



Delegation and ambiguity in correlated equilibrium [☆]

Simon Grant, Ronald Stauber ^{*}

Research School of Economics, Australian National University, ACT 2601, Australia



ARTICLE INFO

Article history:

Received 26 May 2021

Available online 29 January 2022

JEL classification:

C72

D81

Keywords:

Correlated equilibrium

Coarse correlated equilibrium

Ambiguity

Maxmin

ABSTRACT

In the context of normal-form games with complete information, we introduce a notion of correlated equilibrium that allows partial delegation to a mediator and ambiguity in the correlation device. Without ambiguity, the sets of equilibrium action distributions are equivalent to those for coarse correlated equilibrium (Moulin and Vial, 1978). With correlation devices that incorporate ambiguity, any action distribution that Pareto dominates a coarse correlated equilibrium or a correlated equilibrium (Aumann, 1974), can be approximated with an arbitrary degree of precision using the proposed equilibrium notion. These approximations are attained in one-shot, static strategic interactions, and do not require repeated play. We also analyze such equilibria when the set of feasible posteriors is exogenously constrained, which yields, as a special case, a definition and characterization of an “ambiguous correlated equilibrium” that does *not* require delegation to the mediator.

© 2022 Elsevier Inc. All rights reserved.

1. Introduction

In the setting of the classic correlated equilibrium (CE) of Aumann (1974), it is assumed that each player chooses an action after she has observed a private signal generated via an extraneous probabilistic correlation device. Myerson (1982) and Aumann (1987) show that every CE has an equivalent *canonical representation*, where the correlation device is given by a probability distribution over action profiles, so that after an action profile is drawn according to this distribution, each player is informed of her action in the drawn profile. The distribution then constitutes a CE if she never has a strict incentive to deviate from selecting the action the device recommends she choose.¹ An extension of this definition of CE was proposed by Moulin and Vial (1978), who introduced a notion of equilibrium that in the recent literature has been dubbed a *coarse correlated equilibrium* (CCE). A CCE is still defined by a distribution over action profiles, but players receive no additional information beyond their (common) knowledge of the equilibrium distribution. Instead, each player can choose either to allow the mediator to implement an action profile drawn according to this equilibrium distribution, or she can unilaterally

[☆] We are grateful to an advisory editor and two referees, whose comments and recommendations have significantly improved the paper. We also thank Frank Riedel and George Mailath for initial discussions on the subject of this paper, seminar audiences at the Australian National University, the Australasian Economic Theory Workshop 2020, the Graduate School of Business, Stanford University, the AMETS seminar series, RUD 2020 and GAMES2020.1 for comments and suggestions, and Jingni Yang for assistance in producing a number of graphs included in the paper.

^{*} Corresponding author.

E-mail addresses: simon.grant@anu.edu.au (S. Grant), ronald.stauber@anu.edu.au (R. Stauber).

¹ Hence, CE and canonical CE are equivalent in terms of the induced ex-ante equilibrium distributions over action profiles. An epistemic comparison of CE and canonical CE from an interim perspective is presented in Bach and Perea (2020).

deviate by selecting her own action (or mixed strategy) *independently* of this distribution.² It is straightforward to show that the set of CE and CCE coincide for two-action games, however, as illustrated in Moulin and Vial (1978), even for games with just two players and three actions, there exist CCE with expected payoffs that dominate those of every CE.

In this paper, we introduce a generalization of CCE that allows for players with ambiguity-averse preferences, together with ambiguous (that is, set-valued) correlation devices. Analogous to the *universal Bayesian solution* analyzed by Forges (1993) in the context of games with incomplete information, an equilibrium is defined by a set of probability distributions over the Cartesian product of the action space and an (arbitrary) signal space.³ A mediator selects one of the probability distributions according to an arbitrary procedure, and draws a profile of actions and signals according to the chosen distribution. Each player is then informed of the realization of her private signal, but not of the realized action profile, her opponents' signals, or the specific probability distributions on which the draw was made. After learning the realizations of their private signals, players update their beliefs using full Bayesian updating (Jaffray, 1992; Pires, 2002), and on the basis of maxmin preferences (Gilboa and Schmeidler, 1989; Gajdos et al., 2008), each decides whether to allow the mediator to implement the associated action profile, or to deviate by independently selecting her own action or mixed strategy. We call the resulting equilibrium notion a *delegated correlated equilibrium (DCE)*, since players essentially delegate their action choices to the mediator, but maintain the option of revoking that delegation after learning the realization of their private signals—in a DCE, no player has a strict incentive to revoke her delegation to the mediator at any signal realization.

Even though the idea of delegation is not standard in the game-theory literature (except for the notion of CCE), one could envisage applications, such as the coordinated behavior of autonomous self-driving cars, where drivers receive certain signals about the environment, and then must decide whether to intervene and take control of the car, or to allow the vehicle to proceed autonomously, possibly employing a procedure that is coordinated with other vehicles on the road. While the decision-making of current autonomous systems for self-driving vehicles is generally independent across vehicles, the real-time communication between artificial-intelligent (AI) systems that our approach suggests is one of the promises of the new 5G networks that are currently being rolled out worldwide. One might view a key feature for such AI systems to be acceptable to the general public, would be that drivers retain the right to take control of their vehicles, even though in a DCE they will never have a strict incentive to invoke this right. Hence, the notion of DCE should be considered relevant for situations in which players are afforded the opportunity to delegate their action choices to AI agents, who may be endowed with an ability to coordinate strategic decisions between players.

Another potential application of our proposed equilibrium notion is in (finite) routing games (see Shoham and Leyton-Brown, 2009; Roughgarden, 2007, for an introduction), which constitute an example of congestion games, as introduced by Rosenthal (1973). In a congestion game, each player chooses a subset from a set of available resources, and incurs a cost for each selected resource that depends on the total number of players who also choose the same resource. In a network routing game, each player/network user selects a path of links on a directed graph connecting a source node and a target node that are player-specific, and incurs a cost for each chosen link that is a function of the number of players whose chosen path also passes through this link, which captures the delay incurred in transmitting a message via the link. To apply the framework of DCE, the networking protocol controlling the transmissions through the network would need to be configured to act as a trusted mediator, who chooses a usage pattern for the network based on an ambiguous correlation device and informs the players of the corresponding signals, such that the players maintain the option of revoking their delegation to the protocol and independently selecting a transmission path. If the objective of network users is to minmax their costs, which seems plausible in the sense of minimizing the worst-case delay incurred by the chosen transmission paths, the methods we propose apply directly to such games.⁴

When the correlation device defining a DCE involves no ambiguity, the resulting *unambiguous DCE (uDCE)* essentially augments the notion of CCE with an informative signal space. Hence, if each player's signal is the same as her randomly drawn action, then the associated uDCE is equivalent to a canonical CE, where the mediator recommends an action to each player. Similarly, when each player only receives a null signal, the resulting uDCE is identical to a CCE. Clearly, such a null signal yields the weakest incentive constraints among all possible signal spaces, and therefore the largest set of action distributions that are implementable by uDCE. Consequently, the set of action distributions that are implementable by uDCE must be identical to the set of CCE. However, if we consider ambiguity-averse players together with ambiguous correlation devices, allowing for informative signal spaces provides a critical benefit in terms of enlarging the set of distributions over action profiles that can be induced by DCE. Specifically, defining an action distribution σ to be *implemented* by a DCE if every action–signal distribution that defines the DCE induces the same marginal distribution σ over action profiles (so the ambiguity in the correlation device is restricted to signals), we show that every action distribution that Pareto dominates a

² An alternative interpretation of CCE requires players either, to commit to implementing a recommended action (before knowing what the recommendation is), or to choose independently an individual action without receiving a recommendation. Applications of CCE can be found in Gérard-Varet and Moulin (1978), Ray and Gupta (2013), Moulin et al. (2014), Young (2004), Hart and Mas-Colell (2015) and Awaya and Krishna (2019).

³ Our model does not assume, however, that the players' signals provide full information regarding their own recommended actions, as is the case for the universal Bayesian solution. Notice also that while the focus of our paper is on complete information games, so as to emphasize the impact of correlation and ambiguity, by building on the setup in Forges (1993) the general approach we propose could be extended to games with incomplete information.

⁴ Ashlagi et al. (2008) already point out that the (implicit) presence of a mediator for CE makes CE especially appropriate for the study of multi-agent systems in AI, and also include a related analysis of congestion games. Roughgarden (2015) also considers CE and CCE in relation to a certain class of congestion games. We will discuss their work and some related implications of properties of DCE in section 6.3.

		Player 2	
		c	d
Player 1	c	3	4
	d	0	1

Fig. 1. A Prisoner's dilemma.

CCE distribution can be approximately implemented using ambiguous correlation devices, in a precise sense to be defined later. The intuition for this result is illustrated by the following example:

Example 1. Consider a classic prisoner's dilemma, as pictured in Fig. 1. The action distribution that assigns probability one to (c, c) dominates the (unique CCE) distribution that assigns probability one to (d, d) . Letting $\delta_{cc}, \delta_{dd} \in \Delta\{cc, cd, dc, dd\}$ denote the degenerate action distributions that assign probability one to (c, c) and (d, d) , respectively, we now construct, for any $\varepsilon \in (0, 1)$, a DCE that implements the distribution $\sigma^\varepsilon := (1 - \varepsilon)\delta_{cc} + \varepsilon\delta_{dd}$.

Given the symmetry of the game, it suffices to allow both players to observe a common signal from a two-element signal space $S = \{s', s''\}$, and define a DCE comprised of two action–signal distributions $\bar{q}, \tilde{q} \in \Delta(\{cc, dc, cd, cc\} \times \{s', s''\})$. For any $(a, s) \in \{cc, dc, cd, cc\} \times \{s', s''\}$, we can set

$$\bar{q}(a, s) = \sigma^\varepsilon(a)\bar{\pi}(s|a), \text{ and } \tilde{q}(a, s) = \sigma^\varepsilon(a)\tilde{\pi}(s|a),$$

so that both \bar{q} and \tilde{q} have the same marginal σ^ε on the action space, and let the conditionals $\bar{\pi}(\cdot|a)$ and $\tilde{\pi}(\cdot|a)$ be given by

$$\bar{\pi}(s'|cc) = \tilde{\pi}(s''|cc) = 1, \text{ and } \bar{\pi}(s''|dd) = \tilde{\pi}(s'|dd) = 1$$

for $a = cc$ and $a = dd$, and be arbitrary for the remaining action profiles. The sets of player posteriors after observing the signals s' and s'' are then both equal to $\{\delta_{cc}, \delta_{dd}\}$, so that, irrespective of the signal that a player receives, the worst-case payoff from maintaining the delegation is attained at δ_{dd} and is equal to 1, and the maxmin payoff from revoking the delegation is also equal to 1 and is attained at δ_{dd} together with the independent mixed strategy assigning probability one to d . Hence, no player has a strict incentive to revoke her delegation.⁵ ◁

As a further contribution of our paper, we also propose a generalization of CE to strategic environments with ambiguity averse players, which is applicable in situations where players can coordinate through an extraneous correlation device, but are not able to delegate their action choices to a mediator. The intuition for the resulting definition of *ambiguous correlated equilibrium (ACE)* follows from a reinterpretation of CE as a uDCE where each player's set of feasible posteriors over action profiles are constrained to lie on those faces of the corresponding probability simplex that assign probability one to the player's own actions, so that signals can be viewed as action recommendations. Based on an analogous notion of constrained DCE, an ACE then restricts the feasible sets of posteriors associated with any player signal to contain only elements that assign probability one to the same action of this player, so that every player can be entrusted to select their own equilibrium actions. The sets of action distributions that are implementable by an ACE can coincide with those that are implementable by DCE for some games, but may also be identical to the sets of CCE for other games (such as, e.g., the prisoner's dilemma).

The definitions of DCE or ACE implicitly assume that the correlation device associated with an equilibrium is common knowledge between the players, and that all players believe that their opponents conform to their respective equilibrium behavior (i.e., players do not revoke their delegation to the mediator for DCE, or select the recommended action for ACE), so that no player has an incentive to deviate unilaterally. How a particular equilibrium correlation device will be chosen, especially in the case when a game possesses multiple DCE/ACE, will depend on the specific strategic environment under consideration. Presumably, a society or group of players will need to establish appropriate social norms that may be enforced through legislative or technical frameworks, and which enable players to attain the common knowledge required to coordinate on an equilibrium and a corresponding mediator. In some sense, this is analogous to making a selection from multiple Nash equilibria, although, especially in the case when AI systems are used to coordinate behavior, and more so when delegation to an AI system is involved, reaching an appropriate social consensus should be non-trivial, and would probably require the consideration of various philosophical, legal and social issues.⁶ While the question of how a society should select an appropriate correlation device together with a corresponding mediator to which to entrust the delegation is beyond the scope of the present paper, the framework we propose should provide a theoretical structure to support a related analysis and/or discussion, for the case when the players or society members are ambiguity averse.

⁵ We will show later that for some games, it is possible to derive analogous results based on DCE where players have a strict preference to maintain their delegation to the mediator.

⁶ Alternatively, a particular society may also voluntarily acquiesce to the correlation device that is imposed on it by the CEO of a relevant technology company. A survey of issues and challenges related to designing socially responsible AI algorithms is presented in Cheng et al. (2021).

The rest of the paper is organized as follows. Section 2 introduces the model, our main equilibrium definition, as well as some preliminary results. Section 3 presents a simple method to construct DCE, and characterizes action distributions that are implementable by DCE. Section 4 defines and analyzes ACE, based on a notion of constrained DCE. Section 5.1 briefly addresses some matters related to dynamic consistency and welfare, and Section 5.2 discusses the robustness of our analysis to alternative preference specifications. Additional related literature is reviewed in Section 6. Proofs of all results that are not immediate from the discussion in the main text can be found in Appendix A.

2. Preliminaries and equilibrium definition

Consider a finite normal-form game $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$, where

- $N := \{1, \dots, n\}$ denotes the set of n players,

and for each player $i \in N$,

- A_i is the (finite) set of actions available to player i , with generic element a_i , and
- u_i is a function that assigns a payoff (utility) for player i to each action profile $a = (a_1, \dots, a_n) \in A := \times_{i \in N} A_i$.

To construct the correlation devices used in defining a delegated correlated equilibrium (DCE), we add to the description of the game a collection $(S_i)_{i \in N}$ comprised of finite sets S_i of possible signal realizations for players $i \in N$, and let $S := \times_i S_i$ denote the resulting signal space, with generic element given by a profile of signal realizations $s = (s_1, \dots, s_n) \in S$. We identify an ambiguous correlation device with a finite set of distributions Q defined over pairs $(a, s) \in A \times S$, so $Q \subset \Delta(A \times S)$. For each distribution $q \in Q$ and realization $s_i \in S_i$, let q_{s_i} denote the marginal of q on the action space A , conditional on s_i . The interpretation of the correlation device is that the common mediator chooses a distribution q from Q according to an arbitrary unspecified procedure (for example, based on an appropriately designed Ellsberg urn), draws a realization (a, s) using the selected q , and informs each player i of her associated signal s_i .

If player i believes that the actions $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$ of her opponents will be determined according to the profile a drawn by the mediator, then knowing the realization s_i , but not the realized action profile a nor the distribution q on which the draw was based, she must choose one of two options: either she can let the mediator implement the action profile a , or she can unilaterally revoke her delegation to the mediator and choose on her own behalf an action from A_i based on any independent mixed strategy $\rho_i \in \Delta A_i$. Her expected maxmin utility from choosing the former option is given by

$$\min_{q \in Q} \sum_{a \in A} u_i(a) q_{s_i}(a),$$

while from the latter it is given by

$$\min_{q \in Q} \sum_{(a'_i, a_{-i}) \in A} u_i(a'_i, a_{-i}) \rho_i(a'_i) \sum_{a_i \in A_i} q_{s_i}(a_i, a_{-i}),$$

where $\sum_{a_i \in A_i} q_{s_i}(a_i, a_{-i})$ is by construction the marginal of q_{s_i} on A_{-i} .

An equilibrium can then be defined as follows:

Definition 1. A delegated correlated equilibrium (DCE) of a game $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$ consists of a finite signal space $S = \times_i S_i$ and a finite set of distributions $Q \subset \Delta(A \times S)$, such that for every player $i \in N$ and each of her signal realizations $s_i \in S_i$,

$$\min_{q \in Q} \sum_{a \in A} u_i(a) q_{s_i}(a) \geq \max_{\rho_i \in \Delta A_i} \left\{ \min_{q \in Q} \sum_{(a'_i, a_{-i}) \in A} u_i(a'_i, a_{-i}) \rho_i(a'_i) \sum_{a_i \in A_i} q_{s_i}(a_i, a_{-i}) \right\}. \tag{1}$$

If the non-deviation condition from equation (1) is satisfied, we say that player i is compliant at the signal realization s_i . Thus, compliance captures the requirement that a player i who receives signal s_i will not object to the mediator choosing an action in the game on her behalf.⁷

⁷ Compliance is similar to the concept of *obedience* from the literature on Bayes correlated equilibrium and information design (see Bergemann and Morris, 2019, for an overview). Note, however, that while obedience refers to players following the action recommendations provided by a mediator, compliance captures the absence of any strict incentives to revoke the delegations to the mediator.

Given any action–signal distribution $q \in \Delta(A \times S)$, we can write $q(a, s)$ as

$$q(a, s) = \sigma(a)\pi(s|a),$$

so that $\sigma \in \Delta A$ is the marginal of q on A , and $\pi(\cdot|a) \in \Delta S$ is the conditional of q over S , given any $a \in A$.⁸ Defining the marginal of $\pi(\cdot|a)$ on S_i as $\pi(s_i|a) := \sum_{s_{-i} \in S_{-i}} \pi(s_i, s_{-i}|a)$, the conditionals $q_{s_i} \in \Delta A$ can be expressed as

$$q_{s_i}(a) = \frac{\sigma(a)\pi(s_i|a)}{\sum_{a' \in A} \sigma(a')\pi(s_i|a')}.$$

We can then view σ as a prior action distribution associated with q , and q_{s_i} as player i 's posterior arising from the signal s_i .

Consider now, as a special case, a DCE where Q is a singleton, so no ambiguity is present in the respective equilibrium. If $q \in \Delta(A \times S)$ denotes the singleton element of Q defining such an equilibrium, and if $\sigma \in \Delta A$ is the marginal of q over A , we will say that σ is implemented through the particular equilibrium. Since there is no strict benefit to mixing in the absence of ambiguity, the definition of a DCE can then be rewritten as follows:

Definition 2. An unambiguous DCE (uDCE) of a game $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$ is a tuple $\langle S, q \rangle$, consisting of a finite signal space $S = \times_i S_i$ and a distribution $q \in \Delta(A \times S)$ with $q(a, s) = \sigma(a)\pi(s|a)$, such that for every player $i \in N$ and each of her possible signal realizations $s_i \in S_i$,

$$\sum_{a \in A} u_i(a)\sigma(a)\pi(s_i|a) \geq \sum_{a \in A} u_i(a'_i, a_{-i})\sigma(a)\pi(s_i|a) \text{ for all } a'_i \in A_i. \tag{2}$$

Thus, in order to establish compliance at any realization s_i of player i 's signal, it is sufficient to ensure that no deviation to an alternative action a'_i is strictly profitable. If the signals of all players are uninformative, that is, if the cardinality of S is one, then the conditions from equation (2) yield the definition of coarse correlated equilibrium (CCE). When $S_i = A_i$, and $\pi(a|a) = 1$, the requirements from (2) reduce to the canonical definition of a standard correlated equilibrium (CE). Hence, uDCE generalizes both standard notions of correlated equilibrium. Since the constraints defining CCE are weaker than those that define CE, we thus have

$$NE \subset CE \subset CCE \subset \text{uDCE}^*,$$

where uDCE^* denotes the distributions over action profiles that are implementable by a uDCE. Furthermore, summing over $s_i \in S_i$ on both sides of the inequalities defining the equilibrium conditions (2), yields conditions that identify $\sigma \in \Delta A$ as a CCE, and thus implies the inclusion $\text{uDCE}^* \subset CCE$, so $\text{uDCE}^* = CCE$.

One of the key benefits that delegation yields within a DCE, is that only the posteriors q_{s_i} , i.e., the first-order beliefs over the action space that are induced by an action–signal distribution $q \in \Delta(A \times S)$, are relevant for the corresponding equilibrium conditions, even though the respective information structure assigns to every signal realization s_i an infinite belief hierarchy over A and the beliefs and higher-order beliefs of i 's opponents. Thus, when constructing (sets of) action–signal distributions that define a DCE or uDCE, we can focus on the specific combination of posteriors that is associated with each signal realization, which determines whether compliance holds for the particular signal. In the unambiguous case, this implies further that any signal space can be restricted to a (finite) subset of ΔA , such that the signal realizations are identified with their induced posteriors, and where the distribution over signals is independent across players, conditional on action profiles.⁹ Specifically, we have

Lemma 1. Given any action–signal distribution $q \in \Delta(A \times S)$ with arbitrary individual signal spaces S_i , there exists an action–signal distribution $p \in \Delta(A \times T)$ with signal spaces $T_i \subset \Delta A$, such that

- (i) p is equivalent to q in the sense that they are both based on the same marginal/prior $\sigma \in \Delta A$, and induce the same collection of posteriors across their respective signal spaces,
- (ii) for each $t_i \in T_i$, $t_i = p_{t_i}$, i.e., each signal is identical to its induced posterior based on p , and
- (iii) the distribution over signals t_i corresponding to p is independent across players, conditional on action profiles a .

We will refer to an action–signal distribution p as characterized in the lemma as a simple action–signal distribution. The lemma implies that if an action distribution is implemented by some uDCE defined by an action–signal distribution $q \in \Delta(A \times S)$, it can also be implemented by a uDCE defined by a simple action–signal distribution $p \in \Delta(A \times T)$ that is

⁸ If $\sigma(a) = 0$, we can set $\pi(\cdot|a) \in \Delta S$ to be any arbitrary distribution.

⁹ The identification of signal realizations with posteriors is analogous to the reduction that is at the core of the seminal Bayesian persuasion paper of Kamenica and Gentzkow (2011), who show that the sender can restrict the signals that she considers to the set of distributions over states of the world.

equivalent to q according to the lemma. However, as will be discussed further in the next section, characterizing DCE's that involve ambiguity requires the introduction of abstract signal “names” that appropriately connect associated posteriors across the action–signal distributions that define a DCE.

The following two lemmas summarize two properties of simple action–signal distributions that characterize the geometric relation between the associated priors and posteriors, and which will provide some technical background for our subsequent analysis and construction of DCE. While the two lemmas just constitute a restatement of the well-known Splitting Lemma of Aumann and Maschler (1995) in the context of simple action–signal distributions, for completeness, in Appendix A we present proofs based on the notation of our model.

Lemma 2. *Suppose that $p \in \Delta(A \times T)$ is a simple action–signal distribution with signal spaces $T_i \subset \Delta A$, so that $p(a, t) = \sigma(a) \prod_{i \in N} \pi_i(t_i|a)$, and let*

$$\text{supp}(\pi_i) := \{t_i \in \Delta A \mid \exists a \in A \text{ such that } \pi_i(t_i|a) > 0\}.$$

Then, for each $i \in N$, there exists a probability distribution $\gamma_i \in \Delta(\text{supp}(\pi_i))$ such that

$$\sigma = \sum_{t_i \in \text{supp}(\pi_i)} \gamma_i(t_i)t_i,$$

that is, the prior σ is a convex combination of the signals/posteriors $t_i \in \text{supp}(\pi_i)$ of each $i \in N$.

Lemma 3. *Let $\sigma \in \Delta A$, $\{t_i^k\}_{k=1}^m \subset \Delta A$, and suppose that there exist numbers $\{\gamma_i^k\}_{k=1}^m \subset (0, 1)$ such that $\sum_{k=1}^m \gamma_i^k = 1$, and*

$$\sigma = \sum_{k=1}^m \gamma_i^k t_i^k,$$

so that σ is a convex combination of $\{t_i^k\}_{k=1}^m$ for every $i \in N$. Then, for every $a \in A$, there exist numbers $\{\pi_i(t_i^k|a)\}_{k=1}^m \subset [0, 1]$ such that $\sum_{k=1}^m \pi_i(t_i^k|a) = 1$, and

$$t_i^k(a) = \frac{\sigma(a)\pi_i(t_i^k|a)}{\sum_{a' \in A} \sigma(a')\pi_i(t_i^k|a')}.$$

Furthermore, when $\sigma(a) > 0$, we have

$$\pi_i(t_i^k|a) = \frac{\gamma_i^k t_i^k(a)}{\sigma(a)}.$$

3. Ambiguously implementable action distributions

Before extending the definition and characterization of implementable action distribution to DCE (with ambiguous correlation devices), we now briefly characterize uDCE with potentially informative signal spaces that are defined by simple action–signal distributions. While, as previously discussed, such uDCE do not enlarge the set of implementable action distributions relative to CCE, this discussion provides critical context and terminology for the subsequent analysis of DCE, and also yields some technical foundations for the version of CE with ambiguous correlation that we introduce in the following section.

In a uDCE defined by a simple action–signal distribution $p(a, t) = \sigma(a) \prod_{i \in N} \pi_i(t_i|a)$, every player i must be compliant at each signal/posterior $t_i \in \Delta A$ that is in the support of $\pi_i(\cdot|a)$ for some a . Hence, the resulting compliance/equilibrium conditions can be summarized as

$$\sum_{a \in A} u_i(a)t_i(a) \geq \sum_{a \in A} u_i(a', a_{-i})t_i(a) \text{ for all } a' \in A_i.$$

If these inequalities hold, we will say that the posterior t_i is *unambiguously compliant*. The set of potential posteriors that are unambiguously compliant is therefore a convex polytope in $\mathbb{R}^{|\Delta A|}$, as it is defined by a finite collection of linear inequalities. Denote this set by \widehat{T}_i , so

$$\widehat{T}_i := \left\{ t_i \in \Delta A \mid \sum_{a \in A} u_i(a)t_i(a) \geq \sum_{a \in A} u_i(a', a_{-i})t_i(a) \text{ for all } a' \in A_i \right\}.$$

If a uDCE is defined by a simple action–signal distribution p with signal spaces T_i , every feasible signal must yield compliant beliefs, and therefore, $T_i \subset \widehat{T}_i$. Since any prior distribution $\sigma \in \Delta A$ that is implemented through such a uDCE is a convex

combination of the relevant signals/posteriors by Lemma 2, and \widehat{T}_i is a convex polytope, the implemented prior σ must also be compliant. As this holds for every player i , the following characterization of action distributions that can be attained through a uDCE is then immediate:

Proposition 4. Fix a game $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$. Then the following statements are equivalent:

- (i) A distribution over action profiles $\sigma \in \Delta A$ can be implemented through a uDCE;
- (ii) $\sigma \in \bigcap_{i \in N} \widehat{T}_i$.

Since $\bigcap_{i \in N} \widehat{T}_i$ also characterizes the set of all action distributions that constitute a coarse correlated equilibrium (CCE), Proposition 4 yields an alternative proof that $\text{uDCE}^* = \text{CCE}$.

We now turn to delegated correlated equilibria (DCE) characterized by a pair $\langle S, Q \rangle$, where $S = \times_{i \in N} S_i$ is a signal space and $Q \subset \Delta(A \times S)$ is an ambiguous correlation device. In order to characterize the compliance conditions for such a DCE $\langle S, Q \rangle$, we need to consider for each signal realization $s_i \in S_i$, the set of conditional distributions over the action space associated with s_i . Denote this set of induced posteriors by \widehat{S}_i , so

$$\widehat{S}_i := \{q_{s_i} \mid q \in Q\} \subset \Delta A.$$

A critical difference to the unambiguous case is that when we introduce ambiguous correlation devices, signal “names,” i.e., their identities, can matter, so that a given signal realization s_i cannot be identified with its associated posterior belief set \widehat{S}_i . Indeed, as demonstrated in Example 2 below (and more generally in Proposition 6), we can construct DCEs based on multiple signal realizations that each yield the same posterior belief set, and which implement distributions over the action space that are not implementable by unambiguous DCE. Beauchêne et al. (2019) refer to distinct signal realizations that yield the same set of posteriors as *synonyms*, and we will occasionally follow this terminology.

The compliance condition for any signal realization s_i that characterizes a DCE $\langle S, Q \rangle$ can be summarized as

$$\min_{t_i \in \widehat{S}_i} \sum_{a \in A} u_i(a) t_i(a) \geq \max_{\rho_i \in \Delta A_i} \left\{ \min_{t_i \in \widehat{S}_i} \sum_{(a'_i, a_{-i}) \in A} u_i(a'_i, a_{-i}) \rho_i(a'_i) \sum_{a_i \in A_i} t_i(a_i, a_{-i}) \right\}. \tag{3}$$

To characterize compliant belief sets \widehat{S}_i , note that the minimum on the RHS of the above inequality is concave in ρ_i , as the point-wise minimum of a family of functions that are linear in ρ_i . For compliance to hold, the maximum over ρ_i of this concave function cannot exceed $\min_{t_i \in \widehat{S}_i} \sum_{a \in A} u_i(a) t_i(a)$ (which is independent of, and hence constant in, ρ_i). Define

$$\underline{t}_i[\widehat{S}_i] := \arg \min_{t_i \in \widehat{S}_i} \sum_{a \in A} u_i(a) t_i(a).$$

Then, since for every $\rho_i \in \Delta A_i$,

$$\sum_{(a'_i, a_{-i}) \in A} u_i(a'_i, a_{-i}) \rho_i(a'_i) \sum_{a_i \in A_i} \underline{t}_i[\widehat{S}_i](a_i, a_{-i}) \geq \min_{t_i \in \widehat{S}_i} \sum_{(a'_i, a_{-i}) \in A} u_i(a'_i, a_{-i}) \rho_i(a'_i) \sum_{a_i \in A_i} t_i(a_i, a_{-i}),$$

we get the following sufficient condition for compliance of a belief set \widehat{S}_i :

Lemma 5. Suppose $\widehat{S}_i \subset \Delta A$ is a finite belief set. If $\underline{t}_i[\widehat{S}_i] \in \widehat{T}_i$, that is, if $\underline{t}_i[\widehat{S}_i]$ is unambiguously compliant, then \widehat{S}_i satisfies equation (3), and is therefore compliant.

While all belief sets \widehat{S}_i associated with a DCE $\langle S, Q \rangle$ must be compliant, Lemma 5 shows that this does not necessarily require all elements of such compliant sets \widehat{S}_i to be unambiguously compliant themselves. This suggests that a distribution $\sigma \in \Delta A$ that is not unambiguously implementable, may be implementable through an ambiguous correlation device, if σ can be expressed as a convex combination of posterior beliefs that allows each s_i to be associated with a belief set \widehat{S}_i containing at least one unambiguously compliant distribution. Before proving a general result along this line, we illustrate and motivate this result using a game taken from Aumann (1974), which we reinterpret as an example of a simple strategic interaction between two self-driving vehicles.

Example 2. Consider a situation in which two self-driving vehicles are approaching a narrow bridge from opposite sides. Each faces the choice between driving *defensively* (denoted, d) or *egotistically*, that is, aggressively (denoted, e). If both drive defensively then they can safely pass by each other on the bridge proceeding at a reasonable speed. However, if one chooses to drive egotistically while the other adopts a defensive posture, then the defensive driver is forced to wait before venturing onto the bridge until the aggressor has sped past. We model this as a 2×2 game with a payoff matrix pictured in Fig. 2.

This game possesses two Nash equilibria (NE) in pure strategies, (d, e) and (e, d) , and a (symmetric) NE in mixed strategies, where each driver plays d with probability $\frac{2}{3}$, yielding an expected payoff of $4\frac{2}{3}$. In order to provide a simple graphical

		Driver 2	
		d	e
Driver 1	d	6	7
	e	2	0

Fig. 2. Payoff matrix for Example 2.




Fig. 3. The set of CCE distributions $(\sigma_{dd}, \sigma_{de}, \sigma_{ed})$ is represented by the darker-gray-shaded area. The set of unambiguously compliant posteriors for player 1 is bounded by the blue line, and for player 2 by the red line. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

illustration of CCE/uDCE distributions, we restrict our analysis to σ that assign probability zero to the outcome (e, e) , and use σ_{dd} , σ_{de} and σ_{ed} to denote the probabilities assigned to strategy profiles (d, d) , (d, e) and (e, d) , respectively.¹⁰ The compliance conditions describing a CCE $\sigma = (\sigma_{dd}, \sigma_{de}, \sigma_{ed})$ reduce to

$$\sigma_{de} \geq \frac{1}{2}\sigma_{dd}, \text{ and } \sigma_{ed} \geq \frac{1}{2}\sigma_{dd}.$$

Thus, the set of CCE distributions is represented by the darker-gray-shaded area of the simplex in Fig. 3. The set of symmetric CCE with $\sigma_{de} = \sigma_{ed}$ lie on the dashed line bisecting the simplex vertically, with the maximal symmetric payoff given by $\sigma^\bullet = (\sigma_{dd}^\bullet, \sigma_{de}^\bullet, \sigma_{ed}^\bullet) = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$ (the black dot in Fig. 3), and yielding an expected payoff of $5\frac{1}{4}$ to each driver. The set of expected payoffs associated with all the CCE $(\sigma_{dd}, \sigma_{de}, \sigma_{ed})$ characterized above are represented by the gray-shaded area in Fig. 4. In this figure, the set of expected payoffs attainable by any randomization over action profiles is enclosed within the dashed lines, the NE payoffs are represented by the white-filled dots, and the optimal symmetric CCE payoffs by the black dot.

The maximal symmetric payoffs that are attainable in this game are $(6, 6)$, corresponding to the action profile (d, d) . The associated action distribution $\sigma^* = (\sigma_{dd}^*, \sigma_{de}^*, \sigma_{ed}^*) := (1, 0, 0)$ is not unambiguously compliant, and since it is a vertex of the associated simplex, it also cannot be expressed as a convex combination that places positive probability on at least one unambiguously compliant action distribution. However, for every $\varepsilon \in (0, 1)$, the vector

$$\sigma^\varepsilon = \begin{pmatrix} \sigma_{dd}^\varepsilon \\ \sigma_{de}^\varepsilon \\ \sigma_{ed}^\varepsilon \end{pmatrix} := (1 - \varepsilon)\sigma^* + \varepsilon\sigma^\bullet = (1 - \varepsilon) \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + \varepsilon \begin{pmatrix} \frac{1}{2} \\ \frac{1}{4} \\ \frac{1}{4} \end{pmatrix} = \begin{pmatrix} 1 - \frac{\varepsilon}{2} \\ \frac{\varepsilon}{4} \\ \frac{\varepsilon}{4} \end{pmatrix}$$

is a convex combination of σ^* and the optimal symmetric CCE distribution $\sigma^\bullet = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$, and has expected payoffs that, for small ε , are arbitrarily close to $(6, 6)$. We will show that $\sigma^* = (1, 0, 0)$ can be approximately implemented in the sense that, for any $\varepsilon \in (0, 1)$, σ^ε can be implemented through an ambiguous correlation device. Furthermore, the correlation device implementing σ^ε only requires a two-element set $Q = \{\bar{q}, \tilde{q}\}$ of action–signal distributions, where both marginals of \bar{q} and \tilde{q} over $\{dd, de, ed\}$ are equal to σ^ε , so that there is no ambiguity regarding the implemented distribution over action profiles in the associated DCE.

To construct Q , note first that since σ^ε is a convex combination of σ^* and σ^\bullet , Lemma 3 shows that there exists a simple action–signal distribution with prior σ^ε and posteriors σ^* and σ^\bullet . We can thus choose an abstract two-element signal space $S = \{s', s''\}$, and replicate this simple action–signal distribution twice, to get two distributions $\bar{q}, \tilde{q} \in \Delta(\{dd, de, ed\} \times \{s', s''\})$, with

¹⁰ If we interpret a CCE/uDCE as a model of self-driving vehicles that are coordinated by a central AI system, setting the probability of (e, e) to zero would be equivalent to assuming that it is common knowledge that the AI is restricted from choosing the aggressive action for both vehicles, which would presumably lead to an accident.




Fig. 4. The set of payoffs for CCE ($\sigma_{dd}, \sigma_{de}, \sigma_{ed}$) is represented by the gray-shaded area.

$$\bar{q}(a, s) = \sigma^\varepsilon(a)\bar{\pi}(s|a) \text{ and } \tilde{q}(a, s) = \sigma^\varepsilon(a)\tilde{\pi}(s|a),$$

and

$$\bar{\pi}(s'|a) = \tilde{\pi}(s''|a) \text{ and } \bar{\pi}(s''|a) = \tilde{\pi}(s'|a) \text{ for all } a \in \{dd, de, ed\},$$

such that¹¹

$$\bar{q}_{s'} = \tilde{q}_{s''} = \sigma^* \text{ and } \bar{q}_{s''} = \tilde{q}_{s'} = \sigma^\bullet.$$

It follows that the belief sets \hat{s}' and \hat{s}'' associated with the ambiguous correlation device $Q = \{\bar{q}, \tilde{q}\}$ are both equal to $\hat{s}' = \hat{s}'' = \{\sigma^*, \sigma^\bullet\}$, so that s' and s'' are synonyms. While we could construct such signal spaces and beliefs separately and independently across the two players, since the game is symmetric, it suffices to allow both players to observe the same ambiguous signals s' and s'' arising from Q . Since σ^\bullet is unambiguously compliant for each player, and $\underline{t}[\hat{s}'] = \underline{t}[\hat{s}''] = \sigma^\bullet$, Lemma 5 implies that both of these ambiguous signals are compliant for every player, and therefore that Q implements σ^ε . \triangleleft

The derivation of the DCE $\langle S, Q \rangle$ in Example 2 illustrates a simple method that can be used to construct an ambiguous correlation device that implements an action distribution *unambiguously*: If, for a given prior σ (corresponding to σ^ε in the example), there exists for each player, a compliant set of posteriors so that σ is in the interior of the convex hull of every such set, Lemma 3 allows us to construct a simple action–signal distribution with each player’s signal space equal to the player’s compliant set of posteriors, and prior given by σ .¹² We can then choose one abstract signal name for every posterior in the compliant set, and replicate the simple action–signal distribution so that based on the associated conditional probabilities, each posterior is assigned to every signal name across the replicated distributions, and thus, every signal name induces the same compliant set of posteriors. This approach to constructing DCE can also be used to generalize the result establishing the possibility of approximate implementation as seen in the example.

Definition 3. An action distribution $\sigma \in \Delta A$ is *implemented* by a DCE $\langle S, Q \rangle$ if $\sigma = \text{marg}_A[q]$ for every $q \in Q$ (hence, while the correlation device Q may be ambiguous, there is *no ambiguity* regarding the action profile distribution σ induced by the DCE). An action distribution $\sigma \in \Delta A$ can be *approximately implemented* if for every $\varepsilon > 0$, there exists a distribution $\sigma^\varepsilon \in \Delta A$ that lies within an ε -neighborhood of σ , and which can be implemented by a DCE $\langle S^\varepsilon, Q^\varepsilon \rangle$.

Proposition 6. Any distribution over action profiles that Pareto dominates a CCE can be approximately implemented through a DCE. Specifically, suppose that σ is a CCE, and assume that $\sigma^* \in \Delta A$ satisfies $\mathbb{E}_{\sigma^*}[u_i] \geq \mathbb{E}_\sigma[u_i]$ for every $i \in N$. Then, for any $\varepsilon \in (0, 1)$,

$$\sigma^\varepsilon := (1 - \varepsilon)\sigma^* + \varepsilon\sigma$$

¹¹ In this particular case, based on the formulas from Lemma 3, we just need to set $\bar{\pi}(s'|dd) = \tilde{\pi}(s''|dd) = \frac{2-2\varepsilon}{2-\varepsilon}$, $\bar{\pi}(s''|dd) = \tilde{\pi}(s'|dd) = \frac{\varepsilon}{2-\varepsilon}$, $\bar{\pi}(s''|de) = \tilde{\pi}(s'|de) = 1$, and $\bar{\pi}(s''|ed) = \tilde{\pi}(s'|ed) = 1$.

¹² In the example, the same signals are used for both players due to symmetry, but the general approach, illustrated in Example 3 below, allows distinct sets of posteriors across players, as long as the prior σ is in the interior of the convex hull of every player’s posterior set.




Fig. 5. Payoffs in the dark-gray-shaded area can be implemented by a DCE.

can be implemented through a DCE $\langle S, Q \rangle$, such that

- (i) Q is a two-element set $Q = \{\bar{q}, \tilde{q}\}$, where $\text{marg}_A[\bar{q}] = \text{marg}_A[\tilde{q}] = \sigma^\varepsilon$, and
- (ii) the signal space S is a two-element set $S = \{s', s''\}$, with every player receiving the same signal realization, and such that both signals yield identical belief sets $\hat{\sigma}' = \hat{\sigma}'' = \{\sigma^*, \sigma\}$.

Proposition 6 states that every action distribution that Pareto dominates a CCE distribution can be approximately implemented through a DCE. However, many such distributions can themselves be expressed as a convex combination of a CCE distribution and a dominating distribution, which implies that all such distributions will in fact be exactly implementable. Returning to the game with two self-driving vehicles from Example 2, it follows that all payoff profiles in the dark-gray shaded area of Fig. 5, except for those on the (red lined) Pareto frontier, will be exactly implementable.¹³

The ambiguous DCE constructed in Example 1 (the prisoner’s dilemma discussed in the introduction) and Example 2 both involve weak preferences for compliance, such that the compliance constraints (3) hold with equality in equilibrium. A key question is then whether this is a feature of DCE, or whether it is just a property of the particular games or equilibria. We can address this question based on the same two games: In Appendix A, we show that every compliant belief set for the prisoner’s dilemma from Example 1 must be such that the compliance constraints hold with equality, which implies that there can be no strict preference for compliance in any associated DCE. Example 3 below then considers the drivers game from Example 2, and constructs a class of DCE that also approximate the payoffs (6, 6) as those from Example 2, but where the compliance constraints are satisfied with strict inequality. It follows that weak preferences for compliance at equilibrium are *not* a feature of DCE, but that they may necessarily arise in the context of certain games.

Example 3. Consider again the two drivers game from Example 2. We next show that the payoffs (6, 6) associated with the action profile (d, d) can also be approximated using a DCE where compliance is *strictly* preferred to revoking the delegation to the mediator for each signal a player may receive. To construct such a DCE, consider, for players $i \in \{1, 2\}$, the set of posteriors $\{t_{i1}, t_{i2}, t_{i3}\}$ given by

$$\begin{aligned}
 t_{11} &= (1 - \delta, \delta, 0), & t_{12} &= \left(\frac{1}{2}, \frac{1}{4}, \frac{1}{4}\right), & t_{13} &= \left(0, \frac{1}{3}, \frac{2}{3}\right), \\
 t_{21} &= (1 - \delta, 0, \delta), & t_{22} &= \left(\frac{1}{2}, \frac{1}{4}, \frac{1}{4}\right), & t_{23} &= \left(0, \frac{2}{3}, \frac{1}{3}\right),
 \end{aligned}$$

where $\delta > 0$ denotes a small strictly positive number, and the ordering of probability values in each vector corresponds to (d, d) , (d, e) and (e, d) . Given the beliefs set $\{t_{i1}, t_{i2}, t_{i3}\}$ for any player i , the minimum on the LHS of the compliance conditions (3) is attained at t_{i2} , with associated expected payoff $5\frac{1}{4}$, and the maxmin on the RHS of the compliance conditions is

¹³ While the Pareto frontier can thus not be attained by any equilibrium, a society or group of players could aim to select a correlation device that implements an approximating σ^ε based on the smallest ε that is perceived by the players as being probabilistically distinct from 0—an experimental analysis may be able to identify such a value ε .




Fig. 6. The set of action distributions that are implementable by the proposed DCE is represented by the interior of the intersection of the blue and red triangles.

attained at t_{i3} for any $\rho_i \in \Delta A_i$, with associated expected payoff $4\frac{2}{3}$. It follows that the compliance condition holds strictly for the belief sets $\{t_{i1}, t_{i2}, t_{i3}\}$.

To construct a DCE that approximates the payoffs (6, 6), refer first to Fig. 6, and note that in this figure, the convex combinations of $\{t_{i1}, t_{i2}, t_{i3}\}$ are enclosed by the blue triangle for driver 1, and by the red triangle for driver 2. As $\delta \rightarrow 0$, the intersection of these triangles approaches (1, 0, 0) arbitrarily closely, so that we can choose a point in the interior of this intersection that approximates (1, 0, 0) to any desired degree of precision. Denoting such a point by σ , this point can be expressed as a convex combination of each player's belief set $\{t_{i1}, t_{i2}, t_{i3}\}$, with strictly positive weight for every respective distribution, so that we can apply Lemma 3 to construct a simple action–signal distribution with prior σ and posteriors $\{t_{i1}, t_{i2}, t_{i3}\}$ for any player i . By specifying abstract individual signal spaces $S_i = \{s'_i, s''_i, s'''_i\}$ for each player $i \in \{1, 2\}$, we can then replicate the simple action–signal distribution three times by cycling through the respective conditional probabilities across the signal names s'_i, s''_i and s'''_i , so that the induced beliefs sets are equal across signals for each player, and given by $\hat{s}'_i = \hat{s}''_i = \hat{s}'''_i = \{t_{i1}, t_{i2}, t_{i3}\}$. \triangleleft

The approach used to construct the equilibrium from the previous example, which only requires a collection of finite compliant belief sets as a starting point, without the need to specify a priori signal spaces, can be generalized to derive a further characterization of action profile distributions that can be implemented through DCE:

Proposition 7. For every $i \in N$, let $R_i \subset \Delta A$ denote a finite set of action profile distributions such that each R_i constitutes a compliant belief set for player i , and use \tilde{R}_i to denote the relative interior of the convex hull of R_i . Then every $\sigma \in \bigcap_{i \in N} \tilde{R}_i$ can be implemented through a DCE that requires $|R_i|$ signal realizations for each player i , and $\max_{i \in N} |R_i|$ action–signal distributions.

4. Ambiguous correlated equilibrium

The results derived in the previous section show how introducing ambiguity into the correlation devices previously associated with CCE can yield welfare-improving enlargements of the sets of implementable action distributions. One drawback of the resulting notion of DCE is that it requires the availability of a mediator who can implement the action choices recommended by the correlation device. While this need not be a problem in certain strategic settings, it restricts the applicability of DCE to situations where players may have an opportunity to coordinate through some extraneous correlation device, but are not able to delegate their action choices to such a mediator. We will next analyze whether the implementability results for DCE apply to such environments, by introducing an extension of CE to settings with ambiguity averse players, which we call an *ambiguous correlated equilibrium (ACE)*.

Our approach to defining such ACE follows from a reinterpretation of CE as an unambiguous DCE (uDCE) with an informative signal space, where, in addition, the set of feasible posteriors induced by any signal realization is constrained to lie in some exogenously specified subset of ΔA . In standard CE, signals implicitly (or explicitly for canonical CE) involve action recommendations for each player, and hence, a player's associated posteriors must place probability one on the respective optimal actions of the player. Hence, we can view a CE as a uDCE where the posteriors of a player are restricted to belong to the faces of the probability simplex ΔA that assign probability one to this player's actions. For such a uDCE, players can be entrusted to choose their own equilibrium actions after having observed their respective signals, so delegation of action choices to a mediator is not required.¹⁴

¹⁴ Note also that such a uDCE will be equivalent to a canonical CE, if for each action of a player, there is only one signal that yields a posterior assigning probability one to this action.




Fig. 7. Constrained posteriors—blue for driver 1, red for driver 2.

Example 4. As an illustration, consider again the drivers game from Example 2. Referring to Fig. 7, in a CE, any signal that recommends action d to player 1 must yield a posterior that lies on the edge of the simplex that connects the points $\sigma_{dd} = 1$ and $\sigma_{de} = 1$. Similarly, keeping the assumption that outcome (e, e) is assigned probability zero, any signal that recommends action e to player 1 must yield a posterior that assigns probability one to (e, d) . Hence, the set of feasible posteriors for driver 1 is represented by the union of the blue line (connecting $\sigma_{dd} = 1$ and $\sigma_{de} = 1$) and blue dot (representing $\sigma_{ed} = 1$). Analogously, the set of feasible posteriors for driver 2 is represented by the union of the red line (connecting $\sigma_{dd} = 1$ and $\sigma_{ed} = 1$) and red dot (representing $\sigma_{de} = 1$).

Furthermore, the optimality of a recommended action in a CE is equivalent to compliance of the associated posterior for the equivalent constrained uDCE, which restricts the set of posteriors that can be associated with a CE to the thick segments of the blue and red lines in Fig. 7 (satisfying $\sigma_{de} \geq \frac{1}{2}\sigma_{dd}$ for driver 1, and $\sigma_{ed} \geq \frac{1}{2}\sigma_{dd}$ for driver 2), in addition to the blue and red dot, respectively. Following Proposition 4 and Lemma 3, an action distribution σ can then be implemented by a CE if and only if it is a convex combination of posteriors that are both feasible and compliant for both players, which yields the same gray-shaded area of action distributions in the simplex that are also implementable by CCE (as expected, since CE and CCE are equivalent for two-action games). \triangleleft

Motivated by this example, we introduce a general definition of constrained DCE:

Definition 4. For a game $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$, let $T_i^c \subset \Delta A$ denote an exogenously given set of feasible posterior beliefs for any player $i \in N$. Then a DCE $\langle S, Q \rangle$ is a *constrained DCE* with respect to the constraint sets $(T_i^c)_{i \in N}$, if $\widehat{s}_i \in T_i^c$ for every player $i \in N$ and $s_i \in S_i$.¹⁵

For an *unambiguous* constrained DCE (i.e., a constrained uDCE), the requirement that $\widehat{s}_i \in T_i^c$ is equivalent to every posterior associated with a signal realization being an element of the constraint set T_i^c . Hence, every such posterior must be both feasible, in the sense of lying in T_i^c , and compliant, as an element of \widehat{T}_i . Proposition 4 and Lemma 3 then imply that an action distribution σ is implementable as an constrained uDCE with constraint sets $(T_i^c)_{i \in N}$, if and only if $\sigma \in \bigcap_{i \in N} \text{conv}(T_i^c \cap \widehat{T}_i)$, where conv denotes the convex hull of the respective sets.¹⁶

Proposition 7 provided a convenient method to construct DCE that only requires a collection of finite compliant belief sets as a starting point, without any need to specify a priori signal spaces. The associated techniques to characterize DCE also allow us to derive a version of Proposition 6 for constrained DCE, which follows below after an additional definition.

Definition 5. Let $T_i^c \subset \Delta A$ denote the set of feasible posterior beliefs of a player i , and let $\sigma \in \Delta A$. Then a finite set $\overline{R}_i[\sigma] \subset \Delta A$ is a *supporting set* of σ for player i , if $\overline{R}_i[\sigma] \subset T_i^c$ and σ belongs to the relative interior of the convex hull of $\overline{R}_i[\sigma]$.

Proposition 8. Given a game $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$, let $T_i^c \subset \Delta A$ denote the set of feasible posterior beliefs for any player $i \in N$. Suppose $\sigma \in \bigcap_{i \in N} \text{conv}(T_i^c \cap \widehat{T}_i)$ is an unambiguous constrained DCE, and that $\sigma^* \in \Delta A$ satisfies $\mathbb{E}_{\sigma^*}[u_i] \geq \mathbb{E}_{\sigma}[u_i]$ for every $i \in N$. Then σ^* can be approximately implemented through an (ambiguous) constrained DCE if one of the following conditions is satisfied:

¹⁵ While we introduce constrained DCE as a foundation for defining ambiguous CE, exogenous constraints on the structure of the signal space could also arise from various technical or legal requirements that require players to be provided with information of a certain kind, which one might envision to be the case in the context of self-driving vehicles, or other situations where game play is mediated by a central AI system.

¹⁶ Since the constraint sets T_i^c were defined independently of the sets \widehat{T}_i of unambiguously compliant posteriors, their intersection may be empty, and hence a corresponding constrained uDCE need not exist. However, the same argument that identifies $\bigcap_{i \in N} \text{conv}(T_i^c \cap \widehat{T}_i)$ as the set of implementable action distributions, also implies that a constrained uDCE exists if and only if $\bigcap_{i \in N} \text{conv}(T_i^c \cap \widehat{T}_i)$ is non-empty.




Fig. 8. Constrained Pareto dominance.

- (a) for each player i there exist supporting sets $\bar{R}_i[\sigma]$ and $\bar{R}_i[\sigma^*]$ such that $\bar{R}_i[\sigma] \cup \bar{R}_i[\sigma^*]$ is a compliant belief set;
- (b) $\sigma^* \in \bigcap_{i \in N} T_i^c$, so σ^* is itself a feasible posterior for all players.

While in the unambiguous case, CE can equivalently be defined as an unambiguous constrained DCE where any player's feasible posteriors assign probability one to the player's own actions, this characterization does not directly carry over to a definition of ambiguous CE. This is a consequence of the fact that an ambiguous belief set based on such a constraint set may still contain distinct posteriors that assign probability one to different actions of a player, and hence such a belief set may not yield a unique action that the player can select by herself. The following example illustrates how this issue may be resolved, and also shows how the previously used construction of DCE may need to be adjusted as a result, so that not all signals can be synonyms.

Example 5. Consider a constrained version of the drivers game from Example 2, where every feasible posterior of a player assigns probability one to an action of this player. If we define

$$R_1 := \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} \frac{2}{3} \\ \frac{1}{3} \\ 0 \end{pmatrix} \right\}, \text{ and } R_2 := \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \frac{2}{3} \\ 0 \\ \frac{1}{3} \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \right\},$$

then, for each $i \in \{1, 2\}$, the elements of R_i are feasible posteriors for player i , and R_i constitutes a compliant belief set for i . Refer to Fig. 8, and note that each point in the blue-hatched area can be expressed as a convex combination of elements of R_1 , and that each point in the red-hatched area can be expressed as a convex combination of elements of R_2 . Furthermore, points that lie in the red-hatched area but not in the blue-hatched area are constrained unambiguously compliant for player 1, and points that lie in the blue-hatched area but not in the red-hatched area are constrained unambiguously compliant for player 2. It then follows from Proposition 7 or Proposition 8 that all points in the relative interior of the union of the blue-hatched and red-hatched areas are implementable by a constrained DCE, which implies that all action distributions that Pareto dominate a CCE can be approximately implemented through such a constrained DCE.

Even though the individual posteriors associated with the constrained DCE described above assign probability one to a particular action of each player, the resulting belief sets R_1 and R_2 each contain some posteriors that assign probability one to d , and others that assign probability one to e , and therefore, the resulting DCE still require delegation to a mediator. However, we can easily amend the construction of these DCE to show that the same action distributions can also be implemented based on constrained DCE that *do not* require delegation. To see this, consider any $\sigma \in \Delta A$ that also lies in \tilde{R}_1 (the relative interior of the convex hull of R_1), and let $S_1 = \{s_1^d, s_1^{d*}, s_1^e\}$ denote a set of signal names for player 1. Following Lemma 3, we can construct conditional probabilities $\bar{\pi}(\cdot|a)$ and $\tilde{\pi}(\cdot|a)$ over S_1 so that

$$\bar{q}(a, s_1) = \sigma(a)\bar{\pi}(s_1|a) \text{ and } \tilde{q}(a, s_1) = \sigma(a)\tilde{\pi}(s_1|a),$$

and the belief sets derived from $Q = \{\bar{q}, \tilde{q}\}$ for the signals in S_1 are

$$\hat{s}_1^d = \hat{s}_1^{d*} = \left\{ (1, 0, 0), \left(\frac{2}{3}, \frac{1}{3}, 0 \right) \right\} \text{ and } \hat{s}_1^e = \{(0, 0, 1)\}.$$

Each of these belief sets are compliant, and the corresponding posteriors assign probability one to the same action of player 1, so the player can be entrusted to choose her own action after being informed of the respective signal realization, and therefore delegation to the mediator is not required. Analogously, we can define a signal space $S_2 = \{s_2^d, s_2^{d*}, s_2^e\}$ for player 2 that generates belief sets

$$\hat{s}_2^d = \hat{s}_2^{d*} = \left\{ (1, 0, 0), \left(\frac{2}{3}, 0, \frac{1}{3} \right) \right\} \text{ and } \hat{s}_2^e = \{(0, 1, 0)\}$$

for any prior $\sigma \in \tilde{R}_2$. It follows that every action distribution that lies in the relative interior of the union of the blue-hatched and red-hatched areas in Fig. 8 can be implemented based on a DCE where players are allowed to choose their own actions, as recommended by their posterior belief sets.

The DCE described above are constructed so that the convex combinations of the associated posteriors (the union of the blue-hatched and red-hatched areas) exactly generate the action distributions that Pareto dominate the set of CCE/uDCE. A simpler class of constrained DCE can however be derived by analogously assigning to the signals s_1^d, s_1^{d*} and s_1^e of player 1 the belief sets

$$\hat{s}_1^d = \hat{s}_1^{d*} = \{(1, 0, 0), (0, 1, 0)\} \text{ and } \hat{s}_1^e = \{(0, 0, 1)\},$$

and to the signals s_2^d, s_2^{d*} and s_2^e of player 2, the belief sets

$$\hat{s}_2^d = \hat{s}_2^{d*} = \{(1, 0, 0), (0, 0, 1)\} \text{ and } \hat{s}_2^e = \{(0, 1, 0)\}.$$

Since $(0, 1, 0)$ is unambiguously compliant for player 1, and $(0, 0, 1)$ is unambiguously compliant for player 2, these belief sets are compliant, and the respective posteriors span the full simplex for both players, so every action distribution in the interior of the simplex can be implemented by appropriately choosing the conditional probabilities defining the corresponding action–signal distributions \bar{q} and \tilde{q} . The resulting DCE are especially appealing with regard to the interpretation of the game as capturing the interaction between two drivers—when a driver $i \in \{1, 2\}$ receives signals s_i^d or s_i^{d*} , she is instructed to choose action d , i.e., drive defensively, and only knows that her opponent chooses either d or e with probability one. The resulting ambiguity then implies that i 's belief sets \hat{s}_i^d or \hat{s}_i^{d*} are compliant, so that i optimally selects the recommended action d . Similarly, if i receives the signal s_i^e , she knows that her opponent plays d for sure, so that the recommended action e is again optimal. \triangleleft

The construction in the example of constrained DCE that are not reliant on delegation, and where players can be entrusted to choose their own actions, as in standard CE, motivates the following definition:

Definition 6. An *ambiguous correlated equilibrium (ACE)* of a game $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$ is a constrained DCE (S, Q) with constraint sets $(T_i^c)_{i \in N}$, such that

- (i) for every $t_i \in T_i^c$, there exists an action $a_i \in A_i$ so that $\text{marg}_{A_i}[t_i] = \delta_{a_i}$, where $\delta_{a_i} \in \Delta A_i$ denotes the degenerate probability distribution that assigns probability one to a_i , and
- (ii) for every $s_i \in S_i$ and $t_i, t'_i \in \hat{s}_i$, $\text{marg}_{A_i}[t_i] = \text{marg}_{A_i}[t'_i]$.

The discussion from Example 5 shows that the conclusion of Propositions 6 and 8 holds for the drivers game even when the solution concept is restricted to ACE, i.e., every action distribution that Pareto dominates a CE can also be approximated through an ACE. However, if we consider the prisoner's dilemma from Example 1, it is easy to see that any (potentially ambiguous) belief set of a player that places probability one on the player's cooperative action c will yield a best response to play d , and therefore that this result does not hold in general across all games. The following characterization of action distributions that can be implemented by an ACE, can be used to determine whether an action distribution that dominates a CE is implementable for a particular game:

Proposition 9. For any action distribution $\sigma \in \Delta A$ and $i \in N$, let A_i^σ denote the set of $a_i \in A_i$ such that $\text{marg}_{A_i}[\sigma](a_i) > 0$. Then σ can be (unambiguously) implemented by an ambiguous correlated equilibrium if and only if for every $i \in N$, there exists a finite set $\bar{R}_i[\sigma] \subset \Delta A$, such that

- (i) σ belongs to the relative interior of the convex hull of $\bar{R}_i[\sigma]$,
- (ii) there exists a partition $\{\bar{R}_{a_i}\}_{a_i \in A_i^\sigma}$ of $\bar{R}_i[\sigma]$ satisfying $\text{marg}_{A_i}[t_i] = \delta_{a_i}$ for every $t_i \in \bar{R}_{a_i}$, and
- (iii) \bar{R}_{a_i} is a compliant belief set for every $a_i \in A_i^\sigma$.

The characterization from Proposition 9 also applies in a setting where every posterior belonging to a belief set \bar{R}_{a_i} is a degenerate distribution $\delta_a \in \Delta A$, which assigns probability one to an action profile a whose i -th coordinate is equal to a_i , as in the last class of constrained DCE described in Example 5.

An alternative approach to introduce ambiguity to CE would have been to extend the canonical CE and define an ambiguous version as a set of distributions over action profiles, so that each signal of a player informs the player of her recommended action according to the distributions from this set. Such an equilibrium would still be an ACE according to our definition, but would yield a strictly smaller set of equilibria, as it would only be able to introduce ambiguity through a non-singleton set of action distributions, in contrast to ACE which allows ambiguity to be introduced via signals, while keeping the associated distribution over action profiles unambiguous.

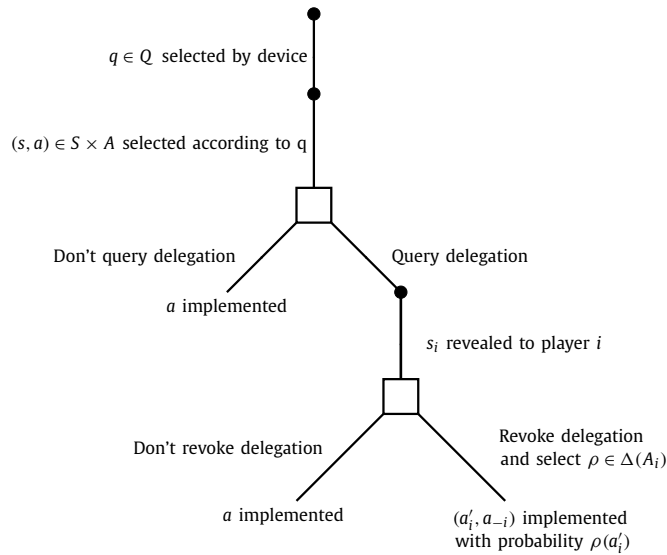


Fig. 9. Timeline from perspective of player i of a delegated correlation device with the addition of the option to first query the delegation.

5. Discussion

5.1. Welfare

In the statements of Proposition 6 and Proposition 8, we compared action distributions using Pareto dominance with respect to associated expected payoffs. This approach follows the convention in the literature on correlated equilibrium, where player welfare induced by a CE $\sigma \in \Delta A$ is defined by the expected utility induced by σ ,

$$U_i(\sigma) := \mathbb{E}_\sigma[u_i] = \sum_{a \in A} u_i(a)\sigma(a).$$

Since our definition of implementation of an action distribution by an (ambiguous) DCE required the induced action distribution to be unambiguous, these expected utilities do not necessitate a representation of the players' ambiguity-sensitive preferences. However, given that our model also involves an interim stage after players observe their signals associated with a DCE, the above expression for expected utility in fact corresponds to an ex-ante evaluation of player welfare, before players observe their respective signals.

One way to motivate such an ex-ante approach to welfare is to consider the perspective of a social planner, who aims to implement a regulatory framework for the interactions modeled by the game that, in equilibrium, yields an ex-ante Pareto efficient outcome. If such a social planner is not concerned by the players' interim expected payoffs (which would generally be lower than their ex-ante payoffs, as a result of the ambiguity they face after observing potentially ambiguous signals), and only worries about the payoffs associated with the (unambiguous) implemented action distribution, then the social planner would choose a regulatory framework/game structure that has the potential to attain the Pareto dominating DCE described in Propositions 6 and 8.

From the point of view of the players, the comparison of equilibria based on the expected payoffs $U_i(\sigma)$ induced by the associated action distributions requires players to evaluate equilibria from an ex-ante perspective. It is well-known that ambiguity averse players are generally not dynamically consistent (see, for example, Siniscalchi, 2011, for a recent discussion and analysis). However, if players have to decide on the type of strategic interaction they would like to participate in before observing any potential signals, the ex-ante perspective is appropriate, and would be equivalent to the standard consistent planning approach to analyzing dynamic decisions proposed by Strotz (1955–1956), and justified axiomatically in Siniscalchi (2011). For example, in a self-driving car setting, a driver may need to decide before venturing on the road whether to use an AI-equipped car that has the ability to engage in correlated play when encountering other self-driving cars, or whether to use a regular car that does not have the ability to employ ambiguous correlation. Since this decision would need to be made before any signals associated with a DCE are observed, the ex-ante point of view is justified by consistent planning, and our results would imply that such a driver would choose to use the AI-equipped car as long as there is a possibility to encounter other cars that are similarly equipped to play the Pareto dominating DCE.

An alternative method that avoids the lower interim payoffs associated with ambiguous signals, is to extend the game under consideration by adding an additional stage that provides an option to each player to query the delegation first. Fig. 9 illustrates the timeline of such an extended game from the perspective of player i while suppressing the decisions of i 's

opponents, assuming that the associated correlation device constitutes a DCE, so that as part of an appropriate extension of the definition of DCE to the resulting extensive game form, player i can expect all players other than herself to be compliant (either by not querying the delegation, or by not revoking it should they choose to query). Notice that player i observes the realization s_i of the ambiguous signal generated by the delegated correlated device only should she choose to exercise her option to query the delegation. The conditions that characterize a DCE would then imply that each player would optimally choose not to query the delegation which would result in the action distribution recommended by the correlation device being played, and the players only having to evaluate the DCE based on the associated ex-ante payoffs. For example, in such an extension of the two drivers game from Example 2, when the AI driving system detected the AI system of another self-driving car, the driver would first be given the option to query the AI's delegation of handling the interaction with the other self-driving car. Only if the driver chose to query the delegation would the AI system reveal to the driver her (ambiguous) signal. But anticipating that her ambiguous signal is part of a DCE, the driver would strictly prefer not to query the delegation, thus leaving the interaction to be handled by the AI systems of the two cars.

Finally, note also that even though our notion of implementation for action distributions requires equilibrium ambiguity to be restricted to signals, and our main results seem to suggest that it is sufficient to focus the analysis of most games to such DCE with unambiguous action distributions, for a general DCE (S, Q) , an appropriate welfare measure would need to take into account that individual distributions in Q could yield distinct action distributions. To define a corresponding notion of player welfare, let $\sigma^q \in \Delta A$ denote the marginal of any $q \in Q$ on A . The (ex-ante) individual welfare of player i associated with (S, Q) could then be defined by

$$U_i(Q) := \min_{q \in Q} \mathbb{E}_{\sigma^q}[u_i] = \min_{q \in Q} \sum_{a \in A} u_i(a) \sigma^q(a),$$

and from a social point of view, different equilibria could be ranked based on a Pareto criterion associated with the equilibrium values for $U_i(Q)$, or a social welfare function that aggregates the players' individual ex-ante welfare.

5.2. Robustness to alternative preference specifications

While our analysis so far has assumed that players possess maxmin preferences based on the full set of posteriors induced by the correlation devices that define an equilibrium, the equilibrium notions we proposed can easily be adapted to alternative preference specifications, and furthermore, the results we derived may still hold, depending on the games under consideration. If we instead assume that players' preferences have an α -maxmin (Jaffray, 1989; Gul and Pesendorfer, 2015) or smooth ambiguity aversion (Klibanoff et al., 2005) representation, a DCE can still be defined by a correlation device (S, Q) , consisting of a signal space $S = \times_{i \in N} S_i$ and a set of action-signal distributions $Q \subset \Delta(A \times S)$, together with appropriately adjusted compliance conditions. Moreover, since the additional constraints introduced to define ACE do not depend on a specific preference representation, they carry over unchanged to the α -maxmin and smooth models.

For any posterior belief set $\hat{s}_i \subset \Delta A$ of player i arising from a correlation device (S, Q) based on full Bayesian updating, the associated compliance condition with α -maxmin preferences is given by

$$\alpha_i \min_{t_i \in \hat{s}_i} \sum_{a \in A} u_i(a) t_i(a) + (1 - \alpha_i) \max_{t_i \in \hat{s}_i} \sum_{a \in A} u_i(a) t_i(a) \geq \max_{\rho_i \in \Delta A_i} \left\{ \alpha_i \min_{t_i \in \hat{s}_i} \sum_{(a'_i, a_{-i}) \in A} u_i(a'_i, a_{-i}) \rho_i(a'_i) \tilde{t}_i(a_{-i}) + (1 - \alpha_i) \max_{t_i \in \hat{s}_i} \sum_{(a'_i, a_{-i}) \in A} u_i(a'_i, a_{-i}) \rho_i(a'_i) \tilde{t}_i(a_{-i}) \right\},$$

where $\alpha_i \in [0, 1]$ is player i 's Hurwicz pessimism/optimism index that measures the player's attitude towards ambiguity, and $\tilde{t}_i(a_{-i}) := \text{marg}_{A_{-i}}[t_i](a_{-i}) = \sum_{a_i \in A_i} t_i(a_i, a_{-i})$. Similarly, the compliance condition for smooth preferences is given by

$$\sum_{t_i \in \hat{s}_i} \mu_{s_i}(t_i) \phi_i \left(\sum_{a \in A} u_i(a) t_i(a) \right) \geq \max_{\rho_i \in \Delta A_i} \left\{ \sum_{t_i \in \hat{s}_i} \mu_{s_i}(t_i) \phi_i \left(\sum_{(a'_i, a_{-i}) \in A} u_i(a'_i, a_{-i}) \rho_i(a'_i) \tilde{t}_i(a_{-i}) \right) \right\},$$

where $\mu_{s_i} \in \Delta(\hat{s}_i)$ measures player i 's subjective assessment of the likelihood of the possible posteriors induced by the signal realization s_i , and $\phi_i : \mathbb{R} \rightarrow \mathbb{R}$ is a continuous and strictly increasing function that captures the player's attitude towards ambiguity.¹⁷

¹⁷ Specifying the subjective beliefs μ_{s_i} separately for each s_i allows for these beliefs to be derived from an initial subjective distribution over Q , but also for the possibility that μ_{s_i} depends only on the posterior belief set induced by s_i , so that synonyms yield identical subjective beliefs. Note also that while ϕ_i , and analogously α_i for the case of α -maxmin preferences, can be viewed as exogenous to a specific equilibrium played, the subjective beliefs μ_{s_i} would arise as a consequence of the correlation device that defines an equilibrium, and could therefore be viewed as an integral component of the associated equilibrium.

It is easy to see that any DCE that involves strict preferences for compliance must be locally robust with respect to α -maxmin preferences, in the sense that the associated belief sets will still be compliant for sufficiently large values of α_i . Returning to the prisoner's dilemma and the drivers game examples, also shows that the strong implementability results we derived previously for DCE and ACE, can still hold for α -maxmin and smooth preferences. While the sufficient condition for compliance of a belief set from Lemma 5 requires maxmin preferences, Proposition 7 still applies with α -maxmin or smooth preferences, so that constructing a DCE just requires a judicious selection of appropriate belief sets that are compliant with respect to the given preference specifications. In particular, the two examples presented in Appendix B show that for the prisoner's dilemma and the drivers game, all the action distributions that are implementable by DCE following Proposition 6, are also implementable with α -maxmin preferences, for any values of $\alpha_i \in (0, 1)$. Furthermore, there exists a large class of parametric specifications of the smooth preference model for which similar implementability properties hold for the drivers game. On the other hand, ACE does impose some restrictions on implementability for the drivers game, either with respect to the specifications of the alternative preference models, or the set of implementable action distributions (i.e., all the action distributions that dominate a CCE can only be approximately implemented via ACE with α -maxmin preferences if each $\alpha_i \geq \frac{1}{3}$, and for $\alpha_i < \frac{1}{3}$, the set of ACE-implementable distributions converges to the set of CCE as all $\alpha_i \rightarrow 0$).

Our initial model based on maxmin preferences, as well as the extensions to α -maxmin and smooth preferences we discussed above, assume that the players' expected payoffs are computed based on the exact "objective" posterior belief sets derived from the correlation device using full Bayesian updating. Gajdos et al. (2008) propose an alternative axiomatic representation of maxmin preferences based on an *objectively given* set of beliefs, but where the corresponding maxmin expected payoffs are computed with respect to a selection from the given objective belief set. Their model is ideally suited to our setting, as we can interpret a posterior belief set \hat{s}_i that is induced by a given signal realization s_i as objectively given, and consider players who compute their maxmin expected payoffs based on a selection $\varphi_i(\hat{s}_i)$, so that the (subjective) operator φ_i can be viewed as embodying player i 's attitude towards ambiguity. As a specific functional form, Gajdos et al. (2008) propose that the operator φ_i may "shrink" the given objective belief sets \hat{s}_i towards a reference distribution, which, in our setting, could be the prior σ . Since the resulting preference representation still has a maxmin form, but now with respect to the subjective belief sets $\varphi_i(\hat{s}_i)$, Lemma 5 still applies, so that the question of implementability now just depends on how "large" the objective posterior sets \hat{s}_i can be constructed so that their shrunken versions $\varphi_i(\hat{s}_i)$ remain compliant.

6. Related literature

6.1. Ellsberg equilibrium

In the ambiguous case, the notion of DCE can be viewed as a correlated version of the *Ellsberg equilibrium* proposed by Riedel and Sass (2014). Riedel and Sass (2014) study a novel class of normal-form games, termed Ellsberg games, in which ambiguity averse players are allowed to choose their actions based on independent ambiguous randomization devices. Thus, in contrast to the standard theory where a player's strategy set is given by the set of probability distributions over her action set, in an Ellsberg game a player's strategy set can be modeled as the set of all non-empty sets of distributions over actions. Assuming that players' preferences can be represented by maxmin expected utility, Riedel and Sass (2014) characterize the resulting set of Ellsberg equilibria, and show that they can yield interesting predictions that are distinct from those associated with Nash equilibrium.

6.2. Ambiguity in information and mechanism design

There is now a growing literature showing that ambiguity can be beneficial in information design problems with ambiguity averse agents. Kellner and Le Quement (2017, 2018) demonstrate this in the context of cheap talk games (which differ from the sender-receiver games from the Bayesian persuasion literature in that the sender is not assumed to be able to commit to a communication device). Somewhat closer to our paper, Beauchêne et al. (2019) introduce the possibility of ambiguous, that is, set-valued communication devices into the classic sender-receiver game with a single receiver of Kamenica and Gentzkow (2011). As for the ambiguous correlation devices we consider, their set-valued communication devices result in a set of probabilistic posterior beliefs (over states of the world) that are derived based on full Bayesian updating, and are thus analogous to the ambiguous belief sets over action profiles in our framework. Beauchêne et al. (2019) also construct ambiguous communication devices where distinct signals are synonyms in the sense that they are associated with identical belief sets, which, while yielding similar features as the ambiguous correlation devices used for the proof of our Proposition 6, rely on a different construction necessitated by the alternative space of uncertainty of the sender-receiver game, and by the fact that the receiver must be induced to choose appropriate actions after receiving the signals, whereas for DCE players must only be prevented from revoking their delegation to the mediator. Given that the sender in a persuasion game aims to design a communication device that maximizes her own payoffs, Beauchêne et al. (2019) also show that any sender payoff that is attained by an ambiguous communication device, can also be attained by one that only contains two elements, analogous to the two-element correlation devices from Proposition 6.

The endogenous introduction of ambiguity through a mechanism designer can also be advantageous in standard mechanism design environments.¹⁸ Di Tillio et al. (2017) allow the seller in a classic screening model to employ ambiguous, set-valued mechanisms, and show that such a seller can attain higher expected payoffs based on mechanisms that are ambiguous in both the allocation and transfer rule. Their results also extend to independent private values auctions when valuations are drawn from an atomless distribution. Guo (2019) shows that first-best surplus extraction and interim individually rational and ex-post budget balanced implementation can be guaranteed through mechanisms that are only ambiguous in the transfer rules, as long as distinct payoff-relevant types of agents are always assumed to have distinct beliefs.

An interesting model that, in a certain sense, combines both mechanism and information design, is presented by Bose and Renou (2014), who consider a classic social choice setting, and allow the mechanism designer to endogenously introduce ambiguous beliefs of agents regarding other agents' types through the use of a dynamic mediated communication game that precedes the play of the actual, unambiguous, allocation mechanism. Bose and Renou (2014) show that if agents are ambiguity averse and update their beliefs using the full Bayesian updating rule, then social choice functions that are not implementable based on unambiguous prior beliefs, may be implementable after agents' beliefs are appropriately manipulated using an ambiguous communication game. This approach to generating ambiguous beliefs through communication is analogous to our method of introducing ambiguous beliefs regarding the action profiles that are implemented by a mediator, through an ambiguous correlation device. The objective in both cases is to induce a certain type of behavior through engineering appropriate ambiguous beliefs under which this behavior is optimal. In our games, delegation plays a crucial role, signals and the resulting beliefs are induced simultaneously with the draw of the action profiles, and the role of the ambiguity in beliefs is to prevent players from revoking their delegation to the mediator. In Bose and Renou (2014), the uncertainty is over the agents' privately-known payoff-relevant types, and the role of the players' ambiguous beliefs regarding other players' types is to induce them to truthfully reveal their own payoff-relevant types, in order to enable the implementation of the given social choice function.

Introducing ambiguous communication can also be beneficial in mechanism design problems that allow for certifiable information, as demonstrated by Ayouni and Koessler (2017). Their paper proposes a model where the mechanism designer has the option of requesting information certification by an agent based on an ambiguous strategy, and shows that the resulting ambiguous communication device eliminates the effect of restrictions on the amount of information that can be certified, in the sense that any allocation rule that is implementable with no limits on certification, is also implementable with limited certification and an ambiguous communication strategy.

6.3. The price of anarchy and related welfare comparisons

The computer science literature has proposed, and extensively analyzed a number of “inefficiency measures” that compare welfare outcomes arising from equilibrium solution concepts in strategic games, to maximal welfare levels that are attainable in such settings when outcomes can be chosen dictatorially without taking into account the strategic nature of the problems. The main objective of this analysis is to establish upper bounds on the ratios of certain equilibrium welfare and maximal welfare levels, which can then be interpreted as bounding the inefficiency arising from “selfish” equilibrium behavior.¹⁹ Various equilibrium solution concepts and selection methods from the associated equilibrium sets have been considered in this context. The first such measure, proposed by Koutsoupias and Papadimitriou (1999) and currently known as the *price of anarchy*, considers the ratio of the maximal overall welfare and the minimal welfare across the set of pure strategy Nash equilibria (the inverse of this ratio is used when the objective is cost minimization). Roughgarden (2015) extends the analysis of the price of anarchy to Nash equilibrium in mixed strategies, CE and CCE. An analogous measure that compares maximal overall welfare levels with maximal equilibrium welfare levels is the *price of stability* when computed with respect to Nash equilibria, or the *enforcement value* when computed with respect to CE.²⁰ Within this literature, congestion games provide a key application for which various inefficiency bounds have been characterized.

The principal connection to the analysis of DCE and ACE follows from the fact that DCE, and for some games also ACE, allow the approximate implementation of all action distributions that Pareto dominate CE and CCE outcomes in a game, and do so without introducing any ambiguity over the induced equilibrium action distributions. As a consequence, when dealing with ambiguity averse players, we can reinterpret the maximal overall welfare levels from the price of anarchy literature as maximal DCE/ACE welfare levels, so that the associated bounds on the inefficiency of equilibria can now equivalently be viewed as measuring the benefits arising from the ability to use ambiguous correlation devices and delegate to a mediator for DCE, or just the benefit of ambiguous correlation for ACE.

¹⁸ An alternative literature considers mechanism design problems where ambiguity of agents' beliefs regarding the state of the world and other agents' types is assumed exogenously. See, for example, Bose et al. (2006), Bose and Daripa (2009), Bodoh-Creed (2012), De Castro et al. (2017) or De Castro and Yannelis (2018).

¹⁹ An introduction to this approach of quantifying the inefficiency of equilibria is presented in Roughgarden and Tardos (2007).

²⁰ See Ashlagi et al. (2008) for appropriate references, as well as an analysis of the value of correlation in CE.

Appendix A. Proofs

Proof of Lemma 1. For each player i , define a mapping $\tau_i : S_i \rightarrow \Delta A$ by $\tau_i(s_i) := q_{s_i}$, and let $T_i := \tau_i(S_i) \subset \Delta A$ and $T := \times_i T_i$. Defining a probability distribution $\tilde{p} \in \Delta(A \times T_1 \times \dots \times T_n)$ by

$$\tilde{p} := q \circ (\text{id}, \tau_1^{-1}, \dots, \tau_n^{-1}),$$

where id denotes the identity function with domain A , then yields $t_i = \tilde{p}_{t_i}$ for every $t_i \in T_i$. Clearly, if σ is the prior/marginal over A associated with q , it must also be the prior for \tilde{p} , and we can write $\tilde{p}(a, t) = \sigma(a)\tilde{\pi}(t|a)$. If we then set

$$\pi_i(t_i|a) := \tilde{\pi}(t_i|a) = \sum_{t_{-i} \in T_{-i}} \tilde{\pi}(t_i, t_{-i}|a),$$

and define a probability distribution $p \in \Delta(A \times T)$ by

$$p(a, t) := \sigma(a) \prod_{i \in N} \pi_i(t_i|a),$$

it follows that $\tilde{p}_{t_i} = p_{t_i}$, and therefore $p_{t_i} = t_i$. Furthermore, the distributions over signals corresponding to p are independent conditional on a , and thus p satisfies all the properties stated in the lemma. \square

Proof of Lemma 2. For every $a \in A$,

$$\begin{aligned} \sigma(a) &= \sigma(a) \sum_{t_i \in \text{supp}(\pi_i)} \pi_i(t_i|a) = \sum_{t_i \in \text{supp}(\pi_i)} \sigma(a)\pi_i(t_i|a) \\ &= \sum_{t_i \in \text{supp}(\pi_i)} \left[\sum_{a' \in A} \sigma(a')\pi_i(t_i|a') \right] \frac{\sigma(a)\pi_i(t_i|a)}{\sum_{a' \in A} \sigma(a')\pi_i(t_i|a')} \\ &= \sum_{t_i \in \text{supp}(\pi_i)} \left[\sum_{a' \in A} \sigma(a')\pi_i(t_i|a') \right] t_i(a). \end{aligned}$$

Hence, if we define $\gamma_i(t_i) := \sum_{a' \in A} \sigma(a')\pi_i(t_i|a')$, then

$$\sum_{t_i \in \text{supp}(\pi_i)} \gamma_i(t_i) = \sum_{a' \in A} \sigma(a') \sum_{t_i \in \text{supp}(\pi_i)} \pi_i(t_i|a') = 1,$$

and $\sigma = \sum_{t_i \in \text{supp}(\pi_i)} \gamma_i(t_i)t_i$. \square

Proof of Lemma 3. If $\sigma(a) = 0$ for some $a \in A$, the assumption that all $\gamma_i^k > 0$ implies that $t_i^k(a) = 0$ for all k , and hence, we can choose any arbitrary conditional distribution $\pi_i(\cdot|a)$. For $a \in A$ with $\sigma(a) > 0$, define

$$\pi_i(t_i^k|a) := \frac{\gamma_i^k t_i^k(a)}{\sigma(a)}.$$

Clearly, $\pi_i(t_i^k|a) \geq 0$, and furthermore, since $\sigma(a) = \sum_{k=1}^m \gamma_i^k t_i^k(a)$ for all a , $\pi_i(t_i^k|a) \leq 1$, and

$$\sum_{k=1}^m \pi_i(t_i^k|a) = \frac{1}{\sigma(a)} \sum_{k=1}^m \gamma_i^k t_i^k(a) = 1.$$

To prove the last property, note that, by definition,

$$t_i^k(a) = \frac{\sigma(a)\pi_i(t_i^k|a)}{\gamma_i^k},$$

and

$$\sum_{a' \in A} \sigma(a')\pi_i(t_i^k|a') = \sum_{a' \in A} \sigma(a') \frac{\gamma_i^k t_i^k(a')}{\sigma(a')} = \gamma_i^k. \quad \square$$

Proof of Proposition 6. Since σ^ε is a convex combination of σ^* and σ , we can apply Lemma 3 to define

$$\begin{aligned} \bar{\pi}(s'|a) &= \tilde{\pi}(s'|a) = \frac{(1-\varepsilon)\sigma^*(a)}{\sigma^\varepsilon(a)}, \\ \bar{\pi}(s''|a) &= \tilde{\pi}(s''|a) = \frac{\varepsilon\sigma(a)}{\sigma^\varepsilon(a)}, \end{aligned}$$

and construct distributions $\bar{q}, \tilde{q} \in \Delta(A \times S)$ by setting

$$\begin{aligned} \bar{q}(a, s') &= \sigma^\varepsilon(a)\bar{\pi}(s'|a) = \sigma^\varepsilon(a)\frac{(1-\varepsilon)\sigma^*(a)}{\sigma^\varepsilon(a)} = (1-\varepsilon)\sigma^*(a), \\ \bar{q}(a, s'') &= \sigma^\varepsilon(a)\bar{\pi}(s''|a) = \sigma^\varepsilon(a)\frac{\varepsilon\sigma(a)}{\sigma^\varepsilon(a)} = \varepsilon\sigma(a), \\ \tilde{q}(a, s') &= \sigma^\varepsilon(a)\tilde{\pi}(s'|a) = \sigma^\varepsilon(a)\frac{\varepsilon\sigma(a)}{\sigma^\varepsilon(a)} = \varepsilon\sigma(a), \\ \tilde{q}(a, s'') &= \sigma^\varepsilon(a)\tilde{\pi}(s''|a) = \sigma^\varepsilon(a)\frac{(1-\varepsilon)\sigma^*(a)}{\sigma^\varepsilon(a)} = (1-\varepsilon)\sigma^*(a). \end{aligned}$$

By construction, we have $\widehat{s}' = \widehat{s}'' = \{\sigma^*, \sigma\}$, and since $\underline{t}_i[\widehat{s}'] = \underline{t}_i[\widehat{s}''] = \sigma$, it follows from Lemma 5 that both signal realizations yield compliant belief sets for all n players. Hence, the defined (S, Q) constitutes a DCE. \square

Proof that strict preferences for compliance are not possible in the prisoner’s dilemma. Consider the prisoner’s dilemma depicted in Fig. 1, and let \widehat{s}_1 denote a compliant belief set for player 1.²¹ For any $t_1 = (t_{11}, t_{12}, t_{13}, t_{14}) \in \widehat{s}_1$, assume that the coordinates of t_1 show the respective probabilities of the four action profiles clockwise, starting with (c, c) . For any mixed strategy $\rho_1 \in \Delta\{c, d\}$ of player 1, the minimum on the RHS of the compliance condition (3) is attained at an element of \widehat{s}_1 that maximizes the marginal probability of player 2 choosing d , i.e., letting $\tilde{t}_1 \in \widehat{s}_1$ denote such a minimizer, we must have

$$\tilde{t}_{12} + \tilde{t}_{13} \geq t_{12} + t_{13}, \text{ for all } t_1 \in \widehat{s}_1.$$

Note also that the maximum on the RHS of condition (3) will always be attained by the mixed strategy ρ_1 that assigns probability 1 to d .

Now let $\underline{t}_1 := \underline{t}_1[\widehat{s}_1]$ (with some abuse of notation). Then there are two possibilities, either $\underline{t}_1 = \tilde{t}_1$, or $\underline{t}_1 \neq \tilde{t}_1$. If $\underline{t}_1 = \tilde{t}_1$, the assumption that \widehat{s}_1 is compliant yields

$$3\tilde{t}_{11} + 0\tilde{t}_{12} + 1\tilde{t}_{13} + 4\tilde{t}_{14} \geq 4(\tilde{t}_{11} + \tilde{t}_{14}) + 1(\tilde{t}_{12} + \tilde{t}_{13}),$$

which implies that $\tilde{t}_{11} = \tilde{t}_{12} = 0$, and therefore that the above inequality holds with equality. If $\underline{t}_1 \neq \tilde{t}_1$, then

$$3\tilde{t}_{11} + 0\tilde{t}_{12} + 1\tilde{t}_{13} + 4\tilde{t}_{14} \geq 3\underline{t}_{11} + 0\underline{t}_{12} + 1\underline{t}_{13} + 4\underline{t}_{14} \geq 4(\underline{t}_{11} + \underline{t}_{14}) + 1(\underline{t}_{12} + \underline{t}_{13}),$$

where the first inequality follows from the definition of \underline{t}_1 , and the second from compliance. As before, this implies that $\underline{t}_{11} = \underline{t}_{12} = 0$, and thus that both of these inequalities must hold with equality. \square

Proof of Proposition 7. Consider any $\sigma \in \bigcap_{i \in N} \widetilde{R}_i$. Then for every $i \in N$, σ can be expressed as a convex combination of elements of R_i , such that each element has strictly positive weight. Following Lemma 3, there exists a simple action–signal distribution with prior σ and posteriors R_i . Thus, we can choose $|R_i|$ signals names s_i for every i , and replicate the simple action–signal distribution $\max_{i \in N} |R_i|$ times, so that the signal names are cycled through to get $\widehat{s}_i = R_i$ for every s_i . \square

Proof of Proposition 8. Assume condition (a) holds. Then since $\overline{R}_i[\sigma]$ is a supporting set of σ , and $\overline{R}_i[\sigma^*]$ a supporting set of σ^* , every $(1-\varepsilon)\sigma^* + \varepsilon\sigma$ for $\varepsilon \in (0, 1)$ belongs to the relative interior of the convex hull of $\overline{R}_i[\sigma] \cup \overline{R}_i[\sigma^*]$. As this holds for every player $i \in N$, Proposition 7 implies that there exist implementable action distributions that are arbitrarily close to σ^* .

To show that condition (b) also yields the stated result, we show that (b) implies (a). Since σ is assumed to be an unambiguous constrained DCE, there exist supporting sets $\overline{R}_i[\sigma]$ of σ for every player $i \in N$. Given that σ^* Pareto dominates σ , Lemma 5 implies that each $\overline{R}_i[\sigma] \cup \{\sigma^*\}$ is a compliant belief set for player i , and thus (a) holds. \square

Proof of Proposition 9. A finite $\overline{R}_i[\sigma] \subset \Delta A$ that satisfies properties (i) and (ii) constitutes a supporting set of σ relative to a constraint set T_i^c satisfying the requirements for an ACE according to Definition 6. As in the proof of Proposition 7, we can

²¹ Using symmetry, an analogous argument applies for player 2.

then construct an ambiguous correlation device such that for each $i \in N$, every \bar{R}_{a_i} for $a_i \in A_i^\sigma$ is the belief set associated with one specific signal realization. The proof then follows directly from the assumption that each \bar{R}_{a_i} is compliant, based on the definition of an ACE as a constrained DCE.

For the converse, consider the belief sets associated with all signal realizations that recommend a fixed action $a_i \in A_i$ as part of the given ACE, and note that the union of any resulting compliant belief sets is also compliant, so that all such signal realizations can be combined into one yielding a belief set \bar{R}_{a_i} . □

Appendix B. Examples of robustness to alternative preference specifications

Example 6 (Robustness in the prisoner’s dilemma). Consider again the prisoner’s dilemma from Example 1. It is easy to see that if $\alpha_i < 1$, then the DCE we initially constructed in the example is not compliant for α -maxmin preferences, which also illustrates that Lemma 5 cannot hold. To construct a class of DCE with α -maxmin preferences that approximately implements all action distributions that Pareto dominate the CCE (d, d) , let $\delta_{a_1 a_2}$ denote the degenerate action distribution that assigns probability one to (a_1, a_2) , where $a_1, a_2 \in \{c, d\}$, and set

$$\tau_1 := \frac{1}{3}\delta_{cc} + \frac{2}{3}\delta_{cd}, \text{ and } \tau_2 := \frac{1}{3}\delta_{cc} + \frac{2}{3}\delta_{dc},$$

so that τ_1 and τ_2 yield the expected payoff profiles $(1, \frac{11}{3})$ and $(\frac{11}{3}, 1)$, respectively. If we then define

$$R_1 := \{\delta_{dc}, \delta_{cc}, \tau_1, \delta_{dd}\}, \text{ and } R_2 := \{\delta_{cd}, \delta_{cc}, \tau_2, \delta_{dd}\},$$

then each R_i is compliant for player i for every $\alpha_i \in [0, 1]$, with associated compliance condition

$$\alpha_i 1 + (1 - \alpha_i)4 \geq \alpha_i 1 + (1 - \alpha_i)4.$$

The approach to constructing DCE characterized in Proposition 7, with all of a player’s signal realizations as synonyms, then shows that every $\sigma \in \bar{R}_1 \cap \bar{R}_2$, i.e., every σ that lies in the intersection of the relative interiors of the convex hulls of R_1 and R_2 , can be implemented by a DCE, which implies that every action distribution that Pareto dominates the CCE (d, d) is approximately implementable with α -maxmin preferences, irrespective of the values of α_i .

If we instead assume that players have smooth preference representations, then the compliance condition for player 1 associated with the posterior belief set R_1 becomes

$$\begin{aligned} &\mu_{s_1}(\delta_{dc})\phi_1(4) + \mu_{s_1}(\delta_{cc})\phi_1(3) + \mu_{s_1}(\tau_1)\phi_1(1) + \mu_{s_1}(\delta_{dd})\phi_1(1) \\ &\geq \mu_{s_1}(\delta_{dc})\phi_1(4) + \mu_{s_1}(\delta_{cc})\phi_1(4) + \mu_{s_1}(\tau_1)\phi_1(2) + \mu_{s_1}(\delta_{dd})\phi_1(1), \end{aligned}$$

which is *not* satisfied for a strictly increasing ϕ_1 if we assume that μ_{s_1} has full support on R_1 , so the DCE we derived above for α -maxmin preferences do not yield equilibria with smooth preferences. This observation is quite general for the prisoner’s dilemma—if we make a full support assumption for the subjective beliefs μ_{s_i} , no action distribution that Pareto dominates the CCE (d, d) can be implemented with smooth preferences.

Concluding this example, note also that the CCE (d, d) is also the unique ACE with α -maxmin or smooth preferences, following the same reasoning as for maxmin preferences. ◁

Example 7 (Robustness in the drivers game). Considering the drivers game from Example 2, it is straightforward to check that the DCE constructed in the example is not compliant with α -maxmin preferences if $\alpha_i < 1$, or with smooth preferences if any associated subjective beliefs have full support. We start by showing that all action distributions that Pareto dominate the equal mixture of (e, d) and (d, e) , i.e., the distribution that results in the worst symmetric CCE payoff in Fig. 4, can be approximately implemented with α -maxmin or smooth preferences. To construct the associated belief sets, let $\delta_{a_1 a_2}$ again denote the degenerate action distribution that assigns probability one to (a_1, a_2) , where $a_1, a_2 \in \{d, e\}$, and set

$$\kappa := \frac{1}{2}\delta_{ed} + \frac{1}{2}\delta_{de}, \tau_1 := \frac{5}{8}\delta_{dd} + \frac{3}{8}\delta_{de}, \text{ and } \tau_2 := \frac{5}{8}\delta_{dd} + \frac{3}{8}\delta_{ed},$$

so that the expected payoff profiles associated with κ , τ_1 and τ_2 are $(4.5, 4.5)$, $(4.5, \frac{51}{8})$ and $(\frac{51}{8}, 4.5)$, respectively. If we define

$$R_1 := \{\delta_{ed}, \delta_{dd}, \tau_1, \kappa\}, \text{ and } R_2 := \{\delta_{de}, \delta_{dd}, \tau_2, \kappa\},$$

the compliance condition for the posterior set R_1 based on α -maxmin preferences becomes

$$\alpha_1 4.5 + (1 - \alpha_1)7 \geq \alpha_1 [p3.5 + (1 - p)4] + (1 - \alpha_1)[p7 + (1 - p)6],$$

where p denotes the probability assigned to e in any mixed deviation of player 1, so that compliance is *strictly* preferred as long as $\alpha_1 > 0$. Analogously, the compliance condition for R_1 with smooth preferences is

$$\begin{aligned} & \mu_{s_1}(\delta_{ed})\phi_1(7) + \mu_{s_1}(\delta_{dd})\phi_1(6) + \mu_{s_1}(\tau_1)\phi_1(4.5) + \mu_{s_1}(\kappa)\phi_1(4.5) \\ & \geq [\mu_{s_1}(\delta_{ed}) + \mu_{s_1}(\delta_{dd})]\phi_1(p7 + (1 - p)6) + \mu_{s_1}(\tau_1)\phi_1\left(p\frac{35}{8} + (1 - p)4.5\right) + \mu_{s_1}(\kappa)\phi_1(p3.5 + (1 - p)4), \end{aligned}$$

which holds for a large class of parametric specifications for μ_{s_1} and ϕ_1 . Since the compliance conditions for player 2 with respect to R_2 are symmetric, Proposition 7 implies that all $\sigma \in R_1 \cap \tilde{R}_2$ are implementable with both α -maxmin and smooth preferences.²²

Now consider implementability via ACE with α -maxmin preferences. Using the setting from Example 5, for player 1, any ACE implementing a distribution that Pareto dominates a CCE must be constructed based on an unambiguous belief set consisting of $\{(0, 0, 1)\}$ that recommends action e , and an ambiguous belief set that recommends action d . While the belief set $\left\{(1, 0, 0), \left(\frac{2}{3}, \frac{1}{3}, 0\right)\right\}$ for player 1 used in the first ACE constructed in Example 5 is not compliant if $\alpha_1 < 1$, the belief set $\{(1, 0, 0), (0, 1, 0)\}$ from the second ACE is compliant as long as

$$\alpha_1 2 + (1 - \alpha_1)6 \geq \alpha_1 0 + (1 - \alpha_1)7 \Leftrightarrow \alpha_1 \geq \frac{1}{3}.$$

Analogously, for player 2, $\{(0, 1, 0)\}$ is unambiguously compliant, and $\{(1, 0, 0), (0, 0, 1)\}$ is compliant as long as $\alpha_2 \geq \frac{1}{3}$. Hence, if both $\alpha_1, \alpha_2 \geq \frac{1}{3}$, all action distributions that Pareto dominate a CCE distribution are still approximately implementable with α -maxmin preferences, and furthermore, since the associated belief sets each contain two elements, the same conclusion holds qualitatively with smooth preferences. If $\alpha_1 < \frac{1}{3}$, the most beneficial (in terms of implementability) belief set recommending action d for player 1 has the form $\{(q, 1 - q, 0), (0, 1, 0)\}$, where q is the largest probability value that satisfies the compliance condition

$$\alpha_1 2 + (1 - \alpha_1)(4q + 2) \geq (1 - \alpha_1)7q \Leftrightarrow q \leq \frac{2}{3(1 - \alpha_1)} \xrightarrow{\alpha_1 \rightarrow 0} \frac{2}{3}.$$

Setting $q = \frac{2}{3(1 - \alpha_1)}$, and using an analogous argument for player 2 to construct a belief set recommending action d given by $\left\{\left(\frac{2}{3(1 - \alpha_2)}, 0, 1 - \frac{2}{3(1 - \alpha_2)}\right), (0, 0, 1)\right\}$, yields the largest corresponding set of ACE-implementable action distributions for $\alpha_1, \alpha_2 < \frac{1}{3}$, which can be seen to converge to the set of CCE as both $\alpha_1, \alpha_2 \rightarrow 0$. ◁

References

Ashlagi, I., Monderer, D., Tennenholtz, M., 2008. On the value of correlation. *J. Artif. Intell. Res.* 33, 575–613.
 Aumann, R.J., 1974. Subjectivity and correlation in randomized strategies. *J. Math. Econ.* 1, 67–96.
 Aumann, R.J., 1987. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55, 1–18.
 Aumann, R.J., Maschler, M., 1995. *Repeated Games with Incomplete Information*. MIT Press.
 Awaya, Y., Krishna, V., 2019. Communication and cooperation in repeated games. *Theor. Econ.* 14, 513–553.
 Ayouni, M., Koessler, F., 2017. Hard evidence and ambiguity aversion. *Theory Decis.* 82, 327–339.
 Bach, C.W., Perea, A., 2020. Two definitions of correlated equilibrium. *J. Math. Econ.* 90, 12–24.
 Beauchêne, D., Li, J., Li, M., 2019. Ambiguous persuasion. *J. Econ. Theory* 179, 312–365.
 Bergemann, D., Morris, S., 2019. Information design: a unified perspective. *J. Econ. Lit.* 57, 44–95.
 Bodoh-Creed, A.L., 2012. Ambiguous beliefs and mechanism design. *Games Econ. Behav.* 75, 518–537.
 Bose, S., Daripa, A., 2009. A dynamic mechanism and surplus extraction under ambiguity. *J. Econ. Theory* 144, 2084–2114.
 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, 1853–1872.
 Cheng, L., Varshney, K.R., Liu, H., 2021. Socially responsible AI algorithms: issues, purposes, and challenges. *J. Artif. Intell. Res.* 71, 1137–1181.
 De Castro, L., Liu, Z., Yannelis, N.C., 2017. Implementation under ambiguity. *Games Econ. Behav.* 101, 20–33.
 De Castro, L., Yannelis, N.C., 2018. Uncertainty, efficiency and incentive compatibility: ambiguity solves the conflict between efficiency and incentive compatibility. *J. Econ. Theory* 177, 678–707.
 Di Tillio, A., Kos, N., Messner, M., 2017. The design of ambiguous mechanisms. *Rev. Econ. Stud.* 84, 237–276.
 Forges, F., 1993. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory Decis.* 35, 277–310.
 Gajdos, T., Hayashi, T., Tallon, J.M., Vergnaud, J.C., 2008. Attitude toward imprecise information. *J. Econ. Theory* 140, 27–65.
 Gérard-Varet, L.A., Moulin, H., 1978. Correlation and duopoly. *J. Econ. Theory* 19, 123–149.
 Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. *J. Math. Econ.* 18, 141–153.
 Gul, F., Pesendorfer, W., 2015. Hurwicz expected utility and subjective sources. *J. Econ. Theory* 159, Part A, 465–488.
 Guo, H., 2019. Mechanism design with ambiguous transfers: an analysis in finite dimensional naive type spaces. *J. Econ. Theory* 183, 76–105.
 Hart, S., Mas-Colell, A., 2015. Markets, correlation, and regret-matching. *Games Econ. Behav.* 93, 42–58.
 Jaffray, J.Y., 1989. Linear utility theory for belief functions. *Oper. Res. Lett.* 8, 107–112.
 Jaffray, J.Y., 1992. Bayesian updating and belief functions. *IEEE Trans. Syst. Man Cybern.* 22, 1144–1152.
 Kamenica, E., Gentzkow, M., 2011. Bayesian persuasion. *Am. Econ. Rev.* 101, 2590–2615.

²² To approximately implement all action distributions that dominate the CCE from Example 2, we can instead redefine $R_1 = R_2 := \left\{\delta_{ed}, \delta_{dd}, \delta_{de}, \frac{2}{9}\delta_{ed} + \frac{2}{9}\delta_{de} + \frac{5}{9}\delta_{ee}\right\}$, which still yields *strict* compliance with α -maxmin preferences for any $\alpha_1, \alpha_2 \in (0, 1)$, but the associated compliance conditions with smooth preferences require more stringent assumptions on the respective parameterization than the previously constructed DCE.

- Kellner, C., Le Quement, M.T., 2017. Modes of ambiguous communication. *Games Econ. Behav.* 104, 271–292.
- Kellner, C., Le Quement, M.T., 2018. Endogenous ambiguity in cheap talk. *J. Econ. Theory* 173, 1–17.
- Klibanoff, P., Marinacci, M., Mukerji, S., 2005. A smooth model of decision making under ambiguity. *Econometrica* 73, 1849–1892.
- Koutsoupias, E., Papadimitriou, C., 1999. Worst-case equilibria. In: *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pp. 404–413.
- Moulin, H., Ray, I., Gupta, S.S., 2014. Improving Nash by coarse correlation. *J. Econ. Theory* 150, 852–865.
- Moulin, H., Vial, J.P., 1978. Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *Int. J. Game Theory* 7, 201–221.
- Myerson, R.B., 1982. Optimal coordination mechanisms in generalized principal-agent problems. *J. Math. Econ.* 10, 67–81.
- Pires, C.P., 2002. A rule for updating ambiguous beliefs. *Theory Decis.* 53, 137–152.
- Ray, I., Gupta, S.S., 2013. Coarse correlated equilibria in linear duopoly games. *Int. J. Game Theory* 42, 541–562.
- Riedel, F., Sass, L., 2014. Ellsberg games. *Theory Decis.* 76, 469–509.
- Rosenthal, R.W., 1973. A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* 2, 65–67.
- Roughgarden, T., 2007. Routing games. In: Nisan, N., Roughgarden, T., Tardos, É., Vazirani, V.V. (Eds.), *Algorithmic Game Theory*. Cambridge University Press, pp. 461–486. chapter 18.
- Roughgarden, T., 2015. Intrinsic robustness of the price of anarchy. *J. ACM* 62, 32.
- Roughgarden, T., Tardos, É., 2007. Introduction to the inefficiency of equilibria. In: Nisan, N., Roughgarden, T., Tardos, É., Vazirani, V.V. (Eds.), *Algorithmic Game Theory*. Cambridge University Press, pp. 443–459. chapter 17.
- Shoham, Y., Leyton-Brown, K., 2009. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press.
- Siniscalchi, M., 2011. Dynamic choice under ambiguity. *Theor. Econ.* 6, 379–421.
- Strotz, R.H., 1955–1956. Myopia and inconsistency in dynamic utility maximization. *Rev. Econ. Stud.* 23, 165–180.
- Young, H.P., 2004. *Strategic Learning and Its Limits*. Oxford University Press.