



Interactive Information Design

Frédéric Koessler, Marie Laclau, Tristan Tomala

► To cite this version:

Frédéric Koessler, Marie Laclau, Tristan Tomala. Interactive Information Design. Mathematics of Operations Research, 2021, 41 (1), pp.153–175. <10.1287/moor.2020.1119>. <halshs-01791918v3>

HAL Id: halshs-01791918

<https://shs.hal.science/halshs-01791918v3>

Submitted on 6 Sep 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.




HAL Authorization

Interactive Information Design

Frédéric Koessler,^a Marie Laclau,^b Tristan Tomala^b

^aParis School of Economics–CNRS, 75014 Paris, France; ^bHEC Paris and GREGHEC–CNRS, 78351 Jouy-en-Josas, France

Contact: frederic.koessler@psemail.eu (FK); laclau@hec.fr (ML); tomala@hec.fr,  <https://orcid.org/0000-0003-3718-7633> (TT)

Received: February 25, 2019

Revised: March 2, 2020; December 17, 2020

Accepted: November 20, 2020

Published Online in Articles in Advance:
June 7, 2021

MSC2010 Subject Classification: Primary:
91A10, 91A27

OR/MS Subject Classification: Primary:
Games/group decisions: noncooperative

<https://doi.org/10.1287/moor.2020.1119>

Copyright: © 2021 INFORMS

Abstract. We study the interaction between multiple information designers who try to influence the behavior of a set of agents. When each designer can choose information policies from a compact set of statistical experiments with countable support, such games always admit subgame-perfect equilibria. When designers produce public information, every equilibrium of the simple game in which the set of messages coincides with the set of states is robust in the sense that it is an equilibrium with larger and possibly infinite and uncountable message sets. The converse is true for a class of Markovian equilibria only. When designers produce information for their own corporation of agents, robust pure-strategy equilibria exist and are characterized via an auxiliary normal-form game in which the set of strategies of each designer is the set of outcomes induced by Bayes correlated equilibria in her corporation.

Funding: F. Koessler acknowledges the support of the Agence Nationale de la Recherche (ANR) through the programs Investissements d’Avenir [Grant ANR-17-EURE-001] and StratCom [Grant ANR-19-CE26-0010-01]. M. Laclau gratefully acknowledges the support of the ANR through the programs Investissements d’Avenir [Grant ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047] and CIGNE [Grant ANR-15-CE38-0007-01]. T. Tomala gratefully acknowledges the support of the HEC Foundation and the ANR program Investissements d’Avenir [Grant ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047].

Keywords: Bayes correlated equilibrium • Bayesian persuasion • information design • splitting games • statistical experiments

1. Introduction

Decision makers often receive information from various interested parties who communicate about diverse pieces of information. For instance, consider competing pharmaceutical firms who aim to obtain approval to release new drugs on the market. Each firm aims to persuade the Food and Drug Administration that its product is effective and would prefer products of other firms not to be approved. In other cases, divisions within an organization or university try to persuade the head of the organization to allocate a position to their department, or managers design the distribution of information in organizations to improve and coordinate agents’ efforts. In all of these examples, interested parties design information to influence the behavior of decision makers. The aim of this paper is to provide a general theoretical framework to analyze such situations. We establish the existence of equilibria, characterize mixed or pure equilibria in special settings, and show when it is without loss of generality to restrict designers to the use of simple information policies.

Our modeling setup is as follows. There are n information designers and k agents. There is an unknown n -dimensional payoff-relevant state parameter. The set of possible payoff-relevant states is finite. Each designer i controls agents’ private information about component i of the state by choosing a statistical experiment that draws messages as a function of the i th component and sends these messages to the agents. Agents observe the chosen statistical experiments, the realized messages, and the outcome of a public randomization device. Finally, each agent chooses an action from some finite set. The payoffs for the agents and the designers depend on the realized state and on the action of every agent.

In such games of information design, designers’ expected utilities induced by the agents’ equilibrium behavior are typically discontinuous (and generally not even upper semicontinuous) in the agents’ beliefs and hence in the profile of statistical experiments. However, we show in Theorem 1 that if designers choose information policies from a compact set of statistical experiments with countable support, then the $(n + k)$ -player game between the designers and the agents admits subgame-perfect equilibria. We first show that, for every profile of statistical experiments, the induced Bayesian game played by the agents has a nonempty and compact set of Bayes-Nash equilibrium outcomes and that the equilibrium correspondence is well behaved.¹ We then apply the existence result of Simon and Zame [38] to show that there exists a selection from continuation equilibria induced by

statistical experiments such that, by backward induction, the induced n -player game played between the designers has a Nash equilibrium.² We provide a simple strictly competitive example with binary states and a single agent (Example 4, adapted from Sion and Wolfe [39]), which shows that mixed (ϵ)-equilibria may fail to exist for some selection of the optimal continuation strategy for the agent.

We study two broad subclasses of information design games for which we investigate additional equilibrium properties.

First, we consider information design games in which designers send *public* messages to the agents. Each profile of experiments chosen by designers induces a public information structure, and the set of continuation equilibria depends only on agents' first-order beliefs over each component of the state. For each designer, choosing a statistical experiment is therefore equivalent to choosing a *splitting* of the prior belief into common posterior beliefs. Concavification methods can be used to characterize designers' best responses. We study "simple" information design games in which the number of messages available to designer i is equal to the number of states in dimension i . In Theorem 2, we show that the set of subgame-perfect equilibrium outcomes of the simple game is robust in the sense that it is included in the set of subgame-perfect equilibrium outcomes of all games with more messages than states (with possibly infinite and uncountable message sets). The inclusion can be strict even with a single designer and a single agent. However, for the class of *Markovian* equilibria in which agents' strategies depend only on their posterior beliefs, the simple game is "canonical" in the sense that the set of Markovian subgame-perfect equilibrium outcomes of all games with more messages than states coincides with the set of subgame-perfect equilibrium outcomes of the simple game (Theorem 4). When there is a continuum of messages, we show in Theorem 3 that it is without loss of generality to focus on pure strategies for the designers: we prove that any mixed-strategy equilibrium can be replicated by an equilibrium in pure strategies. Combined with Theorems 1 and 2, this result implies that there is a pure-strategy equilibrium in the game in which designers are able to choose any statistical experiment together with any (possibly infinite) set of messages.

Second, we consider *rectangular corporation games* in which each designer i controls the (public and private) information of a given corporation of agents (each agent belongs to only one corporation). The utility functions of agents in corporation i depend only on the actions taken by agents in corporation i and on the information controlled by designer i . Hence, from the point of view of agents in corporation i , only the messages received from their own designer matter. Designers, however, interact through how they influence agents' behavior in *all* corporations. In such games, we show that robust equilibria exist in pure strategies with finite sets of messages. We provide a simple characterization of such equilibria via an auxiliary n -player normal-form game. In this game, the set of pure strategies of player i is the set of action profiles (or state-dependent action profiles if the designers' preferences are state-dependent) of the agents in corporation i . We show that the set of equilibria of a mixed extension of this auxiliary game with convex constraints on strategies is included in the set of subgame-perfect equilibrium outcomes of the original information design game. If each corporation is composed of a single agent, then, under a regularity condition, the two sets coincide.

Related Literature

The characterization of optimal information structures with a single designer has been studied and applied in many articles in the economics literature. See, among others, Bergemann and Morris [9], Eliaz and Forges [13], Eliaz and Serrano [14], Jehiel [18], Kamenica and Gentzkow [19], Lewis and Sappington [25], Mathevet et al. [29], and Taneva [42]. When there is a single designer and a single agent, our model of information design exactly corresponds to the model of Kamenica and Gentzkow [19]. In such a setting, a subgame-perfect equilibrium is obtained by concavifying the designer's indirect utility function, which represents her expected utility as a function of the agent's belief. This characterization is analogous to that obtained in the literatures on repeated games with incomplete information, splitting games, and acyclic gambling games (Aumann and Maschler [6], Aumann et al. [7], Laraki [21, 22], Laraki and Renault [23], Mertens et al. [30], Oliu-Barton [34], Sorin [40]).

In Kamenica and Gentzkow [19], the characterization of optimal information structures is obtained by assuming that when an agent is indifferent between several actions, she chooses an optimal action that favors the designer. Under this assumption, the indirect utility function of the designer is upper semicontinuous, and, thus, an optimal solution exists. With multiple designers who do not have the same preferences, there is no analogue to the designer-preferred tie-breaking rule, so the existence of exact best responses is not guaranteed. If there is a single agent and the strategy of the agent depends on her belief only (and not on the precise statistical experiments used by the designers), then an ϵ -best response of a designer can be obtained by taking the concave closure of her payoff function (with the strategies of the other designers fixed). However, this does not guarantee the existence of an ϵ -equilibrium between the designers. Another important difference between the one-designer and the multiple-designer cases concerns the required richness of the message space (or the ability of the

designers to use mixed strategies, i.e., randomizations over statistical experiments). In the single-designer case, the required number of messages is bounded above by the number of actions (by the revelation principle) or the number of states (by a suitable version of Carathéodory's theorem applied to the concave closure of the payoff function). As we illustrate, neither of these results applies with multiple designers, even if there is a single agent.

A few recent articles consider the case of multiple information designers in specific environments. Gentzkow and Kamenica [15] and Li and Norman [26] consider a situation in which the designers can produce the same information for a single agent. This corresponds to our model when the state spaces of all the designers are the same and the states are perfectly correlated. In this case, there is a trivial equilibrium in which designers fully disclose the state. Gentzkow and Kamenica [15] provide conditions under which all pure-strategy equilibria are more informative than the collusive outcome. Albrecht [1], Au and Kawai [4, 5], and Boleslavsky and Cotton [10] show equilibrium existence and provide explicit equilibrium characterizations in an example similar to Example 3 by fixing a symmetric tie-breaking rule for the agent. Levy et al. [24] study competition between multiple information designers (media owners) with a receiver (reader) who suffers from correlation neglect and show that competition sometimes hurts the receiver. Koessler et al. [20] study strictly competitive and multistage information design games between two designers and an agent.

Our analysis is also related to the literature on competing mechanism designers (see, e.g., Martimort and Stole [28], Myerson [33], Peters [35], and Szentes [41]). In this literature, multiple principals control allocations by offering mechanisms to the agents. Then, after the mechanisms have been proposed, agents reveal their private information (including information about the proposed mechanisms) to the principals. By contrast, in our setting, principals (designers) control agents' beliefs by offering statistical experiments to the agents, and agents observe the messages realized by those statistical experiments and then make a decision. Despite these differences, some difficulties identified in that literature are relevant when principals compete in the design of information. For example, the issue of equilibrium existence illustrated by Myerson [33, section 4] with multiple mechanism designers is related to the existence problem in our framework for some selection of agents' continuation equilibrium strategies. The reason is that, in general, there does not exist any continuous selection of equilibrium strategies for the agents. Equilibrium existence is resolved by using the existence theorem of Simon and Zame [38] and by considering endogenous selections from the whole set of continuation equilibria. A similar approach has been used by Carmona and Fajardo [12] to show the existence of equilibrium in common agency games. The identification of conditions under which we can restrict attention to simple games (in which the message spaces coincide with the state spaces in public information design games, or the message spaces coincide with the action spaces in rectangular corporation games) is related to robustness properties of direct mechanisms in competing mechanisms games (see, e.g., Attar et al. [3] and Han [16]).

The next section presents general information design games and an equilibrium existence result. In Section 3, we consider information design games in which designers disclose information publicly. Section 4 considers rectangular corporation problems and contains an application (Example 5) with two designers and four agents, where we compare public and private information design. In the appendix, we provide useful properties of equilibria in Bayesian games with countable states, compact actions, and continuous and bounded payoffs.

2. General Model and Equilibrium Existence

The Environment

There is a set $N = \{1, \dots, n\}$ of information *designers* and a set $K = \{1, \dots, k\}$ of *agents*. The set of *states* $\Theta = \prod_{i \in N} \Theta_i$ is endowed with a common prior probability distribution $p^0 \in \Delta(\Theta)$.³ All players are uninformed about the state. Designer $i \in N$ discloses information about a parameter $\theta_i \in \Theta_i$, and agent $j \in K$ chooses an action $a_j \in A_j$. All sets Θ_i and A_j are nonempty and finite.⁴

The set of action profiles is $A = \prod_{j \in K} A_j$. The payoff of each player depends on the state and on the action profile. The payoff of designer i (respectively, agent j) is denoted $u_i(a; \theta)$ (respectively, $v_j(a; \theta)$).

Information Structures and Statistical Experiments

At the ex ante stage, before states are drawn, each designer i chooses how to disclose information about dimension i of the state by choosing a mapping that specifies the distributions of messages to the agents, conditional on the states in Θ_i . This mapping will be referred to as a *statistical experiment*.

Definition 1. A statistical experiment for designer i is a mapping $x_i : \Theta_i \rightarrow \Delta(\mathbb{N}^K)$ from Θ_i to the Borel probability distributions over the set of profiles of messages $m_i = (m_i^j)_{j \in K} \in \mathbb{N}^K$, where \mathbb{N} is the set of natural numbers.

A statistical experiment induces a discrete⁵ information structure for the agents: it selects a profile of messages $m_i = (m_i^j)_{j \in K}$ from the probability distribution $x_i(m_i | \theta_i)$ and privately delivers message m_i^j to agent j .

This model includes as a particular case an environment in which designers disclose information about a common and arbitrary finite state variable $\omega \in \Omega$. Let $\Theta_i = \Omega$ for every i , and assume that the states are perfectly correlated between dimensions; that is, $p^0(\theta_1, \dots, \theta_n) > 0$ only if $\theta_i = \omega$ for every i . More generally, it covers the case in which the information that each designer i can disclose is represented by a partition Ω_i of Ω and designer i chooses a statistical experiment $x_i : \Omega \rightarrow \Delta(\mathbb{N}^K)$, which is measurable with respect to Ω_i .

In the following, we consider games where each designer i can choose any experiment in an exogenously fixed feasible set X_i . Introducing feasible sets of experiments allows us to model situations where designers face technological constraints on information disclosure policies. For instance, designer i may only be able to use a finite number of messages for each agent. In that case, the cardinality of the support of experiments in X_i is bounded by the size of the message space.

As another example, designer i may be constrained to use deterministic experiments $x_i : \Theta_i \rightarrow \mathbb{N}^K$, in which case at most $|\Theta_i|$ messages from designer i are possible for each agent. As a third example, X_i could be the set of *public* experiments where designer i sends the same message to all agents (i.e., perfectly correlated messages).

The set of all statistical experiments for designer i is $\Delta(\mathbb{N}^K)^{\Theta_i}$. We assume throughout that for each designer i , the set of feasible experiments $X_i \subseteq \Delta(\mathbb{N}^K)^{\Theta_i}$ is compact for the 1-norm on the space \mathcal{Z} of summable families $z = (z(m_i | \theta_i))_{\theta_i, m_i}$ in $(\mathbb{R})^{(\mathbb{N}^K)^{\Theta_i}}$, given by

$$\|z\|_1 = \sum_{\theta_i \in \Theta_i} \sum_{m_i \in \mathbb{N}^K} |z(m_i | \theta_i)|.$$

Compactness is a minimal technical assumption to obtain the existence of equilibria of the game between designers. Note that this norm is the generalization of the distance in total variation over probability distributions. Additionally, X_i is compact for this norm if and only if it is compact for the weak topology over probability distributions (\mathcal{Z} is isomorphic to the space ℓ_1 of summable sequences; see Aliprantis and Border [2, theorem 16.24, p. 537]).

For instance, the set of experiments with bounded support is compact. More generally, if there exists a summable family z such that $x_i(m_i | \theta_i) \leq z(m_i | \theta_i)$, for all x_i in X_i and all θ_i, m_i , then X_i is compact (see the characterization in Aliprantis and Border [2, theorem 16.22, p. 536]).

Information Design Game

Let $X = \prod_i X_i$ be a profile of sets of feasible experiments for all designers. Our main object of interest is the $(n + k)$ -player game G_X , whose timing is as follows:

1. Each designer i chooses a statistical experiment $x_i \in X_i$; these choices are simultaneous.
2. Agents publicly observe (x_1, \dots, x_n) .
3. Nature draws the state $\theta = (\theta_1, \dots, \theta_n)$ according to $p^0 \in \Delta(\Theta)$ and a uniformly distributed public signal $\omega \in [0, 1]$ (called a *public correlation device* hereafter);
4. For each designer i , a profile of messages $m_i = (m_i^1, \dots, m_i^k)$ is drawn with probability $x_i(m_i | \theta_i)$.
5. Each agent j observes the public signal ω and her profile of *private* messages $m^j = (m_i^j)_{i \in N}$, and then chooses an action $a_j \in A_j$.

We allow for mixed strategies in this game: the choice of x_i by designer i at step 1 is possibly random, and the choice of a_j by agent j at step 5 is as well.

Subgame-Perfect Equilibrium

We consider subgame-perfect equilibria of the game G_X , simply referred to as *equilibria* in the rest of the paper. For each profile of experiments $x = (x_i)_{i \in N} \in \prod_{i \in N} X_i$, let $G_X(x)$ be the k -player Bayesian game induced by x . It is the subgame starting from stage 2 in the aforementioned timeline.

Unless stated otherwise, let $M_i = \mathbb{N}^K$ be the set of messages available to designer i , and let $M = \prod_{i \in N} M_i$ be the set of all possible message profiles. Let $\mathcal{E}_X(x) \subseteq \Delta(A)^M$ be the set of Bayes-Nash equilibrium *outcomes* of $G_X(x)$, that is, mappings from message profiles to distributions of action profiles. In Lemma A.1 in the appendix, we show that the set of Bayes-Nash equilibrium outcomes $\mathcal{E}_X(x)$ is nonempty and compact for every $x \in X$. In addition, since $G_X(x)$ is a Bayesian game extended by a public randomization device, $\mathcal{E}_X(x)$ is convex. It coincides with the set of public correlated equilibrium outcomes of $G_X(x)$.

In equilibrium, the agents play a Bayes-Nash equilibrium of $G_X(x)$ for each x . Consider a Borel-measurable mapping $\tau : \prod_{i \in N} X_i \rightarrow \Delta(A)^M$ such that $\tau(x) \in \mathcal{E}_X(x)$ for each x , namely, a Borel-measurable selection from the

equilibrium correspondence $\mathcal{E}_X(\cdot)$. Such a selection induces a simultaneous game between the designers with the payoff function

$$U_i^\tau(x_i, x_{-i}) = \sum_{\theta \in \Theta} \sum_{m \in M} \sum_{a \in A} p^0(\theta) \prod_i x_i(m_i | \theta_i) \tau(x) [a|m] u_i(a; \theta),$$

where $\tau(x)[a|m]$ is the probability of a conditional on m under $\tau(x)$.

It is important to note that the function $x \mapsto U_i^\tau(x)$ is possibly discontinuous, because $x \mapsto \tau(x)$ may be discontinuous. This is true even with a single agent: when the agent is indifferent among several actions at some belief induced by some x , her optimal action may switch from one action to another in the neighborhood of her belief (see Example 2).

Definition 2. An equilibrium of G_X is a Borel-measurable selection τ from the correspondence $\mathcal{E}_X(\cdot)$ and a mixed strategy Nash equilibrium of the game $((U_i^\tau)_i, (X_i)_i)$.

Note that we allow for mixed strategies; that is, each designer i chooses a Borel probability measure over X_i . The main result in this section is that, for every profile of compact sets X , equilibria exist for the information design game G_X .

Theorem 1. For every profile of compact sets X , the game G_X admits an equilibrium.

Proof. For all $x \in \prod_{i \in N} X_i$ and $y \in \Delta(A)^M$, let

$$U_i(x, y) = \sum_{\theta \in \Theta} \sum_{m \in M} \sum_{a \in A} p^0(\theta) \prod_i x_i(m_i | \theta_i) y(a|m) u_i(a; \theta)$$

be the expected utility of designer i given the profile of statistical experiments x and the outcome y of $G_X(x)$. Let $U(x, y) = (U_i(x, y))_{i \in N}$, and define correspondences $U: \prod_{i \in N} X_i \rightarrow \mathbb{R}^N$ by

$$U(x) = \{U(x, y) : y \in \mathcal{E}_X(x)\}.$$

From Lemmas A.1 and A.2 in the appendix, the equilibrium correspondence $\mathcal{E}_X(\cdot)$ has nonempty compact values and is upper hemicontinuous. Since there is a public randomization device, it also has convex values. By continuity of the payoff function $U(x, y)$, $U(\cdot)$ also has compact convex values and is upper hemicontinuous. From the main theorem of Simon and Zame [38], there exists a Borel-measurable selection $U^*(x) \in U(x)$ from U such that the normal-form game $((U_i^*)_i, (X_i)_i)$ admits a Nash equilibrium in mixed strategies. Hence, there exists a Borel-measurable selection τ from the correspondence $\mathcal{E}_X(\cdot)$ such that $(U_i^\tau)_i = (U_i^*)_i$ and $((U_i^\tau)_i, (X_i)_i)$ admits a Nash equilibrium in mixed strategies. \square

Remark 1. Several comments are in order.

1. The correspondence $U(\cdot)$ does not generally admit continuous selections; see Example 2.
2. An (ε) -equilibrium may not exist for some selections $\tau(\cdot) \in \mathcal{E}_X(\cdot)$. This is shown in Example 4 with two designers, one agent, and binary states.
3. In general, a public randomization device is needed to guarantee the existence of equilibria. In the proof of the theorem, this ensures that the set of Bayes-Nash equilibrium outcomes of the game $G_X(x)$ is convex for every profile of experiments x . In the case of one agent, the set of optimal continuation strategies of a single agent is convex; therefore, public correlation is not needed.
4. Suppose that there is a single designer i who can use as many messages as the action profile. That is, let $X_i(M_i)$ be the set of experiments with support included in a finite set $M_i = \prod_{j \in K} M_i^j \subset \mathbb{N}^K$ such that $|M_i^j| \geq |A_j|$ for every $j \in K$; that is, $X_i(M_i) = \Delta(M_i)^{\Theta_i}$. Then, the set $\bigcup_x \mathcal{E}_X(x)$ is the set of *Bayes correlated equilibria* (BCE) (Bergemann and Morris [8]), which is convex and compact. Thus, one equilibrium (but not all) is found by maximizing the designer's payoff over the set of BCE. In such an equilibrium, the public randomization device is not needed, and the designer selects her favorite continuation Bayes-Nash equilibrium along the equilibrium path. Of course, with multiple designers, such a selection does not exist if designers have misaligned preferences.
5. Consider the case where each designer is constrained by the number of available messages. That is, with the notation of the previous point, for each designer i , $X_i = X_i(M_i)$ for some finite set M_i .

Even with one designer, it is not clear how the set of equilibrium outcomes varies with the size of the message spaces. It can be shown that information structures with higher-order beliefs that appear in global games (Carlsson and van Damme [11], Rubinstein [37]) can be used by the designer to induce a unique continuation equilibrium outcome in $\mathcal{E}_X(x)$. Such structures typically use message sets that are larger than the action or the state space.⁶ A fortiori, the set of equilibrium outcomes with multiple designers also depends on the message spaces, even if $|M_i^j| \geq |A_j|$ and $|M_i^j| \geq |\Theta_i|$ for every i and j . We will provide sharper equilibrium characterizations with respect to the message spaces in public information design games (Section 3) and rectangular corporation games (Section 4).

Relatedly, an equilibrium outcome of G_X with $X_i = X_i(M_i)$ for each i may not be an equilibrium outcome of the game where designers can choose *any* experiment, that is, when the size of the message space is a strategic choice.

Again, our equilibrium characterizations in the special classes of public information design games and rectangular corporation games allow us to conclude that equilibrium outcomes in these classes of information design games are robust to such deviations.

3. Public Information Design

In this section, we assume that each designer i sends public messages to all agents, so that every continuation game is one with symmetric information. Precisely, each designer i is restricted to using a *public* statistical experiment, that is, such that the message m_i^j observed by agent j from designer i is the same as the message $m_i^{j'}$ observed by agent j' from designer i , for every $j, j' \in K$.

For each designer i , we fix a finite set of possible public messages $M_i \subset \mathbb{N}$ and let him choose any statistical experiment $x_i : \Theta_i \rightarrow \Delta(M_i)$ from Θ_i to the set of probability distributions over the set of public messages M_i . The set of public experiments $X_i = \Delta(M_i)^{\Theta_i}$ is compact, so the existence result of the previous section applies. We prove the existence of an equilibrium with infinite message spaces in Theorem 2.

To simplify notations, we assume that the prior probability distribution is the product of its marginal distributions: $p^0 = \otimes_{i \in N} p_i^0$, with $p_i^0 \in \Delta(\Theta_i)$. That is, designers produce and disclose independent pieces of information to the agents. Since we make no restriction on the form of the utility functions, this assumption is without loss of generality. Precisely, given any prior $p^0 \in \Delta(\Theta)$ and utility functions $(u_i)_{i \in N}$ and $(v_j)_{j \in K}$, the game is “equivalent” to a game with stochastically independent states by letting $\hat{p}^0(\theta) = \prod_{i \in N} (\sum_{\tilde{\theta}_{-i} \in \Theta_{-i}} p^0(\theta_i, \tilde{\theta}_{-i}))$, $\hat{u}_i(a, \theta) = \frac{p^0(\theta)}{\hat{p}^0(\theta)} u_i(a, \theta)$, and $\hat{v}_j(a, \theta) = \frac{p^0(\theta)}{\hat{p}^0(\theta)} v_j(a, \theta)$ for every θ in the support of p^0 . It is immediate to see that the set $\hat{\mathcal{E}}_X(x)$ of Bayes-Nash equilibrium outcomes of the transformed game (with independent priors) coincides with the original set $\mathcal{E}_X(x)$ for every profile of statistical experiments x and that each designer’s expected payoff $\hat{U}_i^\tau(x)$ of the transformed game is equal to $U_i^\tau(x)$ for every x and τ .

In public information design games, the choice of a statistical experiment $x_i : \Theta_i \rightarrow \Delta(M_i)$ by designer i can be reduced to a choice of a *splitting* of the prior p_i^0 into at most $|M_i|$ common posteriors for the agents. This “belief-based” approach allows us to use concavification techniques in the proofs of Theorems 2 and 4 and to prove equilibrium existence for public information design games with a continuum of messages in Theorem 3. It is also convenient to define sequentially rational continuation equilibria (for beliefs off the equilibrium path) in the proofs of Theorems 2 and 3. Finally, it allows us to define strategies that are Markovian (with respect to common beliefs) in Section 3.3.

From Statistical Experiments to Splittings

Given a public statistical experiment $x_i \in \Delta(M_i)^{\Theta_i}$, a message $m_i \in M_i$ is publicly observed by the agents with total probability

$$\lambda_i(m_i) = \sum_{\theta_i \in \Theta_i} x_i(m_i | \theta_i) p_i^0(\theta_i).$$

If $\lambda_i(m_i) > 0$, then the agents’ posterior beliefs conditional on m_i are derived from Bayes’s rule as follows:

$$\mathbf{p}_i(\theta_i | m_i) = \frac{x_i(m_i | \theta_i) p_i^0(\theta_i)}{\lambda_i(m_i)}, \quad \text{for every } \theta_i \in \Theta_i.$$

Posterior beliefs satisfy $\sum_m \lambda_i(m_i) \mathbf{p}_i(m_i) = p_i^0$, where $\mathbf{p}_i(m_i) = (\mathbf{p}_i(\theta_i | m_i))_{\theta_i \in \Theta_i}$. The set of distributions of posteriors that average on $p_i^0 \in \Delta(\Theta_i)$ will be referred to as the set of *splittings* of p_i^0 :

$$\mathcal{S}_i(p_i^0) = \left\{ \mu_i \in \Delta(\Delta(\Theta_i)) : \int_{p_i \in \Delta(\Theta_i)} p_i d\mu_i(p_i) = p_i^0 \right\},$$

where $\Delta(\Delta(\Theta_i))$ is the set of Borel probability measures over the compact set $\Delta(\Theta_i)$. Throughout, $\Delta(\Delta(\Theta_i))$ and $\mathcal{S}_i(p_i^0)$ are endowed with the weak-* topology (for which they are compact).

We denote by $\mathcal{S}_i^{M_i}(p_i^0)$ the set of splittings with finite support of cardinality at most $|M_i|$. Any such splitting can be represented by a pair $\mu_i = (\lambda_i, \mathbf{p}_i)$ with $\lambda_i \in \Delta(M_i)$, $\mathbf{p}_i : M_i \rightarrow \Delta(\Theta_i)$, and $\sum_m \lambda_i(m_i) \mathbf{p}_i(m_i) = p_i^0$. Every statistical experiment $x_i \in \Delta(M_i)^{\Theta_i}$ induces a splitting in $\mathcal{S}_i^{M_i}(p_i^0)$ satisfying:

$$p_i^0(\theta_i) = \sum_{m_i \in M_i} \lambda_i(m_i) \mathbf{p}_i(\theta_i | m_i).$$

Conversely, from the splitting lemma (Aumann et al. [7]), every splitting $(\lambda_i, \mathbf{p}_i)$ with finite support is generated by the statistical experiment given by $x_i(m_i | \theta_i) = \frac{\lambda_i(m_i) \mathbf{p}_i(\theta_i | m_i)}{p_i^0(\theta_i)}$ for θ_i in the support of p_i^0 .

In public information design games, beliefs are common among agents, and the set of continuation Bayes-Nash equilibrium outcomes given x depends only on the distributions of posteriors (the splittings) and on the realized posteriors induced by x . Hence, the information design game in which each designer i chooses any statistical experiment $x_i \in \Delta(M_i)^{\Theta_i}$ is equivalent to the game in which each designer i chooses any splitting $\mu_i \in \mathcal{S}_i^{M_i}(p_i^0)$ and agents publicly observe $\mu = (\mu_i)_{i \in N}$ and $\mathbf{p}(m) = (\mathbf{p}_i(m_i))_{i \in N}$. We denote this game by PG_M . Similarly, we let PG_∞ be the game where each designer i chooses any splitting $\mu_i \in \mathcal{S}_i(p_i^0)$ with possibly infinite support.

With some abuse of notations, we extend the utility function of each designer $i \in N$ and agent $j \in K$ as follows. For every distribution of action profiles $y \in \Delta(A)$, let $v_j(y; \theta) = \sum_{a \in A} y(a) v_j(a; \theta)$ and $u_i(y; \theta) = \sum_{a \in A} y(a) u_i(a; \theta)$ for $i \in N$. For every $p \in \prod_{i \in N} \Delta(\Theta_i)$, we also denote

$$v_j(y; p) = \sum_{\theta \in \Theta} p_1(\theta_1) \times \dots \times p_n(\theta_n) v_j(y; \theta) \text{ and } u_i(y; p) = \sum_{\theta \in \Theta} p_1(\theta_1) \times \dots \times p_n(\theta_n) u_i(y; \theta).$$

For each $p \in \prod_{i \in N} \Delta(\Theta_i)$, let $Y(p) \subseteq \Delta(A)$ be the set of public correlated equilibrium outcomes (i.e., distributions of action profiles) of the game played between the agents when their common belief is given by p . That is, $Y(p)$ is the convex hull of the set of (mixed) Nash equilibria of the game $((v_j(\cdot; p))_j, (A_j)_j)$. The set $Y(p)$ is nonempty, compact, and convex.

In the public information design game PG_M , agents' strategies induce a continuation outcome denoted by

$$\tau : \prod_{i \in N} \mathcal{S}_i^{M_i}(p_i^0) \times \Delta(\Theta_i) \rightarrow \Delta(A).$$

For every profile of splittings $\mu = (\mu_i)_{i \in N} \in \prod_{i \in N} \mathcal{S}_i^{M_i}(p_i^0)$ and every profile of posteriors $p \in \prod_{i \in N} \Delta(\Theta_i)$, $\tau(\mu, p)$ is the distribution of actions played by the agents in the continuation game following μ and p .

In an equilibrium of PG_M , it must be that $\tau(\mu, p) \in Y(p)$ for every $\mu \in \prod_{i \in N} \mathcal{S}_i^{M_i}(p_i^0)$ and every p in the support of μ . Any such Borel-measurable selection τ induces a game between the designers with payoffs:

$$U_i^\tau(\mu) = \int_p u_i(\tau(\mu, p); p) d\mu(p), \quad i = 1, \dots, n.$$

An equilibrium of PG_M is given by a continuation equilibrium outcome τ , with $\tau(\mu, p) \in Y(p)$ for every $\mu \in \prod_{i \in N} \mathcal{S}_i^{M_i}(p_i^0)$ and p in the support of μ , and by a Nash equilibrium of the game $((U_i^\tau)_i, (\mathcal{S}_i^{M_i}(p_i^0))_i)$. We consider equilibria both in pure strategies $\mu \in \prod_i \mathcal{S}_i^{M_i}(p_i^0)$ and in mixed strategies for the designers, that is, profiles of Borel probability distributions $\zeta \in \prod_i \Delta(\mathcal{S}_i^{M_i}(p_i^0))$ over the compact sets of splittings. A mixed strategy $\zeta_i \in \Delta(\mathcal{S}_i^{M_i}(p_i^0))$ for designer i is called *mixed splitting*.

When the continuation outcome does not depend on μ , we write $y(p) = \tau(\mu, p)$ and

$$U_i^y(p) = u_i(y(p); p), \quad U_i^y(\mu) = \int_p u_i(y(p); p) d\mu(p).$$

Robustness to Message Sets

The first result of this section is that the set of equilibrium outcomes expands with the number of messages when the sets of messages are large enough ($|M_i| \geq |\Theta_i|$ for every $i \in N$). This implies that every equilibrium outcome of the *simple game* PG_Θ , where the message space M_i of each designer i has the same cardinality as Θ_i , is also an equilibrium outcome of every game with larger (possibly infinite) message spaces. It also implies that such an equilibrium is “robust” in the sense that it would still be an equilibrium in the broader game where each designer i is allowed to choose the set of messages M_i strategically.

Equilibrium outcomes of the simple game PG_Θ are also interesting, because, in the single designer case, they include the designer-preferred equilibrium outcome (see Kamenica and Gentzkow [19] and Section 3.4.1). It should be noted that the designer-preferred equilibrium is *Markovian* in the sense that it can be implemented with agents' strategies that depend only on their posterior beliefs. Whereas there is no analogue to the designer-preferred equilibrium when there are multiple designers (except if all designers have the same preferences), we will show in Section 3.3 that equilibrium outcomes of the simple game PG_Θ include all *Markovian equilibrium outcomes* of all games with more messages than states.

Theorem 2. Let M and M' be two profiles of message sets such that $|M'_i| \geq |M_i| \geq |\Theta_i|$ for every $i \in N$. The set of equilibrium outcomes of PG_M is included in the set of equilibrium outcomes of $PG_{M'}$, as well as in the set of equilibrium outcomes of PG_∞ .

This theorem drastically simplifies the study of the equilibria of public information design games since it shuts down the search for equilibrium message spaces. We know that we can find robust equilibria by letting designers use fixed finite message sets that need not contain more messages than states. Combined with Theorem 1, this also implies that the set of equilibrium outcomes of PG_∞ is nonempty. Note that this last result cannot be deduced from Theorem 1, where designers are restricted to choosing from compact sets of experiments with countable support.

The intuition for Theorem 2 is as follows. Take an equilibrium profile (of mixed splittings) ζ^* of PG_Θ , and let designers play the same mixed splittings in PG_M . This is possible because $S_i^{\Theta_i}(p_i^0) \subseteq S_i^{M_i}(p_i^0)$ whenever $|M_i| \geq |\Theta_i|$. Along the equilibrium path in PG_M , consider the same continuation equilibria as in PG_Θ . Off the equilibrium path, if designer i deviates unilaterally from ζ_i^* to a splitting $\mu_i \in S_i^{M_i}(p_i^0)$, consider the *worst* continuation equilibrium for i . If this deviation is profitable in PG_M , then it is also profitable in PG_Θ , because the best response of designer i is obtained as a concavification using $|\Theta_i|$ posteriors, a contradiction. Here is the formal proof.

Proof. Consider an equilibrium $((\zeta_i^*)_{i \in N}, \tau)$ of PG_M , where each ζ_i^* is a mixed strategy in $\Delta(S^{M_i}(p_i^0))$ and $\tau : \prod_i S^{M_i}(p_i^0) \times \Delta(\Theta_i) \rightarrow \Delta(A)$ is a continuation equilibrium outcome for the agents. We will prove the result by extending τ to all splittings in such a way that no designer can profitably deviate to any splitting, with either finite or infinite support.

Recall that, for all (μ, p) , $\tau(\mu, p) \in Y(p)$ is a continuation equilibrium outcome given the public belief p . For each designer i and belief p , fix $y^i(p) \in \arg\min_{y \in Y(p)} u_i(y, p)$. Define the extension $\tilde{\tau} : \prod_i S_i(p_i^0) \times \Delta(\Theta_i) \rightarrow \Delta(A)$ as follows:

- If $\mu_i \in S^{M_i}(p_i^0)$ for all i , then $\tilde{\tau}(\mu, p) = \tau(\mu, p)$ for all p .
- If there is a single designer i such that $\mu_i \notin S^{M_i}(p_i^0)$, then choose $\tilde{\tau}(\mu_i, \mu_{-i}, p) = y^i(p)$.
- For all other profiles μ , let $\tilde{\tau}(\mu, p)$ be arbitrary in $Y(p)$.

We claim that $((\tilde{\zeta}_i^*)_{i \in N}, \tilde{\tau})$ is an equilibrium of PG_∞ . Clearly, no designer i can profitably deviate to $\tilde{\mu}_i \in S^{M_i}(p_i^0)$. Suppose that designer i deviates to $\tilde{\mu}_i \notin S^{M_i}(p_i^0)$. Her expected payoff is

$$\begin{aligned} U_i^{\tilde{\tau}}(\tilde{\mu}_i, \zeta_{-i}^*) &= \int_{\mu_{-i}} \left\{ \int_{p_i, p_{-i}} u_i(y^i(p_i, p_{-i}); p_i, p_{-i}) d\tilde{\mu}_i(p_i) d\mu_{-i}(p_{-i}) \right\} d\zeta_{-i}^*(\mu_{-i}) \\ &= \int_{p_i} \left\{ \int_{\mu_{-i}} \int_{p_{-i}} u_i(y^i(p_i, p_{-i}); p_i, p_{-i}) d\mu_{-i}(p_{-i}) d\zeta_{-i}^*(\mu_{-i}) \right\} d\tilde{\mu}_i(p_i). \end{aligned}$$

Denoting $U_i^{\tilde{\tau}}(p_i, \zeta_{-i}^*) = \int_{\mu_{-i}} \int_{p_{-i}} u_i(y^i(p_i, p_{-i}); p_i, p_{-i}) d\mu_{-i}(p_{-i}) d\zeta_{-i}^*(\mu_{-i})$, we have

$$U_i^{\tilde{\tau}}(\tilde{\mu}_i, \zeta_{-i}^*) = \int_{p_i} U_i^{\tilde{\tau}}(p_i, \zeta_{-i}^*) d\tilde{\mu}_i(p_i) \leq \sup_{\mu_i \in S_i^{\Theta_i}(p_i^0)} \int_{p_i} U_i^{\tilde{\tau}}(p_i, \zeta_{-i}^*) d\mu_i(p_i) = \sup_{\mu_i \in S_i^{\Theta_i}(p_i^0)} \int_{p_i} U_i^{\tilde{\tau}}(p_i, \zeta_{-i}^*) d\mu_i(p_i).$$

Indeed, the right-hand side is the concave closure $\text{cav}_{p_i} U_i^{\tilde{\tau}}(p_i, \zeta_{-i}^*)$ of $U_i^{\tilde{\tau}}(p_i, \zeta_{-i}^*)$ with respect to p_i , and the supremum is achieved with $|\Theta_i|$ posteriors.⁷

If the deviation $\tilde{\mu}_i$ is profitable, then there exists $\mu_i^* \in S_i^{\Theta_i}(p_i^0) \subseteq S^{M_i}(p_i^0)$ such that

$$U_i^{\tilde{\tau}}(\zeta_i^*, \zeta_{-i}^*) = U_i^{\tilde{\tau}}(\zeta_i^*, \zeta_{-i}^*) < U_i^{\tilde{\tau}}(\tilde{\mu}_i, \zeta_{-i}^*) \leq U_i^{\tilde{\tau}}(\mu_i^*, \zeta_{-i}^*) \leq U_i^{\tau}(\mu_i^*, \zeta_{-i}^*),$$

where the last inequality follows from the fact that $\tilde{\tau}(\cdot, p)$ is the least-preferred continuation equilibrium of designer i . This contradicts the fact that ζ^* is an equilibrium of PG_M .

Therefore, if ζ^* is an equilibrium of PG_M , then it is also an equilibrium in the public information design games with more messages, as desired. \square

Theorem 2 shows that the set of equilibrium outcomes increases with the number of messages, that is, with the size of the supports of splittings. In general, the inclusions can be strict, as illustrated in the next example.

Example 1. Consider one designer, one agent, and two states $\Theta = \{0, 1\}$. The prior probability of state $\theta = 1$ is $p^0 = 1/2$. The set of actions of the agent is $A = A^1 \times A^2$, where A^1 is any finite set satisfying $\{0, 1/2, 1\} \subseteq A^1 \subseteq [0, 1]$, and $A^2 = \{0, 1\}$. We choose a utility function $v(a^1, a^2; p)$ of the agent that depends only on $a^1 \in A^1$ and which is such that $a^1 = 0$ is optimal *only* at belief $p = 0$, $a^1 = 1/2$ is optimal *only* at belief $p = 1/2$, and $a^1 = 1$ is optimal *only* at belief $p = 1$. The utility of the designer depends only on the agent's action and is given by $u(a^1, a^2) = \frac{1}{4} - (a^1 - 1/2)^2 - a^2$. Note that a^2 here is only for punishing or rewarding the designer.

Consider the splitting μ^3 of PG_{M^3} with $|M^3| = 3$, which induces the distribution $(1/4, 1/2, 1/4)$ over the posteriors $\{0, 1/2, 1\}$. Consider the continuation equilibrium (i.e., optimal strategy for the agent) τ , which consists of playing the pair of actions $\tau(\mu, p) = (p, 0) \in A$ if $\mu = \mu^3$, and $\tau(\mu, p) = (a^1(p), 1) \in A$ if $\mu \neq \mu^3$, where $a^1(p)$ maximizes $v(a^1, a^2; p)$. That is, the agent plays optimally given her belief, but she punishes with $a^2 = 1$ when the designer deviates from μ^3 (it is optimal for the agent to do so).

Clearly, (μ^3, τ) is an equilibrium of PG_{M^3} , because the designer gets a positive payoff, while she gets a strictly negative payoff if she deviates. As we have observed, τ is also a continuation equilibrium outcome. This equilibrium in PG_{M^3} induces the outcome $\rho \in \Delta(A)^\Theta$, which is such that

$$\rho(0, 0 \mid \theta = 0) = \rho(1/2, 0 \mid \theta = 0) = \rho(1, 0 \mid \theta = 1) = \rho(1/2, 0 \mid \theta = 1) = 1/2.$$

In particular, it induces the ex ante distribution $(\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ over the actions $\{(0, 0), (\frac{1}{2}, 0), (1, 0)\}$.

Consider now the game PG_{M^2} , with $|M^2| = 2$. To induce the same outcome ρ noted earlier, the strategy profile must induce the ex ante distribution $(\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ over the actions $\{(0, 0), (\frac{1}{2}, 0), (1, 0)\}$; thus, the designer must induce the distribution $(\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ over the posteriors $\{0, \frac{1}{2}, 1\}$. The only way to achieve this with splittings with two posteriors is to use a mixed splitting $\zeta \in \Delta(S^{M^2}(p^0))$ that randomizes over the nonrevealing splitting μ_{NR}^2 and the fully revealing splitting μ_{FR}^2 with equal probabilities. In addition, along the equilibrium path, the agent must choose $a^2 = 0$. However, the designer has a profitable deviation from ζ , which consists of playing the splitting μ_{NR}^2 with probability 1. Indeed, playing μ_{NR}^2 gives him a payoff of $1/4$, whereas playing μ_{FR}^2 gives him a payoff of 0.

This example shows that there exists an equilibrium outcome of PG_{M^3} that is not an equilibrium outcome of the simple game PG_{M^2} . This example can be generalized to show that, for every $M' > M \geq 2$, there is an equilibrium outcome of $PG_{M'}$ that is not an equilibrium outcome of PG_M .

The second result in this section states that, in the game PG_∞ , it is without loss of generality to focus on pure strategies for the designers.

Theorem 3. *The set of equilibrium outcomes of PG_∞ coincides with the set of equilibrium outcomes of PG_∞ in which designers use pure strategies.*

This shows that the set of relevant strategies to be considered to obtain all possible equilibria of information design games with public messages (and $|M_i| \geq |\Theta_i|$) is “simply” the set of all splittings, and no randomization over splittings or experiments is required.

The proof of Theorem 3 goes as follows. For each profile of mixed splittings in $\prod_i \Delta(S_i(p_i^0))$, we construct a profile of pure splittings in $\prod_i S_i(p_i^0)$ that induces the same distributions of posterior beliefs. Namely, each mixture of splittings is replaced by the “expected” splitting. Then, it suffices to construct a continuation outcome after the profile of pure splittings in $\prod_i S_i(p_i^0)$, which induces the same expected payoffs for the designers.

The intuitive reason why this is an equilibrium is the following. Each designer i receives the same expected payoff from randomizing over splittings as from playing the pure strategy. Since it was optimal to play the randomization, it should also be optimal to play the equivalent pure strategy, as it does not involve indifference constraints due to randomization.

Combined with Theorem 2, this result implies that every equilibrium outcome of the simple game PG_Θ is an equilibrium outcome of PG_∞ in which designers use pure strategies.

Proof. Consider an equilibrium $((\zeta_i, \zeta_{-i}), \tau)$ of PG_∞ in mixed strategies. The mixed splitting $\zeta_i \in \Delta(S_i(p_i^0))$ for designer i induces a distribution $F_i \in \Delta(S_i(p_i^0) \times \Delta(\Theta_i))$ defined by $dF_i(\mu_i, p_i) = d\mu_i(p_i)d\zeta_i(\mu_i)$. To be precise, for any a Borel set $B \subseteq S_i(p_i^0) \times \Delta(\Theta_i)$, denote $B_{\mu_i} = \{p_i : (\mu_i, p_i) \in B\}$ and let

$$F_i(B) = \int_{\mu_i, p_i} \mathbf{1}_B dF_i(\mu_i, p_i) = \int_{\mu_i} \left(\int_{p_i} \mathbf{1}_{B_{\mu_i}} d\mu_i(p_i) \right) d\zeta_i(\mu_i) = \int_{\mu_i} \mu_i(B_{\mu_i}) d\zeta_i(\mu_i).$$

The marginal distribution of p_i under F_i is a splitting (the “expected splitting”) denoted by μ_{F_i} . The expected payoff of designer i under $((\zeta_i, \zeta_{-i}), \tau)$ is written as:

$$\begin{aligned} U_i^\tau(\zeta_i, \zeta_{-i}) &= \int_{\mu_i, \mu_{-i}} \left\{ \int_{p_i, p_{-i}} u_i(\tau(\mu_i, \mu_{-i}, p_i, p_{-i}); p_i, p_{-i}) d\mu_i(p_i) d\mu_{-i}(p_{-i}) \right\} d\zeta_i(\mu_i) d\zeta_{-i}(\mu_{-i}) \\ &= \int_{\mu_i, p} u_i(\tau(\mu_i, \mu_{-i}, p_i, p_{-i}); p_i, p_{-i}) dF_i(p_i, \mu_i) dF_{-i}(p_{-i}, \mu_{-i}). \end{aligned}$$

Now, we use Fubini's theorem to exchange integration order. Let us write the distribution $dF_i(p_i, \mu_i)$ as the product of its marginal distribution $dF_i(p_i)$ and of a conditional distribution $dF_i(\mu_i | p_i)$ (these are well defined since the sets $\mathcal{S}_i(p_i^0)$ and $\Delta(\Theta_i)$ are compact metric). We obtain

$$\begin{aligned} U_i^\tau(\zeta_i, \zeta_{-i}) &= \int_{\mu, p} u_i(\tau(\mu_i, \mu_{-i}, p_i, p_{-i}); p_i, p_{-i}) dF_i(\mu_i | p_i) dF_i(p_i) dF_{-i}(\mu_{-i} | p_{-i}) dF_{-i}(p_{-i}) \\ &= \int_p u_i \left(\int_{\mu} \tau(\mu_i, \mu_{-i}, p_i, p_{-i}) dF_i(\mu_i | p_i) dF_{-i}(\mu_{-i} | p_{-i}); p_i, p_{-i} \right) dF_i(p_i) dF_{-i}(p_{-i}). \end{aligned}$$

If player i deviates to ξ_i , inducing the distribution G_i , then we have $U_i^\tau(\xi_i, \zeta_{-i}) \leq U_i^\tau(\zeta_i, \zeta_{-i})$ and

$$U_i^\tau(\xi_i, \zeta_{-i}) = \int_p u_i \left(\int_{\mu} \tau(\mu_i, \mu_{-i}, p_i, p_{-i}) dG_i(\mu_i | p_i) dF_{-i}(\mu_{-i} | p_{-i}); p_i, p_{-i} \right) dG_i(p_i) dF_{-i}(p_{-i}).$$

Let F_i, F_{-i} be the representations of mixed splittings as joint distributions, and let $\mu_i^* = \mu_{F_i}$ be the expected splitting of designer i .

Define $\tau^* : \prod_i \mathcal{S}_i(p_i^0) \times \Delta(\Theta_i) \rightarrow \Delta(A)$ as follows:

- $\tau^*(\mu^*, p) = \int_{\mu} \tau(\mu_i, \mu_{-i}, p_i, p_{-i}) dF_i(\mu_i | p_i) dF_{-i}(\mu_{-i} | p_{-i})$,
- $\tau^*(\mu_i, \mu_{-i}^*, p) = y^i(p) \in \arg \min_{y \in Y(p)} u_i(y, p)$,
- for all other profiles μ , let $\tau^*(\mu, p)$ be arbitrary in $Y(p)$.

By construction, the payoff on the equilibrium path is

$$\begin{aligned} U_i^\tau(\zeta_i, \zeta_{-i}) &= \int_p u_i \left(\int_{\mu} \tau(\mu_i, \mu_{-i}, p_i, p_{-i}) dF_i(\mu_i | p_i) dF_{-i}(\mu_{-i} | p_{-i}); p_i, p_{-i} \right) dF_i(p_i) dF_{-i}(p_{-i}) \\ &= \int_p u_i(\tau^*(\mu^*, p), p) d\mu_i^*(p_i) d\mu_{-i}^*(p_{-i}) = U_i^{\tau^*}(\mu_i^*, \mu_{-i}^*). \end{aligned}$$

Deviating to μ_i gives a payoff

$$\begin{aligned} U_i^{\tau^*}(\mu_i, \mu_{-i}^*) &= \int_p u_i(y^i(p), p) d\mu_i(p_i) d\mu_{-i}^*(p_{-i}) \\ &= \int_p u_i(y^i(p), p) dG_i(p_i) dF_{-i}(p_{-i}) \\ &\leq \int_p u_i \left(\int_{\mu} \tau(\mu_i, \mu_{-i}, p_i, p_{-i}) dG_i(\mu_i | p_i) dF_{-i}(\mu_{-i} | p_{-i}); p_i, p_{-i} \right) dG_i(p_i) dF_{-i}(p_{-i}) \\ &= U_i^\tau(\xi_i, \zeta_{-i}) \leq U_i^\tau(\zeta_i, \zeta_{-i}) = U_i^{\tau^*}(\mu_i^*, \mu_{-i}^*), \end{aligned}$$

where the first inequality holds because $y^i(p)$ is the punishing continuation equilibrium outcome. This concludes the proof. \square

Combining Theorems 1, 2, and 3, we obtain the following:

Corollary 1. *The game PG_∞ admits an equilibrium in which designers use pure strategies.*

Another interesting implication of this result is that if designers were able to choose any statistical experiment (with or without countable support) together with any message space, then the game would have an equilibrium (i.e., an equilibrium of PG_∞) in which designers play pure strategies.

Canonical Games and Markovian Equilibria

In Example 1, more messages yield more equilibrium outcomes. This finding relies on equilibria where the action of the agent depends not only on the posterior belief but also on the splitting used by the designer. Let us consider equilibria where the actions of agents depend only on the belief.

Definition 3. An equilibrium $((\zeta_i^*)_{i \in N}, \tau)$ is *Markovian* if $\tau(\mu, p) = y(p) \in Y(p)$ does not depend on μ .

This equilibrium refinement requires strategies to be Markovian⁸ with respect to common beliefs, a natural state variable for this model. The next result is that the set of outcomes of Markovian equilibria of the simple game coincides with the set of outcomes of Markovian equilibria of games with more messages (possibly infinitely many). Thus, the simple game may be called *canonical* for the set of Markovian equilibria.⁹

Theorem 4. Let M be a profile of message sets such that $|M_i| \geq |\Theta_i|$ for every $i \in N$. The set of Markovian equilibrium outcomes of PG_M coincides with the set of Markovian equilibrium outcomes of PG_∞ . This set also coincides with the set of Markovian equilibrium outcomes of PG_∞ in which designers use pure strategies.

Remark 2. It can be noted that those results extend directly to models with continuous actions. The only important assumption in the model is that the correspondence $p \mapsto Y(p)$ is nonempty, convex, compact-valued, and upper hemicontinuous. In particular, when $Y(p)$ is a singleton for each p , all equilibria are Markovian. This is the case, for instance, if there is a single agent with strictly concave preferences and if there are many games with quadratic preferences such as beauty contests, Cournot competition, and network games. For such environments, the simple game PG_Θ is canonical for *all* equilibrium outcomes. That is, the set of equilibrium outcomes of the game PG_Θ coincides with the set of equilibrium outcomes of every game PG_M with $|M_i| \geq |\Theta_i|$ for every $i \in N$, as well as with the set of equilibrium outcomes of PG_∞ .

The proof of this result relies on a lemma that states that every splitting can be replicated by mixing over splittings of the simple game. Let us call $\mathcal{S}_i^{\Theta_i}(p_i^0)$ the set of *canonical splittings* for designer i . A *mixed canonical splitting* of p_i^0 is a Borel probability distribution $\zeta_i \in \Delta(\mathcal{S}_i^{\Theta_i}(p_i^0))$ over the compact set of canonical splittings. A mixed splitting ζ_i induces an “expected” splitting μ_{ζ_i} , defined for each Borel set $B \subseteq \Delta(\Theta_i)$ by

$$\mu_{\zeta_i}(B) = \int \mu_i(B) d\zeta_i(\mu_i).$$

It is easy to see that $\mu_{\zeta_i} \in \mathcal{S}_i^{\Theta_i}(p_i^0)$. We have the following representation result.

Lemma 1.

$$\mathcal{S}_i(p_i^0) = \overline{\text{co}} \mathcal{S}_i^{\Theta_i}(p_i^0) = \left\{ \mu_{\zeta_i} : \zeta_i \in \Delta(\mathcal{S}_i^{\Theta_i}(p_i^0)) \right\},$$

where $\overline{\text{co}}$ denotes the closure of the convex hull.

In words, any splitting is a (limit of a) convex combination of canonical splittings. Equivalently, any splitting is the expected splitting induced by randomizing over canonical splittings.

Now, the intuition that every Markovian equilibrium outcome is an equilibrium outcome of the simple game (Theorem 4) is the following. By Lemma 1, every equilibrium splitting in PG_∞ can be replaced by a mixed canonical splitting, which induces the same distribution of posteriors in PG_Θ . When the equilibrium is Markovian, continuation equilibrium outcomes depend only on agents’ beliefs. Hence, if a deviation from the mixed canonical splitting is profitable in the simple game, then it is also profitable in the original game, which is a contradiction. This argument does not hold if continuation equilibrium outcomes also depend on designers’ splittings, as illustrated in Example 1.

Let us turn to the formal proofs.

Proof of Lemma 1. Since $\mathcal{S}_i(p_i^0)$ is convex and compact, $\overline{\text{co}} \mathcal{S}_i^{\Theta_i}(p_i^0) \subseteq \mathcal{S}_i(p_i^0)$. Let $\mu_i^* \in \mathcal{S}_i(p_i^0)$; we want to show that $\mu_i^* \in \overline{\text{co}} \mathcal{S}_i^{\Theta_i}(p_i^0)$. First, observe that $\mathcal{S}_i^{\Theta_i}(p_i^0)$ is compact in $\Delta(\Delta(\Theta_i))$, and, therefore, the closure of its convex hull $\overline{\text{co}} \mathcal{S}_i^{\Theta_i}(p_i^0)$ is also compact (see, e.g., Aliprantis and Border [2, theorem 5.35, p. 185]). Suppose that $\mu_i^* \notin \overline{\text{co}} \mathcal{S}_i^{\Theta_i}(p_i^0)$. By the separation theorem (Aliprantis and Border [2, theorem 5.79, p. 207]), there exists a continuous function $f : \Delta(\Theta_i) \rightarrow \mathbb{R}$ such that

$$\int_{p_i} f(p_i) d\mu_i^*(p_i) > \sup_{\mu_i \in \overline{\text{co}} \mathcal{S}_i^{\Theta_i}(p_i^0)} \int_{p_i} f(p_i) d\mu_i^*(p_i) \geq \sup_{\mu_i \in \mathcal{S}_i^{\Theta_i}(p_i^0)} \int_{p_i} f(p_i) d\mu_i^*(p_i).$$

This inequality is impossible. The left-hand side is smaller than $\text{cav}f(p_i^0)$, and the right-hand side is $\text{cav}f(p_i^0)$, since the cav of any function is achieved with $|\Theta_i|$ posteriors (Rockafellar [36, corollary 17.1.5, p. 157]). We conclude that $\mu_i^* \in \overline{\text{co}} \mathcal{S}_i^{\Theta_i}(p_i^0)$, and, therefore, $\mathcal{S}_i(p_i^0) = \overline{\text{co}} \mathcal{S}_i^{\Theta_i}(p_i^0)$.

The other set equality in Lemma 1 follows from three observations. First, $\text{co} \mathcal{S}_i^{\Theta_i}(p_i^0)$ is the set of all μ_{ζ_i} for ζ_i in $\Delta(\mathcal{S}_i^{\Theta_i}(p_i^0))$ with finite support. Second, the mapping $\zeta_i \mapsto \mu_{\zeta_i}$ is continuous for the weak-* topology. To see this, suppose that ζ_i^n weak-* converges to ζ_i . For every continuous function $f : \Delta(\Theta_i) \rightarrow \mathbb{R}$,

$$\int_{p_i} f(p_i) d\mu_{\zeta_i^n}(p_i) = \int_{\mu_i} \left(\int_{p_i} f(p_i) d\mu_i(p_i) \right) d\zeta_i^n(\mu_i).$$

The weak-* convergence of ζ_i^n implies that this converges to

$$\int_{\mu_i} \left(\int_{p_i} f(p_i) d\mu_i(p_i) \right) d\zeta_i(\mu_i) = \int_{p_i} f(p_i) d\mu_{\zeta_i}(p_i).$$

Thus, $\mu_{\zeta_i^n}$ weak-* converges to μ_{ζ_i} . Third, the set of probability measures with finite support is dense in $\Delta(\mathcal{S}_i^{\Theta_i}(p_i^0))$ (Aliprantis and Border [2, theorem 15.10, p. 513]). The closure of the convex hull $\overline{\text{co}} \mathcal{S}_i^{\Theta_i}(p_i^0)$ is thus the set of all μ_{ζ_i} s. \square

Proof of Theorem 4. The argument of the proof of Theorem 2 directly applies to show that a Markovian equilibrium outcome of PG_M is a Markovian equilibrium outcome of PG_∞ . Conversely, consider a Markovian equilibrium $((\zeta_i^*)_{i \in N}, y)$ of PG_∞ . First, we reduce it to a Markovian equilibrium $((\mu_i^*)_{i \in N}, y)$ of PG_∞ in pure strategies for the designers. As before, a mixed splitting ζ_i^* is represented by a joint distribution $F_i^* \in \Delta(\mathcal{S}_i(p_i^0) \times \Delta(\Theta_i))$. The equilibrium payoff of designer i is

$$\begin{aligned} U_i^y(\zeta_i^*, \zeta_{-i}^*) &= \int_{\mu} \int_p u_i(y(p_i, p_{-i}); p_i, p_{-i}) d\mu_i(p_i) d\mu_{-i}(p_{-i}) d\zeta_i^*(\mu_i) d\zeta_{-i}^*(\mu_{-i}) \\ &= \int_{\mu, p} u_i(y(p_i, p_{-i}); p_i, p_{-i}) dF_i^*(\mu_i, p_i) dF_{-i}^*(\mu_{-i}, p_{-i}) \\ &= \int_p u_i(y(p_i, p_{-i}); p_i, p_{-i}) \int_{\mu} dF_i^*(\mu_i, p_i) dF_{-i}^*(\mu_{-i}, p_{-i}) \\ &= \int_p u_i(y(p_i, p_{-i}); p_i, p_{-i}) d\mu_{F_i^*}(p_i) d\mu_{F_{-i}^*}(p_{-i}) = U_i^y(\mu_{F_i^*}, \mu_{F_{-i}^*}). \end{aligned}$$

If designer i deviates from the mixed equilibrium to a pure $\tilde{\mu}_i$, then $U_i^y(\tilde{\mu}_i, \zeta_{-i}^*) \leq U_i^y(\zeta_i^*, \zeta_{-i}^*)$ and,

$$U_i^y(\tilde{\mu}_i, \zeta_{-i}^*) = \int_p u_i(y(p_i, p_{-i}); p_i, p_{-i}) d\tilde{\mu}_i(p_i) d\mu_{F_{-i}^*}(p_{-i}) = U_i^y(\tilde{\mu}_i, \mu_{F_{-i}^*}),$$

which shows that $((\mu_{F_i^*}, \mu_{F_{-i}^*}), y)$ is an equilibrium.

Now, let (μ^*, y) be a Markovian equilibrium of PG_∞ in pure strategies for the designers. From Lemma 1, there exists $\zeta_i^* \in \Delta(\mathcal{S}_i^{\Theta_i}(p_i^0))$ such that $\mu_i^* = \mu_{\zeta_i^*}$. We claim that $((\zeta_i^*)_{i \in N}, y)$ is a mixed equilibrium of PG_∞ . We have

$$\begin{aligned} U_i^y(\mu_i^*, \mu_{-i}^*) &= \int_{\mu} \int_p u_i(y(p_i, p_{-i}); p_i, p_{-i}) d\mu_i(p_i) d\mu_{-i}(p_{-i}) d\zeta_i^*(\mu_i) d\zeta_{-i}^*(\mu_{-i}) \\ &= U_i^y(\zeta_i^*, \zeta_{-i}^*). \end{aligned}$$

Suppose that designer i deviates to some pure $\tilde{\mu}_i \in \mathcal{S}_i^{\Theta_i}(p_i^0)$. The expected payoff is

$$\begin{aligned} U_i^y(\tilde{\mu}_i, \zeta_{-i}^*) &= \int_{\mu_{-i}} \int_p u_i(y(p_i, p_{-i}); p_i, p_{-i}) d\tilde{\mu}_i(p_i) d\mu_{-i}(p_{-i}) d\zeta_{-i}^*(\mu_{-i}) \\ &= U_i^y(\tilde{\mu}_i, \mu_{-i}^*) \leq U_i^y(\mu_i^*, \mu_{-i}^*) = U_i^y(\zeta_i^*, \zeta_{-i}^*). \end{aligned}$$

This completes the proof of the theorem. \square

Examples

The One-Designer, One-Agent Case. Public information design games subsume the model of Bayesian persuasion of Kamenica and Gentzkow [19], where there is a single designer and a single agent. When there is a single agent, the set $Y(p)$ of continuation equilibrium outcomes is simply the set of optimal mixed actions of the agent:

$$Y(p) = \Delta(A(p)), \text{ where } A(p) = \arg \max_{a \in A} v(a; p).$$

An equilibrium is a Borel-measurable selection $\tau(\mu, p) \in Y(p)$ for each μ and p and an optimal strategy for the designer that solves $\sup_{\mu \in \mathcal{S}(p^0)} U^\tau(\mu)$. The best equilibrium for the designer is Markovian and is obtained by selecting the designer-preferred optimal strategy τ^* of the agent,

$$\tau^*(\mu, p) = y^*(p) \in \arg \max_{y \in Y(p)} u(y, p).$$

Then, the induced utility for the designer, denoted by $U^*(p) = u(y^*(p); p)$ is upper semicontinuous, and the designer has a best response inducing the payoff:

$$\max_{\mu \in S(p^0)} U^*(\mu) = \text{cav}_p U^*(p^0).$$

Example 2. The designer can design information on $\Theta = \{\theta^1, \theta^2\}$, and $|M| = 2$. The agent has three possible actions, $A = \{a^1, a^2, a^3\}$, and her utility function is given by

		a^1	a^2	a^3
	θ^1	2	3	3
	θ^2	2	1	0

Denote by p the probability of θ^1 . The optimal action of the agent is a^1 when $p < \frac{1}{2}$ and a^2 when $p \in (\frac{1}{2}, 1)$. Every randomization between a_1 and a_2 is optimal at $p = \frac{1}{2}$, and every randomization between a_2 and a_3 is optimal at $p = 1$. Let $y(\frac{1}{2}) \in Y(\frac{1}{2})$ be identified with the probability that the agent plays a^2 at $p = \frac{1}{2}$, and let $y(1) \in Y(1)$ be identified with the probability that the agent plays a^3 at $p = 1$. Consider a designer's utility function that depends only on the action and is given by

$$u(a, \theta) = \begin{cases} 0 & \text{if } a = a^1, \\ 1 & \text{if } a = a^2, \\ 3 & \text{if } a = a^3. \end{cases}$$

Given the strategy $y(p)$ of the agent, the induced utility of the designer as a function of the posterior p is given by the function $u(y(p); p)$ represented in Figure 1. The concavification is given by the dotted curve. Observe that the concavification depends on the tie-breaking rule at $p = 1$ but not at $p = 1/2$. The designer-preferred optimal strategy y^* is such that $y^*(\frac{1}{2}) = y^*(1) = 1$, which yields

$$\text{cav}_p u(y^*(p^0); p^0) = \text{cav}_p U^*(p^0) = 3p^0.$$

Two Strictly Competitive Examples. In the following example, we show that with a finite set of messages, if we arbitrarily fix a continuation equilibrium, then an equilibrium (and even an ε -equilibrium) in pure strategies may not exist. An equilibrium in pure strategies with a finite set of messages exists for *some* continuation equilibrium, which depends both on the posteriors *and* on the splittings used by the designers.


Example 3 (Being Perceived as Better). The agent (a buyer) decides to buy either from firm 1 (action 1) or from firm 2 (action 2). Each designer (firm) wants to maximize the probability of trade. The state (valuations of the buyer) is $\theta = (\theta_1, \theta_2) \in \{0, 1\} \times \{0, 1\}$, and the prior probability that $\theta_i = 1$ is $p_i^0 = \frac{1}{2}$. The agent's payoff is $v(i, \theta) = \theta_i$, $i = 1, 2$, so the set of optimal actions is given by

$$A(p_1, p_2) = \begin{cases} \{1\} & \text{if } p_1 > p_2 \\ \{2\} & \text{if } p_1 < p_2 \\ \{1, 2\} & \text{if } p_1 = p_2, \end{cases}$$

where action $a = i$ means buying from i . Note that, for every strategy of the agent, the game between the designers is zero-sum.

A constant tie-breaking rule is given by a fixed $\alpha \in [0, 1]$ such that the probability of choosing action $a = 1$ is α whenever $p_1 = p_2$. After normalization of the utility functions, we have (see Figure 2)

$$U_1(p_1, p_2) = -U_2(p_1, p_2) = U(p_1, p_2) = \begin{cases} 1 & \text{if } p_1 > p_2, \\ -1 & \text{if } p_2 > p_1, \\ 2\alpha - 1 & \text{if } p_1 = p_2. \end{cases}$$

Figure 1. Concavification in Example 2. $u(y(p); p)$ in solid lines and $\text{cav}_p u(y(p); p)$ in dotted lines.

In game PG_M with $|M_1| = |M_2| = 2$, consider a splitting μ_2 of designer 2 given by the convex combination $\frac{1}{2} = \lambda p_2 + (1 - \lambda)p'_2$ with $p_2 \leq \frac{1}{2} \leq p'_2$. The best-reply payoff of designer 1 is $\text{cav}_{p_1} U(p_1, \mu_2)$, where cav_{p_1} denotes the concavification with respect to the first variable, the second one being fixed. We have

$$U(p_1, \mu_2) = \begin{cases} -1 & \text{if } p_1 < p_2, \\ \lambda(2\alpha - 1) - (1 - \lambda) = 2\alpha\lambda - 1 & \text{if } p_1 = p_2, \\ \lambda - (1 - \lambda) = 2\lambda - 1 & \text{if } p_2 < p_1 < p'_2, \\ \lambda + (1 - \lambda)(2\alpha - 1) = 2\lambda + 2\alpha(1 - \lambda) - 1 & \text{if } p_1 = p'_2, \\ 1 & \text{if } p_1 > p'_2. \end{cases}$$

It is easy to verify that $\text{cav}_{p_1} U(\frac{1}{2}, \mu_2) \geq \frac{\alpha}{2}$ for every μ_2 . In addition, if designer 2 plays fully revealing (i.e., $p_2 = 0$, $p'_2 = 1$), then $\text{cav}_{p_1} U(\frac{1}{2}, \mu_2) = \frac{\alpha}{2}$. Hence,


$$\min_{\mu_2 \in \mathcal{S}_2(\frac{1}{2})} \sup_{\mu_1 \in \mathcal{S}_2(\frac{1}{2})} U(\mu_1, \mu_2) = \frac{\alpha}{2}.$$

The reasoning is symmetric for designer 2, and we obtain

$$\max_{\mu_1 \in \mathcal{S}_2(\frac{1}{2})} \inf_{\mu_2 \in \mathcal{S}_2(\frac{1}{2})} U(\mu_1, \mu_2) = -\frac{(1 - \alpha)}{2}.$$

Therefore, for every fixed $\alpha \in [0, 1]$, the induced two-player game has no value in pure strategies and, thus, no equilibrium in which designers use pure strategies.¹⁰

The game has an equilibrium in pure strategies if the agent's strategy depends on the splittings (hence, this equilibrium is not Markovian according to Definition 3). Indeed, suppose that both designers completely reveal the state, in which case the agent randomizes equally when she is indifferent ($\alpha = \frac{1}{2}$). This yields an expected payoff equal to zero for both designers. If designer 1 deviates, then the agent chooses designer 2's preferred action

Figure 2. Payoff of designer 1 in Example 3 as a function of the agent's beliefs (p_1, p_2) .

($\alpha = 0$), and if designer 2 deviates, then the agent chooses designer 1's preferred action ($\alpha = 1$). It is easy to see that, with this strategy, no designer can achieve an expected probability of trade higher than $\frac{1}{2}$, so this constitutes an equilibrium in which designers use pure strategies.

Now suppose that designers can use mixed strategies, and consider any optimal strategy for the agent. Suppose that designer 2 plays the “uniform strategy”: she randomizes uniformly over all possible symmetric splittings $\frac{1}{2} = \frac{1}{2}p_2 + \frac{1}{2}(1 - p_2)$, with $p_2 \in [0, 1]$ (uniformly distributed).¹¹ Then, the payoff of designer 1 is

$$U(p_1, \mu_2) = \int_0^1 U(p_1, p_2) dp_2 = \int_0^{p_1} 1 dp_2 + \int_{p_1}^1 -1 dp_2 = 2p_1 - 1.$$

This is linear (thus concave), and the value at $\frac{1}{2}$ is 0. Thus, against the uniform strategy, 0 is the best payoff that designer 1 can achieve. Hence, the game has a value (equal to 0), and each designer has an optimal strategy (the uniform one), regardless of the optimal strategy of the agent.

In the previous example, there exists a mixed equilibrium between the designers for every optimal strategy of the agent. The next example, adapted from Sion and Wolfe [39], shows that, for some optimal strategy of the agent, an ε -equilibrium (in pure or mixed strategies) may not exist, regardless of the message spaces of the designers. This justifies our approach of considering the $(n + k)$ -player game involving the agents instead of fixing a continuation equilibrium outcome for the agents and analyzing the induced n -player game between the designers.

Example 4. As in Example 3, the state is $\theta = (\theta_1, \theta_2) \in \{0, 1\} \times \{0, 1\}$, and the prior probability that $\theta_i = 1$ is $p_i^0 = \frac{1}{2}$. The space of posteriors $[0, 1] \times [0, 1]$ is partitioned into nine regions, depending on the optimal action of the agent, who thus has at least nine possible actions. It is easy to construct payoffs for the agent that generate those regions.¹² Her optimal action leads to payoffs $U(p_1, p_2) = U_1(p_1, p_2) = -U_2(p_1, p_2) = 1$ when her belief belongs to the white area of Figure 3, and $U_1(p_1, p_2) = -U_2(p_1, p_2) = -1$ when it belongs to the gray area. The agent is indifferent between several actions at all points of discontinuities, in which case $U_1(p_1, p_2) = -U_2(p_1, p_2)$ can take any value in $[-1, 1]$.

Note first that $\text{cav}_{p_1} U(\frac{1}{2}, p_2) = 1$ for every $p_2 \in [0, 1]$. Therefore, $\text{vex}_{p_2} \text{cav}_{p_1} U(\frac{1}{2}, \frac{1}{2}) = 1$, where $\text{vex}_{p_2}(f) = -\text{cav}_{p_2}(-f)$ is the *convexification* of f with respect to the second variable. This means that, for every optimal strategy of the agent, if designer 1 *could observe* the posterior realized by designer 2, then she would achieve a payoff of 1: if the posterior induced by designer 2 is such that $p_2 \in (0, 1)$, then designer 1 does not reveal any information; otherwise, designer 1 fully discloses θ_1 . In other words, $\text{vex}_{p_2} \text{cav}_{p_1} U(\frac{1}{2}, \frac{1}{2})$ is the value of the zero-sum game where designer 2 plays first and designer 1 observes the realized posterior before choosing her splitting. This is an upper bound of the supinf value of the simultaneous move game.


The value of $\text{vex}_{p_2} U(p_1, \frac{1}{2})$ depends on the strategy of the agent at posteriors $(p_1, p_2) = (\frac{1}{2}, 0)$ and $(p_1, p_2) = (\frac{1}{2}, 1)$. For example, if $U(\frac{1}{2}, 0) = U(\frac{1}{2}, 1) = \bar{u} \in [-1, 1]$, then

$$\text{vex}_{p_2} U\left(p_1, \frac{1}{2}\right) = \begin{cases} -1 & \text{if } p_1 \neq \frac{1}{2} \\ \bar{u} & \text{if } p_1 = \frac{1}{2}, \end{cases}$$

which implies that $\text{cav}_{p_1} \text{vex}_{p_2} U(\frac{1}{2}, \frac{1}{2}) = \bar{u}$. We have

$$\text{cav}_{p_1} \text{vex}_{p_2} U\left(\frac{1}{2}, \frac{1}{2}\right) \leq \max_{\mu_1 \in \mathcal{S}(\frac{1}{2})} \inf_{\mu_2 \in \mathcal{S}(\frac{1}{2})} U(\mu_1, \mu_2) \leq \min_{\mu_1 \in \mathcal{S}(\frac{1}{2})} \sup_{\mu_2 \in \mathcal{S}(\frac{1}{2})} U(\mu_1, \mu_2) \leq \text{vex}_{p_2} \text{cav}_{p_1} U\left(\frac{1}{2}, \frac{1}{2}\right).$$

Figure 3. Payoff of designer 1 in Example 4 as a function of the agent's beliefs (p_1, p_2) .



Thus, for the strategy always favoring designer 1 ($\bar{u} = 1$), the information design game has an equilibrium value:

$$\max_{\mu_1 \in \mathcal{S}(\frac{1}{2})} \inf_{\mu_2 \in \mathcal{S}(\frac{1}{2})} U(\mu_1, \mu_2) = \min_{\mu_1 \in \mathcal{S}(\frac{1}{2})} \sup_{\mu_2 \in \mathcal{S}(\frac{1}{2})} U(\mu_1, \mu_2) = 1.$$

This equilibrium is very simple: designer 1 does not reveal any information, and she gets a payoff of 1 regardless of the strategy of designer 2.

Consider now the “symmetric” optimal strategy for the agent: she randomizes in such a way that the payoff is equal to 0 at all points of discontinuities of $U(p_1, p_2)$ of Figure 3. If the induced game between the designers has a value (supinf = infsup), then, from the previous discussion, it should belong to $[\bar{u}, 1] = [0, 1]$ (with the symmetric strategy, $\bar{u} = 0$). However, we can now follow Sion and Wolfe [39] to show that this game has no value. The example of Sion and Wolfe [39] is a normal-form game in which player 1 chooses an action $x \in [0, \frac{1}{2}]$ and player 2 chooses an action $y \in [\frac{1}{2}, 1]$, and where the payoff is given by our function $U(p_1, p_2)$ with $p_1 = x$, $p_2 = y$, and the symmetric optimal strategy for the agent. It is easy to see that the mixed extension of this game is equivalent to the information design game induced by the symmetric optimal strategy for the agent. Indeed, for every mixed strategy in these authors’ example, there exists a symmetric splitting in the information design game inducing the same payoff. Conversely, for every (symmetric or asymmetric) splitting, there exists an equivalent mixed strategy in their game.

Sion and Wolfe [39] show that the maxmin payoff is $\frac{1}{3}$ and that the minmax payoff is $\frac{3}{7}$. A maxmin strategy for designer 1 is the splitting $\frac{1}{2} = (1/6) \times 0 + (1/6) \times 1 + (2/3) \times \frac{1}{2}$. A minmax strategy for designer 2 is the splitting $\frac{1}{2} = (2/7) \times 0 + (2/7) \times 1 + (1/7) \times \frac{1}{4} + (1/7) \times \frac{3}{4} + (1/14) \times \frac{3}{8} + (1/14) \times \frac{5}{8}$. Hence,

$$\max_{\mu_1 \in \mathcal{S}(\frac{1}{2})} \inf_{\mu_2 \in \mathcal{S}(\frac{1}{2})} U(\mu_1, \mu_2) = \frac{1}{3} < \min_{\mu_1 \in \mathcal{S}(\frac{1}{2})} \sup_{\mu_2 \in \mathcal{S}(\frac{1}{2})} U(\mu_1, \mu_2) = \frac{3}{7}.$$

This shows that, with the symmetric optimal strategy for the agent, the induced information design game between the designers has no value and therefore no ε -equilibrium.

4. Rectangular Corporation Problems

In this section, we study the case in which each designer i discloses information only to a group of agents that we call her *corporation*. Precisely, for each designer i , there is a corporation of n_i agents numbered ij , $j = 1, \dots, n_i$, and corporations are assumed to form a partition of the set of agents. Each agent ij has a finite set of actions A_{ij}^j , and we let $A_i = \prod_{j=1}^{n_i} A_{ij}^j$ denote the set of action profiles for corporation i . The payoff $v_i^j(\theta_i, a_i)$ of each agent ij in corporation i depends on θ_i and a_i , that is, on the state and actions of corporation i . There is a finite set of messages $M_i^j \subset \mathbb{N}$ from each designer i to each agent ij , such that $|M_i^j| \geq |A_{ij}^j|$. We assume that the set of feasible experiments X_i for designer i is the set of all experiments with support in $M_i = \prod_{j=1}^{n_i} M_i^j$. We call such games *rectangular corporation games* and denote them by RG_M with $M = (M_i)_{i \in N}$.

As in the previous section, we assume that the prior probability distribution $p^0 = \otimes_{i \in N} p_i^0$ is the product of its marginal distributions $p_i^0 \in \Delta(\Theta_i)$ (this is without loss of generality; see Remark 3). For simplicity, we consider state-independent preferences for the designers: The payoff $u_i(a)$ of each designer i depends on $a = (a_i)_{i \in N} \in \prod_{i \in N} A_i$, the profile of actions of all agents (from all corporations). Our results extend to state-dependent preferences (see also Remark 3).

Rectangular corporation games fit well with the example of competing pharmaceutical companies mentioned in the introduction, where the agents (e.g., members of the Center for Drug Evaluation and Research of the Food and Drug Administration) review the new drug application of each company to determine if the drug is safe and effective. See also Example 5, where agents in corporations form committees that decide to approve complementary projects of entrepreneurs in new technologies.

Note that every information design game with a single designer is a rectangular corporation game. Additionally, a model with one agent per corporation may represent a common agent who chooses an action for each dimension i and whose utility is separable across dimensions. Note that Examples 3 and 4 in the previous section are *not* rectangular corporation games.

When designer i chooses the information structure for her corporation, she induces a continuation Bayesian game for it. The set of continuation equilibrium outcomes in corporation i induced by all possible experiments of designer i is the set of Bayes correlated equilibria (Bergemann and Morris [8]) in corporation i . A statistical experiment x_i for corporation i is called *direct* if $x_i \in \Delta(A_i)^{\Theta_i}$. That is, x_i is direct if, for all i, j , $|M_i^j| = |A_{ij}^j|$ and messages are one-to-one identified with actions.

Definition 4. A direct statistical experiment x_i^* is a *Bayes correlated equilibrium* (BCE) for corporation i if, for each agent ij and each pair of actions $a_i^j, b_i^j \in A_i^j$,

$$\sum_{\theta_i \in \Theta_i} \sum_{a_i^{-j}} p_i^0(\theta_i) x_i^*(a_i^j, a_i^{-j} | \theta_i) v_i^j(a_i^j, a_i^{-j}; \theta_i) \geq \sum_{\theta_i \in \Theta_i} \sum_{a_i^{-j}} p_i^0(\theta_i) x_i^*(a_i^j, a_i^{-j} | \theta_i) v_i^j(b_i^j, a_i^{-j}; \theta_i).$$

In words, a BCE is a statistical experiment where each agent's message consists of a privately recommended action, with the property that each agent has an incentive to play it if others do so. Denote by $C_i(p_i^0)$ the set of distributions of action profiles of corporation i induced by BCEs. That is, the set of $y_i \in \Delta(A_i)$ for which there exists a BCE x_i^* such that $y_i(a_i) = \sum_{\theta_i \in \Theta_i} p_i^0(\theta_i) x_i^*(a_i | \theta_i)$. The set of BCEs is nonempty and described by finitely many linear inequalities; it is therefore convex and compact, as is $C_i(p_i^0)$. Additionally, any mixture of BCEs is a BCE.

Consider the auxiliary n -player normal-form game $(C_i(p_i^0), u_i)_{i \in N}$. Theorem 5 yields a tractable two-step procedure for solving information design games between corporations: First characterize the set of BCE outcomes $C_i(p_i^0)$ in each corporation i . Second, find the Nash equilibria between designers in the normal-form game $(C_i(p_i^0), u_i)_{i \in N}$ in the same way as in standard multilinear games (see Example 5 for an illustration and a comparison with the belief-based approach used in the previous section).

Theorem 5. The normal-form game $(C_i(p_i^0), u_i)_{i \in N}$ has a nonempty compact set of pure-strategy equilibria $E(p^0)$. Every distribution of actions in $E(p^0)$ is a pure-strategy equilibrium outcome of the rectangular corporation game RG_M .

Proof. The game $(C_i(p_i^0), u_i)_{i \in N}$ is obtained from the mixed extension of the finite game $(A_i, u_i)_{i \in N}$ by considering the nonempty, convex, and compact subsets of feasible strategies $C_i(p_i^0) \subseteq \Delta(A_i)$. Therefore, it admits a pure-strategy equilibrium by Nash's theorem.

Given an equilibrium of $(C_i(p_i^0), u_i)_{i \in N}$, consider the corresponding experiments $(x_i^*)_{i \in N}$ and assume that designers play the profile of pure strategies $(x_i^*)_{i \in N}$. Since these are BCEs, each agent for each corporation has an incentive to play the recommended action with probability 1.

Suppose that designer i deviates from experiment x_i^* to another experiment x_i , and consider continuation equilibria in which agents in other corporations ignore this deviation (this is sequentially rational for them, as their payoffs do not depend on the information and actions of corporation i). In every such equilibrium, agents in corporation i play an equilibrium of the continuation Bayesian game with information structure x_i . A simplified version of the revelation principle in Myerson [33, proposition 2] shows that this continuation equilibrium is also a Bayes correlated equilibrium. Therefore, it induces some distribution in $C_i(p_i^0)$, and the deviation is not profitable for designer i . \square

Remark 3. We make several comments.

1. The proof of Theorem 5 shows that every distribution of actions in $E(p^0)$ can be induced by finite statistical experiments with $|A_i|$ messages for each designer and pure strategies for each designer i and each agent ij . This implies that all rectangular corporation games RG_M with $|M_i^j| \geq |A_i^j|$ for every i, j have a (common) equilibrium in pure strategies.

2. Theorem 5 easily extends to the case where the utilities of the designers depend on the state. Suppose that the utility function of designer i is

$$u_i(a_1, \dots, a_n; \theta_1, \dots, \theta_n).$$

Denote by $\mathcal{B}_i(p_i^0)$ the set of BCEs for corporation i given the prior p_i^0 , and consider the game where each designer i chooses a BCE $b_i \in \mathcal{B}_i(p_i^0)$. Given a strategy profile (b_1, \dots, b_n) in this game, the expected payoff of designer i is

$$U_i(b_1, \dots, b_n) = \sum_{\theta \in \Theta} p^0(\theta) u_i(b_1(\theta_1), \dots, b_n(\theta_n); \theta).$$

The normal-form game $(\mathcal{B}_i(p_i^0), U_i)_i$ admits a Nash equilibrium. Indeed, the set $\mathcal{B}_i(p_i^0)$ is nonempty, convex, and compact, and the mapping $b_i \mapsto U_i(b_i, b_{-i})$ is linear, so Nash's theorem applies. The rest of the proof of Theorem 5 extends directly.

3. Theorem 5 also extends to correlated priors. We can use the same transformation as in Section 3 from a correlated prior to an independent one and set

$$\hat{v}_i^j(a_i, \theta_i) = \frac{p^0(\theta)}{\hat{p}^0(\theta)} v_i^j(a_i, \theta_i) = \frac{p^0(\theta_{-i} | \theta_i)}{\hat{p}^0(\theta_{-i})} v_i^j(a_i, \theta_i).$$

Since agent ij in corporation i does not receive messages from corporations $-i$, her conditional probability on θ_{-i} will remain $p^0(\theta_{-i} | \theta_i)$, even after observing message m_i . Therefore,

$$\sum_{\theta_{-i}} \Pr(\theta_{-i} | \theta_i, m_i^j) v_i^j(a_i, \theta_i) = \sum_{\theta_{-i}} p^0(\theta_{-i} | \theta_i) v_i^j(a_i, \theta_i) = \sum_{\theta_{-i}} \hat{p}^0(\theta_{-i}) \hat{v}_i^j(a_i, \theta_i) = v_i^j(a_i, \theta_i).$$

This shows that the transformed utilities also define a corporation game. Note that, for this reasoning to work, it is important that each designer is restricted to sending messages to her corporation only. If designers can send messages to agents outside of their corporations, then the transformed game might not be a corporation game.

4. With a single designer, the equilibrium of $(C(p^0), u)$ is the designer-preferred equilibrium outcome of the information design game. However, an equilibrium outcome of the information design game may not be an equilibrium of $(C_i(p_i^0), u_i)_{i \in N}$, even with a single agent. To see this, consider the single-designer game of Example 2 with $p^0 = \frac{1}{2}$. There is an equilibrium with no information disclosure in which the agent plays action a^1 if $p < \frac{1}{2}$ and action a^2 if $p \geq \frac{1}{2}$, so action a^2 is played with probability 1 on the equilibrium path. However, game $(C(\frac{1}{2}), u)$ has a unique equilibrium outcome: $E(\frac{1}{2}) = \{\frac{1}{2}\delta_{a^1} + \frac{1}{2}\delta_{a^3}\}$. This example is, however, nongeneric in the sense that there is an action (i.e., a^3) that is not essential: there is no belief under which a^3 is the unique optimal action. Similarly, in Example 1, no action is essential for the agent; the unique equilibrium of the auxiliary game $(C(\frac{1}{2}), u)$ is the first best for the designer (the nonrevealing splitting), but the information design game has many other equilibrium outcomes.

5. With several agents in corporations, $E(p^0)$ might be a strict subset of equilibrium outcomes of RG_M , even for generic games. To see this, consider one designer with a corporation of two agents. The two agents play a game with complete information given by the following matrix in all states:

	<i>L</i>	<i>R</i>
<i>T</i>	1,1	0,0
<i>B</i>	0,0	1,1

Suppose that the utility of the designer is 1 if (T, L) is played and 0 otherwise. In this game, $E(p^0)$ is a singleton obtained by selecting the correlated equilibrium that the designer prefers, namely, (T, L) . However, (B, R) is a Bayes-Nash equilibrium of the matrix game for every statistical experiment (here, this is just a correlation device). Thus, there is an equilibrium of the game of information design where the agents choose (B, R) irrespective of experiments and messages. Thus, $E(p^0)$ is a strict subset of the equilibrium outcomes of the information design game. This latter property also holds for a neighborhood of this game. Indeed, for payoff functions of the agents close enough to the aforementioned matrix, (T, L) and (B, R) , are strict Bayes-Nash equilibria in all states, so the same logic applies.

6. The subset $E(p^0)$ coincides with the set of all equilibrium outcomes of RG_M under the following conditions: (a) there is one agent in each corporation; (b) there is no public randomization device; and (c) for each action $a_i \in A_i$, there is a belief $p_i \in \Delta(\Theta_i)$ such that a_i is the unique optimal action of agent i at p_i .

Condition (b) ensures that the actions $(a_i)_i$ are statistically independent of each other. Then, suppose that there is an equilibrium outcome of RG_M that is not in $E(p^0)$. This means that some designer i has a deviation $\tilde{y}_i \in C_i(p_i^0)$ that is profitable by more than some $\varepsilon > 0$. Consider the splitting that induces \tilde{y}_i . From condition (c), it is possible to perturb the posteriors slightly in such a way that they are still induced by a splitting of p^0 , and, at each posterior, there is a single optimal action for the agent. Moreover, we can do this by selecting the optimal action that designer i prefers. For small perturbations, this would be a deviation profitable by $\varepsilon/2$ in the information design game, which is a contradiction.

The single-designer/single-agent problem is a particular case of a rectangular corporation problem. Hence, the previous remark implies that, under condition (c), the information design game (with at least as many messages as actions) has a unique equilibrium payoff for the designer, given by $\text{cav}_p U^*(p^0)$ and which is achievable by pure strategies.

The following example illustrates our characterization in an investment game with two designers and four agents.

Example 5. Consider two designers who represent entrepreneurs who own projects for complementary new technologies (e.g., new hardware and software or new battery and microprocessor). For each entrepreneur, investing in the new technology must be approved by a committee, and the investment is profitable only if the other entrepreneur also invests in the new technology. The characteristics of the technology of entrepreneur i are

uncertain and denoted by $\theta_i \in \Theta_i = \{\underline{\theta}, \bar{\theta}\}$. The characteristics are good ($\theta_i = \bar{\theta}$) with probability $p_i^0 \in (0, 1)$ and bad ($\theta_i = \underline{\theta}$) with probability $1 - p_i^0$.

For each entrepreneur i , the agents $i1$ and $i2$ in her corporation examine the application and collectively decide whether to approve or not approve the investment of entrepreneur i . Based on their information about θ_i , each agent ij votes “yes” (action $a_i^j = 1$) or “no” (action $a_i^j = 0$). They prefer to coordinate on “yes” when the state is $\theta_i = \bar{\theta}$ and on “no” when the state is $\theta_i = \underline{\theta}$. Their payoffs $(v_i^1(\theta_i, a_i^1, a_i^2), v_i^2(\theta_i, a_i^1, a_i^2))$ are given by the following tables, where the lines correspond to the choices of agent $i1$, the columns to the choices of agent $i2$, and the tables to the states for designer i , and $\gamma \in (0, 2)$ parameterizes an agent’s payoff when she makes the right choice alone:

$\theta_i = \bar{\theta}$	1	0
1	2, 2	$\gamma, 0$
0	$0, \gamma$	0, 0

$\theta_i = \underline{\theta}$	1	0
1	0, 0	$0, \gamma$
0	$\gamma, 0$	2, 2

An entrepreneur’s project is approved if and only if at least one of the members of her corporation chooses action 1.¹³ Her payoff is normalized to 0 if she is not investing. If a single entrepreneur invests, then she incurs a cost equal to $c > 0$. If both entrepreneurs invest, then each obtains a net benefit of 1.¹⁴ To restrict the number of cases to study, we assume the prior that the state is good to be low for both entrepreneurs: $p_i^0 < \min\{\frac{1}{3}, \frac{\gamma}{2+\gamma}\}$ for $i = 1, 2$.

Public Information Design. Assume first, as in Section 3, that designers send public messages to all agents. Consider the Bayesian game with symmetric information played by agents in corporation i when agents’ common belief about θ_i is p_i . Under public information and common belief p_i , agents in corporation i play the following p_i -average game:

	1	0
1	$2p_i, 2p_i$	$\gamma p_i, \gamma(1 - p_i)$
0	$\gamma(1 - p_i), \gamma p_i$	$2(1 - p_i), 2(1 - p_i)$


Given p_i , the Nash equilibria in corporation i are as follows:

- (1, 1) is an equilibrium if and only if $p_i \geq \frac{\gamma}{2+\gamma}$;
- (0, 0) is an equilibrium if and only if $p_i \leq \frac{\gamma}{2+\gamma} := \bar{p}$;
- there is a mixed equilibrium in which action 1 is played with probability $\frac{2(1-p_i)-\gamma p}{2-\gamma}$ if and only if $\frac{\gamma}{2+\gamma} \leq p_i \leq \frac{2}{2+\gamma}$.

For example, if we select the continuation equilibrium that maximizes the probability of approval, then the designers’ payoffs as a function of the agent’s beliefs $(p_1, p_2) \in [0, 1]^2$ are given by Figure 4.

Since $p_i^0 < \frac{\gamma}{2+\gamma}$ for $i = 1, 2$, there is an equilibrium in which entrepreneurs never invest in the new technology by choosing a noninformative policy. We now identify the equilibrium that is Pareto-optimal for the entrepreneurs, that is, the equilibrium that maximizes the probability of joint investment. Such an equilibrium is Markovian, and we know from the analysis of Section 3 that we can restrict attention to splittings with at most two posteriors. The public information policy that maximizes the probability of investment for designer i is the splitting of

Figure 4. Payoff of the designers in Example 5 as a function of agents’ common beliefs (p_1, p_2) when the continuation equilibrium that maximizes the probability of approval is selected.



p_i^0 on the posterior 0 with probability $1 - p_i^0 \bar{p}$ and on the posterior \bar{p} with probability $p_i^0 \bar{p}$. Hence, with such an information policy, the probability of investment for entrepreneur i is equal to $p_i^0 \bar{p} = \frac{p_i^0(2+\gamma)}{\gamma}$. Designer i prefers to invest if and only if the other designer invests with probability of at least $\frac{c}{1+c}$, so we must have $\frac{c}{1+c} \leq \frac{p_i^0(2+\gamma)}{\gamma}$ for $i = 1, 2$. Otherwise, in the unique equilibrium outcome of the public information design game, entrepreneurs never invest in the new technology.

Private Information Design. We now study this example with no restriction on information disclosure policies and use the equilibrium characterization for rectangular corporation games based on BCE outcomes. Clearly, as under public information design, there is always an equilibrium with no investment. However, we will show that the conditions to have an equilibrium in which both entrepreneurs invest with positive probability are weaker than those under public information design whenever $\gamma > 1$. In addition, under these conditions, the probability of joint investment is strictly higher, and, therefore, entrepreneurs are strictly better off when they can use private information policies.

Since a designer's payoff is symmetric with respect to her two agents and since the two agents are symmetric, we can focus on BCE that are symmetric between the two agents. A direct statistical experiment $x_i : \{\bar{\theta}, \underline{\theta}\} \rightarrow \Delta(\{1, 0\}^2)$ is denoted by the following table:

$\theta_i = \bar{\theta}$	1	0
1	\bar{x}_i^{11}	\bar{x}_i^{10}
0	\bar{x}_i^{10}	\bar{x}_i^{00}

$\theta_i = \underline{\theta}$	1	0
1	\underline{x}_i^{11}	\underline{x}_i^{10}
0	\underline{x}_i^{10}	\underline{x}_i^{00}

where $\bar{x}_i^{11} + 2\bar{x}_i^{10} + \bar{x}_i^{00} = \underline{x}_i^{11} + 2\underline{x}_i^{10} + \underline{x}_i^{00} = 1$. The direct statistical experiment x_i is a BCE for corporation i if and only if the following obedient constraints are satisfied:

$$\begin{aligned} p_i^0(2\bar{x}_i^{11} + \gamma\bar{x}_i^{10}) &\geq (1 - p_i^0)(\gamma\underline{x}_i^{11} + 2\underline{x}_i^{10}), \\ (1 - p_i^0)(\gamma\underline{x}_i^{10} + 2\underline{x}_i^{00}) &\geq p_i^0(2\bar{x}_i^{10} + \gamma\bar{x}_i^{00}). \end{aligned}$$

The set $C_i(p_i^0)$ of strategies for designer i in the auxiliary normal-form game is the set of direct statistical experiments satisfying these two inequalities. When feasible, the equilibrium that is Pareto-optimal for designers is the equilibrium that maximizes the probability of investment. That is, it maximizes the probability $p_i^0(1 - \bar{x}_i^{00}) + (1 - p_i^0)(1 - \underline{x}_i^{00})$ that at least one agent chooses action 1 in each corporation. The solution of this program is

$$\bar{x}_i^{00} = 0, \quad \underline{x}_i^{00} = \begin{cases} 1 - \frac{2p^0}{(1-p^0)\gamma} & \text{if } \gamma < 1 \\ 1 - \frac{2p^0}{(1-p^0)} & \text{if } \gamma > 1, \end{cases} \quad \text{and} \quad \underline{x}_i^{11} = \begin{cases} 1 - \underline{x}_i^{00} & \text{if } \gamma < 1 \\ 0 & \text{if } \gamma > 1. \end{cases}$$

Hence, to maximize the probability of investment when $\gamma < 1$, the actions of agents in a corporation must be perfectly correlated ($\bar{x}_i^{10} = \underline{x}_i^{10} = 0$); thus, the information disclosure policy is public. The maximal probability of investment in corporation i is

$$p_i^0(1 - \bar{x}_i^{00}) + (1 - p_i^0)(1 - \underline{x}_i^{00}) = p_i^0 + (1 - p_i^0) \frac{2p^0}{(1-p^0)\gamma} = \frac{p^0(2+\gamma)}{\gamma}.$$

As under public information design, this corresponds to an equilibrium between the two designers if and only if $\frac{c}{1+c} \leq \frac{p_i^0(2+\gamma)}{\gamma}$, $i = 1, 2$.

By contrast, when $\gamma > 1$, to maximize the probability of investment, the actions of agents in a corporation should not be perfectly correlated when the state is $\theta_i = \underline{\theta}$. We have $\underline{x}_i^{10} = \frac{1-\underline{x}_i^{00}}{2} > 0$, so the information disclosure policy is not public. The maximal probability of investment in corporation i is now

$$p_i^0(1 - \bar{x}_i^{00}) + (1 - p_i^0)(1 - \underline{x}_i^{00}) = p_i^0 + (1 - p_i^0) \frac{2p^0}{(1-p^0)} = 3p_i^0 > \frac{p^0(2+\gamma)}{\gamma},$$

which is higher than the probability under public information design. Such an equilibrium between the two designers exists if and only if $\frac{c}{1+c} \leq 3p_i^0$, $i = 1, 2$, which is a weaker condition than that under public information. In particular, when $\frac{p_i^0(2+\gamma)}{\gamma} < \frac{c}{1+c} < 3p_i^0$, the unique equilibrium with public information is the no-investment equilibrium, while there exists an equilibrium with private information policies in which there is investment with probability $3p_i^0$ in each corporation i .

Acknowledgments

The authors thank the associate editor, the anonymous referees, Philippe Bich, Françoise Forges, Keiichi Kawai, Rida Laraki, Francesco Nava, Alessandro Pavan, Jérôme Renault, Ludovic Renou, Vasiliki Skreta, Takuro Yamashita, and seminar participants at Cambridge, EUI Florence, HEC, LSE, Northwestern, PSE, TSE, University of Besançon, University of Bristol, University of Copenhagen, University of Manchester, University of Warwick, the Game Theory Seminar in Paris (IHP), the workshop “Equilibria in Games: Existence, Selection and Dynamics” in Santiago, the Paris symposium in games, and the PGMO days in EDF Laboratory Saclay for useful comments and suggestions. This paper previously circulated under the title “Competitive Information Design.”

Appendix. C3B Bayesian Games

In this appendix, we provide properties of Bayes-Nash equilibria for a class of Bayesian games that includes all possible continuation Bayesian games in our information design environment. There is a finite set of players $K = \{1, \dots, k\}$ and a state space $\Theta \times \prod_{j \in K} M_j$. In state $(\theta, m_1, \dots, m_k)$, player $j \in K$ is privately informed about $m_j \in M_j$. All the sets Θ, M_1, \dots, M_k are assumed to be (finite or) countable. Denote $M = \prod_{j \in K} M_j$, and let $P \in \Delta(\Theta \times M)$ be the common prior. Each player j has a set of actions A_j and a bounded payoff function $v_j : \prod_{j \in K} A_j \times \Theta \times M \rightarrow [0, 1]$. We assume that each set of actions is compact metric and that each payoff function $v_j(\cdot, \theta, m)$ is continuous with respect to the action profile for each (θ, m) .

This game is denoted $G(P)$, and this class of games is called *countable-compact-continuous-bounded* (C3B) to indicate countable states, compact actions, and continuous and bounded payoffs.

Lemma A.1. The set of Bayes-Nash equilibria of a C3B game $G(P)$ is nonempty and compact

This result can be deduced from more general results existing in the literature (see, e.g., Milgrom and Weber [31]). We provide the following elementary proof for the sake of completeness.

Proof. The set of pure strategies of player j in $G(P)$ is $S_j = A_j^{M_j}$, which is metric compact, since it is a countable product of metric compacts (Aliprantis and Border [2, theorem 2.61, p. 52 and theorem 3.36, p. 89]). Given a strategy profile $s \in \prod_{i \in N} S_i$, the payoff of player j is

$$V_j(s) = \sum_{(\theta, m) \in \Theta \times M} P(\theta, m) v_j(s_j(m_j), (s_i(m_i))_{i \neq j}, \theta, m),$$

which is continuous with respect to s . To see this, consider a sequence s^r of strategy profiles that converges to profile s as $r \rightarrow \infty$. Convergence in a countable product of compact sets is equivalent to pointwise convergence, so we have $s^r(m) \rightarrow s(m)$ for all $m \in \prod_{j \in K} M_j$. Since $\Theta \times M$ is countable, the family of numbers $(P(\theta, m))_{(\theta, m) \in \Theta \times M}$ is summable with sum 1. Thus, for each $\varepsilon > 0$, there exists a finite set $B \subseteq \Theta \times M$ such that $P(B) \geq 1 - \frac{\varepsilon}{2}$. Since payoffs are bounded (by 1), we have

$$|V_j(s) - V_j(s^r)| \leq \frac{\varepsilon}{2} + \sum_{(\theta, m) \in B} P(\theta, m) |v_j(s_j(m_j), (s_i(m_i))_{i \neq j}, \theta, m) - v_j(s_j^r(m_j), (s_i^r(m_i))_{i \neq j}, \theta, m)|.$$

The second term is a finite sum that tends to 0 as $r \rightarrow \infty$; thus, $|V_j(s) - V_j(s^r)| \leq \varepsilon$ for large r .

The (normal-form) game $G(P)$ thus has compact sets of pure strategies and continuous payoff functions. By Fan-Glicksberg’s theorem, the game admits an equilibrium in mixed strategies. It is then straightforward that the set of equilibria is compact: by continuity, the limit of a sequence of equilibria is an equilibrium. \square

Consider now the equilibrium correspondence $\mathcal{E} : \Delta(\Theta \times M) \rightarrow \Delta(A^M)$, where $M = \prod_{j \in K} M_j$ denotes the set of message profiles, $A = \prod_{j \in K} A_j$ is the set of action profiles, and, for each $P \in \Delta(\Theta \times M)$, $\mathcal{E}(P)$ is the set of mixed Bayes-Nash equilibria of $G(P)$. Note that by Kuhn’s theorem, $\mathcal{E}(P)$ can equivalently be viewed as the set of Bayes-Nash equilibria in behavior strategies and thus as a subset of $\Delta(A)^M$.

Let $\mathcal{P} \subset \Delta(\Theta \times M)$ be a subset of probability distributions and denote by $\mathcal{E}_{\mathcal{P}}$ the restriction of \mathcal{E} to \mathcal{P} . Since states are countable, a set $\mathcal{P} \subset \Delta(\Theta \times M)$ is compact for the $\|\cdot\|_1$ norm if and only if it is weakly compact (Aliprantis and Border [2, theorem 16.24, p. 537]), and it is henceforth simply called *compact*.

Lemma A.2. The Bayes-Nash equilibrium correspondence \mathcal{E} is Borel-measurable. If \mathcal{P} is compact, then $\mathcal{E}_{\mathcal{P}}$ admits measurable selections and is upper hemicontinuous

Proof. The correspondence \mathcal{E} is clearly Borel-measurable since equilibria are defined by weak inequalities and payoffs are continuous. From the measurable selection theorem (Aliprantis and Border [2, theorem 18.13, p. 600]), \mathcal{E} admits

measurable selections when its domain and image are metric compacts. We consider a compact domain \mathcal{P} by assumption. The countable product A^M of metric compacts is metric compact, and so is $\Delta(A^M)$ with the weak-* topology. To prove that $\mathcal{E}_{\mathcal{P}}$ is upper hemicontinuous, we check that it has a closed graph; that is, if $P^r \rightarrow P$, $s^r \rightarrow s$, and $s^r \in \mathcal{E}_{\mathcal{P}}(P^r)$, then $s \in \mathcal{E}_{\mathcal{P}}(P)$. For this, it is enough to show that, for any sequence $P^r \rightarrow P$ in \mathcal{P} and any sequence of strategy profiles $s^r \rightarrow s$,

$$V_j^r(s^r) := \sum_{(\theta, m) \in \Theta \times M} P^r(\theta, m) v_j(s^r(m), \theta, m) \xrightarrow{r \rightarrow \infty} V_j(s) := \sum_{(\theta, m) \in \Theta \times M} P(\theta, m) v_j(s(m), \theta, m).$$

We have

$$\begin{aligned} |V_j^r(s^r) - V_j(s)| &\leq \sum_{(\theta, m) \in \Theta \times M} |P^r(\theta, m) - P(\theta, m)| v_j(s^r(m), \theta, m) \\ &\quad + \left| \sum_{(\theta, m) \in \Theta \times M} P(\theta, m) (v_j(s^r(m), \theta, m) - v_j(s(m), \theta, m)) \right| \\ &\leq \|P^r - P\|_1 + |V_j(s^r) - V_j(s)|. \end{aligned}$$

This concludes the proof since $P^r \rightarrow P$, $s^r \rightarrow s$, and $V_j(s)$ is continuous in s . \square

Application to Subgames of Information Design Games

In our information design environment, the set of players in a continuation Bayesian game $G_X(x)$ is the set K of agents, A_j is the finite set of actions of agent j , and $M_j = \mathbb{N}^N$ is the set of messages from the designers to agent j . The utility of each agent j does not depend on $m \in M$ and is therefore bounded. For $x = (x_i)_{i \in N} \in X = \prod_{i \in N} X_i$, let $P_x(\theta, m) = p(\theta) \prod_{i \in N} x_i(m_i | \theta_i)$ and $\mathcal{P} = \{P_x : x \in X\}$. Then, $G_X(x) := G(P_x)$ is C3B and \mathcal{P} is compact. Hence, Lemmas A.1 and A.2 imply that the equilibrium correspondence $\mathcal{E}_X(\cdot)$ has nonempty compact values and is upper hemicontinuous, as required in the proof of Theorem 1.

Endnotes

¹In the appendix, we show this property more generally for Bayesian games with countable states, compact actions, and continuous and bounded payoffs.

²This methodology is similar to the one used in the literature on the existence of subgame-perfect equilibria in continuous games. Of particular interest here is the paper of Harris et al. [17] (see also Mariotti [27]), who prove the existence of subgame-perfect equilibria in continuous games with almost perfect information, assuming, as we do, that there is a public randomization device. However, our general model does not have almost perfect information since agents may receive private messages from the designers, and these results do not apply.

³Throughout the paper, for every compact set S , $\Delta(S)$ denotes the set of Borel probability measures over S .

⁴Most of our results generalize to infinite action sets (see the appendix).

⁵Considering experiments with discrete support ensures that the continuation game played by the agents after observing the profile of experiments is well defined and admits a compact set of Bayes-Nash equilibria (see the appendix).

⁶See Moriya and Yamashita [32] for an explicit example involving an infinite number of messages with one designer, two agents, two states, and two actions.

⁷For any function $\varphi : \Delta(\Theta_i) \rightarrow \mathbb{R}$, $\sup_{\mu_i \in S_i(p_i)} \int_{p_i} \varphi(p_i) d\mu_i(p_i)$ is the (pointwise) smallest concave function $\text{cav}\varphi(p_i)$ on $\Delta(\Theta_i)$ satisfying $\text{cav}\varphi(p_i) \geq \varphi(p_i)$ for all $p_i \in \Delta(\Theta_i)$, and the supremum is achieved with $|\Theta_i|$ posteriors. See, e.g., Aumann et al. [7] or Rockafellar [36, corollary 17.1.5, p. 157] for more technical details.

⁸This equilibrium refinement should be more precisely called a *public correlated Markovian equilibrium* whenever public correlation is needed to generate $y(p) \in Y(p)$. With a slight abuse of words, we simply call this refinement a *Markovian equilibrium*.

⁹In particular, Theorem 4 sheds some light on the results from Boleslavsky and Cotton [10]. In their setup, two schools compete to persuade an employer to hire their student by designing grading rules. The authors show that in equilibrium the posterior belief generated by the grading rule is uniform, which implicitly requires that schools use grading policies with a continuum of grades. Theorem 4 implies that the equilibrium can be reinterpreted as a mixed-strategy equilibrium over grading rules that involve only two grades.

¹⁰Assuming $\alpha = 1/2$, Albrecht [1] shows that there is no pure-strategy equilibrium for any finite number of messages.

¹¹A corresponding infinite statistical experiment $x_2 : \Theta_2 \rightarrow [0, 1]$ draws a message in $[0, 1]$ from the conditional density functions $f_{x_2}(m|\theta_2 = 0) = 2m$ and $f_{x_2}(m|\theta_2 = 1) = 2(1 - m)$.

¹²Given any straight line in $[0, 1] \times [0, 1]$, one can find a corresponding decision problem with two actions. It is then possible to combine such problems with additively separable payoffs to generate any polyhedral region.

¹³A similar analysis can be performed by requiring unanimity from the agents in the corporation. In this case, it can be shown that the conditions for the existence of an equilibrium with investment would be the same under public and private information design. We therefore assume that one “yes” vote is enough for investment to illustrate the difference between public and private information design.

¹⁴Similar information design problems, but with a single designer only, have been studied by Bergemann and Morris [9] and Taneva [42].

References

- [1] Albrecht B C (2017) Political persuasion. Accessed May 18, 2021, <https://briancalbrecht.github.io/albrecht-political-persuasion.pdf>.
- [2] Aliprantis C, Border K (2006) *Infinite Dimensional Analysis* (Springer, Berlin).
- [3] Attar A, Campioni E, Piaser G (2018) On competing mechanisms under exclusive competition. *Games Econ. Behav.* 111:1–15.
- [4] Au PH, Kawai K (2019) Competitive disclosure of correlated information. *Econ. Theory*, ePub ahead of print January 9, <https://doi.org/10.1007/s00199-018-01171-7>.
- [5] Au PH, Kawai K (2020) Competitive information disclosure by multiple senders. *Games Econ. Behav.* 119:56–78.
- [6] Aumann RJ, Maschler MB (1967) Repeated games with incomplete information: A survey of recent results. Report, U.S. Arms Control and Disarmament Agency, Washington, DC.
- [7] Aumann RJ, Maschler MB, Stearns R (1995) *Repeated Games of Incomplete Information* (MIT Press, Cambridge, MA).
- [8] Bergemann D, Morris S (2016) Bayes correlated equilibrium and the comparison of information structures in games. *Theoret. Econ.* 11:487–522.
- [9] Bergemann D, Morris S (2019) Information design: A unified perspective. *J. Econom. Lit.* 57(1):44–95.
- [10] Boleslavsky R, Cotton C (2015) Grading standards and education quality. *Amer. Econom. J. Microeconom.* 7(2):248–279.
- [11] Carlsson H, van Damme E (1993) Global games and equilibrium selection. *Econometrica* 61(5):989–1018.
- [12] Carmona G, Fajardo J (2009) Existence of equilibrium in common agency games with adverse selection. *Games Econ. Behav.* 66(2):749–760.
- [13] Eliaz K, Forges F (2015) Information disclosure to Cournot duopolists. *Econ. Lett.* 126:167–170.
- [14] Eliaz K, Serrano R (2014) Sending information to interactive receivers playing a generalized prisoners' dilemma. *Internat. J. Game Theory* 43:245–267.
- [15] Gentzkow M, Kamenica E (2017) Competition in persuasion. *Rev. Econom. Stud.* 84:300–322.
- [16] Han S (2007) Strongly robust equilibrium and competing-mechanism games. *J. Econom. Theory* 137(1):610–626.
- [17] Harris C, Reny P, Robson A (1995) The existence of subgame-perfect equilibrium in continuous games with almost perfect information: A case for public randomization. *Econometrica* 63(3):507–544.
- [18] Jehiel P (2014) On transparency in organizations. *Rev. Econom. Stud.* 82(2):736–761.
- [19] Kamenica E, Gentzkow M (2011) Bayesian persuasion. *Amer. Econom. Rev.* 101(6):2590–2615.
- [20] Koessler F, Laclau M, Renault J, Tomala T (2019) Long information design. *mimeo*.
- [21] Laraki R (2001) The splitting game and applications. *Internat. J. Game Theory* 30:359–376.
- [22] Laraki R (2001) Variational inequalities, system of functional equations and incomplete information repeated games. *SIAM J. Control Optim.* 40(2):516–524.
- [23] Laraki R, Renault J (2020) Acyclic gambling games. *Math. Oper. Res.* 45(4):1237–1257.
- [24] Levy G, de Barra IM, Razin R (2017) Persuasion with correlation neglect. *mimeo*.
- [25] Lewis TR, Sappington DEM (1994) Supplying information to facilitate price discrimination. *Internat. Econom. Rev.* 35(2):309–327.
- [26] Li F, Norman P (2018) On Bayesian persuasion with multiple senders. *Econ. Lett.* 170:66–70.
- [27] Mariotti T (2000) Subgame-perfect equilibrium outcomes in continuous games of almost perfect information. *J. Math. Econom.* 34(1):99–128.
- [28] Martimort D, Stole L (2002) The revelation and delegation principles in common agency games. *Econometrica* 70(4):1659–1673.
- [29] Mathevet L, Perego J, Taneva I (2020) On information design in games. *J. Political Econom.* 128(4):1370–1404.
- [30] Mertens J-F, Sorin S, Zamir S (2015) *Repeated Games* (Cambridge University Press, Cambridge, UK).
- [31] Milgrom PR, Weber RJ (1985) Distributional strategies for games with incomplete information. *Math. Oper. Res.* 10(4):619–632.
- [32] Moriya F, Yamashita T (2020) Asymmetric-information allocation to avoid coordination failure. *J. Econom. Management Strategy* 29(1):173–186.
- [33] Myerson RB (1982) Optimal coordination mechanisms in generalized principal-agent problems. *J. Math. Econom.* 10(1):67–81.
- [34] Oliu-Barton M (2018) The splitting game: Uniform value and optimal strategies. *Dynam. Games Appl.* 8:157–179.
- [35] Peters M (2001) Common agency and the revelation principle. *Econometrica* 69(5):1349–1372.
- [36] Rockafellar R (1970) *Convex Analysis* (Princeton University Press, Princeton, NJ).
- [37] Rubinstein A (1989) The electronic mail game: Strategic behavior under 'almost common knowledge.' *Amer. Econom. Rev.* 79(3):385–391.
- [38] Simon LK, Zame WR (1990) Discontinuous games and endogenous sharing rules. *Econometrica* 58(4):861–872.
- [39] Sion M, Wolfe P (1957) On a game without a value. Dresner M, Tucker AW, Wolfe P, eds. *Contributions to the Theory of Games*, vol. 3 (Princeton University Press, Princeton, NJ), 299–306.
- [40] Sorin S (2002) *A First Course on Zero-Sum Repeated Games* (Springer, Berlin).
- [41] Szentes B (2014) Contractible contracts in common agency problems. *Rev. Econ. Stud.* 82(1):391–422.
- [42] Taneva I (2019) Information design. *Amer. Econom. J. Microeconom.* 11(4):151–185.