Games and Economic Behavior 101 (2017) 20–33

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Implementation under ambiguity ☆

Luciano I. de Castro [a], Zhiwei Liu [b,*], Nicholas C. Yannelis [a]

[a] *Department of Economics, The University of Iowa, Iowa City, USA*
[b] *International School of Economics and Management, Capital University of Economics and Business, Beijing, China*

ARTICLE INFO

Article history:
Received 8 June 2014
Available online 11 November 2015

JEL classification:
D51
D61
D81
D82

Keywords:
Maximin preferences
Maximin efficient allocations
Maximin equilibrium
Implementation

ABSTRACT

We introduce the idea of implementation under ambiguity. In particular, we study maximin efficient notions for an ambiguous asymmetric information economy (i.e., economies where agents' preferences are maximin *à la* Wald, 1950). The interest on the maximin preferences lies in the fact that maximin efficient allocations are always incentive compatible (de Castro and Yannelis, 2009), a result which is false with Bayesian preferences. A noncooperative notion called maximin equilibrium is introduced which provides a noncooperative foundation for individually rational and maximin efficient notions. Specifically, we show that given any arbitrary individually rational and ex-ante maximin efficient allocation, there is a direct revelation mechanism that yields the efficient allocation as its unique maximin equilibrium outcome. Thus, an incentive compatible, individually rational and efficient outcome can be reached by means of noncooperative behavior under ambiguity.

1. Introduction

There is a growing literature on applications of ambiguity aversion and non-expected utility in general to different areas in economics, e.g., Hansen and Sargent (2001) in Macroeconomics; Aryal and Stauber (2014), Bodoh-Creed (2012), Bose and Renou (2014), Bose et al. (2006), de Castro and Yannelis (2009, 2013), in game theory and mechanism design; Angelopoulos and Koutsougeras (2015), de Castro et al. (2011, 2014), He and Yannelis (2015), Liu (2014), Liu and Yannelis (2015) in general equilibrium; Dominiak et al. (2012), Haisley and Weber (2010), Ivanov (2011) in experimental economics; Cohen and Meilijson (2014) and Even and Lehrer (2014) in decision theory, among others. See also the symposium in Ellsberg et al. (2011). Clearly, maximin preferences have been widely used in computer science and engineering, where analyses based on worst-case scenarios are widespread. Although our motivation is not computer science, but the implementation of efficient and incentive compatible allocations, it is worth noticing the increasing interest of maximin preferences in computer science and game theory. A lot of games are nowadays played by robots (in internet ads auctions and in auctions for bandwidth in congested networks, for instance), and these robots are programmed by computer engineers that may incorporate these preferences, since it is a natural way of thinking in engineering. Therefore, such preferences are not so uncommon as we might initially think.

☆ On different occasions, we have benefited from comments, discussions, suggestions and questions by Camelia Bejan, Jean-Marc Bonnisseau, Bernard Cornet, Huiyi Guo, Richard Hislop, Atsushi Kajii, Mamoru Kaneko, John Ledyard, V. Filipe Martins-da-Rocha, Stephen Morris, Rabee Tourky, Walter Trockel and Chi Leung Wong. Thanks are also due to an Editor and two reviewers. Zhiwei Liu gratefully acknowledges funding from the Research and Innovation Centre of Metropolis Economic and Social Development, Capital University of Economics and Business.

* Corresponding author.
E-mail addresses: decastro.luciano@gmail.com (L.I. de Castro), liuzhiwei@cueb.edu.cn (Z. Liu), nicholasyannelis@gmail.com (N.C. Yannelis).

http://dx.doi.org/10.1016/j.geb.2015.10.010

In this paper we study an *ambiguous asymmetric information economy*, i.e., an economy consisting of a finite set of states of nature, a finite set of agents, each of whom is characterized by a *type set*, a *multi-prior* set, a *random initial endowment* and an *ex post utility function*. More importantly, the agents have maximin preferences, *à la* Wald.

In an ambiguous asymmetric information economy, any efficient allocation is incentive compatible with respect to the maximin preferences (de Castro and Yannelis, 2009). In other words, maximin preferences solve the conflict between incentive compatibility and efficiency. Recall that in the standard expected utility (Bayesian) framework, an efficient allocation may not be incentive compatible as it was shown by Holmström and Myerson (1983).

But, could one provide a noncooperative foundation for the maximin efficient allocations in terms of some game theoretic solution concept? In other words, can individually rational and ex-ante maximin efficient allocations be reached by means of noncooperative behavior? What would be the appropriate game theoretic solution concept?

In view of an ambiguous asymmetric information economy, one should not expect to employ any Bayesian Nash type equilibrium notion. To this end, we introduce the idea of a *maximin equilibrium*. Roughly speaking, in a maximin equilibrium, each agent maximizes his payoff lowest bound, that is, each agent maximizes the payoff that takes into account the worst actions of all the other agents against him and also the worst state that can occur.

The main result of the paper is that given any arbitrary individually rational and ex-ante maximin efficient allocation, there is a *direct revelation mechanism* that yields this allocation as its unique maximin equilibrium outcome, i.e., each individually rational and ex-ante maximin efficient allocation is implementable as a maximin equilibrium. A corollary of this result is that each maximin core allocation and each maximin value allocation is implementable as a maximin equilibrium. Therefore, incentive compatible, individually rational and efficient outcomes can be reached by means of noncooperative behavior under ambiguity.

The paper is organized as follows. Section 2 defines an ambiguous asymmetric information economy, and introduces the individually rational and ex-ante maximin efficient notions. In Section 3, we introduce the direct revelation mechanism, the maximin equilibrium, and present the main result of the paper. In Section 4, we also discuss the relationship of our result with de Castro and Yannelis (2009). In Section 5, we discuss the relationship between our paper and the robust implementation of Bergemann and Morris (2005, 2009). Finally, we conclude in Section 6.

2. Ambiguous asymmetric information economy

Let I be a set of N agents, i.e., $I = \{1, \cdots, N\}$. Agent i's type $t_i \in T_i$ summarizes his private information in the interim. The finite set of states of nature is $T = T_1 \times \cdots \times T_N$, and $t \in T$ is a state of nature.

Since each agent observes his type in the interim, it is natural to assume that at ex ante each agent is able to form a probability assessment over his types. That is, there is a measure μ_i generating t_i.

Assumption 1. For each i and for each type t_i, $\mu_i(t_i) > 0$.

Let Δ_i be the set of all probability measures on T that agrees with μ_i,

$$\Delta_i = \left\{ \text{probability measure } \pi_i : 2^T \to [0,1] \mid \sum_{t_{-i}} \pi_i(t_i, t_{-i}) = \mu_i(t_i), \forall t_i \in T_i \right\}.$$

Let P_i, a nonempty, closed and convex subset of Δ_i, be agent i's multi-prior set.

Let $\mathbb{R}_+^\ell$ be the ℓ good commodity space. An agent receives his initial endowment $e_i : T \to \mathbb{R}_+^\ell$ in the interim stage. That is, we have $e_i(t_i, t_{-i}) = e_i\left(t_i, t'_{-i}\right)$ for all t_{-i} and t'_{-i}. In the terminology of auctions, this corresponds to private values endowments, and it is an assumption that was used by Jackson and Swindles (2005) for instance.

Now, let $x_i : T \to \mathbb{R}_+^\ell$ denote *agent i's allocation* (or in short, *i-allocation*). Denote by L_i the *set of all allocations* of agent i, and by $x = (x_1, \cdots, x_N)$ *an allocation* of the economy. An allocation x is said to be *feasible*, if for each $t \in T$, $\sum_{i \in I} x_i(t) = \sum_{i \in I} e_i(t)$.

Let $u_i : \mathbb{R}_+^\ell \times T \to \mathbb{R}$ be agent i's *ex post utility function*, taking the form of $u_i(c_i; t)$, where c_i denotes agent i's consumption. Each agent knows his utility function in the interim, and therefore $u_i(c_i; t_i, t_{-i}) = u_i\left(c_i; t_i, t'_{-i}\right)$ for all t_{-i} and t'_{-i}. We postulate that each agent i's preferences on L_i are maximin (see Gilboa and Schmeidler, 1989).

Definition 1. Take any two allocations of agent i, f_i and h_i, from the set L_i. Agent i prefers f_i to h_i under the maximin preferences (written as $f_i \succeq_i^{MP} h_i$)

$$\min_{\pi_i \in P_i} \sum_{t \in T} u_i(f_i(t); t) \pi_i(t) \geq \min_{\pi_i \in P_i} \sum_{t \in T} u_i(h_i(t); t) \pi_i(t). \tag{1}$$

The general multi-prior model includes both the Bayesian and the Wald-type maximin preferences of de Castro and Yannelis (2009)[1] as special cases. Indeed, if P_i is a singleton set for each agent, then the multi-prior preferences become the Bayesian preferences. If $P_i = \Delta_i$ for each agent, then the multi-prior preferences become the maximin preferences in de Castro and Yannelis (2009), where the following formulation is equivalent to (1),

$$\sum_{t_i \in T_i} \left(\min_{t_{-i} \in T_{-i}} u_i \left(f_i \left(t_i, t_{-i} \right); t_i, t_{-i} \right) \right) \mu_i \left(t_i \right) \geq \sum_{t_i \in T_i} \left(\min_{t_{-i} \in T_{-i}} u_i \left(h_i \left(t_i, t_{-i} \right); t_i, t_{-i} \right) \right) \mu_i \left(t_i \right). \tag{2}$$

Furthermore, agent i strictly prefers f_i to h_i, $f_i \succ_i^{MP} h_i$, if he prefers f_i to h_i but not the reverse, i.e., $f_i \succeq_i^{MP} h_i$ but $h_i \not\succeq_i^{MP} f_i$.

The interest of the preferences, used in de Castro and Yannelis (2009), comes from the fact that under these preferences any efficient allocation is incentive compatible. Furthermore, only these preferences have this property.[2] Indeed, using Bayesian preferences, an efficient allocation may not be incentive compatible as it was shown by Holmström and Myerson (1983). Also, Ledyard (1977) showed that a core selecting mechanism may not be individually incentive compatible in a complete information setting. Furthermore, Ledyard (1978) showed that the introduction of incomplete information in the Bayesian sense may fail to create incentive compatibility, if a mechanism is not incentive compatible under complete information.

The standard notions of individual rationality and efficiency, when applied to agents with maximin preferences, can be stated as follows.

Definition 2. A feasible allocation $x = (x_i)_{i \in I}$ is said to be (maximin) individually rational, if for each $i \in I$, $x_i \succeq_i^{MP} e_i$.

Definition 3. A feasible allocation $x = (x_i)_{i \in I}$ is said to be ex-ante maximin efficient, if there does not exist another feasible allocation $y = (y_i)_{i \in I}$, such that $y_i \succeq_i^{MP} x_i$ for all i, and $y_i \succ_i^{MP} x_i$ for at least one i.

Remark 1. It should be noted that the concepts *maximin core allocations*, *maximin value allocations* and *maximin Walrasian expectations equilibrium allocations* defined in de Castro and Yannelis (2009), Angelopoulos and Koutsougeras (2015), He and Yannelis (2015) are all individually rational and ex-ante maximin efficient.

Could one provide a noncooperative foundation for these notions? That is, can each individually rational and ex-ante maximin efficient allocation be reached by means of noncooperation? We address this question in the next section.

3. Implementation

3.1. The direct revelation mechanism

A direct revelation mechanism, associated with an allocation and its underlying ambiguous asymmetric information economy, is a noncooperative game, in which agents (players) decide what to report after learning their types.

In the interim, each player i privately observes his type t_i, and receives the initial endowment $e_i(t_i)$. Then, each player i reports $t'_i \in T_i$, but the report may not be truthful. A report t'_i is a lie if it differs from the player's type t_i.

Definition 4. A *strategy* of player i is a function $s_i : T_i \to T_i$.

Let S_i denote player i's strategy set. Denote by $S = \times_{i \in I} S_i$ the strategy set, and let $s \in S$ denote a strategy profile. With a slightly abused notation, we use $s(t)$ to denote the players' reports, when they adopt the strategy profile s, and the realized state is t. That is, $s(t) = (s_1(t_1), \cdots, s_N(t_N))$. Clearly, for any $t \in T$, $s(t) \in \times_{i \in I} T_i$.

The players' reports are announced simultaneously. Based on the reports, redistribution takes place. Fig. 1 shows the time line.

A planned redistribution (net transfer) is the adjustments needed to go from the initial endowment e to a planned allocation x.

Definition 5. Let x be the planned allocation. *The planned redistribution of agent i at state t* is given by $x_i(t) - e_i(t)$.

The actual redistribution, on the other hand, depends on the planned redistribution and the players' reports.

[1] See also de Castro et al. (2014).

[2] In addition to the fact that maximin preferences solve the conflict between efficiency and incentive compatibility, it has been shown in Angelopoulos and Koutsougeras (2015), Aryal and Stauber (2014), de Castro and Yannelis (2009), de Castro et al. (2011, 2014), He and Yannelis (2015), Liu (2014) that the adoption of the maximin preferences provides new insights and superior outcomes than the Bayesian preferences. Furthermore, it is known that the maximin preferences solve the Ellsberg Paradox (see Ellsberg, 1961, and for example de Castro and Yannelis, 2013).

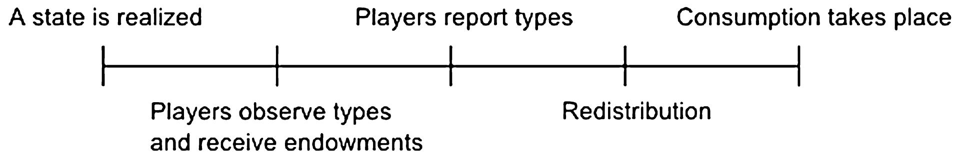

Fig. 1. Time line.

Definition 6. Let $x-e$ denote a planned redistribution, and $t'_1, \cdots, t'_N$ a list of reports. Then *the actual redistribution* of player i is given by

$$D_i\left(x-e,\left(t'_1,\cdots,t'_N\right)\right)=x_i\left(t'_1,\cdots,t'_N\right)-e_i\left(t'_1,\cdots,t'_N\right). \tag{3}$$

Definition 7. Let g_i be the outcome function of player i, which depends on the reports of the players and the realized state of nature, i.e.,

$$g_i\left(\left(t'_1,\cdots,t'_N\right),t\right)=e_i\left(t\right)+D_i\left(x-e,\left(t'_1,\cdots,t'_N\right)\right), \tag{4}$$

where $e_i\left(t\right)+D_i\left(x-e,\left(t'_1,\cdots,t'_N\right)\right)$ is the bundle of the goods, that player i ends up consuming.

The implementation literature often assumes that the set of feasible alternatives is independent of the state of nature. It follows that if the realized state is t, and the players' reported state is $\hat{t}$, then the players end up with the social choice $x\left(\hat{t}\right)$. The implicit assumption is that the social choice $x\left(\hat{t}\right)$ is feasible at the state t. In our context, the players receive initial endowment first, and then redistribute the endowments based on their reports. The relevant feasibility condition is that each player is rich enough to participate in the revelation mechanism. That is, for each i, t and $\hat{t}$, we have $e_i\left(t\right)+x_i\left(\hat{t}\right)-e_i\left(\hat{t}\right)\in\mathbb{R}^{\ell}_{+}$.

Finally, each player i has a final payoff function. It tells us the final payoff that player i ends up, given a list of reports and a realized state of nature. Formally,

Definition 8. Denote by v_i the final payoff function of player i. It takes the form of

$$v_i\left(\left(t'_1,\cdots,t'_N\right);t\right)=u_i\left(e_i\left(t\right)+D_i\left(x-e,\left(t'_1,\cdots,t'_N\right)\right);t\right). \tag{5}$$

A direct revelation mechanism associated with a planned allocation x and its underlying ambiguous asymmetric information economy is a set $\Gamma=\left\langle I,S,x-e,\{g_i\}_{i\in I},\{v_i\}_{i\in I}\right\rangle$.

3.2. *Maximin equilibrium*

An immediate question is that, what would be a reasonable solution concept for Γ? In view of an ambiguous asymmetric information economy, the standard Bayesian Nash solution concept is not suitable here (see Example 1 in Section 3.3). Below we introduce the notion of maximin equilibrium.

Definition 9. In a direct revelation mechanism $\Gamma=\left\langle I,S,x-e,\{g_i\}_{i\in I},\{v_i\}_{i\in I}\right\rangle$, a strategy profile $s^*=\left(s^*_1,\cdots,s^*_N\right)$ constitutes a *maximin equilibrium* ($\mathbb{ME}$), if for each player i, his strategy s^*_i maximizes his interim payoff lower bound, that is, the function $s^*_i: T_i\to T_i$ satisfies, for each t_i,

$$\min_{t'_{-i},t_{-i}\in T_{-i}} v_i\left(s^*_i\left(t_i\right),t'_{-i};t_i,t_{-i}\right)\geq\min_{t'_{-i},t_{-i}\in T_{-i}} v_i\left(\hat{t}_i,t'_{-i};t_i,t_{-i}\right), \tag{6}$$

for all $\hat{t}_i\in T_i$, where t'_{-i} denotes the reports from all the other players, i.e., $t'_{-i}\in T_{-i}=\times_{j\neq i}T_j$.

In other words, each player maximizes the payoff that takes into account the worst reports t'_{-i} of all the other players against him and also the worst state that can occur. Note, a state $t=(t_i,t_{-i})$ is made up with t_i and t_{-i}. In the interim, player i knows t_i, so the worst state is determined by t_{-i}. The maximin equilibrium simply says that every player adopts a criterion *à la* Wald (1950).

Remark 2. In contrast to the restricted maximin equilibrium notion of Dasgupta et al. (1979) and the consistent planning equilibrium of Bose and Renou (2014), our maximin equilibrium notion does not need each player to correctly guess his opponents' strategies to reach an equilibrium. Furthermore, the maximin equilibrium is unique, whenever truth telling is optimal for each player. This is not necessarily the case with the restricted maximin equilibrium notion or the consistent planning equilibrium notion. This is the main reason, we introduce the notion of maximin equilibrium.

Remark 3. Note that in equation (6), we take minima over t_{-i} and t'_{-i}. The first one refers to the uncertainty with respect to the opponents' types and the second one represents the *strategic uncertainty*, that is, it is related to possible false reports by opponents. Therefore, we treat strategic uncertainty in a similar way to the uncertainty with respect to types. This might be considered natural in some settings; for instance, when there is complete ignorance both with respect to types and actions. However, other formulations might be more convenient in different settings. An example of alternative formulation is the *restricted maximin equilibrium* introduced by Dasgupta et al. (1979, p. 207).

We study the implementation of an allocation with the help of its corresponding direct revelation mechanism. In particular, we say an allocation x is implementable, if x can be realized through a maximin equilibrium of the direct revelation mechanism $\Gamma = \langle I, S, x-e, \{g_i\}_{i\in I}, \{v_i\}_{i\in I}\rangle$.

Let $\mathbb{ME}(\Gamma)$ denote the set of maximin equilibria of the mechanism Γ.

Definition 10. Let x be an allocation of an ambiguous asymmetric information economy $\mathcal{E}$, and $\mathbb{ME}(\Gamma)$ the set of maximin equilibria of the mechanism $\Gamma = \langle I, S, x-e, \{g_i\}_{i\in I}, \{v_i\}_{i\in I}\rangle$. We say that *the allocation x is implementable as a maximin equilibrium of the mechanism* Γ if,

$$\exists s^* \in \mathbb{ME}(\Gamma), \text{ such that } g_i\left(s^*(t), t\right) = x_i(t),$$

for each $t \in T$ and for each $i \in I$.

Definition 11. We say that a strategy profile s is *truth telling*, if $s_i(t_i) = t_i$ for all t_i and i. We denote such a strategy profile by s^T.

Remark 4. Clearly, if the truth telling strategy profile s^T constitutes a maximin equilibrium of the mechanism Γ, i.e., $s^T \in \mathbb{ME}(\Gamma)$, then the allocation x is implementable as a maximin equilibrium of the mechanism Γ.

Indeed, under the truth telling strategy profile s^T, the list of reports associated with each state t is $s^T(t) = (t_1, \cdots, t_N)$. That is, the players always tell the truth. As a consequence, we have

$$\begin{aligned} g_i\left(s^T(t), t\right) &= g_i\left((t_1, \cdots, t_N), t\right) \\ &= e_i(t) + D_i\left(x-e, (t_1, \cdots, t_N)\right) \\ &= e_i(t) + x_i(t) - e_i(t) = x_i(t), \end{aligned}$$

for each $t \in T$ and for each $i \in I$ – the requirement of Definition 10. Furthermore, when $s^T \in \mathbb{ME}(\Gamma)$, we say Γ has a *truth telling maximin equilibrium*.

3.3. An example

The example below shows that a maximin individually rational and maximin ex ante efficient allocation is implementable as a maximin equilibrium, while in the same economy a Bayesian individually rational and Bayesian ex ante efficient allocation is not implementable as a Bayesian Nash equilibrium.

Example 1. There are two agents, two commodities, and four possible states of nature $T = \{a, b, c, d\}$. The ex post utility functions of the agents are

$$u_1\left(c_1^1, c_1^2; t\right) = \begin{cases} \sqrt{c_1^1} + \sqrt{c_1^2} & \text{if } t \in \{a, b\} \\ \sqrt{c_1^1} + 1.2\sqrt{c_1^2} & \text{if } t \in \{c, d\} \end{cases}$$

$$u_2\left(c_2^1, c_2^2; t\right) = \begin{cases} \sqrt{c_2^1} + \sqrt{c_2^2} & \text{if } t \in \{a, c\} \\ \sqrt{c_2^1} + 1.2\sqrt{c_2^2} & \text{if } t \in \{b, d\}. \end{cases}$$

The agents' random initial endowments and type sets are:

$$\begin{aligned} &(e_1(a), e_1(b), e_1(c), e_1(d)) = [(8,6); (8,6); (10,4); (10,4)]; \quad T_1 = \{\{a,b\}, \{c,d\}\} \\ &(e_2(a), e_2(b), e_2(c), e_2(d)) = [(4,10); (6,8); (4,10); (6,8)]; \quad T_2 = \{\{a,c\}, \{b,d\}\}. \end{aligned}$$

That is, when the state of nature is a, agent 1 is of type $t_1 = \{a, b\}$ and agent 2 is of type $t_2 = \{a, c\}$. Furthermore, each agent assigns a probability of $\frac{1}{2}$ on each of his type, $\mu_1(\{a,b\}) = \mu_1(\{c,d\}) = \frac{1}{2}$, and $\mu_2(\{a,c\}) = \mu_2(\{b,d\}) = \frac{1}{2}$.

Now, if we assume that the agents have Bayesian preferences and impose a common prior of probability $\frac{1}{4}$ on each state of nature, then, allocation

$$x=\begin{pmatrix} x_1(a), & x_1(b), & x_1(c), & x_1(d) \\ x_2(a), & x_2(b), & x_2(c), & x_2(d) \end{pmatrix}=\begin{pmatrix} (6,8); & (7,5.74); & (7,8.26); & (8,6) \\ (6,8); & (7,8.26); & (7,5.74); & (8,6) \end{pmatrix},$$

is individually rational and ex-ante efficient with respect to the Bayesian preferences, i.e., x is in the Bayesian core. Its corresponding planned redistribution is

$$x_1-e_1=[(-2,2);(-1,-0.26);(-3,4.26);(-2,2)]$$
$$x_2-e_2=[(2,-2);(1,0.26);(3,-4.26);(2,-2)].$$

The direct revelation mechanism $\Gamma=\langle I,S,x-e,\{g_i\}_{i\in I},\{v_i\}_{i\in I}\rangle$ has only one Bayesian Nash equilibrium. The equilibrium is s with $s_1(\{a,b\})=s_1(\{c,d\})=\{c,d\}$, and $s_2(\{a,c\})=s_2(\{b,d\})=\{b,d\}$. That is, regardless of the realized state, the players will agree that state d occurs and exchange 2 units of good 1 with 2 units of good 2. Consequently, the allocation

$$y=\begin{pmatrix} y_1(a), & y_1(b), & y_1(c), & y_1(d) \\ y_2(a), & y_2(b), & y_2(c), & y_2(d) \end{pmatrix}=\begin{pmatrix} (6,8); & (6,8); & (8,6); & (8,6) \\ (6,8); & (8,6); & (6,8); & (8,6) \end{pmatrix}$$

is the unique Bayesian Nash equilibrium outcome. Clearly, the allocation x, which is different from y, is not implemented.

It turns out that if the agents have Wald-type maximin preferences, the allocation y is individually rational and ex-ante efficient with respect to the maximin preferences, i.e., y is in the maximin core. Furthermore, the agents are able to reach the allocation y through the direct revelation mechanism $\Gamma=\langle I,S,y-e,\{g_i\}_{i\in I},\{v_i\}_{i\in I}\rangle$. Indeed, the planned redistribution $y-e$, which is uniform across states, is

$$y_1(a)-e_1(a)=y_1(b)-e_1(b)=y_1(c)-e_1(c)=x_1(d)-e_1(d)=(-2,2);$$
$$y_2(a)-e_2(a)=y_2(b)-e_2(b)=y_2(c)-e_2(c)=y_2(d)-e_2(d)=(2,-2).$$

Regardless of the reports of the players and the realized state, the transfer will be the same, i.e., two units of good 1 for two units of good 2. No player can benefit from lying. The unique[3] maximin equilibrium is

$$s=(s_1(\{a,b\})=\{a,b\},s_1(\{c,d\})=\{c,d\};s_2(\{a,c\})=\{a,c\},s_2(\{b,d\})=\{b,d\}).$$

Clearly, the players reach the allocation y through this equilibrium. That is, y is implemented.

In the next section we show that each maximin individually rational and ex ante maximin efficient allocation is implementable as a maximin equilibrium.

3.4. Implementation

This section presents our main result, i.e., in the mechanism $\Gamma=\langle I,S,x-e,\{g_i\}_{i\in I},\{v_i\}_{i\in I}\rangle$, where x is individually rational and ex-ante maximin efficient, no player has an incentive to lie, and the allocation x is implementable through its corresponding mechanism Γ. Formally:

Theorem 1. *Denote by x an individually rational and ex-ante maximin efficient allocation, and $\mathbb{ME}(\Gamma)$ the set of maximin equilibria of the direct revelation mechanism $\Gamma=\langle I,S,x-e,\{g_i\}_{i\in I},\{v_i\}_{i\in I}\rangle$. Then, there exists a truth telling maximin equilibrium s^T, which is the unique maximin equilibrium of the mechanism Γ (i.e., $\{s^T\}=\mathbb{ME}(\Gamma)$), for which we have $g_i(s^T(t),t)=x_i(t)$, for each $t\in T$ and for each $i\in I$, i.e., the allocation x is implementable as a maximin equilibrium of its corresponding mechanism Γ.*

Remark 5. The implementation shares some similarities with the (truthful) implementation of Dasgupta et al. (1979, p. 189) – an allocation can be (truthfully) implemented, if there exists a direct revelation mechanism (a game in which players report their private information) for which truth telling is its equilibrium (based on some game theoretic solution concept), and the truth telling equilibrium yields the allocation as its outcome.

Remark 6. In contrast to the partial implementation with ambiguity sensitive individuals of Bose and Renou (2014), we have full implementation. But we differ from the full implementation of Jackson (1991), Palfrey and Srivastava's (1989), and Hahn and Yannelis (2001), in that, we do not implement all equilibrium allocations. Instead, we pick an allocation, and fully implement the allocation with a mechanism, *à la* Bergemann and Morris (2009). In particular, we show that given any arbitrary individually rational and ex-ante maximin efficient allocation x, its corresponding mechanism $\Gamma=\langle I,S,x-e,\{g_i\}_{i\in I},\{v_i\}_{i\in I}\rangle$ yields the allocation x as its *unique* maximin equilibrium outcome.

In addition, this paper contributes to the growing literature on efficiency with ambiguity sensitive individuals. de Castro and Yannelis (2009) show that there is no longer a conflict between efficiency and incentive compatibility, if and only if

[3] We assume that a player lies only if he can benefit from doing so.

the agents' preferences are maximin *à la* Wald. Bose et al. (2006) characterize the level of ambiguity needed for each full insurance auction to be optimal. Moreover, they pin down the conditions on the players' multi-prior sets, that make the full insurance auction the unique optimal auction. Bodoh-Creed (2012) develops a payoff equivalence theorem for mechanisms with ambiguity sensitive participants. He also studies the constrained efficient, budget balanced bilateral trade mechanism and shows that increased ambiguity improves the efficiency of the mechanism. This is in the spirit of de Castro and Yannelis' result (de Castro and Yannelis, 2009), who show that the Wald-type maximin preferences provide higher efficiency than Bayesian preferences.

Since maximin core allocations (de Castro et al., 2011), maximin value allocations (Angelopoulos and Koutsougeras, 2015) and maximin Walrasian expectations equilibrium allocations (de Castro et al., 2011) are individually rational and ex-ante efficient, it follows from the main theorem that they are implementable as a maximin equilibrium.

3.5. Proof of the main theorem

Given a direct revelation mechanism $\Gamma = \langle I, S, x-e, \{g_i\}_{i\in I}, \{v_i\}_{i\in I}\rangle$, we show that if the allocation x is a maximin individually rational and ex-ante maximin efficient allocation, then there does not exist a player i, a type t_i, and a lie $\hat{t}_i$ (clearly, $t_i \neq \hat{t}_i$), such that

$$\min_{t'_{-i}, t_{-i}\in T_{-i}} v_i\left(t_i, t'_{-i}; t_i, t_{-i}\right) < \min_{t'_{-i}, t_{-i}\in T_{-i}} v_i\left(\hat{t}_i, t'_{-i}; t_i, t_{-i}\right). \tag{7}$$

Suppose that there exist a player i, a type t_i, and a lie $\hat{t}_i$, such that (7) holds. We will show that the feasible allocation x fails to be ex-ante efficient under the maximin preferences. The argument is along the lines of de Castro and Yannelis (2009) Theorem 3.1.

The left hand side of (7) can be rewritten as

$$\min_{t'_{-i}, t_{-i}\in T_{-i}} v_i\left(t_i, t'_{-i}; t_i, t_{-i}\right) = \min_{t'_{-i}\in T_{-i}} u_i\left(x_i\left(t_i, t'_{-i}\right); t_i, t'_{-i}\right),$$

since the ex post utility function is private valued. Define an i-allocation of player i, $z_i(\cdot)$, such that for each $t'_{-i}\in T_{-i}$, $v_i\left(\hat{t}_i, t'_{-i}; t_i, t'_{-i}\right) = u_i\left(z_i\left(t_i, t'_{-i}\right); t_i, t'_{-i}\right)$, and therefore, the right hand side of (7) can be rewritten as

$$\min_{t'_{-i}, t_{-i}\in T_{-i}} v_i\left(\hat{t}_i, t'_{-i}; t_i, t_{-i}\right) = \min_{t'_{-i}\in T_{-i}} u_i\left(z_i\left(t_i, t'_{-i}\right); t_i, t'_{-i}\right).$$

Clearly, we have $z_i\left(t_i, t'_{-i}\right) = e_i\left(t_i, t'_{-i}\right) + x_i\left(\hat{t}_i, t'_{-i}\right) - e_i\left(\hat{t}_i, t'_{-i}\right)$ for each t'_{-i}. It follows from (7) that

$$\min_{t'_{-i}\in T_{-i}} u_i\left(x_i\left(t_i, t'_{-i}\right); t_i, t'_{-i}\right) < \min_{t'_{-i}\in T_{-i}} u_i\left(z_i\left(t_i, t'_{-i}\right); t_i, t'_{-i}\right). \tag{8}$$

Now, we define an allocation y that Pareto improves x under the maximin preferences. Define for each $j\in I$, the j-allocation $y_j(\cdot)$ by

$$y_j(t') = \begin{cases} z_j\left(t_i, t'_{-i}\right) = e_j\left(t_i, t'_{-i}\right) + x_j\left(\hat{t}_i, t'_{-i}\right) - e_j\left(\hat{t}_i, t'_{-i}\right) & \text{if } t'_i = t_i \text{ and } t'_{-i}\in T_{-i}; \\ x_j\left(t'\right) & \text{otherwise.} \end{cases}$$

Notice that the allocation y is feasible. We only need to check the states at which x and y differ. For each state t' with $t'_i = t_i$ and $t'_{-i}\in T_{-i}$, we have

$$\sum_{j\in I} y_j\left(t'\right) = \sum_{j\in I} z_j\left(t'\right) = \sum_{j\in I} e_j\left(t_i, t'_{-i}\right) + \sum_{j\in I} x_j\left(\hat{t}_i, t'_{-i}\right) - \sum_{j\in I} e_j\left(\hat{t}_i, t'_{-i}\right) = \sum_{j\in I} e_j\left(t_i, t'_{-i}\right),$$

since x is feasible.

From (8) and the definition of y_i, we have

$$\min_{t'_{-i}\in T_{-i}} u_i\left(x_i\left(t_i, t'_{-i}\right); t_i, t'_{-i}\right) < \min_{t'_{-i}\in T_{-i}} u_i\left(y_i\left(t_i, t'_{-i}\right); t_i, t'_{-i}\right) \tag{9}$$

under the type t_i; and for any other type t'_i, we have

$$\min_{t'_{-i}\in T_{-i}} u_i\left(x_i\left(t'_i, t'_{-i}\right); t'_i, t'_{-i}\right) = \min_{t'_{-i}\in T_{-i}} u_i\left(y_i\left(t'_i, t'_{-i}\right); t'_i, t'_{-i}\right).$$

Therefore, combined with the assumption on $\mu_i(\cdot)$ (Assumption 1), we conclude that, for the player i,

$$\sum_{t_i' \in T_i} \left(\min_{t_{-i}' \in T_{-i}} u_i \left(y_i \left(t_i', t_{-i}' \right); t_i', t_{-i}' \right) \right) \mu_i \left(t_i' \right) > \sum_{t_i' \in T_i} \left(\min_{t_{-i}' \in T_{-i}} u_i \left(x_i \left(t_i', t_{-i}' \right); t_i', t_{-i}' \right) \right) \mu_i \left(t_i' \right). \tag{10}$$

That is, player i strictly prefers the i-allocation y_i to the i-allocation x_i under the maximin preferences. Now, it remains to show that for any other player $k \neq i$, we have y_k is preferred to x_k under the maximin preferences.

Fix an arbitrary player $k \neq i$, and an arbitrary type of player k, t_k. Define $X_k = \left\{ x_k \left(t_k, t_{-k}' \right) : t_{-k}' \in T_{-k} \right\}$ and $Y_k = \left\{ y_k \left(t_k, t_{-k}' \right) : t_{-k}' \in T_{-k} \right\}$. We have $Y_k \subset X_k$. Indeed, if $t_i' = t_i$, then $y_k \left(t_k, t_i, t_{-k-i}' \right) = z_k \left(t_k, t_i, t_{-k-i}' \right) = e_k \left(t_k, t_i, t_{-k-i}' \right) + x_k \left(t_k, \hat{t}_i, t_{-k-i}' \right) - e_k \left(t_k, \hat{t}_i, t_{-k-i}' \right) = x_k \left(t_k, \hat{t}_i, t_{-k-i}' \right) \in X_k$; Otherwise, $y_k \left(t_k, t_{-k}' \right) = x_k \left(t_k, t_{-k}' \right) \in X_k$. Therefore, we have that

$$\min_{t_{-k}' \in T_{-k}} u_k \left(x_k \left(t_k, t_{-k}' \right); t_k, t_{-k}' \right) \leq \min_{t_{-k}' \in T_{-k}} u_i \left(y_k \left(t_k, t_{-k}' \right); t_k, t_{-k}' \right).$$

Since the type t_k is arbitrary, we conclude that

$$\sum_{t_k' \in T_k} \left(\min_{t_{-k}' \in T_{-k}} u_k \left(y_k \left(t_k', t_{-k}' \right); t_k', t_{-k}' \right) \right) \mu_k \left(t_k' \right) \geq \sum_{t_k' \in T_k} \left(\min_{t_{-k}' \in T_{-k}} u_k \left(x_k \left(t_k', t_{-k}' \right); t_k', t_{-k}' \right) \right) \mu_k \left(t_k' \right).$$

Also, since player $k \neq i$ is arbitrary, we have for every player $k \neq i$, y_k is preferred to x_k under the maximin preferences.

Thus, the feasible allocation y Pareto improves the allocation x under the maximin preferences, i.e., x fails to be an ex-ante maximin efficient allocation, a contradiction.

Finally, we show that the truth telling maximin equilibrium is the only maximin equilibrium of the mechanism Γ, i.e., $\left\{ s^T \right\} = \mathbb{ME}(\Gamma)$. Suppose otherwise, that is, suppose both s^T and s^* are maximin equilibria of the mechanism Γ, and $s^T \neq s^*$. The truth telling strategy profile s^T is different from the strategy profile s^*, implies that there must exist a player i and a type t_i, such that $s_i^T(t_i) = t_i \neq \hat{t}_i = s_i^*(t_i)$. But $s_i^*(t_i) = \hat{t}_i \neq t_i$ holds, only if lying makes type t_i player i strictly better off,[4] i.e.,

$$\min_{t_{-i}', t_{-i} \in T_{-i}} v_i \left(\hat{t}_i, t_{-i}'; t_i, t_{-i} \right) > \min_{t_{-i}', t_{-i} \in T_{-i}} v_i \left(t_i, t_{-i}'; t_i, t_{-i} \right),$$

a contradiction to the fact that the truth telling strategy profile constitutes a maximin equilibrium of the mechanism.

Clearly, the individually rational and ex-ante maximin efficient allocation x is implemented. Indeed, under the truth telling strategy profile s^T, the list of reports associated to each state t is $s^T(t) = t$. As a consequence, we have

$$\begin{aligned} g_i \left(s^T(t), t \right) &= g_i \left((t_1, \cdots, t_N), t \right) \\ &= e_i(t_i) + D_i \left(x - e, (t_1, \cdots, t_N) \right) \\ &= e_i(t_i) + x_i(t) - e_i(t) = x_i(t), \end{aligned}$$

for each $t \in T$ and for each $i \in I$ – the requirement of Definition 10, and this completes the proof of the Theorem.

4. Relationship with de Castro–Yannelis

We discuss the relationship between this work and the one of de Castro and Yannelis (2009).

Definition 12 *(de Castro–Yannelis).* An allocation x is (maximin) incentive compatible if there is no i, t_i', t_i'' such that

$$\begin{aligned} &\min_{t_{-i}, t_{-i}' \in T_{-i}} u_i \left(e_i \left(t_i', t_{-i}' \right) + x_i \left(t_i'', t_{-i}' \right) - e_i \left(t_i'', t_{-i}' \right); t_i', t_{-i} \right) \\ &\quad > \min_{t_{-i}, t_{-i}' \in T_{-i}} u_i \left(x_i \left(t_i', t_{-i}' \right); t_i', t_{-i} \right). \end{aligned} \tag{11}$$

Definition 13 *(de Castro–Yannelis).* An allocation x is (interim) maximin efficient, if there is no feasible allocation $y = (y_i)_{i \in I}$ such that

$$\min_{t_{-i}, t_{-i}' \in T_{-i}} u_i \left(y_i \left(t_i, t_{-i}' \right); t_i, t_{-i} \right) \geq \min_{t_{-i}, t_{-i}' \in T_{-i}} u_i \left(x_i \left(t_i, t_{-i}' \right); t_i, t_{-i} \right) \tag{12}$$

for every i and $t_i \in T_i$, with strict inequality for some i and t_i.

[4] We assume that a player lies, only if he can benefit from doing so.

Theorem 2 *(de Castro–Yannelis Theorem 3.1). If $x = (x_i)_{i \in I}$ is an (interim) maximin efficient allocation, then x is maximin incentive compatible.*

Notice that de Castro and Yannelis (2009) work with an interim efficiency notion, whereas we work with an ex ante efficiency notion. Example 2 below shows that a maximin individually rational and ex ante maximin efficient allocation, which is maximin incentive compatible in the sense of de Castro and Yannelis (2009), may not be implementable as a maximin equilibrium through its corresponding direct revelation mechanism.[5] Thus, in general an allocation's implementability does not follow from its maximin incentive compatibility.

Example 2. There are two agents, two commodities x and y, and four possible states of nature $T = \{a, b, c, d\}$. The ex post utility function of agent $i = 1, 2$ is

$$u_i\left(c_i^x, c_i^y; t\right) = \begin{cases} \sqrt{c_i^x} + \sqrt{c_i^y} & \text{if } t \in \{a, c\} \\ 10\sqrt{c_i^x} + 10\sqrt{c_i^y} & \text{if } t \in \{b, d\}. \end{cases}$$

The agents' random initial endowments and type sets are:

$$(e_1(a), e_1(b), e_1(c), e_1(d)) = [(10, 4); (8, 2); (8, 2); (3, 4)]; \quad T_1 = \{\{a, d\}, \{b, c\}\}$$
$$(e_2(a), e_2(b), e_2(c), e_2(d)) = [(4, 9); (4, 9); (2, 8); (1, 3)]; \quad T_2 = \{\{a, b\}, \{c, d\}\}.$$

That is, when the state of nature is a, agent 1 (in the interim) knows that his type is $t_1 = \{a, d\}$ and agent 2 knows that her type is $t_2 = \{a, b\}$. Notice that the agents' ex post utility functions and initial endowments are state dependent, but they are not private information measurable. Furthermore, each agent assigns a probability of $\frac{1}{2}$ on each of his or her type, $\mu_1(\{a, d\}) = \mu_1(\{b, c\}) = \frac{1}{2}$, and $\mu_2(\{a, b\}) = \mu_2(\{c, d\}) = \frac{1}{2}$.

A maximin individually rational and ex-ante maximin efficient allocation is

$$x = \begin{pmatrix} x_1(a), & x_1(b), & x_1(c), & x_1(d) \\ x_2(a), & x_2(b), & x_2(c), & x_2(d) \end{pmatrix} = \begin{pmatrix} (7, 6.5); & (5, 5); & (5, 5); & (1, 6) \\ (7, 6.5); & (7, 6); & (5, 5); & (3, 1) \end{pmatrix}.$$

Its corresponding planned redistribution is

$$x_1 - e_1 = [(-3, 2.5); (-3, 3); (-3, 3); (-2, 2)]$$
$$x_2 - e_2 = [(3, -2.5); (3, -3); (3, -3); (2, -2)].$$

We show below that the allocation x is maximin incentive compatible (de Castro and Yannelis, 2009). That is, for every i, t_i', t_i'', we have

$$\begin{aligned} &\min_{t_{-i}, t'_{-i} \in T_{-i}} u_i\left(e_i\left(t_i', t_{-i}'\right) + x_i\left(t_i'', t_{-i}'\right) - e_i\left(t_i'', t_{-i}'\right); t_i', t_{-i}\right) \\ &\quad \le \min_{t_{-i}, t'_{-i} \in T_{-i}} u_i\left(x_i\left(t_i', t_{-i}'\right); t_i', t_{-i}\right). \end{aligned}$$

For $i = 1$, $t_i' = \{a, d\}$, $t_i'' = \{b, c\}$, we have

$$\begin{aligned} &\min_{t_{-i}, t'_{-i} \in T_{-i}} u_i\left(e_i\left(\{a, d\}, t_{-i}'\right) + x_i\left(\{b, c\}, t_{-i}'\right) - e_i\left(\{b, c\}, t_{-i}'\right); \{a, d\}, t_{-i}\right) \\ &\quad = \min\left\{\sqrt{7} + \sqrt{7}, \sqrt{0} + \sqrt{7}, 10\sqrt{7} + 10\sqrt{7}, 10\sqrt{0} + 10\sqrt{7}\right\} \\ &\quad = 2.646 < 3.449 = \\ &\qquad \min\left\{\sqrt{7} + \sqrt{6.5}, \sqrt{1} + \sqrt{6}, 10\sqrt{7} + 10\sqrt{6.5}, 10\sqrt{1} + 10\sqrt{6}\right\} = \\ &\qquad \min_{t_{-i}, t'_{-i} \in T_{-i}} u_i\left(x_i\left(\{a, d\}, t_{-i}'\right); \{a, d\}, t_{-i}\right). \end{aligned}$$

For $i = 1$, $t_i' = \{b, c\}$, $t_i'' = \{a, d\}$, we have

[5] This example is motivated by a referee who raised the question as to how our implementation notion is related to de Castro and Yannelis (2009)'s Theorem 2 above.

$$\min_{t_{-i},t'_{-i}\in T_{-i}} u_i\left(e_i\left(\{b,c\},t'_{-i}\right)+x_i\left(\{a,d\},t'_{-i}\right)-e_i\left(\{a,d\},t'_{-i}\right);\{b,c\},t_{-i}\right)$$
$$=\min\left\{\sqrt{5}+\sqrt{4.5},\sqrt{6}+\sqrt{4},10\sqrt{5}+10\sqrt{4.5},10\sqrt{6}+10\sqrt{4}\right\}$$
$$=4.357<4.472=$$
$$\min\left\{\sqrt{5}+\sqrt{5},\sqrt{5}+\sqrt{5},10\sqrt{5}+10\sqrt{5},10\sqrt{5}+10\sqrt{5}\right\}=$$
$$\min_{t_{-i},t'_{-i}\in T_{-i}} u_i\left(x_i\left(\{b,c\},t'_{-i}\right);\{b,c\},t_{-i}\right).$$

For $i=2$, $t'_i=\{a,b\}$, $t''_i=\{c,d\}$, we have

$$\min_{t_{-i},t'_{-i}\in T_{-i}} u_i\left(e_i\left(\{a,b\},t'_{-i}\right)+x_i\left(\{c,d\},t'_{-i}\right)-e_i\left(\{c,d\},t'_{-i}\right);\{a,b\},t_{-i}\right)$$
$$=\min\left\{\sqrt{6}+\sqrt{7},\sqrt{7}+\sqrt{6},10\sqrt{6}+10\sqrt{7},10\sqrt{7}+10\sqrt{6}\right\}$$
$$=5.095=$$
$$\min\left\{\sqrt{7}+\sqrt{6.5},\sqrt{7}+\sqrt{6},10\sqrt{7}+10\sqrt{6.5},10\sqrt{7}+10\sqrt{6}\right\}=$$
$$\min_{t_{-i},t'_{-i}\in T_{-i}} u_i\left(x_i\left(\{a,b\},t'_{-i}\right);\{a,b\},t_{-i}\right).$$

For $i=2$, $t'_i=\{c,d\}$, $t''_i=\{a,b\}$, we have

$$\min_{t_{-i},t'_{-i}\in T_{-i}} u_i\left(e_i\left(\{c,d\},t'_{-i}\right)+x_i\left(\{a,b\},t'_{-i}\right)-e_i\left(\{a,b\},t'_{-i}\right);\{c,d\},t_{-i}\right)$$
$$=\min\left\{\sqrt{4}+\sqrt{0.5},\sqrt{5}+\sqrt{5},10\sqrt{4}+10\sqrt{0.5},10\sqrt{5}+10\sqrt{5}\right\}$$
$$=2.707<2.732=$$
$$\min\left\{\sqrt{5}+\sqrt{5},\sqrt{3}+\sqrt{1},10\sqrt{5}+10\sqrt{5},10\sqrt{3}+10\sqrt{1}\right\}=$$
$$\min_{t_{-i},t'_{-i}\in T_{-i}} u_i\left(x_i\left(\{c,d\},t'_{-i}\right);\{c,d\},t_{-i}\right).$$

Clearly, the allocation x is maximin incentive compatible.

However, x is not implementable as a maximin equilibrium. Indeed, the direct revelation mechanism $\Gamma=\langle I,S,x-e,\{g_i\}_{i\in I},\{v_i\}_{i\in I}\rangle$ has only one maximin equilibrium. The equilibrium is s with $s_1(\{a,d\})=s_1(\{b,c\})=\{b,c\}$, and $s_2(\{a,b\})=s_2(\{c,d\})=\{a,b\}$. That is, the agents' agreed state is always b, and the redistribution is $x(b)-e(b)$. Consequently, the allocation

$$y=\begin{pmatrix} y_1(a), & y_1(b), & y_1(c), & y_1(d) \\ y_2(a), & y_2(b), & y_2(c), & y_2(d)\end{pmatrix}=\begin{pmatrix}(7,7); & (5,5); & (5,5); & (0,7)\\ (7,6); & (7,6); & (5,5); & (4,0)\end{pmatrix}$$

is the unique maximin equilibrium outcome. Clearly, y is different from x.

Below we outline an alternative proof of our implementation result using the result of de Castro and Yannelis (2009) (Theorem 2 above).

Lemma 1. *If an allocation x is maximin incentive compatible, then it is also implementable as a maximin equilibrium of the mechanism $\Gamma=\langle I,S,x-e,\{g_i\}_{i\in I},\{v_i\}_{i\in I}\rangle$, provided that the initial endowment e_i is private information measurable for each i (i.e., for each i, $e_i(t_i,t_{-i})=e_i(t_i,t'_{-i})$ for all t_{-i} and t'_{-i}).*

Proof. Suppose that the allocation x is maximin incentive compatible for the economy. Furthermore, assume that x is not implementable as a maximin equilibrium through $\Gamma=\langle I,S,x-e,\{g_i\}_{i\in I},\{v_i\}_{i\in I}\rangle$. We will reach a contradiction, with which we conclude the proof.

Since each player's type set is a finite set, the mechanism Γ always has a maximin equilibrium. The allocation x is not implementable implies that for every $s^*\in\mathbb{ME}(\Gamma)$, there exist an $i\in I$ and a $t\in T$ such that $g_i(s^*(t),t)\neq x_i(t)$. It follows that the truth telling strategy profile s^T (Definition 11) is not a maximin equilibrium. That is, there exists a player i, a type t'_i, and a lie t''_i, such that

$$\min_{t'_{-i},t_{-i}\in T_{-i}} v_i\left(t'_i,t'_{-i};t'_i,t_{-i}\right)<\min_{t'_{-i},t_{-i}\in T_{-i}} v_i\left(t''_i,t'_{-i};t'_i,t_{-i}\right). \tag{13}$$

Based on (3), (4) and (5), we can rewrite the left hand side of (13) as

$$\begin{aligned}\min_{t'_{-i},t_{-i}\in T_{-i}} v_i\left(t'_i,t'_{-i};t'_i,t_{-i}\right) &= \min_{t'_{-i},t_{-i}\in T_{-i}} u_i\left(g_i\left(\left(t'_i,t'_{-i}\right),t'_i,t_{-i}\right);t'_i,t_{-i}\right)\\ &= \min_{t'_{-i},t_{-i}\in T_{-i}} u_i\left(e_i\left(t'_i,t_{-i}\right)+D_i\left(x-e,\left(t'_i,t'_{-i}\right)\right);t'_i,t_{-i}\right)\\ &= \min_{t'_{-i},t_{-i}\in T_{-i}} u_i\left(e_i\left(t'_i,t_{-i}\right)+x_i\left(t'_i,t'_{-i}\right)-e_i\left(t'_i,t'_{-i}\right);t'_i,t_{-i}\right)\\ &= \min_{t'_{-i},t_{-i}\in T_{-i}} u_i\left(x_i\left(t'_i,t'_{-i}\right);t'_i,t_{-i}\right).\end{aligned}$$

Also, we can rewrite the right hand side of (13) as

$$\begin{aligned}\min_{t'_{-i},t_{-i}\in T_{-i}} v_i\left(t''_i,t'_{-i};t'_i,t_{-i}\right) &= \min_{t'_{-i},t_{-i}\in T_{-i}} u_i\left(g_i\left(\left(t''_i,t'_{-i}\right),t'_i,t_{-i}\right);t'_i,t_{-i}\right)\\ &= \min_{t'_{-i},t_{-i}\in T_{-i}} u_i\left(e_i\left(t'_i,t_{-i}\right)+D_i\left(x-e,\left(t''_i,t'_{-i}\right)\right);t'_i,t_{-i}\right)\\ &= \min_{t'_{-i},t_{-i}\in T_{-i}} u_i\left(e_i\left(t'_i,t_{-i}\right)+x_i\left(t''_i,t'_{-i}\right)-e_i\left(t''_i,t'_{-i}\right);t'_i,t_{-i}\right).\end{aligned}$$

That is, we have

$$\begin{aligned}&\min_{t_{-i},t'_{-i}\in T_{-i}} u_i\left(e_i\left(t'_i,t_{-i}\right)+x_i\left(t''_i,t'_{-i}\right)-e_i\left(t''_i,t'_{-i}\right);t'_i,t_{-i}\right)\\ &> \min_{t_{-i},t'_{-i}\in T_{-i}} u_i\left(x_i\left(t'_i,t'_{-i}\right);t'_i,t_{-i}\right). \end{aligned}\tag{14}$$

Finally, since the initial endowment e_i is private information measurable, (14) says that the allocation x is not maximin incentive compatible, which is a contradiction. □

Lemma 2. *Suppose the ex post utility functions are private information measurable (i.e., for each i, $u_i(c_i;t_i,t_{-i})=u_i\left(c_i;t_i,t'_{-i}\right)$ for all t_{-i} and t'_{-i}), and Assumption 1 holds. If allocation x is ex ante maximin efficient (as of Definition 3), then it is interim maximin efficient in the sense of de Castro–Yannelis.*

Proof. Suppose that the allocation x is ex ante maximin efficient and not interim maximin efficient. We will reach a contradiction, which concludes the proof.

When the ex post utility functions are private information measurable, the allocation x is not interim maximin efficient implies that there exists a feasible allocation $y=(y_i)_{i\in I}$ such that

$$\min_{t_{-i}\in T_{-i}} u_i\left(y_i(t_i,t_{-i});t_i,t_{-i}\right) \geq \min_{t_{-i}\in T_{-i}} u_i\left(x_i(t_i,t_{-i});t_i,t_{-i}\right) \tag{15}$$

for every i and $t_i\in T_i$, with strict inequality for some i and t_i.

Since $\mu_i(\cdot)>0$ by Assumption 1, we have that, for all i, $y_i\succeq_i^{MP} x_i$,

$$\sum_{t_i\in T_i}\left(\min_{t_{-i}\in T_{-i}} u_i\left(y_i(t_i,t_{-i});t_i,t_{-i}\right)\right)\mu_i(t_i) \geq \sum_{t_i\in T_i}\left(\min_{t_{-i}\in T_{-i}} u_i\left(x_i(t_i,t_{-i});t_i,t_{-i}\right)\right)\mu_i(t_i)$$

and $y_i\succ_i^{MP} x_i$ for at least one i. That is, x is not ex-ante maximin efficient, which is a contradiction. □

Now, we provide an alternative indirect proof of our main theorem using the result of de Castro–Yannelis (Theorem 2 above). Let x be a maximin individually rational and ex-ante maximin efficient allocation. By Lemma 2, x is interim maximin efficient in the sense of de Castro–Yannelis, provided that the agents' ex post utility functions are private information measurable and Assumption 1 holds. Now, by the result of de Castro–Yannelis (Theorem 2 above), x is maximin incentive compatible. Finally, by Lemma 1, x is implementable as a maximin equilibrium, provided that the agents' initial endowments are private information measurable.

5. Relationship with robust implementation (Bergemann and Morris)

The seminal work of Bergemann and Morris (2005) characterizes the environments, in which a social choice function robustly satisfies the interim incentive constraints, i.e., satisfies the interim incentive constraints for any type space, if and only if the ex post incentive constraints are satisfied. In Bergemann and Morris (2005), an agent's type contains a description of his beliefs and his payoff type. The net transfer $x-e$ of Example 3 below is not ex post incentive compatible, where x

is maximin individually rational and ex-ante maximin efficient. Therefore, $x-e$ is not robustly interim incentive compatible in the sense of Bergemann and Morris (2005).

In Bergemann and Morris (2005), a social choice function f satisfies *ex post incentive compatibility* if for all i, t and t_i':

$$\hat{u}_i\left(f\left(t_i,t_{-i}\right),t_i,t_{-i}\right)\geq \hat{u}_i\left(f\left(t_i',t_{-i}\right),t_i,t_{-i}\right).$$

Also, in Bergemann and Morris (2009), a social choice function f satisfies *strict ex post incentive compatibility* if for all i, $t_i'\neq t_i$ and t_{-i}:

$$\hat{u}_i\left(f\left(t_i,t_{-i}\right),t_i,t_{-i}\right)> \hat{u}_i\left(f\left(t_i',t_{-i}\right),t_i,t_{-i}\right).$$

These notions can be applied to our context directly, by letting $f=x-e$, and $\hat{u}_i\left(f\left(t_i',t_{-i}\right),t_i,t_{-i}\right)=u_i\big(e_i\left(t_i,t_{-i}\right)+x\left(t_i',t_{-i}\right)-e\left(t_i',t_{-i}\right),t_i,t_{-i}\big)$.

Example 3. There are two agents 1, 2, two commodities c_1, c_2, and four possible states of nature $T=\{a,b,c,d\}$. The ex post utility functions of the agents are

$$u_1\left(c_1^1,c_1^2;t\right)=\begin{cases}1.22\sqrt{c_1^1}+1.22\sqrt{c_1^2} & \text{if } t\in\{a,b\}\\ \sqrt{c_1^1}+1.05\sqrt{c_1^2} & \text{if } t\in\{c,d\}\end{cases}$$

$$u_2\left(c_2^1,c_2^2;t\right)=\begin{cases}1.02\sqrt{c_2^1}+\sqrt{c_2^2} & \text{if } t\in\{a,d\}\\ 1.2\sqrt{c_2^1}+\sqrt{c_2^2} & \text{if } t\in\{b,c\}.\end{cases}$$

The agents' random initial endowments and type sets are:

$$\begin{aligned}&\left(e_1(a),e_1(b),e_1(c),e_1(d)\right)=\left[(8,6);(8,6);(6,8);(6,8)\right]; && T_1=\{\{a,b\},\{c,d\}\}\\ &\left(e_2(a),e_2(b),e_2(c),e_2(d)\right)=\left[(4,10);(6,6);(6,6);(4,10)\right]; && T_2=\{\{a,d\},\{b,c\}\}.\end{aligned}$$

That is, when the state of nature is a, agent 1 is of type $t_1=\{a,b\}$ and agent 2 is of type $t_2=\{a,d\}$. Furthermore, agent 1 assigns $\mu_1(\{a,b\})=\frac{1}{3}$ and $\mu_1(\{c,d\})=\frac{2}{3}$ to his types, and agent 2 assigns $\mu_2(\{a,d\})=\frac{2}{3}$ and $\mu_2(\{b,c\})=\frac{1}{3}$ to her types.

The allocation

$$\begin{aligned}x&=\begin{pmatrix}x_1(a), & x_1(b), & x_1(c), & x_1(d)\\ x_2(a), & x_2(b), & x_2(c), & x_2(d)\end{pmatrix}\\ &=\begin{pmatrix}(6.8,7.51); & (7.116,7.178); & (5.8,8.4); & (4.901,9.439)\\ (5.2,8.49); & (6.884,4.822); & (6.2,5.6); & (5.099,8.561)\end{pmatrix}\end{aligned}$$

is individually rational and ex-ante efficient with respect to the maximin preferences. Its corresponding planned redistribution is

$$\begin{aligned}x-e&=\begin{pmatrix}x_1(a)-e_1(a), & x_1(b)-e_1(b), & x_1(c)-e_1(c), & x_1(d)-e_1(d)\\ x_2(a)-e_2(a), & x_2(b)-e_2(b), & x_2(c)-e_2(c), & x_2(d)-e_2(d)\end{pmatrix}\\ &=\begin{pmatrix}(-1.2,1.51); & (-0.884,1.178); & (-0.2,0.4); & (-1.099,1.439)\\ (1.2,-1.51); & (0.884,-1.178); & (0.2,-0.4); & (1.099,-1.439)\end{pmatrix},\end{aligned}$$

which is not ex post incentive compatible. Indeed, at state a, agent 1's type is $\{a,b\}$. Reporting the truth (i.e., reporting $\{a,b\}$) gives him a payoff of

$$\begin{aligned}u_1\left(g\left(\{a,b\},\{a,d\},a\right);a\right)&=u_1\left(e_1(a)+x_1(a)-e_1(a);a\right)\\ &=u_1\left((6.8,7.51);a\right)\\ &=1.22\sqrt{6.8}+1.22\sqrt{7.51}=6.525,\end{aligned}$$

which is strictly smaller than his payoff of telling a lie (i.e., reporting $\{c,d\}$)

$$\begin{aligned}u_1\left(g\left(\{c,d\},\{a,d\},a\right);a\right)&=u_1\left(e_1(a)+x_1(d)-e_1(d);a\right)\\ &=u_1\left((6.901,7.439);a\right)\\ &=1.22\sqrt{6.901}+1.22\sqrt{7.439}=6.532.\end{aligned}$$

Furthermore, the readers may wonder if there is any relationship of our ambiguous implementation with Bergemann and Morris (2009) on robust implementation. We will show that the maximin implementation of this paper and the robust implementation are different.

The main difference is that we use the maximin preferences and the maximin equilibrium concept, whereas robust implementation uses the Bayesian preferences and the idea of iterative elimination of never best responses. In our framework, each player takes into account the worst actions of all the other agents against him and also the worst state that can occur. The method of iterative elimination of never best responses is defined for players with Bayesian preferences. Each player has a set of probability assessments on the unknowns. A player considers *each* probability in the set in the process of deletion. A social choice function x is robustly implemented if at each state of nature t, the survived messages, through the outcome function, give us the correct consumption bundle, i.e., $x(t)$. In other words, x has to be the unique outcome of the mechanism. See Bergemann and Morris (2009, 2011).

Robust implementation may not be possible in our context. In Example 3, the maximin individually rational and maximin ex-ante efficient allocation x is implementable as a maximin equilibrium. However, the responsive social choice function $x-e$ in Example 3 is not strictly ex post incentive compatible, and by Theorem 2 of Bergemann and Morris (2009), $x-e$ is not robustly implementable.

6. Concluding remarks

We showed that each individually rational and ex-ante maximin efficient allocation is implementable by means of noncooperative behavior under ambiguity. That is, given any arbitrary individually rational and ex-ante maximin efficient allocation, its corresponding direct revelation mechanism yields the allocation as its unique maximin equilibrium outcome. As a consequence, any maximin core allocation, maximin value allocation and maximin Walrasian equilibrium allocation is implementable.

The new equilibrium notion (maximin equilibrium) takes into account the agents' information constraints – the inability to assign a probability to every state of nature, and to each possible action of his opponents. In a maximin equilibrium, each agent maximizes his payoff lowest bound, i.e., each agent maximizes the payoff that takes into account the worst actions of all the other agents against him and also the worst state that can occur. It turns out that, such a noncooperative behavior (i.e., the maximin equilibrium) enables agents to reach a desirable outcome, i.e., an individually rational and ex-ante maximin efficient allocation, which is also incentive compatible.

The method of iterative elimination as used in robust implementation (Bergemann and Morris, 2009) does not give the same result in our context. Since this method is defined for players with Bayesian preferences, it is an open question if the introduction of maximin preferences on the method of iterative deletion will provide similar results with ours. We conjecture that one way to proceed is to regard the unknowns as the realized state of nature and the actions of all the other players. Thus, having maximin preferences means that each player chooses the best responses taking into account the *worst* probability assessment on the unknowns. When taking into account all possible probabilities, this is equivalent to each player maximizing his payoff taking into account the worst actions of all the other agents against him and also the worst state that can occur, i.e., the maximin equilibrium notion. However, this is a topic for further research.

It is not yet known whether or not the result of this paper holds in the presence of infinitely many states. The difficulty arises from the fact that the minimum of the utility over even countably many states may not exist. Also, the implementation of interim efficient notions seem to be an open question. Some progress in this direction was made in Liu (2015), who showed that the maximin rational expectations equilibrium is implementable as a maximin equilibrium.

References

Angelopoulos, A., Koutsougeras, L.C., 2015. Value allocation under ambiguity. Econ. Theory 59 (1), 147–167.
Aryal, G., Stauber, R., 2014. Trembles in extensive games with ambiguity averse players. Econ. Theory 57, 1–40.
Bergemann, D., Morris, S., 2005. Robust mechanism design. Econometrica 73 (6), 1771–1813.
Bergemann, D., Morris, S., 2009. Robust implementation in direct mechanisms. Rev. Econ. Stud. 76, 1175–1204.
Bergemann, D., Morris, S., 2011. Robust implementation in general mechanisms. Games Econ. Behav. 71 (2), 261–281.
Bodoh-Creed, A., 2012. Ambiguous beliefs and mechanism design. Games Econ. Behav. 75, 518–537.
Bose, S., Ozdenoren, E., Pape, A., 2006. Optimal auctions with ambiguity. Theor. Econ. 1, 411–438.
Bose, S., Renou, L., 2014. Mechanism design with ambiguous communication devices. Econometrica 82 (5), 1853–1872.
Cohen, M., Meilijson, I., 2014. Preference for safety under the Choquet model: in search of a characterization. Econ. Theory 55, 619–642.
Dasgupta, P.S., Hammond, P.J., Maskin, E.S., 1979. The implementation of social choice rules: some general results on incentive compatibility. Rev. Econ. Stud. 46 (2), 185–216.
de Castro, L.I., Yannelis, N.C., 2009. Ambiguity aversion solves the conflict between efficiency and incentive compatibility. Working paper.
de Castro, L.I., Yannelis, N.C., 2013. An interpretation of Ellsberg's Paradox based on information and incompleteness. Econ. Theory Bull. 1 (2), 139–144.
de Castro, L.I., Pesce, M., Yannelis, N.C., 2011. Core and equilibria under ambiguity. Econ. Theory 48, 519–548.
de Castro, L.I., Pesce, M., Yannelis, N.C., 2014. Rational expectations under ambiguity. Working paper.
Dominiak, A., Duersch, P., Lefort, J.-P., 2012. A dynamic Ellsberg urn experiment. Games Econ. Behav. 75 (2), 625–638.
Ellsberg, D., 1961. Risk, ambiguity, and the Savage axioms. Quart. J. Econ. 75 (4), 643–669.
Ellsberg, D., Machina, M.J., Ritzberger, K., Yannelis, N.C., 2011. Symposium on the 50th anniversary of the Ellsberg Paradox. Econ. Theory 48 (2), 219–548.
Even, Y., Lehrer, E., 2014. Decomposition-integral: unifying Choquet and the concave integrals. Econ. Theory 56, 33–58.
Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. J. Math. Econ. 18 (2), 141–153.
Hahn, G., Yannelis, N.C., 2001. Coalitional Bayesian Nash implementation in differential information economies. Econ. Theory 18, 485–509.

Haisley, E.C., Weber, R.A., 2010. Self-serving interpretations of ambiguity in other-regarding behavior. Games Econ. Behav. 68 (2), 614–625.
Hansen, L.P., Sargent, T.J., 2001. Acknowledging misspecification in macroeconomic theory. Rev. Econ. Dynam. 4, 519–535.
He, W., Yannelis, N.C., 2015. Equilibrium theory under ambiguity. J. Math. Econ. 61, 86–95.
Holmström, B., Myerson, R.B., 1983. Efficient and durable decision rules with incomplete information. Econometrica 51 (6), 1799–1819.
Ivanov, A., 2011. Attitudes to ambiguity in one-shot normal-form games: an experimental study. Games Econ. Behav. 71 (2), 366–394.
Jackson, M.O., 1991. Bayesian implementation. Econometrica 59 (2), 461–477.
Jackson, M., Swindles, J., 2005. Existence of equilibrium in single and double private value auctions. Econometrica 73 (1), 93–139.
Ledyard, J.O., 1977. Incentive compatible behavior in core-selecting organizations. Econometrica 45 (7), 1607–1621.
Ledyard, J.O., 1978. Incentive compatibility and incomplete information. J. Econ. Theory 18, 171–189.
Liu, Z., 2014. A note on the welfare of the maximin rational expectations. Econ. Theory Bull. 2, 213–218.
Liu, Z., 2015. Implementation of maximin rational expectations equilibrium. Econ. Theory. http://dx.doi.org/10.1007/s00199-015-0932-5.
Liu, Z., Yannelis, N.C., 2015. Ambiguous implementation: the partition model. Working paper.
Palfrey, T.R., Srivastava, S., 1989. Implementation with incomplete information in exchange economies. Econometrica 57 (1), 115–134.
Wald, A., 1950. Statistical Decision Functions. John Wiley and Sons, New York.