal{F}_i$. Let

$$t_i\left(\hat{E}_1, \cdots, \hat{E}_N\right) = \begin{cases} x_i\left(\hat{\omega}\right) - e_i\left(\hat{\omega}\right) \text{ if} & \bigcap_{j\in I} \hat{E}_j = \left\{\hat{\omega}\right\} \\ D_i \quad \text{if} & \bigcap_{j\in I} \hat{E}_j = \emptyset, \end{cases} \tag{4}$$

where $D_i \in \mathbb{R}^{\ell}$ denotes a punishment, i.e., it is the transfer of agent i when agents' reports are incompatible.[5] Clearly, if the realized state is ω and the agents report the true events, then $\hat{E}_i = E_i(\omega)$ for each i and $\cap_{i\in I} E_i(\omega) = \{\omega\}$. Thus, the transfers $t\left(\hat{E}_1, \cdots, \hat{E}_N\right) = x(\omega) - e(\omega)$ are correct and the agents reach $x(s)$ after these transfers, i.e., the agents end up with $e(\omega) + t\left(\hat{E}_1, \cdots, \hat{E}_N\right) = e(\omega) + x(\omega) - e(\omega) = x(\omega)$.[6] For simplicity, we use $E_{-i}(\omega)$ to denote $(E_1(\omega), \cdots, E_{i-1}(\omega), E_{i+1}(\omega), \cdots, E_N(\omega))$. Also, we use $\hat{E}_i \cap E_{-i}(\omega)$ to denote $\hat{E}_i \cap E_1(\omega) \cap \cdots \cap E_{i-1}(\omega) \cap E_{i+1}(\omega) \cap \cdots \cap E_N(\omega)$.

3.1 Mixed maxmin incentive compatibility

Let $x_i^{\hat{E}_i}(\omega)$ denote agent i's consumption, when the realized state is ω, her report is $\hat{E}_i$ and all other agents report truthfully. Thus,

$$\begin{aligned} x_i^{\hat{E}_i}(\omega) &= e_i(\omega) + t_i\left(\hat{E}_i, E_{-i}(\omega)\right) \\ &= \begin{cases} e_i(\omega) + x_i\left(\hat{\omega}\right) - e_i\left(\hat{\omega}\right) \text{ if} & \hat{E}_i \cap E_{-i}(\omega) = \left\{\hat{\omega}\right\} \\ e_i(\omega) + D_i \quad \text{if} & \hat{E}_i \cap E_{-i}(\omega) = \emptyset. \end{cases} \end{aligned} \tag{5}$$

An allocation is *maxmin incentive compatible* if no agent can improve her interim Wald's maxmin payoff by unilaterally deviating from reporting the true event to reporting another event. That is, an allocation x is maxmin incentive compatible, if for each agent i and for each $E_i \in \mathcal{F}_i$,

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

[5] Agents are punished if and only if their reports are incompatible. The Designer chooses the value of $D_i \in \mathbb{R}^{\ell}$. Some choices are no transfer (i.e., every agent keeps her endowment), imposing the worst possible transfer, or randomly assigning a transfer (see for example, Glycopantis et al. (2001), Liu (2016) and De Castro et al. (2020, 2017b)).

[6] As in De Castro et al. (2017a, b), Moreno-García and Torres-Martínez (2020) and Liu and Yannelis (2021), we impose the following condition: every transfer under x is feasible, i.e., $e_i(\omega) + x_i\left(\hat{\omega}\right) - e_i\left(\hat{\omega}\right) \in \mathbb{R}_+^{\ell}$, for each $i, \omega, \hat{\omega}$. Clearly, if each e_i is constant, then the feasibility condition above is automatically satisfied.

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

We say that an allocation is *mixed maxmin incentive compatible*, if no agent can improve her interim Wald's maxmin payoff by unilaterally deviating from reporting the true event. That is, no agent can become strictly better off by misreporting the true event with a strictly positive probability. Formally, let α_i be a probability distribution over $\mathcal{F}_i$ and $\alpha_i\left(\hat{E}_i\right)$ the probability of reporting the event $\hat{E}_i \in \mathcal{F}_i$. Now, agents face two types of uncertainties. These are uncertainty from the state of nature and uncertainty from α_i. Depending on an agent's subjective belief of which uncertainty resolves first, there are different ways to formulate the mixed maxmin incentive compatibility notion (Liu et al. 2020). We assume that each agent i believes that nature draws the state ω first, then agent i learns $E_i(\omega)$ and the realization $\hat{E}_i$ of the lottery α_i. Then, we have the following definition, which is the strongest mixed maxmin incentive compatibility notion in Liu et al. (2020).

Definition 1 An allocation x is mixed maxmin incentive compatible, if for each agent i and for each $E_i \in \mathcal{F}_i$,

$$\min_{\omega \in E_i} u_i\left(x_i(\omega), \omega\right) \geq \min_{\omega \in E_i} \sum_{\hat{E}_i \in \mathcal{F}_i} u_i\left(x_i^{\hat{E}_i}(\omega), \omega\right) \alpha_i\left(\hat{E}_i\right), \tag{7}$$

for all α_i.

Remark 1 Liu et al. (2020) showed that if an allocation satisfies Definition 1, then it satisfies every mixed maxmin incentive compatible notion in their paper. That is, if an allocation x satisfies Definition 1, then it is mixed maxmin incentive compatible under the Wald's maxmin preferences, regardless of agents' subjective beliefs on which uncertainty resolves first. Furthermore, they showed that Definition 1 is strictly stronger than the maxmin incentive compatible notion.

3.2 Implementation

When a Designer asks agents to report their privately observed events, no agent knows the reports of other agents. Facing the state of nature uncertainty and the strategic uncertainty, agents may adopt a criterion *à la* Wald (1950). That is, each agent maximizes her worst case payoff, i.e., the payoff that takes into account not only the worst state that can occur, but also the worst strategy of all the other agents against her.[7] We say that a feasible allocation x is implementable, if agents end up with x through maximizing their worst case payoffs.

Formally, a direct revelation mechanism associated with an allocation is a non-cooperative game. In this game, when state ω is realized, each agent i privately observes the event $E_i(\omega)$. Then, the agents report events simultaneously.[8]

Definition 2 A *mixed strategy* of agent i is a function $\sigma_i : \mathcal{F}_i \to \Delta\left(\mathcal{F}_i\right)$.

[7] Decerf and Riedel (2020) refer to this as maxmin strategies.

[8] We assume that an agent lies, only if she can benefit from doing so.

For every $E_i \in \mathcal{F}_i$, $\sigma_i(E_i)$ is a lottery over her partition $\mathcal{F}_i$. That is, after observing E_i, agent i reports the realization $\hat{E}_i$ of the lottery $\sigma_i(E_i)$. Let Σ_i denote agent i's strategy set. Denote by $\Sigma = \times_{i \in I} \Sigma_i$ the strategy set. Let $\sigma \in \Sigma$ denote a strategy profile, that is, $\sigma = (\sigma_1, \cdots, \sigma_N)$, then $\sigma(\omega) = (\sigma_1(E_1(\omega)), \cdots, \sigma_N(E_N(\omega)))$. As usual, let Σ_{-i} denote $\times_{j \neq i} \Sigma_j$ and let σ_{-i} denote $(\sigma_1, \cdots, \sigma_{i-1}, \sigma_{i+1}, \cdots, \sigma_N)$. Furthermore, let $\sigma_i(E_i)\left[\hat{E}_i\right]$ denote the probability of reporting $\hat{E}_i$ when agent i adopts the strategy σ_i and her observed event is E_i. If agents adopt the strategy profile σ and the state is ω, then the agents get the transfers $t\left(\hat{E}_1, \cdots, \hat{E}_N\right)$ with probability $\sigma_1(E_1(\omega))\left[\hat{E}_1\right] \times \cdots \times \sigma_N(E_N(\omega))\left[\hat{E}_N\right]$, for each $\left(\hat{E}_1, \cdots, \hat{E}_N\right) \in \times_{i \in I} \mathcal{F}_i$. In state ω, the *outcome* of adopting a strategy profile σ is $g(\sigma, \omega)$: the agents get $e(\omega) + t\left(\hat{E}_1, \cdots, \hat{E}_N\right)$ with probability $\sigma_1(E_1(\omega))\left[\hat{E}_1\right] \times \cdots \times \sigma_N(E_N(\omega))\left[\hat{E}_N\right]$, for each $\left(\hat{E}_1, \cdots, \hat{E}_N\right) \in \times_{i \in I} \mathcal{F}_i$. That is, in state ω, agent i's payoff is

$$\sum_{\hat{E}_i \in \mathcal{F}_i} \sum_{\hat{E}_{-i} \in \mathcal{F}_{-i}} u_i\left(e_i(\omega) + t_i\left(\hat{E}_i, \hat{E}_{-i}\right), \omega\right) \sigma_i(E_i(\omega))\left[\hat{E}_i\right] \sigma_{-i}(E_{-i}(\omega))\left[\hat{E}_{-i}\right],$$

where $\sigma_{-i}(E_{-i}(\omega))\left[\hat{E}_{-i}\right]$ denotes $\times_{j \neq i} \sigma_j\left(E_j(\omega)\right)\left[\hat{E}_j\right]$.

A *direct revelation mechanism* associated with an allocation x is a set $\Gamma = \left\langle I, \Sigma, x, t, g, \{u_i\}_{i \in I}\right\rangle$. In a maxmin equilibrium, each agent chooses a mixed strategy to maximize her payoff that takes into account the worst state that can occur and also the worst mixed strategy of all the other agents against her.[9] In other words, the maxmin equilibrium simply says that every agent adopts a criterion *à la* Wald (1950).

Definition 3 In a direct revelation mechanism Γ, a strategy profile $\sigma^* = \left(\sigma_1^*, \cdots, \sigma_N^*\right)$ constitutes a *maxmin equilibrium*, if for each agent i, her strategy σ_i^* maximizes the lower bound of her interim payoff, that is, the function $\sigma_i^* : \mathcal{F}_i \to \Delta(\mathcal{F}_i)$ satisfies that for each $E_i \in \mathcal{F}_i$,

$$\begin{aligned} &\min_{\omega \in E_i, \sigma_{-i} \in \Sigma_{-i}} \sum_{\hat{E}_i \in \mathcal{F}_i} \sum_{\hat{E}_{-i} \in \mathcal{F}_{-i}} \\ &\quad u_i\left(e_i(\omega) + t_i\left(\hat{E}_i, \hat{E}_{-i}\right), \omega\right) \sigma_i^*(E_i)\left[\hat{E}_i\right] \sigma_{-i}(E_{-i}(\omega))\left[\hat{E}_{-i}\right] \\ &\geq \min_{\omega \in E_i, \sigma_{-i} \in \Sigma_{-i}} \sum_{\hat{E}_i \in \mathcal{F}_i} \sum_{\hat{E}_{-i} \in \mathcal{F}_{-i}} \\ &\quad u_i\left(e_i(\omega) + t_i\left(\hat{E}_i, \hat{E}_{-i}\right), \omega\right) \sigma_i(E_i)\left[\hat{E}_i\right] \sigma_{-i}(E_{-i}(\omega))\left[\hat{E}_{-i}\right], \end{aligned} \tag{8}$$

for all $\sigma_i \in \Sigma_i$.

[9] As in De Castro et al. (2017a, b), Decerf and Riedel (2020), agents adopt maxmin strategies. We differ from Guo and Yannelis (2021) in both solution concepts and strategy sets.

That is, agent i evaluates each lottery choice by taking minima over ω and σ_{-i}. The first one refers to the *state uncertainty*, and the second one represents the *strategic uncertainty*. Let $ME(\Gamma)$ denote the set of maxmin equilibria of the mechanism Γ.

Definition 4 An allocation x is implementable as a maxmin equilibrium of the mechanism Γ, if there exists a maxmin equilibrium σ^*, such that for every ω, the outcome $g(\sigma^*, \omega) = x(\omega)$.

A strategy profile σ is *truth telling*, if $\sigma_i(E_i)[E_i] = 1$ for all E_i and i. We denote such a strategy profile by σ^T. Clearly, if the truth telling strategy profile σ^T constitutes a maxmin equilibrium of the mechanism $\Gamma = \langle I, \Sigma, x, t, g, \{u_i\}_{i\in I}\rangle$, then $g(\sigma^T, \omega) = e(\omega) + t(E_1(\omega), \cdots, E_N(\omega)) = e(\omega) + x(\omega) - e(\omega) = x(\omega)$, for each $\omega \in \Omega$. That is, the allocation x is implementable as a maxmin equilibrium of the mechanism Γ.

4 Implementation of incentive compatible allocations

We show in Example 1 below that if an allocation is maxmin incentive compatible (i.e., it satisfies (6)), then it may not be implementable in a direct revelation mechanism that takes into account that agents may randomize their choices. A stronger incentive compatibility notion (Definition 1), i.e., mixed maxmin incentive compatibility, is needed for implementation in a direct revelation mechanism with mixed strategies (Theorem 1 and Corollary 1 below).

Example 1 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, ex post efficient and maxmin incentive compatible allocation (please see (6)):

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

However, x is not implementable as a maxmin equilibrium when we take into account that agents may randomize their choices. Indeed, suppose that agent 1 observes the event $\{a, b\}$. If she reports the true event $\{a, b\}$, the worst state that can occur is state b and the worst strategy of agent 2 against agent 1 is reporting $\{b, c\}$. Hence, her worst case payoff is $\sqrt{2}$. If she reports the true event $\{a, b\}$ with probability 0.5, then the worst state that can occur is state b and the worst strategy of agent 2 against agent 1 is reporting $\{b, c\}$. Now, her worst case payoff is $\left(\sqrt{2}+\sqrt{3}\right)\cdot 0.5$, which is strictly higher than the worst case payoff of reporting the true event. It turns out that we have a unique maxmin equilibrium, in which each agent always mixes the true event and the lie with equal probabilities. Then, in each state of nature $\omega \in \Omega$, the agents get $x(a)$ with probability 0.25. Clearly, the allocation x is not implemented.

In Example 1, the agents' reports are always compatible. It follows that the Designer never has a chance to punish the agents. However, in general, agents' reports may be incompatible, i.e., $\bigcap_{j\in I}\hat{E}_j = \emptyset$. Thus, the Designer needs to specify transfers $D_i \in \mathbb{R}^{\ell}$, $i \in I$ for the case of incompatible reports. That is, each agent i gets a punishment D_i, when their reports are incompatible. Recall that punishments affect the incentive compatibility and the implementation of an allocation through (5) and (4) respectively.

Clearly, if punishments are too small, an agent may strictly prefer to lie. However, it is not true that very high punishments can induce the agents to report true events in $\Gamma = \left\langle I, \Sigma, x, t, g, \{u_i\}_{i\in I}\right\rangle$. The reason is that agents' reports may be incompatible even if agent i reports truthfully. Now, if the punishment D_i is too high and agent i can avoid incompatible reports through reporting a lie, then she may strictly prefer to lie. Thus, in order to remove agents' incentives to lie in Γ, punishments must not make "the worst case payoff of truth telling" strictly less than "the worst case payoff of lying with a strictly positive probability". Formally,

Definition 5 We say that punishments $D_i \in \mathbb{R}^{\ell}$, $i \in I$ induce truth telling with respect to mixed strategy deviations, if the following does not hold: there exists $E_i \in \mathcal{F}_i$, $\omega^* \in E_i$, $\hat{E}^*_{-i} \in \mathcal{F}_{-i}$ and an $\alpha_i \in \Delta(\mathcal{F}_i)$, such that $E_i \bigcap \hat{E}^*_{-i} = \emptyset$, and

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

We show in Theorem 1 below that each mixed maxmin incentive compatible allocation is implementable as a maxmin equilibrium if and only if punishments D_i, $i \in I$ induce truth telling with respect to mixed strategies deviations.

Theorem 1 *Let $\mathcal{E} = \{\Omega, (u_i, e_i, \mathcal{F}_i, P_i) : i \in I\}$ be an economy where agents have the Wald's maxmin preferences, e_i and u_i are $\mathcal{F}_i$-measurable for each i. Let $\Gamma = \left\langle I, \Sigma, x, t, g, \{u_i\}_{i\in I}\right\rangle$ be a direct revelation mechanism associated with a mixed maxmin incentive compatible allocation x in the economy $\mathcal{E}$. There exists a unique truth telling maxmin equilibrium σ^T of the mechanism Γ (i.e., $\left\{\sigma^T\right\} = ME(\Gamma)$) for which we have $g\left(\sigma^T, \omega\right) = x(\omega)$, for each $\omega \in \Omega$, if and only if punishments D_i, $i \in I$ induce truth telling with respect to mixed strategy deviations.*

Below we provide the intuition. For simplicity, we refer "the payoff that takes into account the worst state" (i.e., (7)) as "the IC payoff" and "the payoff that takes into account not only the worst state that can occur, but also the worst strategy of all the other agents against her" (i.e., (8)) as "the worst case payoff". Clearly, regardless of an agent's report, her worst case payoff is no higher than her IC payoff. If she reports the true event, we have two cases. The "first case" is that "the worst case outcome of truth telling" is $e_i(\omega) + x_i(\hat{\omega}) - e_i(\hat{\omega})$, where $\hat{\omega}$ is a state in the true event $E_i(\omega)$. It turns out that now her worst case payoff is the same as her IC payoff as both e_i and u_i are $\mathcal{F}_i$-measurable. Furthermore, since the allocation x is mixed maxmin incentive compatible, then her IC payoff of using any lottery α_i is lower than her IC payoff of reporting the true event. Thus, we have that her worst case payoff of reporting the true event is higher than her worst case payoff of using any lottery α_i. That is, the "mixed maxmin incentive compatibility condition" guarantees that reporting the true event in Γ is each agent's best choice. The "second case" is that "the worst case outcome of truth telling" is $e_i(\omega) - D_i$. That is, her truthful report is incompatible with other agents' reports. Now, the "inducing truth telling with respect to mixed strategy deviations" condition guarantees that reporting the true event in Γ is each agent's best choice. Furthermore, we assume that an agent lies, only if she can benefit from doing so. It follows that if σ^T is a maxmin equilibrium of Γ, then there is no other maxmin equilibrium. Finally, the "only if" part of the proof is straight forward. Indeed, if the "inducing truth telling with respect to mixed strategy deviations" condition fails, then "the worst case payoff of truth telling" is strictly lower than "the worst case payoff of a lottery". Thus, σ^T cannot be a maxmin equilibrium of Γ. The proof of Theorem 1 is in the Appendix.

De Castro et al. (2017a) adopts the standard type model of the implementation literature in which agents' reports are always compatible. As in Example 1, when agents' reports are always compatible, the Designer never has a chance to punish the agents. Clearly, if there is no chance to punish the agents, then no punishment can cause agents to lie. It follows that any punishment satisfies Definition 5. Thus, we have the following corollary.

Corollary 1 *Let $\mathcal{E} = \{\Omega, (u_i, e_i, \mathcal{F}_i, P_i) : i \in I\}$ be an economy with compatible reports, where agents have the Wald's maxmin preferences, e_i and u_i are $\mathcal{F}_i$-measurable for each i. Let $\Gamma = \langle I, \Sigma, x, t, g, \{u_i\}_{i\in I}\rangle$ be a direct revelation mechanism associated with a mixed maxmin incentive compatible allocation x in the economy $\mathcal{E}$. There exists a unique truth telling maxmin equilibrium σ^T of the mechanism Γ (i.e., $\{\sigma^T\} = ME(\Gamma)$) for which we have $g(\sigma^T, \omega) = x(\omega)$, for each $\omega \in \Omega$.*

Remark 2 Liu et al. (2020) showed that each *efficient allocation* of De Castro and Yannelis (2018); De Castro et al. (2017a, b) and each *interim maxmin value allocation* of Angelopoulos and Koutsougeras (2015) are mixed maxmin incentive compatible under the Wald's maxmin preferences. We extend these work further. Indeed, by Theorem 1 and Corollary 1, we know that these allocations are implementable as a maxmin equilibrium, even if we take into account that agents may randomize their choices.[10]

[10] Recall that many efficient allocations may not be incentive compatible or implementable under the Bayesian preferences, see for example, Holmström and Myerson (1983), Glycopantis and Yannelis (2018), Pram (2020), De Castro et al. (2011), Lombardi and Yoshihara (2020), Qin and Yang (2020), Guo and Yannelis (2020), just to name a few.

5 Concluding remarks

If we have mixed strategies in a direct revelation mechanism, then incentive compatibility has to be in mixed strategies. Otherwise, as Example 1 above indicates that if an allocation is just maxmin incentive compatible, then it may not be implementable. In general, agents' reports may be incompatible and the Designer needs to specify punishments for the case of incompatible reports. We show in Theorem 1 that each mixed maxmin incentive compatible allocation is implementable as a maxmin equilibrium if and only if the punishments do not discourage agents from truth telling.

Appendix

Proof of Theorem 1

Proof **The "if" direction:** Suppose that the truth telling strategy profile σ^T is not a maxmin equilibrium of Γ. Then there exists an agent i, an E_i, and a strategy $\sigma_i \in \Sigma_i$, such that

$$\begin{aligned}&\min_{\omega\in E_i,\sigma_{-i}\in\Sigma_{-i}}\sum_{\hat{E}_{-i}\in\mathcal{F}_{-i}}u_i\left(e_i(\omega)+t_i\left(E_i,\hat{E}_{-i}\right),\omega\right)\sigma_{-i}(E_{-i}(\omega))\left[\hat{E}_{-i}\right]\\&<\min_{\omega\in E_i,\sigma_{-i}\in\Sigma_{-i}}\sum_{\hat{E}_i\in\mathcal{F}_i}\sum_{\hat{E}_{-i}\in\mathcal{F}_{-i}}\\&u_i\left(e_i(\omega)+t_i\left(\hat{E}_i,\hat{E}_{-i}\right),\omega\right)\sigma_i(E_i)\left[\hat{E}_i\right]\sigma_{-i}(E_{-i}(\omega))\left[\hat{E}_{-i}\right].\end{aligned}\tag{9}$$

The inequality (9) can be rewritten as

$$\begin{aligned}&\min_{\omega\in E_i,\hat{E}_{-i}\in\mathcal{F}_{-i}}u_i\left(e_i(\omega)+t_i\left(E_i,\hat{E}_{-i}\right),\omega\right)\\&<\min_{\omega\in E_i,\hat{E}_{-i}\in\mathcal{F}_{-i}}\sum_{\hat{E}_i\in\mathcal{F}_i}u_i\left(e_i(\omega)+t_i\left(\hat{E}_i,\hat{E}_{-i}\right),\omega\right)\sigma_i(E_i)\left[\hat{E}_i\right].\end{aligned}\tag{10}$$

Clearly,

$$\begin{aligned}&\min_{\omega\in E_i,\hat{E}_{-i}\in\mathcal{F}_{-i}}\sum_{\hat{E}_i\in\mathcal{F}_i}u_i\left(e_i(\omega)+t_i\left(\hat{E}_i,\hat{E}_{-i}\right),\omega\right)\sigma_i(E_i)\left[\hat{E}_i\right]\\&\le\min_{\omega\in E_i}\sum_{\hat{E}_i\in\mathcal{F}_i}u_i\left(e_i(\omega)+t_i\left(\hat{E}_i,E_{-i}(\omega)\right),\omega\right)\sigma_i(E_i)\left[\hat{E}_i\right],\end{aligned}$$

therefore by (10) we have that

$$\min_{\omega\in E_i,\hat{E}_{-i}\in\mathcal{F}_{-i}} u_i\left(e_i(\omega)+t_i\left(E_i,\hat{E}_{-i}\right),\omega\right) < \min_{\omega\in E_i}\sum_{\hat{E}_i\in\mathcal{F}_i} u_i\left(e_i(\omega)+t_i\left(\hat{E}_i,E_{-i}(\omega)\right),\omega\right)\sigma_i(E_i)\left[\hat{E}_i\right]. \tag{11}$$

To ease the explanation, denote the left hand side of (11) by

$$u_i\left(e_i\left(\omega^*\right)+t_i\left(E_i,\hat{E}^*_{-i}\right),\omega^*\right)=\min_{\omega\in E_i,\hat{E}_{-i}\in\mathcal{F}_{-i}} u_i\left(e_i(\omega)+t_i\left(E_i,\hat{E}_{-i}\right),\omega\right), \tag{12}$$

where $\omega^*\in E_i$ and $\hat{E}^*_{-i}\in\mathcal{F}_{-i}$ solve the minimization problem above.

Case one: $E_i\bigcap\hat{E}^*_{-i}=\left\{\hat{\omega}\right\}$ for some $\hat{\omega}$ in E_i. Since u_i and e_i are $\mathcal{F}_i$-measurable, we have

$$\begin{aligned} u_i\left(e_i\left(\omega^*\right)+t_i\left(E_i,\hat{E}^*_{-i}\right),\omega^*\right) &= u_i\left(e_i\left(\omega^*\right)+x_i\left(\hat{\omega}\right)-e_i\left(\hat{\omega}\right),\omega^*\right)\\ &= u_i\left(e_i\left(\hat{\omega}\right)+x_i\left(\hat{\omega}\right)-e_i\left(\hat{\omega}\right),\omega^*\right)\\ &= u_i\left(x_i\left(\hat{\omega}\right),\hat{\omega}\right). \end{aligned} \tag{13}$$

Notice that $u_i\left(e_i\left(\omega^*\right)+t_i\left(E_i,\hat{E}^*_{-i}\right),\omega^*\right)=u_i\left(x_i\left(\hat{\omega}\right),\hat{\omega}\right)$ implies that

$$u_i\left(x_i\left(\hat{\omega}\right),\hat{\omega}\right)=\min_{\omega\in E_i,\hat{E}_{-i}\in\mathcal{F}_{-i}} u_i\left(e_i(\omega)+t_i\left(E_i,\hat{E}_{-i}\right),\omega\right) \le \min_{\omega\in E_i} u_i\left(e_i(\omega)+t_i\left(E_i,E_{-i}(\omega)\right),\omega\right)\le u_i\left(x_i\left(\hat{\omega}\right),\hat{\omega}\right).$$

That is, we have

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

Now, from (11), we have

$$\min_{\omega\in E_i} u_i\left(e_i(\omega)+t_i\left(E_i,E_{-i}(\omega)\right),\omega\right) < \min_{\omega\in E_i}\sum_{\hat{E}_i\in\mathcal{F}_i} u_i\left(e_i(\omega)+t_i\left(\hat{E}_i,E_{-i}(\omega)\right),\omega\right)\sigma_i(E_i)\left[\hat{E}_i\right].$$

That is, the allocation x is not mixed maxmin incentive compatible, which is a contradiction. Thus, the truth telling strategy profile σ^T is a maxmin equilibrium of Γ.

<u>Case two:</u> $E_i \bigcap \hat{E}^*_{-i} = \emptyset$. We know that the punishment D_i induces truth telling with respect to mixed strategy deviations. It follows that

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

for every $\alpha_i \in \Delta\left(\mathcal{F}_i\right)$. Now, by (12), we have that for every $\alpha_i \in \Delta\left(\mathcal{F}_i\right)$,

$$\min_{\omega\in E_i,\hat{E}_{-i}\in\mathcal{F}_{-i}} u_i\left(e_i\left(\omega\right)+t_i\left(E_i,\hat{E}_{-i}\right),\omega\right)$$
$$\geq \min_{\omega\in E_i,\hat{E}_{-i}\in\mathcal{F}_{-i}}\sum_{\hat{E}_i\in\mathcal{F}_i}u_i\left(e_i\left(\omega\right)+t_i\left(\hat{E}_i,\hat{E}_{-i}\right),\omega\right)\alpha_i\left(\hat{E}_i\right). \tag{14}$$

That is, inequality (10) cannot hold, which is a contradiction. Thus, the truth telling strategy profile σ^T is a maxmin equilibrium of Γ.

Finally, there is no other maxmin equilibrium. Indeed, an agent lies, only if she can benefit from doing so. Thus, if σ^* is a maxmin equilibrium with $\sigma^*_i\left(E_i\right) \neq \sigma^T_i\left(E_i\right)$, then lying (with a strictly positive probability) makes agent i strictly better off upon observing the event E_i, i.e.,

$$\min_{\omega\in E_i,\sigma_{-i}\in\Sigma_{-i}}\sum_{\hat{E}_{-i}\in\mathcal{F}_{-i}}u_i\left(e_i\left(\omega\right)+t_i\left(E_i,\hat{E}_{-i}\right),\omega\right)\sigma_{-i}\left(E_{-i}\left(\omega\right)\right)\left[\hat{E}_{-i}\right]$$
$$<\min_{\omega\in E_i,\sigma_{-i}\in\Sigma_{-i}}\sum_{\hat{E}_i\in\mathcal{F}_i}\sum_{\hat{E}_{-i}\in\mathcal{F}_{-i}} \tag{15}$$
$$u_i\left(e_i\left(\omega\right)+t_i\left(\hat{E}_i,\hat{E}_{-i}\right),\omega\right)\sigma^*_i\left(E_i\right)\left[\hat{E}_i\right]\sigma_{-i}\left(E_{-i}\left(\omega\right)\right)\left[\hat{E}_{-i}\right],$$

which contradicts to the fact that the truth telling strategy profile σ^T constitutes a maxmin equilibrium of the mechanism. We can conclude that the truth telling maxmin equilibrium is the only maxmin equilibrium of the mechanism Γ, i.e., $\left\{\sigma^T\right\} = ME\left(\Gamma\right)$.

The "only if" direction: We prove by contradiction. Suppose that the punishments D_i, $i \in I$ do not induce truth telling with respect to mixed strategy deviations. That is, there exists $E_i \in \mathcal{F}_i$, $\omega^* \in E_i$, $\hat{E}^*_{-i} \in \mathcal{F}_{-i}$ and an $\alpha_i \in \Delta\left(\mathcal{F}_i\right)$, such that $E_i \bigcap \hat{E}^*_{-i} = \emptyset$ and

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

Since

$$\min_{\omega \in E_i, \hat{E}_{-i} \in \mathcal{F}_{-i}} u_i \left(e_i(\omega) + t_i \left(E_i, \hat{E}_{-i} \right), \omega \right) \leq u_i \left(e_i(\omega^*) + t_i \left(E_i, \hat{E}^*_{-i} \right), \omega^* \right),$$

we have that

$$\begin{aligned} &\min_{\omega \in E_i, \hat{E}_{-i} \in \mathcal{F}_{-i}} u_i \left(e_i(\omega) + t_i \left(E_i, \hat{E}_{-i} \right), \omega \right) \\ &< \min_{\omega \in E_i, \hat{E}_{-i} \in \mathcal{F}_{-i}} \sum_{\hat{E}_i \in \mathcal{F}_i} u_i \left(e_i(\omega) + t_i \left(\hat{E}_i, \hat{E}_{-i} \right), \omega \right) \alpha_i \left(\hat{E}_i \right). \end{aligned} \tag{16}$$

Thus, the worst case payoff of truth telling is strictly lower than the worst case payoff of using a lottery α_i. That is, truth telling σ^T is not a maxmin equilibrium of Γ. We can conclude that if σ^T is a maxmin equilibrium of Γ, then the punishments $D_i, i \in I$ induce truth telling with respect to mixed strategy deviations. □

References

Angelopoulos, A., Koutsougeras, L.C.: Value allocation under ambiguity. Econ. Theory **59**, 147–167 (2015)

De Castro, L., Yannelis, N.C.: Uncertainty, efficiency and incentive compatibility: ambiguity solves the conflict between efficiency and incentive compatibility. J. Econ. Theory **177**, 678–707 (2018)

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

De Castro, L., Liu, Z., Yannelis, N.C.: Implementation under ambiguity. Games Econ. Behav. **101**, 20–33 (2017a)

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

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

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

Gilboa, I., Schmeidler, D.: Maxmin expected utility with non-unique prior. J. Math. Econ. **18**, 141–153 (1989)

Glycopantis, D., Yannelis, N.C.: The maximin equilibrium and the PBE under ambiguity. Econ. Theory Bull. **6**, 183–199 (2018)

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

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

Guo, H., Yannelis, N.C.: Full implementation under ambiguity. Am. Econ. J. **13**, 148–178 (2021)

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

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

Liu, Z.: A note on the welfare of the maximin rational expectations. Econ. Theory Bull. **2**, 213–218 (2014)

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

Liu, Z., Yannelis, N.C.: Persuasion in an asymmetric information economy: a justification of Wald's maxmin preferences. Economic Theory (2021). https://doi.org/10.1007/s00199-021-01364-7

Liu, Z., Song, X., Yannelis, N.C.: Randomization under ambiguity: efficiency and incentive compatibility. J. Math. Econ. **90**, 1–11 (2020)

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

Moreno-García, E., Torres-Martínez, J.P.: Information within coalitions: risk and ambiguity. Econ. Theory **69**, 125–147 (2020)
Pram, K.: Weak implementation. Econ. Theory **69**, 569–594 (2020)
Qin, C.Z., Yang, X.: On the equivalence of rational expectations equilibrium with perfect Bayesian equilibrium. Econ. Theory **69**, 1127–1146 (2020)
Raiffa, H.: Risk, ambiguity, and the Savage axioms: Comment. Q. J. Econ. **75**, 690–694 (1961)
Saito, K.: Preferences for flexibility and randomization under uncertainty. Am. Econ. Rev. **105**, 1246–1271 (2015)
Wald, A.: Statistical Decision Functions. John Wiley and Sons, New York (1950)