MECHANISM DESIGN WITH AMBIGUOUS COMMUNICATION DEVICES

BY SUBIR BOSE AND LUDOVIC RENOU¹

This paper considers mechanism design problems in environments with ambiguity-sensitive individuals. The novel idea is to introduce ambiguity in mechanisms so as to exploit the ambiguity sensitivity of individuals. Deliberate engineering of ambiguity, through ambiguous mediated communication, can allow (partial) implementation of social choice functions that are not incentive compatible with respect to prior beliefs. We provide a complete characterization of social choice functions partially implementable by ambiguous mechanisms.

KEYWORDS: Mechanism design, ambiguity, communication device, multiple priors.

1. INTRODUCTION

THIS PAPER STUDIES MECHANISM DESIGN PROBLEMS in environments with ambiguity-sensitive individuals. The central and novel aspect of our analysis is to consider *ambiguous* mechanisms. That is, we give the designer the possibility to *endogenously* engineer ambiguity in its mechanisms. After all, if individuals are ambiguity-sensitive, why should we preclude the use of ambiguous mechanisms if it helps the designer in achieving its goals? As an example, consider a tax evasion problem.² Tax authorities can be ambiguous about the likelihood of auditing taxpayers, for example, the tax authority can audit a taxpayer based on the draw of a ball from an urn containing a number of balls of different colors with unknown proportions (an “Ellsberg” urn). Would this dose of ambiguity help? As another example, consider the main refinancing operations of the European Central Bank (henceforth, ECB). To provide short-term liquidities, the ECB organizes weekly tender auctions, typically held every Tuesday. The ECB regulations carefully and meticulously specify all aspects of the tender auctions, so that there is little scope for ambiguity at this stage. However, prior to each weekly auction, the ECB communicates with the eligible counterparts about their liquidity needs. And the communication can be ambiguous. Again, would such ambiguity help the ECB in steering the monetary policy? These questions echo the main theme of this paper: does the introduction of ambiguity in mechanism design help the designer in achieving its desired goals?

A natural starting point is to introduce ambiguity in the allocation rule. While simple and indeed appealing, the introduction of ambiguity in the allocation rule does not enlarge the set of (partially) implementable social choice

¹We are particularly indebted to Matthew O. Jackson and three referees for perceptive comments and suggestions; the paper has considerably improved thanks to them. We are grateful to Ehud Lehrer for an insightful exchange that led us to Theorem 2. We also thank Vince Crawford, Peter Klibanoff, Rida Laraki, Vina Liu, Sujoy Mukerji, Jérôme Renault, and Tristan Tomala.

²See Lang and Wambach (2013) for a treatment of this problem.

functions.³ The intuition is simple. While one can design an allocation rule that *formally* depends on some extraneous random variables (e.g., the draw from an “Ellsberg” urn), the (equilibrium) outcome has to be invariant across the realized values of these random variables, if the mechanism is to implement the given social choice function. Hence, a mechanism with an ambiguous allocation rule cannot (in itself) implement a social choice function that is not implementable by a classical (i.e., unambiguous) mechanism.

Nevertheless, we show that, through the use of appropriately constructed ambiguous mechanisms, it is possible to implement social choice functions that are not implementable by using a standard, unambiguous mechanism. The critical ingredient in such construction is to add an *ambiguous* mediated communication stage prior to the *unambiguous* allocation stage. Classical mediated communication consists of sets of messages that agents can send to the designer, sets of messages that agents can receive from the designer, and a (i.e., *single*) mapping from messages received by the designer to probabilities over messages sent by the designer. Ambiguous mediated communication differs by having instead a *set* of mappings from messages received by the designer to probabilities over messages sent. In essence, when an agent receives a message from the designer, he is ambiguous about the exact communication device that has been used by the designer. Our ambiguous mechanism is thus a mediated extension of a classical mechanism. However, unlike a classical mediated extension (Myerson (1986), Forges (1986, 1990)), the extensions we consider are ambiguous. And the basic idea behind the usefulness of ambiguous communication devices is quite simple: even when a social choice function is not implementable given (possibly ambiguous) prior beliefs, the posterior beliefs resulting from the communication may be such that it becomes implementable.

The main contribution of this paper is to provide characterization results for the class of social choice functions implementable by ambiguous mechanisms. An ambiguous “direct” mechanism is a direct revelation mechanism for the allocation stage, preceded by one round of ambiguous mediated communication where the set of possible messages the agents can send to the communication device is restricted to the set of types. Theorem 1 shows that a social choice function is implementable by an ambiguous mechanism if and only if it is implementable by an ambiguous “direct” mechanism. In a nutshell, for a social choice function to be implementable, it must be incentive compatible with respect to some set of (possibly ambiguous) beliefs. Furthermore, the messages sent and received at the communication stage must be such that the required posterior belief sets are indeed generated. Theorem 2 continues the characterization result: it provides conditions as to when a target set of posteriors can be obtained from a given set of priors through the use of some ambiguous communication devices.

³In the sequel, we omit the qualifier “partial” and write “implementation” for “partial implementation.”

This paper contributes to the growing literature on mechanism design with ambiguity-sensitive individuals, for example, Bodoh-Creed (2012), Bose, Ozdenoren, and Pape (2006), Bose and Daripa (2009), De Castro and Yannelis (2012), Lopomo, Rigotti, and Shannon (2010), Salo and Weber (1995), to name a few. This literature differs from our analysis in one fundamental aspect, however. While we give the designer the possibility to endogenously engineer ambiguity, this literature constrains the designer to offer classical unambiguous mechanisms (often static, moreover). In particular, if there is *no* exogenous ambiguity, a social choice function is implementable by an unambiguous mechanism if and only if it is (classically) incentive compatible (Myerson (1979)). In contrast, even when there is no exogenous ambiguity, the set of implementable social choice functions can enlarge when ambiguous mechanisms are allowed.⁴

2. AN INTRODUCTORY EXAMPLE

This section illustrates our main results with the help of a simple example. There are two players, labeled 1 and 2, two (payoff-relevant) types θ and θ' for each player, and two alternatives x and y . Types are private information. Throughout, we write “ i ” for “ $i \in \{1, 2\}$ ” and omit the qualifier “ $j \neq i$ ” whenever no confusion arises.

We assume that players have multiple-prior preferences (Gilboa and Schmeidler (1989)) and prior-by-prior updating (full Bayesian updating); see Pires (2002) for an axiomatization. Let P_i be the set of priors of player i about player j 's types. Player i 's payoff function u_i is defined by $u_i(y, \theta) = u_i(x, \theta') = 1$ and $u_i(x, \theta) = u_i(y, \theta') = 0$.

The designer aims at (partially) implementing the social choice function f given by $f(\theta, \theta) = f(\theta', \theta') = x$, and $f(\theta, \theta') = f(\theta', \theta) = y$.

To crisply illustrate the role of ambiguity in mechanism design, we consider the starkest possible case whereby there is no prior ambiguity, that is, P_i is the singleton $\{(p_i, 1 - p_i)\}$, with p_i the belief of player i about the type of player j being θ . For concreteness, assume that $p_i = 2/3$.

Clearly, no classical mechanism implements the social choice function f since it fails to be incentive compatible. Yet, we argue that introducing some ambiguity in an otherwise classical mechanism makes it possible to implement f .

⁴A notable exception is Di Tillio, Kos, and Messner (2011). These authors also explored the effects of introducing ambiguity in mechanisms. There are important differences, though. First, they focused on the revenue maximization problem in a seller–buyer setup, whereas our objective is to characterize the entire set of social choice functions implementable by ambiguous mechanisms. Second, they introduced ambiguity in the allocation rule. As mentioned earlier, this has no effect for the implementation of social choice functions; hence our focus on the introduction of ambiguity at the pre-allocation stage. Nonetheless, as their work shows, introducing ambiguity at the allocation stage can expand the set of social choice *sets* that are implementable.

The central idea of the paper is to add a communication stage prior to the allocation stage and to introduce ambiguity at the communication stage. To see how this is achieved, consider first a *hypothetical* situation where player i has ambiguous beliefs about the type of player j . Specifically, suppose the set of beliefs of player i is $\Pi_i = \{(0, 1), (1, 0)\}$, that is, he either believes that player j 's type is θ with probability 1, or is θ' with probability 1.⁵ With such beliefs, truth-telling is a max-min equilibrium of the direct revelation game and, therefore, the social choice function f would be implementable if the players had the set of beliefs $\{(0, 1), (1, 0)\}$.

Given that the players' initial beliefs are, however, not $\{(0, 1), (1, 0)\}$, the task is to see how, starting from the prior belief $(2/3, 1/3)$, we can generate the set of posterior beliefs $\{(0, 1), (1, 0)\}$. We now construct an ambiguous communication device that generates the appropriate posteriors. Assume that each player can send message θ or θ' to the designer, can receive message ω or ω' from the designer, and that there are two possible probability systems λ and λ' . Denote $\lambda_i(\omega_i|\theta_j)$ the probability that player i receives the message $\omega_i \in \{\omega, \omega'\}$ conditional on player j 's message $\theta_j \in \{\theta, \theta'\}$ and let $\lambda((\omega_1, \omega_2)|(\theta_1, \theta_2)) = \lambda_1(\omega_1|\theta_2)\lambda_2(\omega_2|\theta_1)$ for each possible (ω_1, ω_2) and (θ_1, θ_2) . Assuming $\lambda_i(\omega|\theta) = 1$ and $\lambda_i(\omega'|\theta') = 1$ thus fully defines the first probability system λ . Similarly, the second probability system λ' is fully specified by $\lambda'_i(\omega|\theta') = 1$ and $\lambda'_i(\omega'|\theta) = 1$. Clearly, if the probability system is λ , player i 's posterior belief is $(1, 0)$ if he receives the message ω , and $(0, 1)$ otherwise. Alternatively, if the probability system is λ' , player i 's posterior belief is $(0, 1)$ if he receives the message ω , and $(1, 0)$ otherwise. Thus, *regardless* of the message received, player i 's set of posteriors is $\{(0, 1), (1, 0)\}$, as required.

It thus remains to show that a stage of mediated and ambiguous communication prior to the allocation stage indeed implements f . Formally, the ambiguous mechanism is as follows. In the first stage, players communicate with the designer through the ambiguous communication device constructed above. In the second stage, players announce their types to the designer, which then implements an alternative according to f . It is important to stress that the allocation depends only on the second-stage reports. The implementing mechanism is thus a classical direct mechanism extended by one round of mediated ambiguous communication.

We now argue that there exists an equilibrium implementing f . Roughly speaking, at an equilibrium of such a dynamic game, a player chooses a strategy at the communication stage among the strategies that he will actually follow at the allocation stage. Consider the following strategies. A player truthfully reports his type to the communication device at the first stage and, regardless of the messages sent and received in the first stage, he again truthfully reports his type at the second stage. It follows from the construction of the communication device that player i 's beliefs at the second stage are given by Π_i , regardless

⁵Alternatively, consider the convex hull of Π_i .

of the messages he has sent and received in the first stage. Thus, player i has an incentive to truthfully reveal his type at the second stage. Lastly, since player i expects the social choice function f to be implemented in the second stage, regardless of the messages sent and received, he has no incentive to lie at the first stage. The social choice function f is thus implementable by an ambiguous mechanism.

We end this section by remarking that our construction continues to apply even with ambiguous priors. For instance, suppose that player i has the set of priors $\text{co}\{(\underline{p}_i, 1 - \underline{p}_i), (\bar{p}_i, 1 - \bar{p}_i)\}$, where $\text{co } Y$ denotes the convex hull of Y . If $\underline{p}_i \leq 1 - \bar{p}_i$, the social choice function is incentive compatible and can be implemented using a classical mechanism. Alternatively, if $\underline{p}_i > 1 - \bar{p}_i$, the social choice function is not incentive compatible, but the ambiguous communication device constructed above still generates the set of posteriors Π_i , so that the social choice function remains implementable. The rest of the paper characterizes the social choice functions that can be implemented with ambiguous communication devices.

3. PRELIMINARIES

Notations

For any collection of n sets, $Y_1, \dots, Y_n$, we let $Y := \times_{j \in \{1, \dots, n\}} Y_j$ and $Y_{-i} := \times_{j \in \{1, \dots, n\} \setminus \{i\}} Y_j$, with generic element y and y_{-i} , respectively. For any measurable space $(Y, \mathbb{B}_Y)$, we denote $\Delta(Y)$ the set of probability measures over Y . Let $\mu: Y \rightarrow \Delta(Y^*)$ be a measurable function between $(Y, \mathbb{B}_Y)$ and $(\Delta(Y^*), \mathbb{B}_{\Delta(Y^*)})$. For any event $E \in \mathbb{B}_{Y^*}$, we write $\mu(y)[E]$ for the probability of the event E according to $\mu(y)$. Throughout the paper, we assume that most sets (types, alternatives, messages, etc.) are finite, so as to avoid unnecessary technicalities (e.g., measurability of strategies, conditioning on events with measure zero, etc.). Most of our results extend naturally to economic problems with continuous spaces. The reader interested in the formal treatment of such extension is referred to the Supplemental Material (Bose and Renou (2014)).

Environments

An environment is a tuple $\langle N, X, (\Theta_i)_{i \in N} \rangle$, where $N := \{1, \dots, n\}$ is a finite set of n players, X a set of alternatives, and for each player $i \in N$, Θ_i is a finite set of payoff-relevant types. Types are privately known. A social choice function $f: \Theta \rightarrow X$ assigns an alternative to each profile of types.

Preferences and Updating

Let $\mathcal{H}_i := \{h_i: \Theta_{-i} \rightarrow X\}$ be the set of player i 's acts. Player i of type θ_i has type-dependent preferences over $\mathcal{H}_i$; preferences are *ambiguity-sensitive*.

For specificity, we assume that preferences have the max–min expected utility representation (Gilboa and Schmeidler (1989)). More precisely, there exist a payoff function $u_i: X \times \Theta \rightarrow \mathbb{R}$ and a nonempty, convex, and closed set $P_i \subseteq \Delta(\Theta_{-i})$ such that player i of type θ_i evaluates the act h_i as

$$\min_{p_i \in P_i} \sum_{\theta_{-i} \in \Theta_{-i}} u_i(h_i(\theta_{-i}), \theta_i, \theta_{-i}) p_i[\theta_{-i}].$$

We assume full support, that is, P_i is in the relative interior of $\Delta(\Theta_{-i})$, and prior-by-prior updating (full Bayesian updating).

Ambiguous Mechanisms

Ambiguous mechanisms have two essential components: allocation mechanisms and ambiguous communication devices. An *allocation* mechanism is a pair $\langle (M_i)_{i \in N}, g \rangle$ where, for each player i , M_i is a finite set of messages and $g: \times_{i \in N} M_i \rightarrow X$ is an allocation rule. An allocation mechanism is thus a classical static mechanism.

A communication device is a tuple $\langle (\hat{\Omega}_i, \Omega_i)_{i \in N}, \lambda \rangle$, where $\hat{\Omega}_i$ is a (finite) set of messages that player i can send to the communication device, Ω_i is a (finite) set of messages that player i can receive from the communication device, and $\lambda: \times_{i \in N} \hat{\Omega}_i \rightarrow \Delta(\times_{i \in N} \Omega_i)$ is a probability system: $\lambda(\hat{\omega})[\omega]$ is the probability that the profile of messages ω is sent to the players by the communication device, conditional on the profile of messages received $\hat{\omega}$. An *ambiguous* communication device is a tuple $\langle (\hat{\Omega}_i, \Omega_i)_{i \in N}, \Lambda \rangle$, where Λ is a finite set of probability systems.

Given an allocation mechanism $\langle M, g \rangle$, a mediated extension of $\langle M, g \rangle$ is a mechanism in which one round of ambiguous mediated communication is allowed before $\langle M, g \rangle$ is played. Communication is private and simultaneous. We call such a mechanism an *ambiguous mechanism*. It is important to stress that the alternative implemented depends only on the messages reported in the allocation mechanism at the second stage; the *entire* role of the communication stage is to affect players' beliefs.⁶

We use standard concepts of *strategies* and *assessments*, with occasional modifications needed for our purpose. The two restrictions on assessments are that they (i) are consistent with the *set* of initial priors, and (ii) are derived via full Bayesian updating whenever possible.⁷

⁶The communication stage consists of a *single* round of communication; in the Supplemental Material, we consider an extension consisting of multiple rounds of communication and explain in what sense the restriction to a single round is without loss of generality.

⁷In particular, suppose that the strategy of player i of type θ_i requires him to send the message $\hat{\omega}_i$ to the communication device at the first stage, but when he has to send a message m_i at the second stage, player i has observed the history $(\hat{\omega}'_i, \omega_i)$ with $\hat{\omega}'_i \neq \hat{\omega}_i$. This history is obviously

Next, we define the notion of *equilibrium*, called *consistent planning equilibrium*, used in this paper. Intuitively, a consistent planning equilibrium is defined inductively by “backward induction.” It captures the idea that, at the communication stage, a player chooses a strategy among the strategies that he will actually follow at the allocation stage.⁸ For a detailed discussion on the decision-theoretic issues behind consistent planning and dynamic consistency, see the Supplemental Material (Bose and Renou (2014)).

Fix an ambiguous mechanism $\langle (\widehat{\Omega}_i, \Omega_i)_{i \in N}, \Lambda, \langle (M_i)_{i \in N}, g \rangle \rangle$. Denote $\{\emptyset\}$ the initial history and $H_i := \widehat{\Omega}_i \times \Omega_i$ the set of possible histories of messages sent and received by player i during the communication phase. A (pure) strategy for player i is a mapping $s_i^* : (\{\emptyset\} \cup H_i) \times \Theta_i \rightarrow \widehat{\Omega}_i \cup M_i$, with $s_i^*(\{\emptyset\}, \theta_i) \in \widehat{\Omega}_i$ and $s_i^*(h_i, \theta_i) \in M_i$, for all θ_i and h_i . In the sequel, we write $\widehat{\omega}_i^*(\theta_i)$ for the message sent at the initial history by player i of type θ_i under s_i^* , that is, $s_i^*(\{\emptyset\}, \theta_i)$. Let $\Pi_i^{H, \Theta}(h_i, \theta_i) \subseteq \Delta(H_{-i} \times \Theta_{-i})$ denote the set of beliefs of player i , conditional on type θ_i and history h_i , about the types θ_{-i} of his opponents and their histories h_{-i} .

DEFINITION 1: A (pure) consistent planning equilibrium is a strategy-assessment pair $(s_i^*, \Pi_i^{H, \Theta})_{i \in N}$ such that, for all $i \in N$, for all $\theta_i \in \Theta_i$, for all $h_i \in H_i$,

$$\begin{aligned} (1) \quad & \min_{\pi_i(h_i, \theta_i) \in \Pi_i^{H, \Theta}(h_i, \theta_i)} \sum_{h_{-i}, \theta_{-i}} u_i(g(s_i^*(h_i, \theta_i), s_{-i}^*(h_{-i}, \theta_{-i})), \theta_i, \theta_{-i}) \\ & \times \pi_i(h_i, \theta_i)[h_{-i}, \theta_{-i}] \\ & \geq \min_{\pi_i(h_i, \theta_i) \in \Pi_i^{H, \Theta}(h_i, \theta_i)} \sum_{h_{-i}, \theta_{-i}} u_i(g(m_i, s_{-i}^*(h_{-i}, \theta_{-i})), \theta_i, \theta_{-i}) \\ & \times \pi_i(h_i, \theta_i)[h_{-i}, \theta_{-i}] \end{aligned}$$

for all $m_i \in M_i$, and

$$\begin{aligned} (2) \quad & \min_{\lambda \in \Lambda, p_i \in P_i} \left(\sum_{\omega_i, \omega_{-i}, \theta_{-i}} u_i(g(s_i^*(\widehat{\omega}_i^*(\theta_i), \omega_i), \theta_i), \right. \\ & \left. s_{-i}^*(\widehat{\omega}_{-i}^*(\theta_{-i}), \omega_{-i}), \theta_{-i}) \right) \\ & \times \lambda((\widehat{\omega}_i^*(\theta_i), \widehat{\omega}_{-i}^*(\theta_{-i}))) [\omega_i, \omega_{-i}] p_i[\theta_{-i}] \end{aligned}$$

“off-path” for player i of type θ_i . Yet, if the history $(\widehat{\omega}_i, \omega_i)$ has “positive probability,” player i of type θ_i can still use Bayes’s rule to make inference about the private histories and types of his opponents. We require him to do so.

⁸For more on consistent planning, see Strotz (1955) and Siniscalchi (2011).

$$\geq \min_{\lambda \in \Lambda, p_i \in P_i} \left(\sum_{\omega_i, \omega_{-i}, \theta_{-i}} u_i \left(g \left(s_i^* \left((\widehat{\omega}_i, \omega_i), \theta_i \right), s_{-i}^* \left((\widehat{\omega}_{-i}^*(\theta_{-i}), \omega_{-i}), \theta_{-i} \right) \right), \theta_i, \theta_{-i} \right) \right. \\ \left. \times \lambda(\widehat{\omega}_i, \widehat{\omega}_{-i}^*(\theta_{-i})) [\omega_i, \omega_{-i}] p_i[\theta_{-i}] \right)$$

for all $\widehat{\omega}_i \in \widehat{\Omega}_i$.

In essence, a consistent planning equilibrium is a perfect Bayesian equilibrium of a game between multiple selves of the same players with max–min expected utility. Equation (1) states that player i has no profitable deviation at the allocation stage, while Equation (2) states that player i has no profitable deviation at the communication stage *given* the strategy he will actually follow at the allocation stage (and thus satisfying Equation (1)). It is worth stressing that the concept of consistent planning equilibrium does not rule out dynamic inconsistency. It is indeed possible that at *every* history h_i (and corresponding posterior beliefs), the equilibrium choices at the allocation stage are suboptimal when evaluated from the viewpoint of the initial history (i.e., according to the prior beliefs).

We say that player i 's history $h_i = (\widehat{\omega}_i, \omega_i)$ has positive probability under the profile of strategies (s_i^*, s_{-i}^*) if there exist a type $\theta_i \in \Theta_i$, a prior $p_i \in P_i$, and a probability system $\lambda \in \Lambda$ such that $s_i^*(\{\emptyset\}, \theta_i) = \widehat{\omega}_i$ and $\sum_{\theta_{-i}, \omega_{-i}} \lambda(\widehat{\omega}_i, s_{-i}^*(\{\emptyset\}, \theta_{-i})) [(\omega_i, \omega_{-i})] p_i[\theta_{-i}] > 0$. The history (h_i, h_{-i}) has positive probability if each h_i has positive probability. We are now ready to define the notion of implementation by ambiguous mechanism.

DEFINITION 2: The ambiguous mechanism $\langle (\widehat{\Omega}_i, \Omega_i)_{i \in N}, \Lambda \rangle, \langle (M_i)_{i \in N}, g \rangle$ (partially) implements the social choice function f if there exists a pure equilibrium $(s^*, \Pi^{H, \Theta})$ such that

$$g(s_i^*(h_i, \theta_i), s_{-i}^*(h_{-i}, \theta_{-i})) = f(\theta_i, \theta_{-i})$$

for all (θ_i, θ_{-i}) , for all (h_i, h_{-i}) having positive probability under s^* .

According to Definition 2, for a social choice function to be implementable, there must exist an ambiguous mechanism and a pure equilibrium such that, for *all* histories of messages sent to and received from the communication device, and for all messages sent at the allocation stage following these histories, the designer implements the correct outcome. In particular, even if player i of type θ_i does not send the message $\widehat{\omega}_i$ at the communication stage (but another of his type does), his (equilibrium) message at the allocation stage following any history $(\widehat{\omega}_i, \omega_i)$ with positive probability must still implement f . We refer the reader to Section 5 and the Supplemental Material for an alternative definition of implementation.

4. CHARACTERIZATION RESULTS

Incentive Compatibility

We first define incentive compatibility.

DEFINITION 3: Let $\Pi_i \subseteq \Delta(\Theta_{-i})$ be a nonempty set of beliefs. The social choice function f is Π_i -incentive compatible if, for each $\theta_i \in \Theta_i$,

$$\begin{aligned} & \min_{\pi_i \in \Pi_i} \sum_{\theta_{-i} \in \Theta_{-i}} u_i(f(\theta_i, \theta_{-i}), \theta_i, \theta_{-i}) \pi_i[\theta_{-i}] \\ & \geq \min_{\pi_i \in \Pi_i} \sum_{\theta_{-i} \in \Theta_{-i}} u_i(f(\theta'_i, \theta_{-i}), \theta_i, \theta_{-i}) \pi_i[\theta_{-i}] \end{aligned}$$

for all $\theta'_i \in \Theta_i$. The social choice function f is incentive compatible if it is Π_i -incentive compatible for some nonempty set of belief Π_i , for each player $i \in N$.

Incentive compatibility is equivalent to truth-telling being a (max–min) equilibrium of the direct revelation game, in which each player i has the set of beliefs Π_i about his opponents' types. When the prior beliefs P_i are not Π_i , we would want one round of (ambiguous) mediated communication to *generate* the required set of beliefs Π_i . In other words, players should have an incentive not only to truthfully reveal their types at the allocation stage, but also to communicate in the “right” way at the communication stage so as to generate the required sets of beliefs.

Before presenting the main results, let us briefly discuss the issue of mixed strategies. The definitions of incentive compatibility and of an equilibrium only invoke pure strategies; deviations to mixed strategies are not even considered. While this is without loss of generality when players are subjective expected utility maximizers, this is *not* so when players are max–min expected utility maximizers. Indeed, with such preferences, a player can be indifferent between two pure strategies and yet strictly prefer a mixture over either of them. Deviations to mixed strategies may therefore be strictly profitable, even when no deviations to pure strategies are.⁹ Thus, in situations where players have access to commitment devices to implement the chosen randomizations (e.g., when they can use websites like stickk.com or delegate to a trusted third party), mixed strategies ought to be considered explicitly. To keep the exposition simple, however, we have chosen to not consider mixed strategies. Extending the definitions to allow for mixed strategies is a straightforward exercise and does not change the qualitative nature of our characterization results. The interested reader is referred to Section S.2 of the Supplemental Material.

⁹For instance, in the introductory example, a player would benefit strictly by randomizing uniformly between θ and θ' , when the set of beliefs is $\{(0, 1), (1, 0)\}$.

The Main Results

Consider the ambiguous communication device $((\Theta_i, \Omega_i)_{i \in N}, \Lambda)$. For each player $i \in N$ and each $\lambda \in \Lambda$, define $\lambda_i: \Theta_i \times \Theta_{-i} \rightarrow \Delta(\Omega_i)$ by $\lambda_i(\theta_i, \theta_{-i})[\omega_i] := \sum_{\omega_{-i}} \lambda(\theta_i, \theta_{-i})[(\omega_i, \omega_{-i})]$, for all $(\theta_i, \theta_{-i}, \omega_i)$. The probability $\lambda_i(\theta_i, \theta_{-i})[\omega_i]$ is the probability of player i receiving the message ω_i from the communication device when $\lambda \in \Lambda$ is the probability system used and (θ_i, θ_{-i}) the profile of types reported by the players. Let Λ_i be the set of such probability systems.

For any $\lambda_i \in \Lambda_i$, for any prior $p_i \in P_i$, let $\zeta_i(p_i, \theta_i, \omega_i, \lambda_i)$ be player i 's posterior belief about his opponents' types, when he sends message θ_i to, and receives message ω_i from, the communication device, that is,

$$\zeta_i(p_i, \theta_i, \omega_i, \lambda_i)[\theta_{-i}] = \frac{\lambda_i(\theta_i, \theta_{-i})[\omega_i] p_i[\theta_{-i}]}{\sum_{\theta_{-i}} \lambda_i(\theta_i, \theta_{-i})[\omega_i] p_i[\theta_{-i}]},$$

if $\sum_{\theta_{-i}} \lambda_i(\theta_i, \theta_{-i})[\omega_i] p_i[\theta_{-i}] > 0$. If $\sum_{\theta_{-i}} \lambda_i(\theta_i, \theta_{-i})[\omega_i] p_i[\theta_{-i}] = 0$, we let $\{\zeta_i(p_i, \theta_i, \omega_i, \lambda_i)\} = \emptyset$.

The following results—Theorems 1 and 2—constitute our characterization results. All proofs are relegated to the [Appendix](#).

THEOREM 1: *The following statements are equivalent:*

1. *The social choice function f is implementable by an ambiguous mechanism.*
2. *For each player $i \in N$, there exists a finite family $(\Pi_i^k)_{k \in K_i}$ of nonempty sets of beliefs such that:*

Condition (IC). For each $k \in K_i$, the social choice function f is Π_i^k -incentive compatible.

Condition (B). There exist a set of messages Ω_i , a set of probability systems Λ_i , and a partition $\{\Phi_i^1, \dots, \Phi_i^k, \dots, \Phi_i^{|K_i|}\}$ of $\Theta_i \times \Omega_i$ such that for all $k \in K_i$, for all $(\theta_i, \omega_i) \in \Phi_i^k$,

$$\bigcup_{p_i \in P_i} \bigcup_{\lambda_i \in \Lambda_i} \{\zeta_i(p_i, \theta_i, \omega_i, \lambda_i)\} = \Pi_i^k.$$

Theorem 1 has the flavor of a revelation principle; it states necessary and sufficient conditions for a social choice function to be implementable by an ambiguous mechanism. In particular, it is without loss of generality to consider direct mechanisms as allocation mechanisms, and type spaces as the sets of messages players can send to the communication device. However, the set of messages that the communication device can send to the players need not be equal to the sets of players' types. According to Theorem 1, for a social choice function to be implementable, it must be incentive compatible with respect to some sets of beliefs (condition (IC)). Moreover, those sets of beliefs must result from the ambiguous communication and, therefore, must satisfy

condition (B). Conversely, to implement a social choice function that satisfies conditions (IC) and (B), consider an ambiguous mechanism consisting of the communication device $\langle (\Theta_i, \Omega_i), \Lambda \rangle$, where each $\lambda \in \Lambda$ is the product of probability systems $\lambda_i \in \Lambda_i$, that is, $\lambda := \bigotimes_{i \in N} \lambda_i$, and the direct allocation mechanism $\langle (\Theta_i)_{i \in N}, f \rangle$. We argue that truth-telling at both the communication and the allocation stages is an equilibrium. To see this, note that regardless of the message sent and received (θ_i, ω_i) at the communication stage, player i 's set of beliefs at the allocation stage is one of the Π_i^k . Since f is Π_i^k -incentive compatible, it is clearly optimal for player i to tell the truth at the allocation stage (given that his opponents are telling the truth). Also, since f is implemented at the allocation stage, regardless of the message sent at the communication stage, it is optimal to tell the truth at the communication stage.

Theorem 1 requires generating a family of posteriors $(\Pi_i^k)_k$, one for each subset Φ_i^k .¹⁰ The next result, Theorem 2, provides conditions as to when the “target” family of posteriors $(\Pi_i^k)_k$ (for which the social choice function is incentive compatible) can be generated from the given set of priors P_i through the use of some ambiguous communication device.

In the sequel, for any posterior π_i and prior p_i , we write π_i/p_i for the posterior-to-prior ratio.

THEOREM 2: *Let $P_i \subseteq \Delta(\Theta_{-i})$ be a nonempty set of priors and $(\Pi_i^k)_{k \in K_i}$ a finite family of posteriors, with $\emptyset \neq \Pi_i^k \subseteq \Delta(\Theta_{-i})$ for each $k \in K_i$. The following statements are equivalent:*

1. Condition (B). *There exist a set of messages Ω_i , a set of probability systems Λ_i , and a partition $\{\Phi_i^1, \dots, \Phi_i^k, \dots, \Phi_i^{|K_i|}\}$ of $\Theta_i \times \Omega_i$ such that, for all $k \in K_i$, for all $(\theta_i, \omega_i) \in \Phi_i^k$,*

$$\bigcup_{p_i \in P_i} \bigcup_{\lambda_i \in \Lambda_i} \{\zeta_i(p_i, \theta_i, \omega_i, \lambda_i)\} = \Pi_i^k.$$

2. *There exist a finite set A_i , a partition $\{A_i^1, \dots, A_i^k, \dots, A_i^{|K_i|}\}$ of A_i into $|K_i|$ subsets, $(\kappa_\alpha : P_i \rightarrow \bigcup_k \Pi_i^k)_{\alpha \in A_i}$ and $(\mu_\alpha : P_i \rightarrow \mathbb{R}_{++})_{\alpha \in A_i}$ such that*

- (i) $\bigcup_{p_i \in P_i} \bigcup_{\alpha \in A_i^k} \{\kappa_\alpha(p_i)\} = \Pi_i^k$ for each k .
- (ii) For each $\alpha \in A_i$, for each pair $(p_i, p'_i) \in P_i \times P_i$,

$$\mu_\alpha(p_i)(\kappa_\alpha(p_i)/p_i) = \mu_\alpha(p'_i)(\kappa_\alpha(p'_i)/p'_i).$$

- (iii) For each $p_i \in P_i$, $p_i = \sum_{\alpha \in A_i} \mu_\alpha(p_i) \kappa_\alpha(p_i)$ and $\sum_{\alpha \in A_i} \mu_\alpha(p_i) = 1$.

The following observation helps in understanding Theorem 2. Condition (B) necessitates the existence of a family Λ_i of probability systems that generate the

¹⁰Observe that each Π_i^k is the finite union of convex sets.

required posteriors, *regardless* of the messages sent and received. It is, however, equivalent to the existence of a *single* probability system that generates the required posteriors by conditioning over *all* messages sent and received.¹¹ This simple observation makes it possible to use classical properties of Bayesian updating, for example, priors are convex combinations of all posteriors, to provide conditions on the set of priors P_i and posteriors (Π_i^k) for condition (B) to hold. And, of course, messages sent and received are nothing else than informative signals about one's opponents' types (and private histories).

So, think of $\alpha \in A_i$ as the realization of a statistical experiment informative about the random variable $\tilde{\theta}_{-i}$, with posterior $\kappa_\alpha(p_i)$ when the prior is p_i . Since any posterior-to-prior ratio $\kappa_\alpha(p_i)/p_i$ must be proportional to the likelihood of $\tilde{\theta}_{-i}$, *all* posterior-to-prior ratios $(\kappa_\alpha(p_i)/p_i)_{p_i \in P_i}$ must be collinear.¹² Condition (ii) reflects this collinearity condition. It has important consequences for the set of posteriors one can generate. Indeed, fix any prior p_i and posterior $\kappa_\alpha(p_i)$. Then, all other posteriors $\kappa_\alpha(p'_i)$ are uniquely determined by the collinearity condition through the system of linear equations: For all $(\theta_{-i}, \theta'_{-i})$,¹³

$$\kappa_\alpha(p'_i)[\theta_{-i}] = \frac{\kappa_\alpha(p_i)[\theta_{-i}]}{\kappa_\alpha(p_i)[\theta'_{-i}]} \frac{p_i[\theta'_{-i}]}{p_i[\theta_{-i}]} \frac{p'_i[\theta_{-i}]}{p'_i[\theta'_{-i}]} \kappa_\alpha(p'_i)[\theta'_{-i}] \quad \text{and}$$

$$\sum_{\theta_{-i}} \kappa_\alpha(p'_i)[\theta_{-i}] = 1.$$

Condition (iii) states that each prior must be a convex combination of some posteriors; this is the classical martingale property of posterior beliefs. In particular, it implies that the set of priors must be in the relative interior of the convex hull of the set of posteriors, that is, $P_i \subseteq \text{relint co}(\bigcup_{k \in K_i} \Pi_i^k)$. This condition explains the need for multiple sets of beliefs $(\Pi_i^k)_{k \in K_i}$, for which a social choice function is incentive compatible. Indeed, suppose that the social choice function f is incentive compatible with respect to Π_i^1 and Π_i^2 . It may well be that P_i is included neither in $\text{relint co} \Pi_i^1$ nor in $\text{relint co} \Pi_i^2$, but is included in $\text{relint co}(\Pi_i^1 \cup \Pi_i^2)$. In such instances, we truly need to generate multiple sets of beliefs. Finally, condition (i) guarantees that all the posteriors are spanned.

An important special case, stated as Corollary 1 below, is when there is a unique prior p_i .

¹¹The formal statement, Lemma 1, is in the [Appendix](#) and the proof is relegated to the Supplemental Material.

¹²Recall that the likelihood of θ_{-i} is proportional to the probability of α given θ_{-i} , $\mathbb{P}(\alpha|\theta_{-i})$. From Bayes's rule, it follows that $\kappa_\alpha(p_i)/p_i \propto \mathbb{P}(\alpha|\cdot)$ for all p_i .

¹³Note that there are $|\Theta_{-i}|^2 + 1$ equations in $|\Theta_{-i}|$ unknowns $(\kappa_\alpha(p'_i)[\theta_{-i}])_{\theta_{-i}}$, but only $|\Theta_{-i}|$ of them are independent. See [Laclau and Renou \(2013\)](#) for a related characterization and additional results.

COROLLARY 1: Let $P_i = \{p_i\}$ and $(\Pi_i^k)_{k \in K_i}$ a finite family of posteriors, with $\emptyset \neq \Pi_i^k \subseteq \Delta(\Theta_{-i})$ for each $k \in K_i$. The following statements are equivalent:

1. Condition (B) holds.
2. $p_i \in \text{relint co}(\bigcup_{k \in K_i} \Pi_i^k)$.

The result follows immediately from the fact that when P_i is a singleton, conditions (ii) and (iii) of Theorem 2 are equivalent to p_i being in the relative interior of the convex hull of $\bigcup_k \Pi_i^k$, and that condition (i) trivially holds.¹⁴ This result is known as the splitting lemma in the literature on repeated games with incomplete information; see [Aumann, Maschler, and Stearns \(1995\)](#). Given a collection $(\Pi_i^k)_k$ for which a social choice function is incentive compatible, it is therefore fairly straightforward to check whether the social choice function is implementable when there is no prior ambiguity.

Another important special case is when the social choice function is $\Delta(\Theta_{-i})$ -incentive compatible, that is, when player i is “maximally” ambiguous about the types of his opponents.¹⁵ (See [De Castro and Yannelis \(2012\)](#) for a study of implementation under maximal ambiguity.)

COROLLARY 2: Any $\Delta(\Theta_{-i})$ -incentive compatible social choice function is implementable by an ambiguous mechanism.

To see this, first note that if a social choice function is $\Delta(\Theta_{-i})$ -incentive compatible, then it is incentive compatible with respect to any subset $\Pi_i \subseteq \Delta(\Theta_{-i})$ that contains the extreme points of $\Delta(\Theta_{-i})$. Now, to apply Theorem 2, choose A_i with the same cardinality as the number of extreme points of $\Delta(\Theta_{-i})$ and fix any $p_i^* \in P_i$. Next, associate an extreme point of $\Delta(\Theta_{-i})$ with each $\alpha \in A_i$; call it $\kappa_\alpha(p_i^*)$. By construction, we have that $p_i^* = \sum_\alpha \mu_\alpha(p_i^*) \kappa_\alpha(p_i^*)$ for some weights $(\mu_\alpha(p_i^*))_\alpha$. All other posteriors $(\kappa_\alpha(p_i))_{\alpha, p_i \neq p_i^*}$ are then fully determined by the collinearity condition and satisfy condition (iii). Finally, simply let $\Pi_i = \bigcup_{p_i} \bigcup_\alpha \{\kappa_\alpha(p_i)\}$.

In the construction above, the cardinality of the set A_i is $|\Theta_{-i}|$, the number of extreme points of $\Delta(\Theta_{-i})$. Similarly, when there is a single prior, we can construct each A_i^k with the same cardinality as Π_i^k .¹⁶ More generally, assume that P_i has finitely many extreme points. Condition (ii) of Theorem 2 (the collinearity condition) implies that each extreme point of P_i must be associated with an extreme point of Π_i^k , for each $\alpha \in A_i^k$. It follows that the minimum number of

¹⁴To see this, note that $\bigcup_k \Pi_i^k$ is finite when P_i is a singleton and thus one can choose any A_i with the same cardinality as $\bigcup_k \Pi_i^k$.

¹⁵Examples of social choice functions that are $\Delta(\Theta_{-i})$ -incentive compatible include the efficient allocation of private and public goods.

¹⁶In fact, since a Π_i^k -incentive compatible social choice function is also $\text{ext}(\Pi_i^k)$ -incentive compatible, we can construct A_i^k with the same cardinality as the set of extreme points of Π_i^k . (For a set Y , $\text{ext}(Y)$ denotes the set of extreme points of Y .)

messages required to span Π_i^k is the number of extreme points of Π_i^k divided by the number of extreme points of P_i .

In general, however, while it is relatively easy to check whether P_i belongs to the relative interior of $\bigcup_k \Pi_i^k$, the spanning and collinearity conditions seem “much harder” to check.

Before leaving this section, let us note that Theorem 2 may prove useful to researchers interested in knowing which sets of posteriors can be obtained from a set of priors through the use of some—but not necessarily ambiguous—communication device. For instance, one may be interested in how the beliefs of a population of individuals change with different policies of public announcements, for example, about macroeconomic variables.

5. CONCLUDING DISCUSSION

We conclude by offering a critical discussion of some of the salient features of our analysis. We indicate certain generalizations but relegate the full development of these to the Supplemental Material.

Other Preferences

While the formal analysis in this paper assumes max–min preferences, most of our results generalize to a wider class of ambiguity-sensitive preferences like α -max–min, variational, Bewley, to name a few. For every such preference, there will be a corresponding set of beliefs reflecting the respective incentive constraints and the task is to generate the required sets of beliefs. The central idea of the paper is belief engineering through the use of ambiguous communication devices and is independent of the exact nature of the ambiguity-sensitive preferences.

A Weaker Definition

A weaker definition of implementation would require the social choice function to be implemented only at histories consistent with truth-telling at the first stage. In other words, this weaker definition would not require the social choice function to be implemented following “untruthful” reports at the communication stage, unlike the present definition.

Our analysis extends to this weaker definition, but with some complications. Intuitively, the difficulty is that the private information of a player at the allocation stage is not only his type, but also his private history, that is, the message he has sent and received. In contrast with our definition, this weaker definition does not require f to be implemented at “untruthful” histories and, therefore, private histories now play a critical role.¹⁷ In a nutshell, the characterization of

¹⁷With the present definition, f is implemented regardless of the private histories, so that private histories are essentially irrelevant. In addition, it implies that players are indifferent between telling the truth and lying at the communication stage.

implementable social choice functions then becomes: There must exist sets of messages $(\Omega_i)_i$, a set of transition probabilities Λ (with $\lambda: \times_i \Theta_i \rightarrow \Delta(\times_i \Omega_i)$), beliefs over private histories and types, and a social choice function f^* mapping private histories and types to alternatives, such that f^* coincides with f at “truthful” histories and is incentive compatible at the interim and ex ante stage, that is, players must have an incentive to truthfully reveal their private information both at the allocation stage and at the communication stage. A reader interested in the formal treatment may consult the Supplemental Material.

Beyond Ambiguous Mechanisms

Natural extensions of ambiguous mechanisms might involve multiple rounds of communication and non-direct (multistage) allocation mechanisms. The reader is referred to Sections S.3 and S.4 of the Supplemental Material for discussions on these.

APPENDIX

PROOF OF THEOREM 1: (2) $\Rightarrow$ (1). The proof is constructive. Suppose that for each player $i \in N$, there exists a unique nonempty valued set $\Pi_i \subseteq \Delta(\Theta_{-i})$ such that conditions (IC) and (B) hold. (The general case of a family of sets is treated similarly and left to the reader.)

We consider a two-stage ambiguous mechanism, where players announce a type to the communication device in the first stage, receive messages in the first stage conditional on their announcements, and again report a type in the second stage. Formally, for each $i \in N$, choose $\lambda_i \in \Lambda_i$ and define $\lambda: \times_{i \in N} \Theta_i \rightarrow \Delta(\times_{i \in N} \Omega_i)$ with $\lambda((\theta_i)_{i \in N})[(\omega_i)_{i \in N}] := \times_{i \in N} \lambda_i(\theta_i, \theta_{-i})[\omega_i]$. The set Λ is the collection of all such λ . The mechanism consists of the ambiguous communication device $\langle \times_{i \in N} \Theta_i, \times_{i \in N} \Omega_i, \Lambda \rangle$ and the allocation mechanism $\langle \times_{i \in N} \Theta_i, f \rangle$.

Consider the following profile of strategies. At the initial history, players truthfully report their types to the ambiguous communication device and, for any history (θ'_i, ω_i) , players again report their type truthfully at the allocation stage.

Notice that if there exists a message ω_i such that, for all $\lambda \in \Lambda$, all $p_i \in P_i$,

$$\sum_{\theta_{-i}, \omega_{-i}} \lambda(\theta_i, \theta_{-i})[(\omega_i, \omega_{-i})] p_i[\theta_{-i}] = 0,$$

then $\{\zeta_i(p_i, \theta_i, \omega_i, \lambda_i) = \emptyset$ for all (p_i, λ_i) , a contradiction with condition (B) (since Π_i is nonempty). Therefore, assessments are well-defined and obtained by Bayes's rule.

By construction of the ambiguous communication device, at the allocation stage, the set of beliefs of player i about the types θ_{-i} of his opponents is Π_i

(since condition (B) holds). Since the social choice function is Π_i -incentive compatible for each player i , it follows that no player has a profitable deviation from truth-telling at any history $(\theta'_i, \omega_i) \in \Theta_i \times \Omega_i$. Moreover, no player has a profitable deviation in the first stage. Indeed, regardless of his first-stage announcement, each player expects the social choice function to be implemented in the second stage and thus has no profitable deviation in the first stage.

(1) $\Rightarrow$ (2). Suppose that the social choice function f is implementable by the ambiguous mechanism

$$\langle\langle \widehat{\Omega}_i, \Omega_i \rangle_{i \in N}, \Lambda \rangle, \langle (M_i)_{i \in N}, g \rangle.$$

There exists a pure equilibrium $(s^*, \Pi^{H, \Theta})$ such that

$$g(s_i^*(h_i, \theta_i), s_{-i}^*(h_{-i}, \theta_{-i})) = f(\theta_i, \theta_{-i})$$

for all (θ_i, θ_{-i}) , for all (h_i, h_{-i}) having positive probability under s^* .

Consider player i of type θ_i and any history $h_i = (\widehat{\omega}_i, \omega_i)$ having positive probability, that is, there exist a type $\theta'_i \in \Theta_i$ of player i , a probability system $\lambda' \in \Lambda$, and a prior $p'_i \in P_i$ such that $s_i^*(\{\emptyset\}, \theta'_i) = \widehat{\omega}_i$ and $\sum_{\theta_{-i}, \omega_{-i}} \lambda'(\widehat{\omega}_i, s_{-i}^*(\{\emptyset\}, \theta_{-i}))[(\omega_i, \omega_{-i})]p'_i[\theta_{-i}] > 0$.

Recall that $\Pi_i^{H, \Theta}(h_i, \theta_i) \subseteq \Delta(H_{-i} \times \Theta_{-i})$ is the set of beliefs of player i about the private histories and types of his opponents, conditional on the private history and type (h_i, θ_i) . It is obtained as follows. Fix any $\lambda \in \Lambda$ and any $p_i \in P_i$. From Bayes's rule, the probability of $((\widehat{\omega}_{-i}, \omega_{-i}), \theta_{-i})$, given the history $(\widehat{\omega}_i, \omega_i)$, the type θ_i , the probability system λ , and the prior p_i , is

$$\begin{cases} \frac{\lambda(\widehat{\omega}_i, s_{-i}^*(\{\emptyset\}, \theta_{-i}))[(\omega_i, \omega_{-i})]p_i[\theta_{-i}]}{\sum_{\theta_{-i}, \omega_{-i}} \lambda(\widehat{\omega}_i, s_{-i}^*(\{\emptyset\}, \theta_{-i}))[(\omega_i, \omega_{-i})]p_i[\theta_{-i}]}, & \text{if } s_{-i}^*(\{\emptyset\}, \theta_{-i}) = \widehat{\omega}_{-i}, \\ 0, & \text{if } s_{-i}^*(\{\emptyset\}, \theta_{-i}) \neq \widehat{\omega}_{-i}, \end{cases}$$

if $\sum_{\theta_{-i}, \omega_{-i}} \lambda(\widehat{\omega}_i, s_{-i}^*(\{\emptyset\}, \theta_{-i}))[(\omega_i, \omega_{-i})]p_i[\theta_{-i}] > 0$, and undefined otherwise.

The set $\Pi_i^{H, \Theta}(h_i, \theta_i)$ is then obtained by taking the union over all $\lambda \in \Lambda$ and $p_i \in P_i$. Since the history $(\widehat{\omega}_i, \omega_i)$ has positive probability, it follows that $\Pi_i^{H, \Theta}(h_i, \theta_i)$ is nonempty and all beliefs follow from Bayes's rule. Note that we apply Bayes's rule at the history $(\widehat{\omega}_i, \omega_i)$, even though it may be that $s_i^*(\{\emptyset\}, \theta_i) \neq \widehat{\omega}_i$. Indeed, it is possible to apply Bayes's rule at the history $(\widehat{\omega}_i, \omega_i)$ and our equilibrium definition requires to apply it whenever possible.

By definition of an equilibrium, we have that

$$\begin{aligned} & \min_{\pi_i(h_i, \theta_i) \in \Pi_i^{H, \Theta}(h_i, \theta_i)} \sum_{(h_{-i}, \theta_{-i})} \pi_i(h_i, \theta_i)[h_{-i}, \theta_{-i}] \\ & \times u_i(g(s_i^*(h_i, \theta_i), s_{-i}^*(h_{-i}, \theta_{-i})), \theta_i, \theta_{-i}) \end{aligned}$$

$$\geq \min_{\pi_i(h_i, \theta_i) \in \Pi_i^{H, \Theta}(h_i, \theta_i)} \sum_{(h_{-i}, \theta_{-i})} \pi_i(h_i, \theta_i)[h_{-i}, \theta_{-i}] \\ \times u_i(g(s_i^*(h_i, \theta'_i), s_{-i}^*(h_{-i}, \theta_{-i})), \theta_i, \theta_{-i})$$

for all θ'_i . Note that if $\pi_i(h_i, \theta_i)[h_{-i}, \theta_{-i}] > 0$, then the history h_{-i} has positive probability. Therefore, the above inequalities are equivalent to

$$\min_{\pi_i(h_i, \theta_i) \in \Pi_i^{H, \Theta}(h_i, \theta_i)} \sum_{(h_{-i}, \theta_{-i})} \pi_i(h_i, \theta_i)[h_{-i}, \theta_{-i}] u_i(f(\theta_i, \theta_{-i}), \theta_i, \theta_{-i}) \\ \geq \min_{\pi_i(h_i, \theta_i) \in \Pi_i^{H, \Theta}(h_i, \theta_i)} \sum_{(h_{-i}, \theta_{-i})} \pi_i(h_i, \theta_i)[h_{-i}, \theta_{-i}] u_i(f(\theta'_i, \theta_{-i}), \theta_i, \theta_{-i})$$

for all θ'_i .

Since the beliefs of player i about the types of his opponents are independent of his own types, we have that $\Pi_i^{H, \Theta}(h_i, \theta_i) = \Pi_i^{H, \Theta}(h_i, \theta'_i)$ for any (θ_i, θ'_i) . Thus, we can suppress the dependence on θ_i in the sequel. Write

$$\Pi_i^\Theta(h_i) := \left\{ \pi_i \in \Delta(\Theta_{-i}) : \right. \\ \left. \pi_i[\cdot] = \sum_{h_{-i}} \pi_i(h_i)[(h_{-i}, \cdot)], \text{ for some } \pi_i(h_i) \in \Pi_i^{H, \Theta}(h_i) \right\}.$$

Finally, consider the set of player i 's private histories having positive probabilities and partition it into K_i subsets such that $\Pi_i^\Theta(h_i) = \Pi_i^\Theta(h'_i)$ for any (h_i, h'_i) in the k th subset. Call Π_i^k the common set of beliefs in the k th subset.

We now construct a set $\Omega_i^* \subseteq \Omega_i$ and a partition $\{\Phi_i^1, \dots, \Phi_i^k, \dots, \Phi_i^{K_i}\}$ of $\Theta_i \times \Omega_i^*$ such that the set of posteriors is Π_i^k given any $(\theta_i, \omega_i) \in \Phi_i^k$. First, note that since s_i^* is a pure strategy, any history $h_i = (\hat{\omega}_i, \omega_i)$ with positive probability corresponds to pairs (θ_i, ω_i) such that $s_i^*(\{\emptyset\}, \theta_i) = \hat{\omega}_i$. Second, denote Ω_i^* the subset of Ω_i such that, for each $\omega_i \in \Omega_i^*$, there exist a type θ'_i , a prior p'_i , and a probability system λ' with

$$\sum_{\omega_{-i}} \lambda'(s_i^*(\{\emptyset\}, \theta'_i), s_{-i}^*(\{\emptyset\}, \theta_{-i}))[(\omega_i, \omega_{-i})] p'_i[\theta_{-i}] > 0.$$

Note that $\{\hat{\omega}_i \in \hat{\Omega}_i : \hat{\omega}_i = s_i^*(\{\emptyset\}, \theta'_i), \text{ for some } \theta'_i\} \times \Omega_i^* \subseteq \hat{\Omega}_i \times \Omega_i$ is the set of player i 's histories with positive probability. Therefore, we can partition the set $\Theta_i \times \Omega_i^*$ into K_i subsets $\{\Phi_i^1, \dots, \Phi_i^k, \dots, \Phi_i^{K_i}\}$ such that the set of posteriors is Π_i^k given any $(\theta_i, \omega_i) \in \Phi_i^k$. To see this, consider two histories $(\hat{\omega}_i, \omega_i)$ and $(\hat{\omega}'_i, \omega'_i)$ with positive probabilities such that the set of posteriors is Π_i^k , conditional on either of them. Since these histories have positive probabilities, there exist θ_i and θ'_i such that $s_i^*(\{\emptyset\}, \theta_i) = \hat{\omega}_i$ and $s_i^*(\{\emptyset\}, \theta'_i) = \hat{\omega}'_i$. Moreover,

both ω_i and ω'_i belong to Ω_i^* . Therefore, (θ_i, ω_i) and (θ'_i, ω'_i) are associated with the same set of posteriors Π_i^k , hence to the same cell Φ_i^k of the partition.

From the above, it follows that

$$\begin{aligned} & \min_{\pi_i^k \in \Pi_i^k} \sum_{\theta_{-i}} \pi_i^k[\theta_{-i}] u_i(f(\theta_i, \theta_{-i}), \theta_i, \theta_{-i}) \\ & \geq \min_{\pi_i^k \in \Pi_i^k} \sum_{\theta_{-i}} \pi_i^k[\theta_{-i}] u_i(f(\theta'_i, \theta_{-i}), \theta_i, \theta_{-i}) \end{aligned}$$

for all θ'_i , that is, f is Π_i^k -incentive compatible. Lastly, since the beliefs are obtained by full Bayesian updating, condition (B) clearly holds, with

$$\lambda_i(\theta_i, \theta_{-i})[\omega_i] = \sum_{\omega_{-i}} \lambda(s_i^*(\{\emptyset\}, \theta_i), s_{-i}^*(\{\emptyset\}, \theta_{-i}))[(\omega_i, \omega_{-i})]$$

for each (θ_i, θ_{-i}) , for each λ . For further reference, note that, for each $(\theta_i, \omega_i) \in \Phi_i^k$, there exist θ_{-i} and λ such that $\lambda_i(\theta_i, \theta_{-i})[\omega_i] > 0$. *Q.E.D.*

PROOF OF THEOREM 2: As mentioned in the main text, Theorem 2 can be stated in terms of a unique probability system. Lemma 1 below states this important equivalence; the proof can be found in the Supplemental Material.

LEMMA 1: *The following statements are equivalent:*

1. *There exist a set of messages Ω_i^* , a probability system $\bar{\lambda}_i: \Theta_i \times \Theta_{-i} \rightarrow \Delta(\Omega_i)$, and a finite partition $\{\Phi_i^1, \dots, \Phi_i^k, \dots, \Phi_i^{|K_i|}\}$ of $\Theta_i \times \Omega_i^*$ such that, for all $k \in K_i$,*

$$\bigcup_{p_i \in P_i} \bigcup_{(\theta_i, \omega_i) \in \Phi_i^k} \{\zeta_i(p_i, \theta_i, \omega_i, \bar{\lambda}_i)\} = \Pi_i^k.$$

2. *There exist a set of messages Ω_i , a set of probability systems λ_i , and a finite partition $\{\Phi_i^1, \dots, \Phi_i^k, \dots, \Phi_i^{|K_i|}\}$ of $\Theta_i \times \Omega_i$ such that, for all $k \in K_i$, for all $(\theta_i, \omega_i) \in \Phi_i^k$,*

$$\bigcup_{p_i \in P_i} \bigcup_{\lambda_i \in \Lambda_i} \{\zeta_i(p_i, \theta_i, \omega_i, \lambda_i)\} = \Pi_i^k.$$

See the Supplemental Material for the proof.

Given Lemma 1, to prove Theorem 2, it thus suffices to show that condition 1 of Lemma 1 and condition 2 of Theorem 2 are equivalent.

We first prove that (1) $\Rightarrow$ (2). Suppose that there exist a set of messages Ω_i , a probability system $\lambda_i: \Theta_i \times \Theta_{-i} \rightarrow \Delta(\Omega_i)$, and a finite partition $\{\Phi_i^1, \dots, \Phi_i^k, \dots, \Phi_i^{|K_i|}\}$ of $\Theta_i \times \Omega_i$ such that, for all $k \in K_i$,

$$\bigcup_{p_i \in P_i} \bigcup_{(\theta_i, \omega_i) \in \Phi_i^k} \{\zeta_i(p_i, \theta_i, \omega_i, \lambda_i)\} = \Pi_i^k.$$

For each $(\theta_i, \omega_i) \in \Theta_i \times \Omega_i$ and each $p_i \in P_i$, define

$$\mu_{(\theta_i, \omega_i)}(p_i) := \sum_{\theta_{-i}} \lambda_i(\theta_i, \theta_{-i})[\omega_i] p_i[\theta_{-i}] / |\Theta_i|.$$

From the proof of Lemma 1 (see the last sentence), we have that $\mu_{(\theta_i, \omega_i)}(p_i) > 0$ and $\sum_{(\theta_i, \omega_i) \in \Theta_i \times \Omega_i} \mu_{(\theta_i, \omega_i)}(p_i) = 1$.

Similarly, for each (θ_i, ω_i) , for each p_i , define $\kappa_{(\theta_i, \omega_i)}(p_i)$ as $\zeta_i(p_i, \theta_i, \omega_i, \lambda_i)$. By construction, we have that $\bigcup_{p_i} \bigcup_{(\theta_i, \omega_i) \in \Phi_i^k} \{\kappa_{(\theta_i, \omega_i)}(p_i)\} = \Pi_i^k$ for all k .

For all (θ_i, ω_i) , for all θ_{-i} , it follows from the definition of ζ_i that

$$\lambda_i(\theta_i, \theta_{-i})[\omega_i] = \left(\sum_{\theta_{-i}} \lambda_i(\theta_i, \theta_{-i})[\omega_i] p_i[\theta_{-i}] \right) \frac{\zeta_i(p_i, \theta_i, \omega_i, \lambda_i)[\theta_{-i}]}{p_i[\theta_{-i}]}.$$

Therefore, we must have that

$$\mu_{(\theta_i, \omega_i)}(p_i) (\kappa_{(\theta_i, \omega_i)}(p_i) / p_i) = \mu_{(\theta_i, \omega_i)}(p'_i) (\kappa_{(\theta_i, \omega_i)}(p'_i) / p'_i)$$

for any (p_i, p'_i) , for any (θ_i, ω_i) .

By summing over (θ_i, ω_i) , we clearly have that

$$p_i = \sum_{(\theta_i, \omega_i)} \mu_{(\theta_i, \omega_i)}(p_i) \kappa_{(\theta_i, \omega_i)}(p_i)$$

for each p_i .

To conclude this direction of the proof, simply let $A_i := \Theta_i \times \Omega_i$ and partition A_i into $|K_i|$ subsets A_i^k , with $A_i^k := \Phi_i^k$, for each k .

We now prove (2) $\Rightarrow$ (1). Let Ω_i be any set for which there exists a bijection ρ between A_i and Ω_i , and define $\lambda_i: \Theta_{-i} \rightarrow \Delta(\Omega_i)$ as follows:

$$\lambda_i(\theta_{-i})[\omega_i] := \mu_{\rho^{-1}(\omega_i)}(p_i) \frac{\kappa_{\rho^{-1}(\omega_i)}(p_i)[\theta_{-i}]}{p_i[\theta_{-i}]}$$

for all $\theta_{-i} \in \Theta_{-i}$, for all $\omega_i \in \Omega_i$, for some arbitrarily chosen $p_i \in P_i$. Note that $\lambda_i(\theta_{-i})[\omega_i] \geq 0$ and $\sum_{\omega_i \in \Omega_i} \lambda_i(\theta_{-i})[\omega_i] = 1$ (since condition (iii) states that $\sum_{\alpha \in A_i} \mu_\alpha(p_i) \kappa_\alpha(p_i) = p_i$ and $\mu_\alpha(p_i) > 0$ for all p_i).

Choose any $p'_i \in P_i$ and any $\omega_i \in \Omega_i$. For each θ_{-i} , the posterior $\zeta_i(p'_i, \theta_i, \omega_i, \lambda_i)[\theta_{-i}]$ of p'_i given ω_i and θ_i is

$$\begin{aligned} & \frac{\mu_{\rho^{-1}(\omega_i)}(p_i) (\kappa_{\rho^{-1}(\omega_i)}(p_i)[\theta_{-i}] / p_i[\theta_{-i}]) p'_i[\theta_{-i}]}{\sum_{\theta_{-i}} \mu_{\rho^{-1}(\omega_i)}(p_i) (\kappa_{\rho^{-1}(\omega_i)}(p_i)[\theta_{-i}] / p_i[\theta_{-i}]) p'_i[\theta_{-i}]} \\ &= \frac{\mu_{\rho^{-1}(\omega_i)}(p'_i) \kappa_{\rho^{-1}(\omega_i)}(p'_i)[\theta_{-i}]}{\sum_{\theta_{-i}} \mu_{\rho^{-1}(\omega_i)}(p'_i) \kappa_{\rho^{-1}(\omega_i)}(p'_i)[\theta_{-i}]} = \kappa_{\rho^{-1}(\omega_i)}(p'_i)[\theta_{-i}], \end{aligned}$$

where the first equality follows from the collinearity condition (ii).

Finally, define the partition $\{\Phi_i^1, \dots, \Phi_i^k, \dots, \Phi_i^{|\mathcal{K}_i|}\}$ of $\Theta_i \times \Omega_i$ with $\Phi_i^k := \Theta_i \times \rho^{-1}(A_i^k)$ for each k . It then follows from the spanning condition (i) that $\bigcup_{p_i \in P_i} \bigcup_{(\theta_i, \omega_i) \in \Phi_i^k} \{\zeta_i(p_i, \theta_i, \omega_i, \lambda_i)\} = \Pi_i^k$, the desired result. *Q.E.D.*

REFERENCES

- AUMANN, R., M. MASCHLER, AND R. STEARNS (1995): *Repeated Games With Incomplete Information*. Cambridge: MIT Press. [1865]
- BODOH-CREED, A. (2012): "Ambiguous Beliefs and Mechanism Design," *Games and Economic Behavior*, 75, 518–537. [1855]
- BOSE, S., AND A. DARIPA (2009): "A Dynamic Mechanism and Surplus Extraction Under Ambiguity," *Journal of Economic Theory*, 144, 2084–2114. [1855]
- BOSE, S., AND L. RENOU (2014): "Supplement to 'Mechanism Design With Ambiguous Communication Devices'," *Econometrica Supplemental Material*, 82, http://www.econometricsociety.org/ecta/supmat/10285_extensions.pdf. [1857, 1859]
- BOSE, S., E. OZDENOREN, AND A. PAPPE (2006): "Optimal Auctions With Ambiguity," *Theoretical Economics*, 1, 411–438. [1855]
- DE CASTRO, L., AND N. C. YANNELIS (2012): "Uncertainty, Efficiency and Incentive Compatibility," Report. [1855, 1865]
- DI TILLIO, A., N. KOS, AND M. MESSNER (2011): "The Design of Ambiguous Mechanisms," Report. [1855]
- FORGES, F. (1986): "An Approach to Communication Equilibria," *Econometrica*, 54, 1375–1385. [1854]
- (1990): "Universal Mechanisms," *Econometrica*, 58, 1341–1364. [1854]
- GILBOA, I., AND D. SCHMEIDLER (1989): "Maxmin Expected Utility With a Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141–153. [1855, 1858]
- LACLAU, M., AND L. RENOU (2013): "Persuasion Games With Ambiguity," Report. [1864]
- LANG, M., AND A. WAMBACH (2013): "The Fog of Fraud—Mitigating Fraud by Strategic Ambiguity," *Games and Economic Behavior*, 81, 255–275. [1853]
- LOPOMO, G., L. RIGOTTI, AND C. SHANNON (2010): "Uncertainty in Mechanism Design," Report. [1855]
- MYERSON, R. (1979): "Incentive Compatibility and the Bargaining Problem," *Econometrica*, 47, 61–73. [1855]
- (1986): "Multistage Games With Communication," *Econometrica*, 54, 323–358. [1854]
- PIRES, C. P. (2002): "A Rule for Updating Ambiguous Beliefs," *Theory and Decision*, 53, 137–152. [1855]
- SALO, A., AND M. WEBER (1995): "Ambiguity Aversion in First Price Sealed Bid Auctions," *Journal of Risk and Uncertainty*, 11, 123–137. [1855]
- SINISCALCHI, M. (2011): "Dynamic Choice Under Ambiguity," *Theoretical Economics*, 6, 379–421. [1859]
- STROTZ, R. (1955): "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 23, 165–180. [1859]

Dept. of Economics, University of Leicester, Leicester LE1 7RH, U.K.; sb345@le.ac.uk

and

Dept. of Economics, University of Essex, Colchester, Essex CO4 3SQ, U.K.; lrenou@essex.ac.uk.

Manuscript received September, 2011; final revision received March, 2014.