

Long Information Design^{*}

Frederic Koessler[†] Marie Laclau[‡] Jérôme Renault[§] Tristan Tomala[¶]

July 14, 2021

Abstract

We analyze information design games between two designers with opposite preferences and a single agent. Before the agent makes a decision, designers repeatedly disclose public information about persistent state parameters. Disclosure continues until no designer wishes to reveal further information. We consider environments with general constraints on feasible information disclosure policies. Our main results characterize equilibrium payoffs and strategies of this long information design game and compare them with the equilibrium outcomes of games where designers move only at a single predetermined period. When information disclosure policies are unconstrained, we show that at equilibrium in the long game, information is revealed right away in a single period; otherwise, the number of periods in which information is disclosed might be unbounded. As an application, we study a competition in product demonstration and show that more information is revealed if each designer could disclose information at a predetermined period. The format that provides the buyer with most information is the sequential game where the last mover is the ex-ante favorite seller.

KEYWORDS: Bayesian persuasion; concavification; convexification; information design; Mertens-Zamir solution; product demonstration; splitting games; statistical experiments; stochastic games.

JEL CLASSIFICATION: C72; D82.

^{*}We thank James Best, Françoise Forges, Marina Halac, Rida Laraki, Laurent Miclo, Alessandro Pavan, Xavier Venel, anonymous referees, and seminar participants at HEC Paris (May 2019), University of Southern Denmark (May 2019, Odense), CREST (June 2019, Palaiseau), University of Strasbourg (June 2019), Frontiers in Design (June 2019, London), Information Design and Splitting Games (June 2019, Paris), the Lisbon Meetings in Game Theory and Applications (November 2019, Lisbon), Aix-Marseille School of Economics (November 2019), Bar-Ilan and Technion University (May 2020), One World Mathematical Game Theory Seminar (June 2020), and Journées SMAI MODE 2020 (September 2020) for useful comments and suggestions.

[†]Paris School of Economics–CNRS, 48 boulevard Jourdan, 75014 Paris, France; frederic.koessler@psemail.eu. Frederic Koessler acknowledges the support of the ANR (Investissements d’Avenir program ANR-17-EURE-001 and StratCom ANR-19-CE26-0010-01).

[‡]HEC Paris and GREGHEC–CNRS, 1 rue de la Libération, 78351 Jouy-en-Josas, France; laclau@hec.fr. Marie Laclau gratefully acknowledges the support of the ANR through the program Investissements d’Avenir (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047) and under grant ANR CIGNE (ANR-15-CE38-0007-01).

[§]TSE (Université Toulouse 1 Capitole), 21 allée de Brienne, 31000 Toulouse, France; jerome.renault@tse-fr.eu. Jérôme Renault gratefully acknowledges funding from ANR-3IA Artificial and Natural Intelligence Toulouse Institute, grant ANR-17-EUR-0010 (Investissements d’Avenir program) and ANR MaSDOL.

[¶]HEC Paris and GREGHEC, 1 rue de la Libération, 78351 Jouy-en-Josas, France; tomala@hec.fr. Tristan Tomala gratefully acknowledges the support of the HEC foundation and ANR/Investissements d’Avenir under grant ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047.

1 Introduction

In many environments, economic agents spend time acquiring information before making decisions, and information often comes from multiple and interested sources. This paper analyzes strategic interactions between two information providers (called the *designers*) who compete for influencing the action of a decision-maker (called the *agent*). The designers have opposite preferences and control the access to independent pieces of information. The first key feature of our model is that designers are able to disclose information in multiple stages, as long as they want. Specifically, at each stage, designers disclose information publicly and choose new disclosure policies at the next stage. Hence, in each stage, a designer can react to the information disclosed by himself and by the other designer. Once both designers have finished releasing information, the agent chooses an action. The payoffs of the designers and of the agent depend on the realized state parameters and on the action taken. The second feature of our model is that it introduces general technological restrictions on available information disclosure policies. The model encompasses the standard information design setting in which a designer can choose any public information policy, but it also covers more realistic environments in which a designer may be constrained to choosing information policies from an exogenous subset. For example, in most real world information design problems, a designer may be constrained to choose deterministic information structures, may run experimental tests with unavoidable false-positive or false-negative results, or may rely on imperfectly reliable experts and reviewers.

As in the literature on Bayesian persuasion initiated by Kamenica and Gentzkow (2011), the term “information designer” refers to a player who is uninformed about a state parameter but is able to choose an information disclosure policy (a statistical experiment) about this parameter to modify agents’ information. Contrary to cheap-talk communication, there is no issue of credibility in information disclosure: the statistical experiment is publicly observable and verifiable. In addition, since an information designer is not privately informed about the state, the choice of the disclosure policy has no signaling effect. In practice, information designers can represent competing sellers who release information about a new product to influence buyers’ valuations. For example, sellers can offer free samples or trial periods to social media influencers, control the access to and content of online, press or magazine reviews, organize product testing and trade fairs, or make announcements about future product development. Alternatively, information designers could be lobbyists who control the informativeness of some studies to influence a policymaker. In such contests, the value of releasing additional information depends on the information publicly revealed by the competing designer.

The aim of the paper is to study designers’ equilibrium information disclosure policies and payoffs in multistage information design problems in which there is a priori no deadline, i.e., no bound on the number of disclosure stages. We refer to the model with no deadline as “long” information design even though, under some conditions on the primitives of the game (see Proposition 2), information disclosure ends very quickly in equilibrium. We compare equilibrium disclosure strategies and players’ welfare of such long information design game, with those of information design games in which each designer can disclose information only at a predetermined stage, either simultaneously (the one-stage simultaneous-move game) or sequentially (the two-stage sequential-move game). Beyond the theoretical interest, our analysis is motivated by the question of how a decision maker should choose to acquire

information from competing designers. For example, a committee could decide in which order (i) evidences should be presented, (ii) experts and reviewers should be consulted and whether to set a deadline to communication. Similarly, a buyer could decide how to acquire information from competing sellers by committing to a purchasing period or by choosing in which order to visit them.

Contributions The main result of this paper is a characterization of equilibrium payoffs and strategies of long information design games in which designers have opposite preferences, noninformative policies are always feasible, and feasible policies are closed under iteration (that is, a distribution of beliefs that can be obtained by a two-step combination of experiments can also be obtained by a single feasible experiment).¹ We show that the long information design game admits a stationary equilibrium, i.e., such that designers’ strategies depend only on the current beliefs. In addition, if there is no constraint on the set of available information disclosure policies, there exists an equilibrium in which information is disclosed at the first stage only. That is, along the equilibrium path, at most one designer discloses some information at the first stage, information disclosure policies are uninformative thereafter, and the agent therefore takes his decision at the end of the second stage. If information disclosure policies are constrained, we provide an example in which, in every equilibrium, the number of disclosure stages is unbounded, even under the maintained assumption that feasible policies are closed under iteration (see Section 3.4.2).

As in the usual case of one designer of Kamenica and Gentzkow (2011), the value of the game (the equilibrium payoff of Designer 1) is derived from the expected payoff of Designer 1, denoted by $u(p_1, p_2)$, as a function of the beliefs (p_1, p_2) , where p_i is the public belief about the information controlled by Designer i . Equilibrium strategies can be directly backed out from the characterization of the value of the game. First, assume that only Designer 1 is active (e.g., because Designer 2 is constrained to be silent); then, according to Kamenica and Gentzkow (2011), Designer 1 would “concavify” the payoff function with respect to p_1 , obtaining what we denote by $\text{cav}_{p_1} u(p_1, p_2)$. Similarly, if only Designer 2 were active, we would obtain the “convexification” with respect to p_2 , which we denote by $\text{vex}_{p_2} u(p_1, p_2)$.

Consider now the two-stage sequential information design game in which Designer 1 can disclose information in the first stage only, and Designer 2 in the second stage. By backward induction, the value of the game is $\text{cav}_{p_1} \text{vex}_{p_2} u(p_1, p_2)$. Similarly, if Designer 2 moves first and Designer 1 last, then the value of the game is $\text{vex}_{p_2} \text{cav}_{p_1} u(p_1, p_2)$. In the one-stage simultaneous move game in which designers move simultaneously in a single stage, the value, which we call *the splitting game value*, is in between $\text{cav}_{p_1} \text{vex}_{p_2} u(p_1, p_2)$ and $\text{vex}_{p_2} \text{cav}_{p_1} u(p_1, p_2)$. In the long information design game, the values and equilibrium strategies are the same for games with simultaneous or alternating moves, and this value is the unique function $v(p_1, p_2)$, which we call *the Mertens-Zamir function*, that satisfies the following system:

$$v(p_1, p_2) = \text{cav}_{p_1} \min\{(u(p_1, p_2), v(p_1, p_2))\} = \text{vex}_{p_2} \max\{(u(p_1, p_2), v(p_1, p_2))\}.$$

¹Under this latter condition, which holds in standard Bayesian persuasion models, multiple stages of information disclosure is irrelevant if there is only one designer. However, with more than one designer, multiple stages of information disclosure become relevant because they allow each designer to react to the information released by the competitor. In a companion paper (Koessler, Laclau, Renault, and Tomala, 2021), we study particular cases where the assumption that feasible policies are closed under iteration can be relaxed.

This system is key to the study of discounted zero-sum repeated games with incomplete information on both sides (Mertens and Zamir, 1971, 1977) and of zero-sum dynamic gambling games (Laraki and Renault, 2020). It allows simple optimal strategies to be derived directly: Designer 1 plays non-informatively if $u(p_1, p_2) \geq v(p_1, p_2)$ and Designer 1 concavifies $\min(u, v)$ if $u(p_1, p_2) < v(p_1, p_2)$; Designer 2 plays noninformatively if $u(p_1, p_2) \leq v(p_1, p_2)$ and Designer 2 convexifies $\max(u, v)$ if $u(p_1, p_2) > v(p_1, p_2)$.

An intuition for this result is as follows. When a Mertens-Zamir function exists, for any profile of current posterior beliefs (p_1, p_2) , there exists an information disclosure policy for Designer 1 such that, if the game would end after his move, then his payoff would be greater than or equal to the value of the Mertens-Zamir function at (p_1, p_2) . Similarly, there exists an information disclosure policy for Designer 2 such that, if the game would end after his move, then the continuation payoff of Designer 1 would be less than or equal to the value of the Mertens-Zamir function at (p_1, p_2) . Proceeding backward, this implies that at the prior beliefs, Designer 1 has an information disclosure strategy that gives him a continuation payoff greater than or equal to the value of the Mertens-Zamir function at the priors, and Designer 2 has an information disclosure strategy that gives Designer 1 a continuation payoff less than or equal to the value of the Mertens-Zamir function at the priors. Thus, the Mertens-Zamir value at the priors can be guaranteed by both designers; since the game is zero-sum between the two designers, it is the unique equilibrium payoff of the game.

The value of the long information design game is between $\text{cav}_{p_1} \text{vex}_{p_2} u(p_1, p_2)$ and $\text{vex}_{p_2} \text{cav}_{p_1} u(p_1, p_2)$, so if these two quantities are equal, all variants of the multistage information design games have the same value. Indeed in such a situation, the Mertens-Zamir value at the priors can be guaranteed by Designer 1 even if he moves first and only once, and it can be defended by Designer 2 even if he moves first and only once, so the order of moves and the time horizon are irrelevant. Hence in this case, the information revealed to the agent in equilibrium is the same for any information disclosure protocol. When $\text{cav}_{p_1} \text{vex}_{p_2} u(p_1, p_2) = \text{vex}_{p_2} \text{cav}_{p_1} u(p_1, p_2)$, it doesn't matter how the agent acquires information from competing designers.²

We provide examples and economic applications in which $\text{cav}_{p_1} \text{vex}_{p_2} u$, $\text{vex}_{p_2} \text{cav}_{p_1} u$, the splitting game value and the Mertens-Zamir function are all different. In such situations, the protocol of information disclosure determines the information revealed to the agent in equilibrium. Hence, our equilibrium characterizations allow to derive directly the agent's preference over disclosure protocols. This comparison is especially relevant if the agent can commit to an information acquisition strategy, or if the disclosure protocol can be regulated.

Application to a Product Demonstration by Two Sellers We apply our results and methodology to a stylized model of competitive and public product demonstration in which two sellers disclose information about their respective products to a representative buyer. All players are initially uncertain about the buyer's valuation (or "match") for each product.³ The buyer decides to buy from the seller

²The way the agent acquires information when the two designers have access to the same information (i.e., states are perfectly correlated) and all information disclosure policies are feasible is also irrelevant because full information disclosure is always an equilibrium in this case (see Section 5.3 where we discuss the general correlated case).

³For a literature review of match advertising see, for example, Renault (2015).

for whom the expected match conditional on the public information is highest (see Boleslavsky and Cotton, 2015 for the analysis of the one-stage simultaneous information design game). We show that regardless of the prior expected match values, sellers' equilibrium strategies are always less informative in the long information design game than the one-stage simultaneous and sequential information design games. In the long information design game, only one designer discloses information, and the expected payoff of the buyer is the same as without information disclosure. The best disclosure protocol for the buyer is the two-stage sequential game in which the seller with the highest ex-ante value moves last.

The intuition for this result is as follows. When seller i knows that he is the last mover, he has an incentive to choose an information policy that has the highest chance to strictly improve his expected match value compared to his competitor $j \neq i$. In turn, this incentivizes seller j to reveal additional match information, especially if seller j has a lower ex-ante expected match. On the contrary, in the long information design game, where each seller has the opportunity to react to the information policy of his competitor by disclosing additional information later, no seller has an incentive to strictly improve his expected match compared to his competitor. Indeed, if seller i chooses an information policy that induces a strictly higher expected match than j with positive probability, then seller j is able to react with an another information disclosure policy that improves his expected match value compared to seller i 's expected match value with positive probability. Such a deviation by seller i is actually strictly detrimental because it provides an informational advantage to seller j : if seller i 's information policy generates a bad signal, then seller j takes over the market without revealing any information. Formally, in this application, the Mertens-Zamir function is strictly concave in the prior about seller 1 and strictly convex in the prior about seller 2.

Related Literature The methodology and results of Bayesian persuasion (Kamenica and Gentzkow, 2011) and information design (e.g., Bergemann and Morris, 2016b,a, Mathevet, Perego, and Taneva, 2020, and Taneva, 2019) are deeply related to repeated games with incomplete information on one side (Aumann and Maschler, 1966, 1967, Aumann, Maschler, and Stearns, 1995) and to the literature on generalized principal-agent problems, correlated and communication equilibria (Aumann, 1974, Myerson, 1982, 1986, Forges, 1986, 1993). See, e.g., the literature reviews in Kamenica (2018), Bergemann and Morris (2019) and Forges (2020). Bayesian persuasion models with constraints on information disclosure policies appear in Perez-Richet (2014), in Boleslavsky and Kim (2018) and Salamanca (2021), where the designer chooses from Bayes-plausible distributions that satisfy some incentive constraints, in Le Treust and Tomala (2019), where the designer is constrained to sending noisy messages, in Wu (2020), where the accuracy of feasible test designs is bounded due to exogenous false-negative errors, and in Matyskova and Montes (2021), where the receiver is able to gather additional information.

The strategic interaction between multiple information designers has been studied under the assumption of simultaneous and one-stage information disclosure by, among others, Gentzkow and Kamenica (2017), Albrecht (2017), Au and Kawai (2020, 2019), Boleslavsky and Cotton (2015, 2018), and Koessler, Laclau, and Tomala (2020). Gentzkow and Kamenica (2017) consider the case in which each designer is

able to choose an information policy that is more informative than that of the other designer.⁴ Albrecht (2017), Au and Kawai (2020, 2019), and Boleslavsky and Cotton (2015, 2018) consider the case in which designers control independent pieces of information in applied examples. Koessler, Laclau, and Tomala (2020) provide existence results and properties of equilibria in games with multiple designers and multiple agents. This latter paper assumes that designers disclose information simultaneously, followed by agents making decisions simultaneously as well.

Multistage information design with a single designer has been studied in dynamic decision problems by, among others, Doval and Ely (2020), Ely (2017), Renault, Solan, and Vieille (2017) and Makris and Renou (2018). Since we assume that the state of nature is persistent and the decision problem is static (the agent makes a decision only once), multistage information design would be irrelevant in our model if there were only one designer. The dynamics of information design is interesting in our setting precisely because there are multiple designers. Sequential information design with multiple information designers has been studied by Li and Norman (2021) (see also Wu, 2020), but there are important differences with our work. First in Li and Norman (2021), the time horizon is finite and commonly known. Second, as in Gentzkow and Kamenica (2017), Li and Norman (2021) assume that designers disclose information about a common state and therefore each designer is able to choose an information policy that is more informative than the one of the competitor. Under these assumptions, they show that the sequential game cannot generate a more informative equilibrium than the simultaneous game. This property is not true in our model.⁵

Our methodology and results are closely related to the contributions in the literature on repeated games with incomplete information on *both* sides, splitting games and acyclic gambling games. The Mertens-Zamir function has been introduced by Mertens and Zamir (1971) (see also Mertens and Zamir, 1977, Sorin, 2002 and Mertens, Sorin, and Zamir, 2015) for the value function $u(p_1, p_2)$ of the one-stage incomplete information zero-sum game, which we replace by the indirect utility function of Designer 1. Mertens and Zamir (1971) have shown that the Mertens-Zamir function is the limit of the value of the infinitely repeated and discounted game as the discount factor tends to one.⁶ It is also the value of zero-sum splitting games studied in Laraki (2001a,b) and Oliu-Barton (2017). Laraki and Renault (2020) have recently extended these results from splitting games to more general stochastic games, called acyclic gambling games. We consider the indirect utility function of designers, given the beliefs and sequentially rational actions of the agent, and obtain our results by adapting the methodology of Renault and Venel (2017) and Laraki and Renault (2020). One main difference with Laraki and Renault (2020) is that we consider terminal payoffs (when the agent makes the decision), while payoffs in splitting games and acyclic gambling games are cumulated and discounted. We also consider and compare various possible timings of the game. Finally, a technical contribution with respect to Laraki and Renault (2020) is that we introduce a continuity condition on the feasible disclosure policies correspondences which is adapted to information design and is different from the non-expansivity condition of Laraki and Renault (2020).

⁴In our model, if designers can both reveal all the information about the payoff-relevant state (i.e., if their private states are perfectly correlated and all information disclosure policies are available), then the value of every multistage information design game coincides with the expected payoff under full information for the agent. See Section 5.3.

⁵See, for example, Section 4.

⁶It is also the limit of the value of the undiscounted N -stage repeated game as $N \rightarrow \infty$. If $\text{cav}_{p_1} \text{vex}_{p_2}$ differs from $\text{vex}_{p_2} \text{cav}_{p_1}$, the undiscounted infinitely repeated game has no value (see Aumann et al., 1995).

The timing of our long information design game is inspired by that of long cheap-talk games (Forges, 1990, Aumann and Hart, 2003). In one of our examples, the equilibrium martingale of posteriors does not reach its limit within a bounded number of disclosure stages, similarly to the “four frogs” example in Aumann and Hart (1986) and Forges (1984, 1990).

The one-stage simultaneous game of our application to a competitive product demonstration by two sellers has been studied, among others, by Albrecht (2017), Au and Kawai (2020) and Boleslavsky and Cotton (2015, 2018). Independently to our work, Whitmeyer (2020) considers a multistage extension of this example with a finite horizon N and discounted payoffs. Differently from our setting, he assumes that there is a sequence of short lived buyers, each buyer takes a decision once at a pre-determined period. He directly solves the game by backward induction and, similarly to us, shows that less information is revealed by the designers in the long information design game than in the static game. Precisely, the solution of Whitmeyer (2020) of this example, converges to the solution of our long information design game when the time horizon N tends to infinity and the discount factor tends to one.

Structure of the Paper In Section 2, we present the model. The main results are in Section 3. The application to competition in product demonstration is developed in Section 4. Several generalizations and extensions are studied in Section 5: we provide sufficient conditions under which our results apply with discontinuous indirect utility functions (Section 5.1); we consider alternative extensive form games with no or exogenous stopping rules (Section 5.2), allow for correlated private states (Section 5.3), and consider the case in which the messages generated by information disclosure policies are unobservable by the designers (Section 5.4). All proofs that are not in the text are in the Appendix.

2 Model

2.1 Environment

There are two *information designers* and a single *agent*. There is a finite set of *states* $\Omega_1 \times \Omega_2$ that is endowed with a common prior probability distribution. At the start of the game, no player is informed about the state. For each $i \in \{1, 2\}$, Designer i is able to design public information about Ω_i . To simplify the exposition, we assume that the prior probability distribution is the product of its marginal distributions, $p_1^0 \otimes p_2^0$, where $p_i^0 \in \Delta(\Omega_i)$ for each i (in all the paper, $\Delta(X)$ denotes the set of Borel probability measures over a compact set X). The extension to correlated priors, which includes the extreme case in which designers can choose information disclosure policies on a common state space, is studied in Section 5.3.

Designers produce and disclose publicly independent pieces of information; thus, each Designer i controls the public belief $p_i \in \Delta(\Omega_i)$ in a Bayes-plausible way.

2.2 Information Disclosure and Admissible Splittings

An *information disclosure policy* for Designer i is a public information structure *à la* Blackwell for Ω_i . When the public belief about Ω_i is $p_i \in \Delta(\Omega_i)$, the policy induces a probability distribution over

posterior beliefs with expectation p_i . Such a mean-preserving spread is called a *splitting* of p_i . The set of splittings of $p_i \in \Delta(\Omega_i)$ for Designer i is denoted:

$$\mathbb{S}_i(p_i) = \left\{ s \in \Delta(\Delta(\Omega_i)) : \int_{\tilde{p}_i \in \Delta(\Omega_i)} \tilde{p}_i ds(\tilde{p}_i) = p_i \right\}.$$

Our analysis covers situations in which designers are able to induce any splitting, i.e., any Bayes-plausible distribution of posteriors, as in the benchmark information design models. We also consider restricted sets of feasible information policies (for both or just one designer), this is important for the model to represent more realistic environments; restricting the set of feasible information policies also helps to illustrate our results with tractable examples. For instance, a simple restriction is to allow designers to choose between either fully revealing or non-revealing policies, as in the simple illustrative example of Section 3.4.1. In Section 2.6, we discuss richer cases of constrained information policies, including restrictions found in the literature.

If designers face technological constraints when choosing information disclosure policies, then the sets of splittings of beliefs that designers are able to induce are constrained as well. We model those constraints as follows. For each i , let $P_i \subseteq \Delta(\Omega_i)$ be a compact set, with $p_i^0 \in P_i$. We call $P_1 \times P_2$ the set of *admissible posteriors*. For every $p_i \in P_i$, we let $S_i(p_i) \subseteq \Delta(P_i) \cap \mathbb{S}_i(p_i)$ be the set of *admissible splittings*, where the correspondence $S_i : P_i \rightrightarrows \Delta(P_i)$ is continuous (i.e., upper hemi-continuous and lower hemi-continuous) with nonempty convex and compact values. Note that since the set $S_i(p_i)$ is convex, randomizing over splittings does not extend the set of feasible information policies.

Throughout the paper, we make two assumptions on admissible splittings (which are satisfied in the unconstrained case). First, we assume that for every $p_i \in P_i$, $\delta_{p_i} \in S_i(p_i)$, where δ_x denotes the Dirac measure at x . I.e., each designer is always able to choose a noninformative disclosure policy. Second, we assume that iterating admissible splittings does not enlarge splitting possibilities. This is the case, e.g., if all splittings are admissible. More generally, consider p_i in P_i , $s_i \in S_i(p_i)$ and let $f : P_i \rightarrow \Delta(P_i)$ be a measurable selection of S_i , $f(p'_i) \in S_i(p'_i)$ for each $p'_i \in P_i$. Define a splitting $f * s_i$ where Designer i first draws a posterior p'_i from s_i , and then a second posterior p_i'' from $f(p'_i)$. Specifically, $f * s_i$ is the probability distribution on $\Delta(P_i)$ defined by $f * s_i(B) = \int f(B|p'_i) ds_i(p'_i)$ for each Borel set $B \subseteq \Delta(P_i)$, where $f(B|p'_i)$ denotes the probability of Borel set B for distribution $f(p'_i)$. Observe that $f * s_i \in \mathbb{S}_i(p_i)$ (i.e., $\int \int p_i'' df(p_i''|p'_i) ds_i(p'_i) = p_i$). Let

$$S_i^2(p_i) = \{f * s_i : s_i \in S_i(p_i), f \text{ measurable selection of } S_i\},$$

be the set of splittings of p_i obtained by iterating S_i twice. We then assume that $\forall i, \forall p_i \in P_i$, $S_i^2(p_i) = S_i(p_i)$.⁷

2.3 Information Design Game

In the long information design game, each designer discloses information at stages $n = 1, 2, \dots$, and the agent makes a decision whenever *both* designers choose a non-revealing splitting at some stage denoted

⁷We relax this assumption when the set P_i is finite in Koessler et al. (2021).

by N^* (if no such stage exists, then $N^* = +\infty$). For instance, in a debate or a committee discussion, participants are asked if they have additional evidence at each stage, and if they do not, the discussion stops.

The initial public belief is (p_1^0, p_2^0) . At each stage $n = 1, 2, \dots$ and for each current belief (p_1^{n-1}, p_2^{n-1}) , each Designer i chooses an admissible splitting $s_i^n \in S_i(p_i^{n-1})$. Each new posterior p_i^n is drawn according to s_i^n , all players observe s_i^n and p_i^n for all i , and the game proceeds to stage $n + 1$. The first stage N^* in which $s_i^{N^*}$ is noninformative for each Designer i ends the game, and the agent chooses an action after stage N^* . If no such stage exists, then both designers choose admissible splittings until $N^* = +\infty$ and the agent chooses an action after stage $N^* = +\infty$.

A n -stage history is a sequence of splittings and posteriors as follows:

$$h^n = (p_1^0, p_2^0, s_1^1, s_2^1, \dots, p_1^{n-1}, p_2^{n-1}, s_1^n, s_2^n, p_1^n, p_2^n).$$

We denote by H^n the set of such histories. A strategy of each Designer i is a sequence σ_i of measurable functions $(\sigma_i^n)_n$, where for every n , $\sigma_i^n : H^{n-1} \rightarrow \Delta(P_i)$ and for every history $h^{n-1} \in H^{n-1}$, $\sigma_i^n(h^{n-1}) \in S_i(p_i^{n-1})$, where p_i^{n-1} is the $(n-1)$ -stage posterior on Ω_i .

Remark 1. Our results also apply to the alternative game format where designers cannot disclose information simultaneously, for example if Designer 1 is active at odd stages and Designer 2 at even stages (see Section 5.2). In that case, the agent makes a decision whenever both designers play non-revealingly consecutively. In Section 5.2, we also show that our results apply to the game with no deadline which never stops, and also to games with random duration which continues at each stage with probability $\delta \in (0, 1)$ or terminates with probability $(1 - \delta)$, when $\delta \rightarrow 1$.

2.4 Agent's Decision Problem and Players' Payoffs

The agent observes all splittings chosen by designers and all induced posteriors. After every play path h^{N^*} of the designers, the agent chooses an action z from a metric compact set Z . The payoff of each player depends on the state $(\omega_1, \omega_2) \in \Omega_1 \times \Omega_2$ and on the action $z \in Z$. The payoff of the agent is denoted by $\tilde{u}_A(z; \omega_1, \omega_2)$, and the payoff of Designer 1 is denoted by $\tilde{u}(z; \omega_1, \omega_2)$. We assume that the game is strictly competitive between the designers, so the payoff of Designer 2 is $-\tilde{u}(z; \omega_1, \omega_2)$. For every state $(\omega_1, \omega_2) \in \Omega_1 \times \Omega_2$, the utility functions $\tilde{u}(z; \omega_1, \omega_2), \tilde{u}_A(z; \omega_1, \omega_2)$ are assumed to be continuous in z . If the agent has a finite set of actions A , we let $Z = \Delta(A)$ be the set of mixed actions and define payoffs linearly by taking the expectation. We extend \tilde{u} and \tilde{u}_A as usual as follows: for z in Z and for each $p_i \in P_i$,

$$\tilde{u}(z; p_1, p_2) = \sum_{(\omega_1, \omega_2) \in \Omega_1 \times \Omega_2} p_1(\omega_1) p_2(\omega_2) \tilde{u}(z; \omega_1, \omega_2) \text{ and,}$$

$$\tilde{u}_A(z; p_1, p_2) = \sum_{(\omega_1, \omega_2) \in \Omega_1 \times \Omega_2} p_1(\omega_1) p_2(\omega_2) \tilde{u}_A(z; \omega_1, \omega_2).$$

For every pair of admissible posteriors (p_1, p_2) , we assume that the decision problem of the agent $\max_{z \in Z} \tilde{u}_A(z; p_1, p_2)$ has a unique solution $z(p_1, p_2)$, which is continuous in (p_1, p_2) according to the

Maximum Theorem. This assumption is satisfied, e.g., if $\tilde{u}_A(z; \omega_1, \omega_2)$ is strictly concave in z . If $\max_{z \in Z} \tilde{u}_A(z; p_1, p_2)$ has multiple solutions, our results also apply as long as there is a continuous selection $z(p_1, p_2)$ of the argmax. Such a selection typically does not exist if the agent has a finite set of actions, Z is the set of mixed actions, and all posteriors are admissible (i.e., $P_i = \Delta(\Omega_i)$ for each i). This is the case in the example presented in Section 4. Importantly, our results and methods can be extended to analyze this example (proofs have to be amended to account for the discontinuities, see Section 5.1), so the study of continuous games also gives insights for discontinuous models.

We assume that the agent's decision is sequentially rational, i.e., the agent chooses $z(p_1, p_2)$ for all possible terminal posteriors $(p_1, p_2) \in P_1 \times P_2$. Thus, for each play path h^{N^*} of the designers, there is a uniquely defined decision of the agent. Let

$$u(p_1, p_2) = \sum_{(\omega_1, \omega_2) \in \Omega_1 \times \Omega_2} p_1(\omega_1) p_2(\omega_2) \tilde{u}(z(p_1, p_2); \omega_1, \omega_2),$$

be Designer 1's expected payoff induced by the optimal action of the agent when public beliefs are given by $(p_1, p_2) \in P_1 \times P_2$. Note that $u(p_1, p_2)$ is continuous in (p_1, p_2) since $\tilde{u}(z; \omega_1, \omega_2)$ is continuous in z and $z(p_1, p_2)$ is continuous in (p_1, p_2) . We extend u and denote by $u(s_1, s_2)$ the expected value of u with respect to $s_i \in \Delta(P_i)$ for each i ,

$$u(s_1, s_2) = \iint u(p_1, p_2) ds_1(p_1) ds_2(p_2).$$

We use a similar notation for the expected utility $u_A(p_1, p_2)$ of the agent.

2.5 Games Between Designers, Equilibria and Values

Our main object of interest is the two-player long information design game induced by the sequentially rational decisions of the agent at the terminal stage, we denote this game by $G(p_1^0, p_2^0)$. A pair of strategies (σ_1, σ_2) of the designers induces a distribution $\mathbb{P}_{(\sigma_1, \sigma_2)}$ over play paths, for which the random sequence of posteriors $(p_1^n, p_2^n)_n$ is a martingale:

$$\forall n, \forall h^n = (p_1^0, p_2^0, s_1^1, s_2^1, \dots, p_1^{n-1}, p_2^{n-1}, s_1^n, s_2^n, p_1^n, p_2^n), \forall i, \mathbb{E}_{\sigma_1, \sigma_2}[p_i^{n+1} | h^n] = p_i^n.$$

The martingale convergence theorem ensures that the terminal beliefs $(p_1^{N^*}, p_2^{N^*})$ are well defined (if $N^* = +\infty$, $p_i^{N^*} = \lim_n p_i^n$ which exists almost surely).

The expected payoff of Designer 1 is given by $U(\sigma_1, \sigma_2) = \mathbb{E}_{\sigma_1, \sigma_2} u(p_1^{N^*}, p_2^{N^*})$ and the expected payoff of Designer 2 is $-U(\sigma_1, \sigma_2)$. A pair of strategies (σ_1^*, σ_2^*) is an ε -equilibrium of the information design game if

$$U(\sigma_1, \sigma_2^*) - \varepsilon \leq U(\sigma_1^*, \sigma_2^*) \leq U(\sigma_1^*, \sigma_2) + \varepsilon,$$

for every σ_1 and σ_2 , it is an equilibrium if it is an ε -equilibrium for $\varepsilon = 0$. The information design game has value V if the sup inf and inf sup payoffs coincide:

$$V = \sup_{\sigma_1} \inf_{\sigma_2} U(\sigma_1, \sigma_2) = \inf_{\sigma_2} \sup_{\sigma_1} U(\sigma_1, \sigma_2).$$

As is well known, the zero-sum information design game has a value if and only if it has ε -equilibria for every $\varepsilon > 0$ and (σ_1^*, σ_2^*) is an equilibrium if and only if $V = \min_{\sigma_2} U(\sigma_1^*, \sigma_2) = \max_{\sigma_1} U(\sigma_1, \sigma_2^*)$. In this case, σ_1^* and σ_2^* are called optimal strategies of Designers 1 and 2, respectively. A strategy σ_i of Designer i is called *stationary* if it is Markovian, i.e., in each stage n , σ_i^n depends only on the posteriors (p_1^{n-1}, p_2^{n-1}) of the previous stage, and σ_i^n does not depend on n . An equilibrium (σ_1^*, σ_2^*) is a *stationary equilibrium* if σ_1^* and σ_2^* are stationary.

2.6 Examples of feasible information policies

We describe below some natural constraints on information policies available to a designer and show how these constraints translate into constraints on the set of admissible posteriors and admissible splittings.

When information policies are unconstrained, Designer i has access to any (statistical or Blackwell) experiment $x : \Omega_i \rightarrow \Delta(M)$, where M is some set of publicly observed messages. In this case, he is able to induce any splitting of any $p_i \in \Delta(\Omega_i)$, the sets of admissible posteriors and splittings are then the entire sets: $P_i = \Delta(\Omega_i)$, $S_i(p_i) = \mathbb{S}_i(p_i)$.

In most concrete applications, information policies are constrained because only a limited number of statistical experiments are available. For example, a medical test usually has unavoidable false-negative or false-positive results; asking a social media influencer to test a new product, or a referee to evaluate a research project or a candidate, is always subject to some errors or predictable biases in the distribution of generated messages. Alternatively, a designer may be constrained to provide hard evidence about the state from an exogenous set of evidences. In addition, the information policy might be required to be ex-post verifiable, in which case an admissible information policy should be a deterministic mapping from states to messages.

To be more explicit about such constrained information disclosure policies, consider an exogenous compact subset of experiments $X_i \subseteq \{x : \Omega_i \rightarrow \Delta(M)\}$, where M is a finite set of messages. Designer i might be able to repeatedly run experiments in X_i as many times as he wants, or he might be able to run a single experiment from X_i only once.

For example, if only deterministic experiments are available, then $P_i = \{p_i \in \Delta(\Omega_i) : \exists E \subseteq \Omega_i, \text{ s.t. } p_i = p_i^0(\cdot|E)\}$ is finite and $S_i(p_i) = \Delta(P_i) \cap \mathbb{S}_i(p_i)$. An interesting subset of deterministic experiments is obtained by following the literature on disclosure games with evidences (e.g., Milgrom, 1981) and mechanism design with evidences (e.g., Green and Laffont, 1986). Precisely, for each state ω_i , Designer i can only choose messages from a nonempty set $M(\omega_i) \subseteq M$, interpreted as the set of evidences available in state ω_i . In such a setting, a natural restriction on information policies is to assume that the designer chooses any deterministic experiment $x : \Omega_i \rightarrow M$ such that $x(\omega_i) \in M(\omega_i)$ for every $\omega_i \in \Omega_i$.⁸ Then, the assumption that the designer is able to choose a noninformative disclosure policy ($\delta_{p_i} \in S_i(p_i)$) is satisfied iff there exists a message $m_0 \in M$ such that $m_0 \in M(\omega_i)$ for every ω_i , i.e., there exists a message m_0 that provides trivial evidence. Note that in this model, the assumption that $S_i^2(p_i) = S_i(p_i)$ is the analogue of the standard “normality” condition in disclosure games and mechanism design with evidences, according to which the collection of sets $\{M^{-1}(m) : m \in M\}$ is

⁸Stochastic evidence production could be added to this setting as well, by allowing stochastic experiments $x : \Omega_i \rightarrow \Delta(M)$ such that $x(m | \omega_i) = 0$ if $m \notin M(\omega_i)$.

closed under intersection.⁹ This assumption is satisfied, e.g., if there is an evidence for every subset $\Omega'_i \subseteq \Omega_i$, or in the classical disclosure model of Dye (1985), where there is a set of states $\Omega_i^E \subseteq \Omega_i$ in which there exists either full or no evidence ($M(\omega_i) = \{m_0, m_{\omega_i}\}$ for every $\omega_i \in \Omega_i^E$) and in the remaining states there is no evidence ($M(\omega_i) = \{m_0\}$ for every $\omega_i \in \Omega_i \setminus \Omega_i^E$).

As another example, take $\Omega_i = \{\underline{\omega}, \bar{\omega}\}$, $X_i = \{x_\varepsilon, x_0\}$, $M = \{\underline{m}, \bar{m}\}$, where x_0 is a non informative experiment, $x_\varepsilon(\underline{m} \mid \underline{\omega}) = 1 - \varepsilon$, $x_\varepsilon(\bar{m} \mid \bar{\omega}) = 1$, and $\varepsilon \in (0, 1)$. Let $p_i^0 = \frac{1}{2}$ be the prior probability of state $\underline{\omega}$. For instance, let ω_i be the characteristic of a product which could be inappropriate for a consumer ($\omega_i = \underline{\omega}$) or appropriate ($\omega_i = \bar{\omega}$) and ε is the probability that the experiment x_ε produces a false-positive result. First, assume that Designer i can run the experiment x_ε only once. Then from Bayes' rule, the set of admissible posterior beliefs on $\underline{\omega}$ is simply $P_i = \{\varepsilon/(\varepsilon + 1), \frac{1}{2}, 1\}$ and the sets of admissible splittings are given by

$$S_i\left(\frac{1}{2}\right) = \Delta(P_i) \cap \mathbb{S}_i\left(\frac{1}{2}\right) \quad \text{and} \quad S_i(p_i) = \delta_{p_i} \text{ for every } p_i \neq \frac{1}{2}.$$

Alternatively, assume that Designer i can repeatedly run x_ε as many times as he wants. Then, the set of admissible posteriors is $P_i = \{0, \varepsilon^n/(\varepsilon^n + 1), 1 : n = 0, 1, 2, \dots\}$ and for every $p_i \in P_i$ the set of admissible splittings is

$$S_i(p_i) = \Delta(\{p'_i \in P_i : p'_i = 1 \text{ or } p'_i \leq p_i\}) \cap \mathbb{S}_i(p_i).$$

Our model applies more generally when each Designer i have access to an arbitrary compact subset X_i of experiments (see Proposition 4 in Section 3.3 for more details).

3 Main Results

In this section, we characterize equilibrium strategies and values for the long information design game presented in the previous section. We also provide several examples to illustrate the described properties.

3.1 Definitions and Preliminary Results

It is well known from the theory of repeated games (Aumann and Maschler, 1966; Aumann et al., 1995) and of Bayesian persuasion (Kamenica and Gentzkow, 2011) that the concavification (or concave closure) of a function is a key concept, as it captures the optimal splitting for a single designer. In our setting, payoffs depend on two variables p_1, p_2 and are zero-sum between designers. Thus, ideally, Designer 1 would like to concavify $u(p_1, p_2)$ with respect to p_1 and Designer 2 would like to concavify $-u(p_1, p_2)$, i.e., convexify $u(p_1, p_2)$, with respect to p_2 . Before studying the interplay between these notions, we need to define what concave and convex mean in our setting since P_1 and P_2 are not necessarily convex sets and splittings are constrained.

In the following Definitions 1 and 2, all functions defined on P_i (or $P_1 \times P_2$) are assumed to be measurable and bounded, and extended to elements of $\Delta(P_i)$ (or $\Delta(P_1) \times \Delta(P_2)$) by (multi)-linearity.

⁹See, for example, Green and Laffont (1986), Forges and Koessler (2005), Bull and Watson (2007).

Definition 1. A function $w : P_1 \rightarrow \mathbb{R}$ is S_1 -concave if for all $p_1 \in P_1$ and all $s_1 \in S_1(p_1)$, $w(p_1) \geq w(s_1)$, where $w(s_1) = \int w(p'_1) ds_1(p'_1)$. A function $w : P_2 \rightarrow \mathbb{R}$ is S_2 -convex if for all $p_2 \in P_2$ and all $s_2 \in S_2(p_2)$, $w(p_2) \leq w(s_2)$. A function $w : P_1 \times P_2 \rightarrow \mathbb{R}$ is S_1 -concave if $w(\cdot, p_2)$ is S_1 -concave for all p_2 , and is S_2 -convex if $w(p_1, \cdot)$ is S_2 convex for all p_1 .

If P_1 is a convex set and $S_1(p_1) = \Delta(P_1) \cap \mathbb{S}_1(p_1)$ for each p_1 , w is concave in the usual sense: $w(p_1)$ is greater than or equal to the expectation of w under any distribution with mean p_1 . Thus, the definition is the generalization to all *admissible* distributions with mean p_1 . For functions of two variables $w(p_1, p_2)$, S_1 -concave means concave with respect to variable p_1 given the admissible splittings S_1 .

Definition 2. The concavification of $w : P_1 \rightarrow \mathbb{R}$ is the smallest S_1 -concave function that is pointwise greater than or equal to w . We denote it by the following:

$$\text{cav}_{p_1} w = \inf\{g : g \geq w, g \text{ is } S_1\text{-concave}\}.$$

The convexification of $w : P_2 \rightarrow \mathbb{R}$ is the largest S_2 -convex function that is pointwise less than or equal to w . We denote it by the following:

$$\text{vex}_{p_2} w = \sup\{g : g \leq w, g \text{ is } S_2\text{-convex}\}.$$

For a function $w : P_1 \times P_2 \rightarrow \mathbb{R}$, $\text{cav}_{p_1} w$ denotes the concavification with respect to p_1 of $w(\cdot, p_2)$ for fixed p_2 , and $\text{vex}_{p_2} w$ denotes the convexification with respect to p_2 of $w(p_1, \cdot)$ for fixed p_1 .

The next lemma shows that the concavification with respect to p_1 (the convexification with respect to p_2) of a continuous function is continuous and corresponds to the optimal choice of information policy for Designer 1 (Designer 2) when the posterior belief about the other designer is fixed.

Lemma 1. For a continuous function $w : P_1 \times P_2 \rightarrow \mathbb{R}$, the functions $\text{cav}_{p_1} w$ and $\text{vex}_{p_2} w$ are continuous on $P_1 \times P_2$ and for every $(p_1, p_2) \in P_1 \times P_2$,

$$\text{cav}_{p_1} w(p_1, p_2) = \max\{w(s_1, p_2) : s_1 \in S_1(p_1)\} \text{ and } \text{vex}_{p_2} w(p_1, p_2) = \min\{w(p_1, s_2) : s_2 \in S_2(p_2)\}.$$

3.2 Benchmark models

In this section, consider benchmark models in which each designer only moves once. We denote by $G^{sim}(p_1^0, p_2^0)$ the one-stage simultaneous game where the set of strategies is $S_i(p_i^0)$ for each Designer i , and the payoff is $u(s_1, s_2)$. Because the strategy sets are convex and compact and the payoff function $(s_1, s_2) \mapsto u(s_1, s_2)$ is continuous and bilinear, by the minmax theorem $G^{sim}(p_1^0, p_2^0)$ has a value, denoted by $\text{SV}(u)(p_1^0, p_2^0)$ and both designers have optimal strategies:

$$\text{SV}(u)(p_1^0, p_2^0) = \max_{s_1 \in S_1(p_1^0)} \min_{s_2 \in S_2(p_2^0)} u(s_1, s_2) = \min_{s_2 \in S_2(p_2^0)} \max_{s_1 \in S_1(p_1^0)} u(s_1, s_2).$$

We call $\text{SV}(u)$ the *splitting value* function of u .

Lemma 2. *The functions $\text{SV}(u)$, $\text{cav}_{p_1} \text{vex}_{p_2} u$ and $\text{vex}_{p_2} \text{cav}_{p_1} u$ are continuous, S_1 -concave and S_2 -convex. In addition, for every $(p_1^0, p_2^0) \in P_1 \times P_2$,*

$$\text{vex}_{p_2} u(p_1^0, p_2^0) \leq \text{cav}_{p_1} \text{vex}_{p_2} u(p_1^0, p_2^0) \leq \text{SV}(u)(p_1^0, p_2^0) \leq \text{vex}_{p_2} \text{cav}_{p_1} u(p_1^0, p_2^0) \leq \text{cav}_{p_1} u(p_1^0, p_2^0).$$

Notice that $\text{cav}_{p_1} \text{vex}_{p_2} u(p_1^0, p_2^0)$ is the value of the 2-stage sequential information design game in which Designer 1 plays first and Designer 2 second, and $\text{vex}_{p_2} \text{cav}_{p_1} u(p_1^0, p_2^0)$ is the value of the sequential game in which Designer 2 plays first and Designer 1 second. Indeed consider the sequential game in which Designer 1 plays first and Designer 2 second. At the first stage, Designer 1 chooses a splitting $s_1 \in S_1(p_1^0)$ that maximizes $u(s_1, p_2^0)$. For any strategy of Designer 2, the expected payoff is $\mathbb{E} \text{cav}_{p_1} u(p_1^0, p_2^0) \geq \text{vex}_{p_2} \text{cav}_{p_1} u(p_1^0, p_2^0)$. Hence, Designer 1 has a strategy that guarantees a payoff of at least $\text{vex}_{p_2} \text{cav}_{p_1} u(p_1^0, p_2^0)$ irrespective of the strategy of Designer 2. Define a strategy σ_2 of Designer 2 as follows. At the second stage, Designer 2 chooses a splitting $s_2 \in S_2(p_2^0)$ that minimizes $\mathbb{E}_{s_2} \text{cav}_{p_1} u(p_1^1, p_2)$. Regardless of the strategy of Designer 1, the expected payoff is at most $\mathbb{E} \text{vex}_{p_2} \text{cav}_{p_1} u(p_1^1, p_2^0) \leq \text{vex}_{p_2} \text{cav}_{p_1} u(p_1^0, p_2^0)$ by S_1 -concavity of $\text{vex}_{p_2} \text{cav}_{p_1} u(p_1, p_2)$. Hence, Designer 2 has a strategy that guarantees that Designer 1's payoff is at most $\text{vex}_{p_2} \text{cav}_{p_1} u(p_1^0, p_2^0)$ for any strategy of Designer 1. Therefore, $\text{vex}_{p_2} \text{cav}_{p_1} u(p_1^0, p_2^0)$ is the value of the game. The argument is analogous for the sequential game in which Designer 2 plays first and Designer 1 second. More generally, it can be shown that $\text{cav}_{p_1} \text{vex}_{p_2} u(p_1^0, p_2^0)$ is the value of any multistage information design game in which Designer 2 moves last, and $\text{vex}_{p_2} \text{cav}_{p_1} u(p_1^0, p_2^0)$ is the value of any multistage information design game in which Designer 1 moves last. In particular, the previous lemma implies that $\text{vex}_{p_2} \text{cav}_{p_1} \text{vex}_{p_2} u(p_1^0, p_2^0) = \text{cav}_{p_1} \text{vex}_{p_2} u(p_1^0, p_2^0)$ and $\text{cav}_{p_1} \text{vex}_{p_2} \text{cav}_{p_1} u(p_1^0, p_2^0) = \text{vex}_{p_2} \text{cav}_{p_1} u(p_1^0, p_2^0)$.

3.3 Values and equilibria of the long information design game

The following definition is due to Mertens and Zamir (1971).

Definition 3. *Let $u : P_1 \times P_2 \rightarrow \mathbb{R}$ be a continuous function. A Mertens-Zamir (MZ) function for u is a function $v : P_1 \times P_2 \rightarrow \mathbb{R}$ such that for every $(p_1, p_2) \in P_1 \times P_2$,*

$$v(p_1, p_2) = \text{cav}_{p_1} \min\{u(p_1, p_2), v(p_1, p_2)\} = \text{vex}_{p_2} \max\{u(p_1, p_2), v(p_1, p_2)\}.$$

By definition, if v is a MZ function for u , then v is S_1 -concave and S_2 -convex. Section 3.4 provides two simple illustrations of this definition. Below we show that when the Mertens-Zamir function exists, then its value at the prior beliefs is equal to the value of the long information design (Theorem 1) and lies in between the values of the 2-stage sequential games in which each designer plays only once (Lemma 3). We also provide an alternative formulation of the MZ function (Proposition 1) that allows simple optimal strategies to be derived directly. Then, we provide general conditions on the splitting correspondances for the Mertens-Zamir function to exist.

Lemma 3. *If v is a MZ function for u , then for every $(p_1, p_2) \in P_1 \times P_2$,*

$$\text{cav}_{p_1} \text{vex}_{p_2} u(p_1, p_2) \leq v(p_1, p_2) \leq \text{vex}_{p_2} \text{cav}_{p_1} u(p_1, p_2).$$

Proposition 1. *A continuous function v is a MZ function for u if and only if v is S_1 -concave, S_2 -convex and satisfies the following properties for all $(p_1, p_2) \in P_1 \times P_2$:*

(C1) *There exists $s_1 \in S_1(p_1)$ such that $v(p_1, p_2) = v(s_1, p_2)$ and $v(p'_1, p_2) \leq u(p'_1, p_2)$, $\forall p'_1 \in \text{supp}(s_1)$.*

(C2) *There exists $s_2 \in S_2(p_2)$ such that $v(p_1, p_2) = v(p_1, s_2)$ and $v(p_1, p'_2) \geq u(p_1, p'_2)$, $\forall p'_2 \in \text{supp}(s_2)$.*

This result is very useful since it allows deriving intuitive strategies from the Mertens-Zamir function (see the discussion after the next theorem).

Theorem 1. *If there exists a continuous MZ function of u , denoted by $\text{MZ}(u)(p_1^0, p_2^0)$, then the long information design game $G(p_1^0, p_2^0)$ has a stationary equilibrium and the value is $V = \text{MZ}(u)(p_1^0, p_2^0)$.*

An informal sketch of the proof is as follows. From properties (C1) and (C2), we know that, setting $v = \text{MZ}(u)$, there exists $s_1 \in S_1(p_1)$ such that $v(p_1, p_2) = v(s_1, p_2)$ and $v(p'_1, p_2) \leq u(p'_1, p_2)$ for all $p'_1 \in \text{supp}(s_1)$, and there exists $s_2 \in S_2(p_2)$ such that $v(p_1, p_2) = v(p_1, s_2)$ and $v(p_1, p'_2) \geq u(p_1, p'_2)$ for all $p'_2 \in \text{supp}(s_2)$. These properties define naturally stationary strategies for both designers that turn out to form an equilibrium. Assume that $v(p_1, p_2)$ is the value of the continuation game starting at (p_1, p_2) . From the point of view of Designer 1, (C1) means that it is possible to choose a splitting that preserves the expected continuation value. Hence, assume that $v(p_1, p_2) < u(p_1, p_2)$; then, Designer 1 can play non-revealingly without reducing the equilibrium continuation payoff. Intuitively, if $v(p_1, p_2) < u(p_1, p_2)$, Designer 1 would be content if the game stopped, as Designer 1 would receive more than the value. If $v(p_1, p_2) > u(p_1, p_2)$, then Designer 1 does not want the game to stop and chooses a splitting s_1 such that $v(p'_1, p_2) \leq u(p'_1, p_2)$ with probability one to reach a point where the realized payoff u would potentially be greater than the value. The symmetry of (C1) and (C2) implies that $v(p_1, p_2)$ can be enforced by both designers. Thus, this is the equilibrium payoff of this zero-sum game.

Proof of Theorem 1. Let $v = \text{MZ}(u)$, and define a strategy σ_1 of Designer 1 as follows. Given posteriors $(p_1, p_2) \in P_1 \times P_2$ at stage n , Designer 1 chooses the non-revealing splitting δ_{p_1} if $u(p_1, p_2) \geq v(p_1, p_2)$, and otherwise chooses a splitting $s_1 \in S_1(p_1)$ such that $v(p_1, p_2) = v(s_1, p_2)$ and $v(p'_1, p_2) \leq u(p'_1, p_2)$, $\forall p'_1 \in \text{supp}(s_1)$. According to Proposition 1, this strategy is well defined. It has the property that for any strategy of Designer 2, $u(p_1^{n+1}, p_2^n) \geq v(p_1^{n+1}, p_2^n)$ almost surely. This implies that $u(p_1^{N^*}, p_2^{N^*}) \geq v(p_1^{N^*}, p_2^{N^*})$ almost surely: either $N^* < +\infty$ and then $p_i^{N^*-1} = p_i^{N^*}$, or $N^* = +\infty$ and we consider the limit. Therefore, $\mathbb{E}u(p_1^{N^*}, p_2^{N^*}) \geq \mathbb{E}v(p_1^{N^*}, p_2^{N^*})$. By S_2 -convexity of v , $\mathbb{E}[v(p_1^{n+1}, p_2^{n+1}) | h^n] \geq \mathbb{E}[v(p_1^{n+1}, p_2^n) | h^n] = v(p_1^n, p_2^n)$ by construction of σ_1 . It follows that $\mathbb{E}v(p_1^{n+1}, p_2^{n+1}) \geq \mathbb{E}v(p_1^n, p_2^n)$, and by induction (and taking limit if $N^* = +\infty$), $\mathbb{E}v(p_1^{N^*}, p_2^{N^*}) \geq \mathbb{E}v(p_1^n, p_2^n) \geq v(p_1^0, p_2^0)$. Thus, there is a strategy of Designer 1 such that for any strategy of Designer 2, $\mathbb{E}u(p_1^{N^*}, p_2^{N^*}) \geq v(p_1^0, p_2^0)$. By symmetry, this is the value of the game. We have constructed a stationary equilibrium. \square

It follows from the above proof that under the equilibrium strategies we have $u(p_1^{N^*}, p_2^{N^*}) = v(p_1^{N^*}, p_2^{N^*})$ almost surely. Thus, if the martingale stops, it must be at points where $u(p_1, p_2) = v(p_1, p_2)$. Furthermore, if designers can always reach points where $u(p_1, p_2) = v(p_1, p_2)$, the martingale actually stops after the first stage ($N^* = 2$), after which the agent takes his decision. This is the case under the conditions below.

Proposition 2. *Assume that both sets P_i are convex, that all splittings are admissible, and that there exists a continuous MZ function $v = \text{MZ}(u)$. There is an equilibrium such that:*

- *If $u(p_1^0, p_2^0) < v(p_1^0, p_2^0)$, then Designer 1 plays revealing at the first stage, and both designers play non-revealing at the second stage;*
- *If $u(p_1^0, p_2^0) > v(p_1^0, p_2^0)$, then Designer 2 plays revealing at the first stage, and both designers play non-revealing at the second stage;*
- *If $u(p_1^0, p_2^0) = v(p_1^0, p_2^0)$, then both designers play non-revealing at the first stage.*

This is a direct consequence of the following lemma.

Lemma 4. *Assume that both P_i are convex and that all splittings are admissible, $S_1(p_1) = \Delta(P_1) \cap \mathbb{S}_1(p_1)$, $S_2(p_2) = \Delta(P_2) \cap \mathbb{S}_2(p_2)$. Then, for any $(p_1, p_2) \in P_1 \times P_2$,*

- *if $u(p_1, p_2) \leq v(p_1, p_2)$, there exists $s_1 \in S_1(p_1)$ such that $v(p'_1, p_2) = u(p'_1, p_2)$, $\forall p'_1 \in \text{supp}(s_1)$,*
- *if $u(p_1, p_2) \geq v(p_1, p_2)$, there exists $s_2 \in S_2(p_2)$ such that $v(p_1, p'_2) = u(p_1, p'_2)$, $\forall p'_2 \in \text{supp}(s_2)$.*

This lemma is a direct extension of a result in Heuer (1992) and in Oliu-Barton (2017). For the sake of completeness, the proof is recalled in the Appendix. It follows that in equilibrium both designers play non-revealing and the martingale is constant if $u(p_1, p_2) = v(p_1, p_2)$. If $u(p_1, p_2) < v(p_1, p_2)$, Designer 1 splits to a point p'_1 such that $u(p'_1, p_2) = v(p'_1, p_2)$ and the martingale is constant thereafter and if $u(p_1, p_2) > v(p_1, p_2)$, Designer 2 splits to a point p'_2 such that $u(p_1, p'_2) = v(p_1, p'_2)$ and the martingale is constant thereafter. Thus, if the sets of admissible posteriors are convex and all splittings are admissible, information disclosure lasts one stage at most, after which the agent takes his decision.¹⁰ In Section 3.4.2, we provide an example with finite sets of posteriors in which the number of disclosure stages is unbounded.

Existence of MZ function. A natural question arising from Theorem 1 is whether an MZ function exists. To answer this, we introduce the following assumption, which strengthens continuity of splitting correspondences. A related condition of “non-expansivity” appears in Renault and Venel (2017).

Assumption 1. *The splitting correspondence of each Designer i satisfies the following. There exists $\theta \in (0, 1]$ such that:*

$$\forall p_i, p'_i \in P_i, \forall s_i \in S_i(p_i), \forall a, b \geq 0, \exists s'_i \in S_i(p'_i) \text{ s.t. } \forall f \in D_\theta^i, |af(s_i) - bf(s'_i)| \leq \|ap_i - bp'_i\|_1, \quad (1)$$

where $D_\theta^i = \{f : P_i \rightarrow [-\theta, \theta], \forall p_i, p'_i \in P_i, \forall a, b \geq 0, |af(p_i) - bf(p'_i)| \leq \|ap_i - bp'_i\|_1\}$.

This condition should be thought of as one possible definition of a Lipschitz correspondence, similarly to the non-expansivity condition in Renault and Venel (2017). The idea is that an existence result requires a fixed point argument. Since the value of the game is a function of beliefs, we need to find a

¹⁰A similar property also appears in Li and Norman (2021, Proposition 6) for non-zero sum sequential information design games with perfectly correlated states and a fixed deadline.

set of functions where the fixed point argument will be applied. This set of functions must be compact, i.e., not too large, but also stable by concavification and convexification, thus not too small. Consider for instance the special case where Designer 2 has a trivial correspondence with only noninformative policies. Then, the solution would be the concavification by Designer 1. Thus, the value functions should at least lie in sets which are stable under concavification or convexification.

As it turns out, Assumption 1 can be reformulated as follows.

Lemma 5. *Assumption 1 is equivalent to: For each Designer i , there exists $\theta \in (0, 1]$ such that for each $f \in D_\theta^i$, the concavification and the convexification of f belong to D_θ^i .*

This shows that the set of functions D_θ^1 (resp. D_θ^2) is chosen to be closed under the concavification (resp. convexification) operator. It is apparent in our proofs that the existence of the value relies on the set of functions closed under those operators, see in particular the proof of Proposition 7 in the appendix.

Here are important cases where the assumption holds.

Proposition 3. *If for each i , $P_i = \Delta(\Omega_i)$ and all splittings are allowed, or P_i is finite, then Assumption 1 is satisfied.*

We now consider the interesting case of splitting correspondences induced by constraints on experiments. Suppose that Designer i is given a finite set of messages M and a compact subset of experiments $X \subseteq \{x : \Omega_i \rightarrow \Delta(M)\}$, which contains a non-revealing experiment. Consider all the splittings that Designer i can generate by repeatedly choosing experiments in X . For instance, if X contains only a non-revealing experiment x_0 and a noisy experiment x_1 , the designer can run x_1 as many time as he wants, depending on the outcomes of previous trials.

More generally, define an “auxiliary strategy” for Designer i as a sequence of measurable mappings $\sigma_i = (\sigma_{it})$, where for each $t \geq 1$, $\sigma_{it} : (X \times M)^{t-1} \rightarrow \Delta(X)$. For each $x \in X$, $\sigma_{it}(x|x_1, m_1, \dots, x_{t-1}, m_{t-1})$ is the probability that the designer runs experiment x conditionally on previous experiments and message realizations $x_1, m_1, \dots, x_{t-1}, m_{t-1}$. A prior p_i and an auxiliary strategy σ_i induce a probability distribution over sequences of experiments and messages: conditional on state ω_i , x_1 is selected according to σ_{i1} , then m_1 according to $x_1(\cdot|\omega_i)$, then x_2 according to $\sigma_{i2}(x_1, m_1)$, then m_2 according to $x_2(\cdot|\omega_i)$, and so on. Denote by $\mu_T(p_i, \sigma_i)$ the induced distribution of posteriors after T trials and $S_{iT}(p_i)$ the set of all such distributions as σ_i varies. Finally let $S_i(p_i)$ be the weak-* closure of $\cup_{T \geq 1} S_{iT}(p_i)$. It is easy to see that $S_i(p_i)$ is a convex compact set and the correspondence S_i satisfies $S_i^2 = S_i$.

Proposition 4. *The splitting correspondence generated by a set of experiments as above satisfies Assumption 1.*

We then have the following existence result.

Theorem 2. *Under Assumption 1, there exists a continuous MZ function of u , denoted by $\text{MZ}(u)$ which is therefore the value of the long information design game $G(p_1^0, p_2^0)$.*

The logic of the proof is the following. Introduce a “discounted game” where Designers 1 and 2 move simultaneously and in each stage, the game ends with exogenous probability $1 - \delta$ and the receiver

takes an action, or the game continues to the next stage with probability $\delta < 1$. Let $v_\delta(p_1^0, p_2^0)$ be the value of this game. The main argument is to show that the family of functions $(v_\delta)_\delta$ is equicontinuous and therefore admits a limit point v as $\delta \rightarrow 1$. This uses the fact that function v_δ is the fixed point of an appropriate Bellmann equation. We then show that v satisfies all the conditions ensuring that $v = \text{MZ}(u)$. The proof follows the steps of the proofs of Propositions 2 and 3 in Laraki and Renault (2020), albeit using a different “continuity” condition (Assumption 1).

Laraki and Renault (2020) use a “non expansivity” assumption which guarantees also the existence of a continuous MZ function. Their assumption is satisfied if $P_i = \Delta(\Omega_i)$ and all splittings allowed or if P_i is finite. However, the splitting correspondence generated by a set of experiments need not be non-expansive. We provide a counter-example in Appendix A.2.

3.4 Examples

In this section we provide two simple examples in which we characterize the values and optimal strategies. In the first example each designer can either fully reveal the state or reveal nothing, and illustrates that the values and equilibrium strategies differ in the simultaneous, sequential and long information design games. In the second example the sets of admissible posteriors are finite and all splittings on those sets are admissible, and the equilibrium martingale of posteriors is unbounded.

3.4.1 Illustrative Example

Consider the following situation. The designers are opposed lobbyists or NGO. The agent is a journalist who would like to write an article or talk about a lobbyist’s case in a TV show only if this lobbyist has disclosed relevant information about his case. The second lobbyist would like to appear in the journalist’s article or TV show only if the first lobbyist does, while the first lobbyist would like the reverse. Formally, for each i , let $\Omega_i = \{0, 1\}$, identify $p_i \in \Delta(\Omega_i)$ with $p_i(1) \in [0, 1]$, and let $p_i^0 = \frac{1}{2}$. Suppose that each designer has only two available disclosure policies: non-revealing or fully revealing, or equivalently that each designer can only use deterministic experiments. The possible posteriors are thus $P_1 = P_2 = \{0, \frac{1}{2}, 1\}$, and all splittings are available on those sets. For each pair of feasible posteriors, the agent takes some optimal action that induces the following payoff for Designer 1:

$$u = \begin{array}{c} \begin{array}{|c|c|c|c|} \hline p_1 = 1 & 0 & 1 & 0 \\ \hline p_1 = 1/2 & 1 & 0 & 1 \\ \hline p_1 = 0 & 0 & 1 & 0 \\ \hline & p_2 = 0 & p_2 = 1/2 & p_2 = 1 \\ \hline \end{array} \end{array}$$

Designer 1 would like to fully reveal his own state when Designer 2 is silent at $p_2 = 1/2$, and Designer 2 would like to reveal his own state if Designer 1 has already revealed.

The payoff function u is neither S_1 -concave nor S_2 -convex. The concavification and convexification

of u are given by the following:

$$\text{cav}_{p_1} u = \begin{array}{|c|c|c|} \hline 0 & 1 & 0 \\ \hline 1 & 1 & 1 \\ \hline 0 & 1 & 0 \\ \hline \end{array} \quad \text{vex}_{p_2} u = \begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 1 & 0 & 1 \\ \hline 0 & 0 & 0 \\ \hline \end{array}$$

If designers can disclose information simultaneously at a single stage, then Designer 1 can guarantee an expected payoff of $\frac{1}{2}$ by revealing the state with probability $\frac{1}{2}$ and remaining silent with probability $\frac{1}{2}$. Indeed, in this case the posterior belief of the agent on Ω_1 is $p_1 = 0$ with probability $\frac{1}{4}$, $p_1 = 1$ with probability $\frac{1}{4}$ and $p_1 = \frac{1}{2}$ with probability $\frac{1}{2}$. Hence, the probability of obtaining the payoff of 1 is equal to $\frac{1}{2}$ for any disclosure policy of Designer 2. Similarly, Designer 2 can guarantee that Designer 1's expected payoff is not higher than $\frac{1}{2}$ by revealing the state with probability $\frac{1}{2}$. This is the splitting value of the model, i.e., the equilibrium payoff of Designer 1 of the one-stage simultaneous move game.

Suppose now that designers play sequentially once and that Designer 1 moves last. In that case, Designer 1 can guarantee a payoff of 1 by playing the opposite of what Designer 2 did (Designer 1 discloses if Designer 2 has not disclosed at stage 1, and does not disclose if Designer 2 has disclosed). This is the *vex cav* value. Similarly, if Designer 2 moves second, Designer 2 can guarantee that Designer 1's payoff is 0 by playing the same way as Designer 1 (Designer 2 discloses if Designer 1 has disclosed before, and does not disclose otherwise). This is the *cav vex* value.

What is the equilibrium if the long information design game? The game is not symmetric. Designer 2 can still apply the strategy above: when Designer 1 discloses, Designer 2 discloses right after, and does not disclose otherwise. Clearly, the resulting payoff is 0 regardless of what Designer 1 does. This is the Mertens-Zamir value of this example. Summarizing, we have the following:¹¹

$$\text{vex}_{p_2} \text{cav}_{p_1} u = \begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 1 & 1 & 1 \\ \hline 0 & 0 & 0 \\ \hline \end{array} \quad \text{cav}_{p_1} \text{vex}_{p_2} u = \begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 1 & 0 & 1 \\ \hline 0 & 0 & 0 \\ \hline \end{array}$$

and

$$\text{SV}(u) = \begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 1 & \frac{1}{2} & 1 \\ \hline 0 & 0 & 0 \\ \hline \end{array} \quad \text{MZ}(u) = \begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 1 & 0 & 1 \\ \hline 0 & 0 & 0 \\ \hline \end{array}$$

3.4.2 Unbounded Disclosure Stages

As in the previous example, for each i , let $\Omega_i = \{0, 1\}$ and identify $p_i \in \Delta(\Omega_i)$ with $p_i(1) \in [0, 1]$. Let $P_i = \{0, 1/3, 2/3, 1\}$ and assume that all splittings on P_i are admissible. Consider the following “matching pennies” utility function for Designer 1:

¹¹In Koessler et al. (2021), we present an extension of this example where $P_1 = \Delta(\Omega_1)$, $P_2 = \Delta(\Omega_2)$, all splittings are admissible, and $\text{SV}(u) \neq \text{MZ}(u)$.

$$u =$$

$p_1 = 1$	0	0	1	1
$p_1 = 2/3$	0	0	1	1
$p_1 = 1/3$	1	1	0	0
$p_1 = 0$	1	1	0	0
	$p_2 = 0$	$p_2 = 1/3$	$p_2 = 2/3$	$p_2 = 1$

Consider, for example, the situation in which $(p_1, p_2) = (1, \frac{2}{3})$. If Designer 2 does not disclose any information, then the utility of Designer 1 is $u(1, \frac{2}{3}) = 1$. The optimal information policy of Designer 2 is to induce the posteriors $p'_2 = 1$ and $p''_2 = \frac{1}{3}$ with the same probability and we have $\text{vex}_{p_2} u(1, \frac{2}{3}) = \frac{1}{2}u(1, \frac{1}{3}) + \frac{1}{2}u(1, 1) = \frac{1}{2}0 + \frac{1}{2}1 = \frac{1}{2}$. More generally, we have the following:

$$\begin{array}{l} \text{vex}_{p_2} u = \begin{array}{|c|c|c|c|} \hline 0 & 0 & \frac{1}{2} & 1 \\ \hline 0 & 0 & \frac{1}{2} & 1 \\ \hline 1 & \frac{1}{2} & 0 & 0 \\ \hline 1 & \frac{1}{2} & 0 & 0 \\ \hline \end{array} \\ \text{cav}_{p_1} \text{vex}_{p_2} u = \begin{array}{|c|c|c|c|} \hline 0 & 0 & \frac{1}{2} & 1 \\ \hline \frac{1}{2} & \frac{1}{4} & \frac{1}{2} & 1 \\ \hline 1 & \frac{1}{2} & \frac{1}{4} & \frac{1}{2} \\ \hline 1 & \frac{1}{2} & 0 & 0 \\ \hline \end{array} \end{array} \quad \begin{array}{l} \text{cav}_{p_1} u = \begin{array}{|c|c|c|c|} \hline 0 & 0 & 1 & 1 \\ \hline \frac{1}{2} & \frac{1}{2} & 1 & 1 \\ \hline 1 & 1 & \frac{1}{2} & \frac{1}{2} \\ \hline 1 & 1 & 0 & 0 \\ \hline \end{array} \\ \text{vex}_{p_2} \text{cav}_{p_1} u = \begin{array}{|c|c|c|c|} \hline 0 & 0 & \frac{1}{2} & 1 \\ \hline \frac{1}{2} & \frac{1}{2} & \frac{3}{4} & 1 \\ \hline 1 & \frac{3}{4} & \frac{1}{2} & \frac{1}{2} \\ \hline 1 & \frac{1}{2} & 0 & 0 \\ \hline \end{array} \end{array}$$

Using $\text{cav}_{p_1} \text{vex}_{p_2} u \leq \text{MZ}(u) \leq \text{vex}_{p_2} \text{cav}_{p_1} u$ and the symmetries of the example, the MZ value function can be written as follows:

$$\text{MZ}(u) = \begin{array}{|c|c|c|c|} \hline 0 & 0 & \frac{1}{2} & 1 \\ \hline \frac{1}{2} & x & y & 1 \\ \hline 1 & y & x & \frac{1}{2} \\ \hline 1 & \frac{1}{2} & 0 & 0 \\ \hline \end{array}$$

where $\frac{1}{4} \leq x \leq \frac{1}{2}$ and $\frac{1}{2} \leq y \leq \frac{3}{4}$. It is easy to see that the solution of the Mertens-Zamir system $\text{MZ}(u)(p_1, p_2) = \text{cav}_{p_1} \min(u, \text{MZ}(u))(p_1, p_2) = \text{vex}_{p_2} \max(u, \text{MZ}(u))(p_1, p_2)$ gives $x = \frac{1}{3}$ and $y = \frac{2}{3}$.¹² From this solution, we immediately obtain the following equilibrium strategies:

- If $p_1 = \frac{2}{3}$ and $p_2 \in \{0, \frac{1}{3}\}$, or if $p_1 = \frac{1}{3}$ and $p_2 \in \{\frac{2}{3}, 1\}$, then Designer 1 splits the belief p_1 to the two neighborhood posteriors with the same probability; otherwise, he discloses no information.
- If $p_2 = \frac{2}{3}$ and $p_1 \in \{\frac{2}{3}, 1\}$, or if $p_2 = \frac{1}{3}$ and $p_1 \in \{0, \frac{1}{3}\}$, then Designer 2 splits the belief p_2 to the two neighborhood posteriors with the same probability; otherwise, he discloses no information.

The equilibrium martingale of posteriors is represented by Figure 1. Note that, as in the ‘‘four frogs’’ example of Forges (1984, 1990), the number of equilibrium disclosure stages in the long information design game is unbounded, but disclosure stops with probability one in finite time, at a profile of posterior beliefs (p_1, p_2) in $\{(0, 1), (0, \frac{2}{3}), (0, 0), (\frac{1}{3}, 0), (\frac{2}{3}, 1), (1, 1), (1, \frac{1}{3}), (1, 0)\}$, represented by a ‘‘*’’ symbol in Figure 1. It is important to emphasize that an unbounded number of stages is required

¹²An algorithm for computing the solution more generally is provided in Koessler et al. (2021).

for this equilibrium, even though every designer is able to induce in a single stage all distributions of posteriors that he would be able to induce by iterating information policies in multiple stages; that is, $S_1^2(p_1) = S_1(p_1)$ and $S_2^2(p_2) = S_2(p_2)$ is satisfied in the example. However, Proposition 2, which guarantees that all the equilibrium information is disclosed in a single stage, does not apply because the sets of admissible posteriors P_1 and P_2 are not convex.




Figure 1: Equilibrium martingale of posteriors in the example of Section 3.4.2.

It is also easy to check that the one-stage splitting value is as follows:

$$SV(u) = \begin{array}{|c|c|c|c|} \hline 0 & 0 & \frac{1}{2} & 1 \\ \hline \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & 1 \\ \hline 1 & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \\ \hline 1 & \frac{1}{2} & 0 & 0 \\ \hline \end{array}$$

Indeed, if a designer splits an interior belief to the two closest posteriors with the same probability when his payoff is different from his first best payoff then, whatever the belief induced by the other designer, the expected utility is $\frac{1}{2}$. For example, when $(p_1, p_2) = (\frac{1}{3}, \frac{2}{3})$, if Designer 1 splits $p_1 = \frac{1}{3}$ to $p_1' = \frac{2}{3}$ with probability $\frac{1}{2}$ and to $p_1'' = 0$ with probability $\frac{1}{2}$, then the expected utility of Designer 1 is $\frac{1}{2}$, and if Designer 2 plays a non-revealing strategy, then the expected utility of Designer 1 is at most $\frac{1}{2}$. Hence, in these situations, the value of the game is $SV(u)(p_1, p_2) = \frac{1}{2}$. Otherwise, the value is 0 or 1. For example, when $(p_1, p_2) = (\frac{2}{3}, 1)$, Designer 2 cannot modify p_2 and Designer 1 gets his first best, so $SV(u)(\frac{2}{3}, 1) = 1$ and no information is disclosed. Similarly, when $(p_1, p_2) = (1, \frac{1}{3})$, Designer 1 cannot modify p_1 and Designer 2 gets his first best, so $SV(u)(1, \frac{1}{3}) = 0$ and no information is disclosed. In this example, for every interior posteriors (p_1, p_2) we have $\text{cav}_{p_1} \text{vex}_{p_2} u(p_1, p_2) < \text{MZ}(u)(p_1, p_2) \neq SV(u)(p_1, p_2) < \text{vex}_{p_2} \text{cav}_{p_1} u(p_1, p_2)$.

4 Competition in Product Demonstration

In this section, we study the case of two designers acting as sellers of products with uncertain match value. The agent is a buyer who chooses the seller from whom to buy the product. The agent prefers to

buy the product with the highest expected match. The static simultaneous-move game $G^{sim}(p_1^0, p_2^0)$ is studied in Boleslavsky and Cotton (2015), who characterize the equilibrium and the value of the game. The model is as follows.

The match value for the product of each seller can be either high (H) or low (L): $\Omega_1 = \Omega_2 = \{L, H\}$. Let $p_1^0 \in (0, 1)$ be the prior probability that the state is H for Designer 1 and $p_2^0 \in (0, 1)$ be the prior probability that the state is H for Designer 2. Denote by $p_1 \in [0, 1]$ and $p_2 \in [0, 1]$ the corresponding posteriors.

The agent must choose either Designer 1 or Designer 2. The agent's payoff is 1 if the chosen designer's state is H and is 0 otherwise. Hence, Designer 1 is chosen if $p_1 > p_2$; Designer 2 is chosen if $p_1 < p_2$, and we assume that the agent randomizes uniformly if $p_1 = p_2$. Designer 1's payoff is equal to 1 if that designer is chosen and is -1 otherwise. Given the posteriors (p_1, p_2) , the expected payoff of Designer 1 is thus as follows:

$$u(p_1, p_2) = \mathbf{1}\{p_1 > p_2\} - \mathbf{1}\{p_1 < p_2\}.$$

Note that $u(p_1, p_2)$ is discontinuous at $p_1 = p_2$. The expected payoff of the agent is as follows:

$$u_A(p_1, p_2) = \max\{p_1, p_2\}.$$

This example can fit the following economic scenario. Designers are two competing firms (e.g., Canon and Nikon) that are about to release a new product or system (e.g., a new camera technology together with plans for lens development and compatibility). The agent is a representative consumer (e.g., a professional photographer) who plans to switch in the near future to one of these two new products. Ex-ante, the agent and the firms do not know the match value for the products (it can be high (H) or low (L) for each product). Suppose that the agent already owns a similar product from an older product generation and is ready to spend some time acquiring match information before making a decision. Firms can release public match information to consumers about their respective products through product demonstration (e.g., press events, reviews, trade fairs, product testing or gradual announcements of a lens roadmap). When choosing how to disclose information, the firm does not know the consumer's match value for the product. Additionally, it does not know in advance the public information feedback (e.g., the ratings of reviews) its product will receive. Finally, firms cannot significantly adjust their prices but are able to adjust their information policies to signals generated by their competitors.¹³ It follows that they have strictly opposite preferences, and both want to attract the consumer.

Below, we characterize and compare the equilibrium strategies and values for the simultaneous, sequential and long information design games. We use those results to analyze the informativeness of firms' strategies and the resulting welfare of the consumer. If we consider finite sets of admissible posteriors, the example fits all our assumptions. For instance, this represents situations in which the consumer reads reviews that apply some rating system (e.g., from one to five stars, or scores

¹³See, e.g., Boleslavsky, Cotton, and Gurnani (2016) for more details on product demonstration and on price flexibility in a similar scenario with a single information designer.

corresponding to relevant features of the product). To compare our results with the equilibria of the static simultaneous game $G^{sim}(p_1^0, p_2^0)$ studied in Boleslavsky and Cotton (2015), we perform the equilibrium analysis assuming that *all* splittings of *all* possible posteriors are admissible. Note that if all splittings are admissible, designers' utility functions are discontinuous on the diagonal $p_1 = p_2$.¹⁴ To deal with this discontinuity, we extend some of the proofs of the previous section to this particular example.¹⁵ In the static simultaneous-move game and in the long information design game, we show that equilibria exist. The 2-stage sequential games admit ε -equilibria for all $\varepsilon > 0$ but not exact equilibria.

Before proceeding to equilibrium analysis, we first consider the fully revealing (FR) and non-revealing (NR) benchmarks. If information is fully revealed to the agent, then the payoff of the latter is 1 if the state is H for Designer 1 or 2, and 0 if the state is L for both designers. The ex-ante expected payoff of the agent is therefore equal to the following:

$$U_A^{FR} = 1 - (1 - p_1^0)(1 - p_2^0) = p_1^0 + p_2^0(1 - p_1^0).$$

This is an upper bound on the ex-ante expected payoff of the agent. A lower bound is obtained with no revelation of information. When no information is revealed, the agent's ex-ante expected payoff is as follows:

$$U_A^{NR} = u_A(p_1^0, p_2^0) = \max\{p_1^0, p_2^0\} < U_A^{FR}.$$

One-Stage Simultaneous Product Demonstration

The equilibrium strategies and payoffs for the one-stage simultaneous-move game have been characterized by Boleslavsky and Cotton (2015). Assume without loss of generality that $p_1^0 \geq p_2^0$. The equilibrium strategies $s_i \in \mathbb{S}_i(p_i^0)$ are unique. For example, if $p_1^0 \leq \frac{1}{2}$,

$$s_1 = U[0, 2p_1^0] \quad \text{and} \quad s_2 = \begin{cases} U[0, 2p_1^0] & \text{with prob. } \frac{p_2^0}{p_1^0}, \\ 0 & \text{with prob. } 1 - \frac{p_2^0}{p_1^0}, \end{cases}$$

where $U[0, 2p_1^0]$ is the uniform distribution over $[0, 2p_1^0]$. The splitting value exists and is given by $SV(u)(p_1^0, p_2^0) = 1 - p_2^0/p_1^0$. If $p_1^0 \leq \frac{1}{2}$, the equilibrium ex-ante expected utility of the agent is as follows:

$$U_A^{SV} = \left(1 - \frac{p_2^0}{p_1^0}\right)p_1^0 + \frac{p_2^0}{p_1^0} \left(\frac{4}{3}p_1^0\right) = p_1^0 + \frac{1}{3}p_2^0.$$

If $p_1^0 > \frac{1}{2}$, the equilibrium ex-ante expected utility of the agent is as follows:

$$U_A^{SV} = \left(\frac{1}{p_1^0} - 1\right) \left[\left(1 - \frac{p_2^0}{p_1^0}\right)(1 - p_1^0) + \frac{p_2^0}{p_1^0} \left(\frac{1}{p_1^0} - 1\right) \frac{4}{3}(1 - p_1^0) + \frac{p_2^0}{p_1^0} \left(2 - \frac{1}{p_1^0}\right) \right] + \left(2 - \frac{1}{p_1^0}\right).$$

¹⁴The designers' utility functions are discontinuous at $p_1 = p_2$ for any tie-breaking rule adopted by the agent.

¹⁵An alternative approach would be to assume a continuous utility by letting the agent tremble and choose Designer 1 with probability that is a continuous function of $p_1 - p_2$, increasing from 0 to 1 (e.g., following a logit rule).

Long Product Demonstration

For every p_1 and p_2 , let

$$v(p_1, p_2) = \frac{p_1 - p_2}{\max(p_1, p_2)}.$$

It is easy to verify that for every p_1 and p_2 we have $v(p_1, p_2) = \text{cav}_{p_1} \min(u, v)(p_1, p_2) = \text{vex}_{p_2} \max(u, v)(p_1, p_2)$ (see Figure 2). Hence, despite the discontinuity of u , there exists a solution to the Mertens-Zamir system. We have the following:

Proposition 5. *The value of the long product demonstration game $G(p_1^0, p_2^0)$ is $v(p_1^0, p_2^0) = (p_1^0 - p_2^0) / \max(p_1^0, p_2^0)$.*

Proof. See Section 5.1, where we derive a more general result in Proposition 6, dealing with discontinuous utilities. \square

The equilibrium strategies are quite intuitive: if the product of a seller has a better perceived match than the product of his competitor, then the latter is forced to react by revealing additional match information. On the contrary, if the match of a seller is perceived as worse than the match of his competitor, then the latter is better off not revealing additional information. Precisely, the equilibrium strategy σ_1 of Designer 1 is such that, given posteriors (p_1, p_2) at stage n , he plays the non-revealing splitting δ_{p_1} if $p_1 \geq p_2$, and the splitting $s_1 = (p_1/p_2)\delta_{p_2} + (1 - p_1/p_2)\delta_0$ if $p_1 < p_2$. The equilibrium strategy σ_2 of Designer 2 is such that, given posteriors (p_1, p_2) at stage n , he plays the non-revealing splitting δ_{p_2} if $p_1 \leq p_2$, and the splitting $s_2 = (p_2/p_1)\delta_{p_1} + (1 - p_2/p_1)\delta_0$ if $p_1 > p_2$. Therefore, σ_1 is as in the proof of Theorem 1 and satisfies $v(s_1, p_2) = v(p_1, p_2) \leq u(s_1, p_2)$.

Even though the values of games $G^{sim}(p_1^0, p_2^0)$ and $G(p_1^0, p_2^0)$ are the same for the designers, the induced equilibrium outcomes are very different. In particular, designers' strategies are always more informative in the one-stage simultaneous-move game than in the long game. The equilibrium payoff of the agent in $G(p_1^0, p_2^0)$ is actually equal to p_1^0 , as in the non-revealing case: For $p_1^0 \geq p_2^0$, $U_A^{MZ} = U_A^{NR} = p_1^0$.

Two-Stage Sequential Product Demonstration

Now, we consider the 2-stage sequential-move games and calculate $\text{cav}_{p_1} \text{vex}_{p_2} u$ and $\text{vex}_{p_2} \text{cav}_{p_1} u$. The concavification of $u(p_1, p_2)$ with respect to p_1 and its convexification with respect to p_2 are given by:

$$\text{cav}_{p_1} u(p_1, p_2) = \begin{cases} 1 & \text{if } p_1 > p_2 \\ -1 + \frac{2p_1}{p_2} & \text{if } p_1 \leq p_2, p_2 \in (0, 1) \\ 0 & \text{if } p_1 = p_2 = 0 \\ -1 + p_1 & \text{if } p_2 = 1 \end{cases} \quad \text{vex}_{p_2} u(p_1, p_2) = \begin{cases} -1 & \text{if } p_2 > p_1 \\ 1 - \frac{2p_2}{p_1} & \text{if } p_2 \leq p_1, p_1 \in (0, 1) \\ 0 & \text{if } p_1 = p_2 = 0 \\ 1 - p_2 & \text{if } p_1 = 1. \end{cases}$$

Note that $\text{cav}_{p_1} u$ is discontinuous at $p_1 = p_2 = 0$ and $p_2 = 1$, and $\text{vex}_{p_2} u$ is discontinuous at $p_1 = p_2 = 0$ and $p_1 = 1$. If only Designer 1 can disclose information at (p_1^0, p_2^0) , the equilibrium is non-revealing. If only Designer 2 can disclose information, that designer splits p_2^0 to p_1^0 (more precisely, to $p_1^0 + \varepsilon$, with $\varepsilon \rightarrow 0$) with probability p_2^0/p_1^0 and to 0 with probability $1 - p_2^0/p_1^0$. In both cases, the agent's expected




Figure 2: MZ function $v(\cdot, p_2)$ for a given p_2 .

payoff is the non-revealing payoff:

$$U_A^{\text{cav}_{p_1}} = U_A^{\text{vex}_{p_2}} = U_A^{NR} = p_1^0.$$

To compute $\text{cav}_{p_1} \text{vex}_{p_2} u$, fix p_2 and consider the optimal splitting of p_1 for Designer 1 when the utility of the latter is given by $\text{vex}_{p_2} u(p_1, p_2)$. We get (see Figure 3):

$$\text{cav}_{p_1} \text{vex}_{p_2} u(p_1, p_2) = \begin{cases} 0 & \text{if } p_1 = p_2 = 0 \\ -1 + p_1(2 - p_2) & \text{if } 2 - \sqrt{2} \leq 2p_2 \\ 1 - \frac{2p_2}{p_1} & \text{if } 2p_2 \leq p_1 \leq 2 - \sqrt{2} \text{ and } p_1 > 0 \\ -1 + \frac{p_1}{2p_2} & \text{if } p_1 \leq 2p_2 < 2 - \sqrt{2} \\ 1 - \frac{p_2(4 - p_1 - 2\sqrt{2})}{3 - 2\sqrt{2}} & \text{if } 2p_2 \leq 2 - \sqrt{2} \leq p_1. \end{cases}$$




Figure 3: Concavification of $\text{vex}_{p_2} u(p_1, p_2)$.

Similarly, we get:

$$\text{vex}_{p_2} \text{cav}_{p_1} u(p_1, p_2) = \begin{cases} 0 & \text{if } p_1 = p_2 = 0 \\ 1 - p_2(2 - p_1) & \text{if } 2 - \sqrt{2} \leq 2p_1 \\ -1 + \frac{2p_1}{p_2} & \text{if } 2p_1 \leq p_2 \leq 2 - \sqrt{2} \text{ and } p_2 > 0 \\ 1 - \frac{p_2}{2p_1} & \text{if } p_2 \leq 2p_1 < 2 - \sqrt{2} \\ -1 + \frac{p_1(4 - p_2 - 2\sqrt{2})}{3 - 2\sqrt{2}} & \text{if } 2p_1 \leq 2 - \sqrt{2} \leq p_2. \end{cases}$$

These computations allow to derive ε -equilibrium strategies as $\varepsilon \rightarrow 0$, and the following expected equilibrium payoff of the agent in the 2-stage sequential game as a function of the priors:

$$U_A^{\text{cav}_{p_1} \text{vex}_{p_2}} = \begin{cases} U_A^{FR} & \text{if } p_2^0 > \frac{2-\sqrt{2}}{2} \\ U_A^{NR} & \text{if } p_2^0 < \frac{2-\sqrt{2}}{2} \text{ and } p_1^0 > 2p_2^0 \\ \frac{p_1^0}{2} + p_2^0 \in (U_A^{NR}, U_A^{SV}) & \text{if } p_2^0 < \frac{2-\sqrt{2}}{2} \text{ and } \frac{4}{3}p_2^0 < p_1^0 < 2p_2^0 \\ \frac{p_1^0}{2} + p_2^0 \in (U_A^{SV}, U_A^{FR}) & \text{if } p_2^0 < \frac{2-\sqrt{2}}{2} \text{ and } p_1^0 < \frac{4}{3}p_2^0, \end{cases}$$

$$U_A^{\text{vex}_{p_2} \text{cav}_{p_1}} = \begin{cases} U_A^{FR} & \text{if } p_1^0 > \frac{2-\sqrt{2}}{2} \\ U_A^{NR} & \text{if } p_1^0 < \frac{2-\sqrt{2}}{2} \text{ and } p_2^0 > 2p_1^0 \\ \frac{p_2^0}{2} + p_1^0 \in (U_A^{NR}, U_A^{SV}) & \text{if } p_1^0 < \frac{2-\sqrt{2}}{2} \text{ and } \frac{4}{3}p_1^0 < p_2^0 < 2p_1^0 \\ \frac{p_2^0}{2} + p_1^0 \in (U_A^{SV}, U_A^{FR}) & \text{if } p_1^0 < \frac{2-\sqrt{2}}{2} \text{ and } p_2^0 < \frac{4}{3}p_1^0. \end{cases}$$

If p_1^0 and p_2^0 are greater than $(2 - \sqrt{2})/2$, the agent receives the fully revealing payoff U_A^{FR} that is higher than the agent's payoff U_A^{SV} in one-stage simultaneous-move game, which itself is higher than the agent's expected payoff in the long information design game. The sequential game also gives a higher expected payoff to the agent than does the simultaneous game if p_1^0 and p_2^0 are less than $(2 - \sqrt{2})/2$ but close to each other (if $\frac{3}{4}p_2^0 < p_1^0 < \frac{4}{3}p_2^0$). Otherwise, if p_1^0 and p_2^0 are less than $(2 - \sqrt{2})/2$, the agent prefers the sequential game in which Designer 1 plays last to the simultaneous game if $p_1^0 > \frac{3}{4}p_2^0$, and prefers the sequential game in which Designer 2 plays last to the simultaneous game if $p_1^0 > \frac{4}{3}p_2^0$.

We summarize the above comparisons as follows:

1. *The buyer prefers that sellers can disclose information at a predetermined stage.* Indeed, designers' equilibrium strategies are less informative in the long information design game: only one designer discloses information, and the expected payoff of the agent is the same as that without information disclosure.
2. *The buyer prefers that the seller with the highest ex-ante expected match value is allowed to disclose last.* Indeed, the agent is better off in the sequential game than in the simultaneous game whenever the designer with the highest ex-ante value has the last-mover advantage.

5 Extensions and Further Results

5.1 Discontinuous Utilities

In the example studied in Section 4, the utility function is discontinuous. Nevertheless, we are able to characterize the value of the game and determine equilibrium strategies using the intuition from Theorem 1. The next result gives sufficient conditions under which the existence of an MZ function implies that it is the value of the game.

Proposition 6. *Assume that there exists a function $v : P_1 \times P_2 \rightarrow \mathbb{R}$ that is S_1 -concave, S_2 -convex and satisfies (C1) and (C2) for all $(p_1, p_2) \in P_1 \times P_2$. Suppose further that v is u.s.c. with respect to p_2 and l.s.c. with respect to p_1 , that the sets*

$$\{(p_1, p_2) \in P_1 \times P_2 : v(p_1, p_2) \leq u(p_1, p_2)\} \text{ and } \{(p_1, p_2) \in P_1 \times P_2 : v(p_1, p_2) \geq u(p_1, p_2)\},$$

are closed and that v is l.s.c. on the closure of $\{(p_1, p_2) \in P_1 \times P_2 : v(p_1, p_2) < u(p_1, p_2)\}$ and u.s.c. on the closure of $\{(p_1, p_2) \in P_1 \times P_2 : v(p_1, p_2) > u(p_1, p_2)\}$. Then, $v(p_1^0, p_2^0)$ is the value of the long information design game $G(p_1^0, p_2^0)$, and the equilibrium is as constructed in the proof of Theorem 1.

There, we show that the stationary strategies given by (C1) and (C2) form an equilibrium of the game under the above weakening of continuity. These conditions are met in the example of Section 4.

5.2 Alternative Extensive Forms

We can extend our model to alternating-move information design games in which designers cannot disclose information simultaneously. The description of the game is the same except that Designer 1 can move only at odd stages, and Designer 2 can move only at even stages. If there are two consecutive stages $N^* - 1$ and N^* such that Designer i chooses a non-revealing splitting at stage $N^* - 1$ and Designer $j \neq i$ chooses a non-revealing splitting at stage N^* , then the agent chooses an action after stage N^* . Otherwise, both designers choose admissible splittings until $N^* = +\infty$ and the agent chooses an action after stage N^* . It is easy to see that Theorem 1 extends directly to this version, and the value of this information design game is also $\text{MZ}(u)(p_1^0, p_2^0)$.

In another version of the game, the game does not stop even if there is no revelation from both designers at some stage. The timing of such game is similar to that of a “long cheap talk” game as defined in Aumann and Hart (2003). This formulation of the game is formally equivalent to our model, and Theorem 1 extends.

In a last version of the game, the game terminates after each stage with exogenous probability $1 - \delta \in (0, 1)$, and the agent takes an action. With probability δ , the game continues to the next stage. This game is equivalent to a discounted game in which a short-lived agent makes a decision at each stage and designers maximize/minimize the discounted average payoff. This game is a discounted stochastic game called a splitting game in Laraki (2001a) and a gambling game in Laraki and Renault (2020). Lemmas 9 and 10 in Appendix A.1 show that this game has value V_δ , and $\lim_{\delta \rightarrow 1} V_\delta = \text{MZ}(u)(p_1^0, p_2^0)$. Thus, our game $G(p_1^0, p_2^0)$ can be approximated by a discounted game with a high discount factor.

5.3 Correlated Information

Our results also extend to correlated priors. The first approach consists of analyzing a modified game with independent priors. Specifically, given any prior $\mu^0 \in \Delta(\Omega_1 \times \Omega_2)$ and utility functions $\tilde{u}_i(z; \omega_1, \omega_2)$, where i denotes Designer 1, Designer 2 or the agent, we define a modified game with stochastically independent states by letting $\hat{\mu}^0(\omega_1, \omega_2) = (\sum_{\tilde{\omega}_2 \in \Omega_2} \mu^0(\omega_1, \tilde{\omega}_2)) \times (\sum_{\tilde{\omega}_1 \in \Omega_1} \mu^0(\tilde{\omega}_1, \omega_2))$ and for every i ,

$$\hat{u}_i(z; \omega_1, \omega_2) = \frac{\mu^0(\omega_1, \omega_2)}{\hat{\mu}^0(\omega_1, \omega_2)} \tilde{u}_i(z; \omega_1, \omega_2),$$

for every (ω_1, ω_2) in the support of μ^0 .¹⁶ This modification preserves the continuity of players’ utility functions and the zero-sum property between the two designers. The equilibria of this modified multistage information design game (without correlation) coincide with the equilibria of the original multistage information design game (with correlation). The second approach consists of defining concavification and convexification for correlated distributions and extending the notions of cav_{p_1} , vex_{p_2} and MZ . The reader is referred to Ponsard and Sorin (1980), Sorin (1984), Mertens et al. (2015) and Orlu-Barton (2017).

Consider now the particular case in which designers disclose information about a common payoff-relevant state; i.e., their private states are perfectly correlated, as in Gentzkow and Kamenica (2017).

¹⁶Such transformations can be found in Aumann et al. (1995, Chap. 2, section 4.2, pages 100–104) and in Myerson (1985).

Denote by $\Omega = \Omega_1 = \Omega_2$ the common set of states and assume that all information disclosure policies are available. In all versions of the information design game, full revelation is an equilibrium. Indeed, when one designer fully discloses the state, the other is indifferent and may fully disclose as well. Thus, the value of each version of the game is that obtained by full revelation. This value coincides with the $\text{cav}_{p_1} \text{vex}_{p_2}$ and $\text{vex}_{p_2} \text{cav}_{p_1}$ values, and thus also with the SV and MZ values. Indeed, all of those functions have to be both concave and convex with respect to the prior. Therefore, they must be linear and equal to $\sum_{\omega} \mu(\omega)u(z(\omega), \omega)$, where $z(\omega)$ is the optimal action of the agent when the agent knows that the state is ω .

5.4 Unobserved Messages

Consider an alternative model in which posteriors are not observed at each stage. More precisely, designers only observe splittings chosen at each stage, until the game between the designers stops. Then, the sequence of splittings chosen is drawn and posteriors are finally privately observed before the agent chooses an action. Contrary to the case where posteriors are publicly observed at each stage by the designers, all the versions of this game, one-stage simultaneous, two-stage sequential or long, have the same value and optimal strategies, which are the splitting value and the one-stage optimal strategies. In particular, when posteriors are privately observed by the agent, it does not matter for him how he would acquire information from competing designers. This observation directly follows from the fact that

$$\text{SV}(u)(p_1, p_2) = \max_{s_1 \in S_1(p_1)} \min_{s_2 \in S_2(p_2)} u(s_1, s_2) = \min_{s_2 \in S_2(p_2)} \max_{s_1 \in S_1(p_1)} u(s_1, s_2).$$

That is, in every version of the game with unobserved messages, Designer 1 can guarantee that his expected utility is at least $\text{SV}(u)(p_1, p_2)$ by playing his one-stage equilibrium strategy in his first move, and Designer 2 can guarantee that the expected utility of Designer 1 is at most $\text{SV}(u)(p_1, p_2)$ by playing his one-stage equilibrium strategy in his first move. Hence, $\text{SV}(u)(p_1, p_2)$ is the value.

It is worth noting that the previous observation does not apply if one considers more general utility functions. To see this, consider the following non-zero sum example, where the set of admissible posteriors is $P_1 = P_2 = \{0, 1/2, 1\}$, all splittings on P_1 and P_2 are admissible, the priors are $p_1^0 = p_2^0 = \frac{1}{2}$, and the utilities of the designers are given by:

$p_1 = 1$	1, 1	4, 0	1, 1
$p_1 = 1/2$	0, 4	3, 3	0, 4
$p_1 = 0$	1, 1	4, 0	1, 1
	$p_2 = 0$	$p_2 = 1/2$	$p_2 = 1$

The one-stage game has similarities with the prisoner dilemma: it is a dominant strategy for both designers to disclose information and, by backward induction, it is the unique equilibrium whatever the finite and fixed number of stages. However, it is easy to see that the infinite horizon game (whether the realized posteriors are observed or not) has a (subgame perfect) equilibrium sustaining no information revelation. Indeed, along the equilibrium path both designers play non-revealing, and they punish by

disclosing information if they observe a deviation from the non-revealing splitting.

A Appendix

A.1 Proofs

Proof of Lemma 1 We show that $\text{cav}_{p_1} w(p_1, p_2) = \max\{w(s_1, p_2) : s_1 \in S_1(p_1)\}$. Fix p_2 and let $h(p_1) := \max\{w(s_1, p_2) : s_1 \in S_1(p_1)\}$. Consider a function $g : P_1 \rightarrow \mathbb{R}$, S_1 -concave such that $g(\cdot) \geq w(\cdot, p_2)$. This pointwise inequality implies that for any $s_1 \in S_1(p_1)$, $g(s_1) \geq w(s_1, p_2)$ and since g is S_1 -concave,

$$g(p_1) \geq g(s_1) \geq w(s_1, p_2).$$

Since this holds for any $s_1 \in S_1(p_1)$, this implies $g(p_1) \geq h(p_1)$, and since this holds for any $g(\cdot) \geq w(\cdot, p_2)$, we obtain $\text{cav}_{p_1} w(p_1, p_2) \geq h(p_1)$.

To obtain the converse inequality, since $h(p_1) \geq w(p_1, p_2)$, it is enough to prove that h is S_1 -concave. Hence, take $s_1 \in S_1(p_1)$ and consider $h(s_1) = \int h(p'_1) ds_1(p'_1)$. Define a measurable selection of S_1 , $f : P_1 \rightarrow \Delta(P_1)$ such that $f(p'_1) \in \arg \max\{u(s'_1, p'_1) : s'_1 \in S_1(p'_1)\}$. We have for each $p'_1 \in P_1$ that $h(p'_1) = w(f(p'_1), p_2) = \int w(p''_1, p_2) df(p''_1|p'_1)$. We then obtain the following:

$$h(s_1) = \iint w(p''_1, p_2) df(p''_1|p'_1) ds_1(p'_1) = \int w(\tilde{p}_1, p_2) d(f * s_1)(\tilde{p}_1) = w(f * s_1, p_2) \leq h(p_1),$$

since $f * s_1 \in S_1(p_1)$ by assumption. We thus have $\text{cav}_{p_1} w(p_1, p_2) = \max\{w(s_1, p_2) : s_1 \in S_1(p_1)\}$ for all p_1 and p_2 . \square

Proof of Lemma 2

1. *Continuity.* From Lemma 1, $\text{vex}_{p_2} u$, $\text{cav}_{p_1} u$, $\text{cav}_{p_1} \text{vex}_{p_2} u$, $\text{vex}_{p_2} \text{cav}_{p_1} u$ are continuous. The correspondence $(p_1, p_2) \mapsto S_1(p_1) \times S_2(p_2)$ is both upper and lower hemicontinuous; hence, by the Maximum Theorem $\text{SV}(u)$ is continuous.

2. *The function $\text{cav}_{p_1} \text{vex}_{p_2} u$ is S_2 -convex.* Denote $F(p_1, p_2) = \text{vex}_{p_2} u(p_1, p_2)$ and take $s_1 \in S_1(p_1)$ such that

$$\text{cav}_{p_1} \text{vex}_{p_2} u(p_1, p_2) = F(s_1, p_2) = \int F(p'_1, p_2) ds_1(p'_1).$$

Since F is S_2 -convex, $F(p'_1, p_2) \leq F(p'_1, s_2)$ for each $s_2 \in S_2(p_2)$. Thus,

$$\begin{aligned} \text{cav}_{p_1} \text{vex}_{p_2} u(p_1, p_2) &= \int F(p'_1, p_2) ds_1(p'_1) \\ &\leq \int F(p'_1, s_2) ds_1(p'_1) \\ &\leq \int \text{cav}_{p_1} \text{vex}_{p_2} u(p'_1, s_2) ds_1(p'_1) \\ &\leq \text{cav}_{p_1} \text{vex}_{p_2} u(p_1, s_2), \end{aligned}$$

where the last inequality holds due to $\text{cav}_{p_1} \text{vex}_{p_2} u$ being S_1 -concave. By symmetry, $\text{vex}_{p_2} \text{cav}_{p_1} u$ is S_1 -concave.

3. *The function $\text{SV}(u)$ is S_1 -concave.* Fix (p_1, p_2) in $P_1 \times P_2$ and $s_1 \in S_1(p_1)$; we have to show that $\text{SV}(u)(p_1, p_2) \geq \text{SV}(u)(s_1, p_2)$. Let $f : P_1 \rightarrow \Delta(P_1)$ be a measurable selection of S_1 such that for each p'_1 in p_1 , $f(p'_1) \in S_1(p'_1)$ is an optimal strategy of Designer 1 in game $G^{\text{sim}}(p'_1, p_2)$. We have for each p'_1 in p_1 that $\forall s_2 \in S_2(p_2)$, $u(f(p'_1), s_2) \geq \text{SV}(u)(p'_1, p_2)$. Taking expectation with respect to s_1 implies the following:

$$\forall s_2 \in S_2(p_2), \int u(f(p'_1), s_2) ds_1(p'_1) \geq \int \text{SV}(u)(p'_1, p_2) ds_1(p'_1) = \text{SV}(u)(s_1, p_2).$$

We get $\forall s_2 \in S_2(p_2)$, $u(f * s_1, s_2) \geq \text{SV}(u)(s_1, p_2)$. Since $f * s_1 \in S_1(p_1)$, we obtain that the value $\text{SV}(u)(p_1, p_2)$ of $G^{\text{sim}}(p_1, p_2)$ is at least $\text{SV}(u)(s_1, p_2)$. Hence, $\text{SV}(u)$ is S_1 -concave.

4. *For each $(p_1, p_2) \in P_1 \times P_2$, $\text{cav}_{P_1} \text{vex}_{P_2} u(p_1, p_2) \leq \text{SV}(u)(p_1, p_2)$.* It suffices to show that there exists $s_1 \in S_1(p_1)$ such that for all $s_2 \in S_2(p_2)$, $u(s_1, s_2) \geq \text{cav}_{P_1} \text{vex}_{P_2} u(p_1, p_2)$. Choose s_1 such that $\text{cav}_{P_1} \text{vex}_{P_2} u(p_1, p_2) = \text{vex}_{P_2} u(s_1, p_2)$. Then,

$$\forall s_2 \in S_2(p_2), u(s_1, s_2) \geq \text{vex}_{P_2} u(s_1, p_2) = \text{cav}_{P_1} \text{vex}_{P_2} u(p_1, p_2).$$

The other inequalities are either trivial or deduced by symmetry. \square

Proof of Lemma 3 To prove the inequality $\text{cav}_{P_1} \text{vex}_{P_2} u(p_1, p_2) \leq \text{MZ}(u)(p_1, p_2)$, observe that $\max(u, \text{MZ}(u)) \geq u$; thus,

$$\text{MZ}(u) = \text{vex}_{P_2} \max(u, \text{MZ}(u)) \geq \text{vex}_{P_2} u.$$

Since $\text{MZ}(u)$ is S_1 -concave, this implies $\text{MZ}(u) \geq \text{cav}_{P_1} \text{vex}_{P_2} u$. The other inequality is obtained by symmetry. \square

Proof of Proposition 1 We first prove the following lemma that states a useful property of concavification and optimal splittings: at an optimal splitting s_1 such that $\text{cav}_{P_1} w(p_1) = w(s_1)$, it must be that $\text{cav}_{P_1} w(p'_1) = w(p'_1)$ on the support of s_1 . This implies that

$$\text{cav}_{P_1} w(s_1) = \int \text{cav}_{P_1} w(p'_1) ds_1(p'_1) = \int w(p'_1) ds_1(p'_1) = \text{cav}_{P_1} w(p_1).$$

Therefore, $\text{cav}_{P_1} w$ is “linear” on the support of S_1 .

Lemma 6. *Let $w : P_1 \rightarrow \mathbb{R}$ be a continuous function. For each $p_1 \in P_1$ and each $s_1 \in S_1(p_1)$ such that $\text{cav}_{P_1} w(p_1) = w(s_1) = \max\{w(s'_1) : s'_1 \in S_1(p_1)\}$, we have the following:*

$$s_1(\{p'_1 \in P_1 : w(p'_1) = \text{cav}_{P_1} w(p'_1)\}) = 1.$$

Proof. Assume the contrary that $\text{cav}_{P_1} w(p_1) = w(s_1)$ and $s_1(\{p'_1 \in P_1 : \text{cav}_{P_1} w(p'_1) > w(p'_1)\}) > 0$. Then, there exist $\varepsilon, \alpha > 0$ such that $s_1(\{p'_1 \in P_1 : \text{cav}_{P_1} w(p'_1) \geq w(p'_1) + \varepsilon\}) = \alpha > 0$. Let $B = \{p'_1 \in P_1 : \text{cav}_{P_1} w(p'_1) \geq w(p'_1) + \varepsilon\}$; define a measurable selection f of S_1 such that $\text{cav}_{P_1} w(p'_1) = w(f(p'_1))$

for all $p'_1 \in B$ and $f(p'_1) = \delta_{p'_1}$ for all $p'_1 \notin B$, and consider the splitting $f * s_1$. We have the following:

$$w(f * s_1) = \int_B \text{cav}_{p_1} w(p'_1) ds_1(p'_1) + \int_{P_1 \setminus B} w(p'_1) ds_1(p'_1) \geq \varepsilon \alpha + \int w(p'_1) ds_1(p'_1).$$

Since $f * s_1 \in S_1(p_1)$, this contradicts $w(s_1) = \max\{w(s'_1) : s'_1 \in S_1(p_1)\}$. \square

To complete the proof of the proposition, we first show that if $v = \text{MZ}(u)$, then v is S_1 -concave and S_2 -convex, and (C1) and (C2) hold. $\text{MZ}(u)$ is S_1 -concave and S_2 -convex; by symmetry, it is enough to prove (C1). Suppose that $v = \text{cav}_{p_1} \min(u, v)$ and consider some $(p_1, p_2) \in P_1 \times P_2$. According to Lemma 6, there exists $s_1 \in S_1(p_1)$ such that $v(p_1, p_2) = \min(u, v)(s_1, p_2)$ and $v(p'_1, p_2) = \min(u, v)(p'_1, p_2)$, $\forall p'_1 \in \text{supp}(s_1)$. It follows that $v(p'_1, p_2) \leq u(p'_1, p_2)$, $\forall p'_1 \in \text{supp}(s_1)$ and $v(p_1, p_2) = v(s_1, p_2)$, so (C1) holds.

Second, we show that if v is S_1 -concave and S_2 -convex and (C1) and (C2) hold, then $v = \text{MZ}(u)$. Suppose that v is S_1 -concave and S_2 -convex, and let $w(p_1, p_2) = \min(u, v)(p_1, p_2)$. Since $w \leq v$, $\text{cav}_{p_1} w \leq \text{cav}_{p_1} v$. According to (C1), there exists $s_1 \in S_1(p_1)$ such that $v(p_1, p_2) = v(s_1, p_2)$ and $w(p'_1, p_2) = v(p'_1, p_2) \leq u(p'_1, p_2)$ for all p'_1 in the support of S_1 . Hence, $v(p_1, p_2) = v(s_1, p_2) = w(s_1, p_2) \leq \text{cav}_{p_1} w(p_1, p_2)$. Thus, $v = \text{cav}_{p_1} \min(u, v)$ and by symmetry, $v = \text{vex}_{p_2} \max(u, v)$. \square

For the proofs of the next four results we let $i = 1$. The proofs are similar for $i = 2$.

Proof of Lemma 4 Consider a continuous function $w : P_1 \rightarrow \mathbb{R}$. For each p_1 such that $w(p_1) < \text{cav}_{p_1} w(p_1)$, there exists a splitting with finite support $s_1 = \sum_m \lambda_m \delta_{p_{1m}}$ such that $\text{cav}_{p_1} w(p_1) = \sum_m \lambda_m w(p_{1m})$ and $w(p'_1) < \text{cav}_{p_1} w(p'_1)$ for each p'_1 in the relative interior of the convex hull of $\{p_{1m} : m\}$. The reason is that $(p_1, \text{cav}_{p_1} w(p_1))$ lies on a face of the convex hull hypograph of the function w if $w(p_1) < \text{cav}_{p_1} w(p_1)$. One just has to obtain this point as a “minimal” convex combination of extreme points of the hypograph (minimal in the sense that the convex hull of its extreme points is minimal for inclusion) among the splittings s_1 such that $\text{cav}_{p_1} w(p_1) = w(s_1)$. Fix p_2 and apply this logic to $v(p_1, p_2) = \text{cav}_{p_1} \min(u, v)(p_1, p_2)$. We obtain $v(p_1, p_2) = \sum_m \lambda_m \min(u, v)(p_{1m}, p_2)$, with $v(p_{1m}, p_2) \leq u(p_{1m}, p_2)$ for all m and $\min(u, v)(p'_1, p_2) < v(p'_1, p_2)$ for all p'_1 in the relative interior of the convex hull of $\{p_{1m} : m\}$. By continuity, this implies that $v(p_{1m}, p_2) = u(p_{1m}, p_2)$ for all m . \square

Proof of Lemma 5 \implies Take f in D_θ^1 , p_1, p'_1 in P_1 , $a, b \geq 0$. Also, let $s_1 \in S_1(p_1)$ be such that $f(s_1) = \text{cav}_{p_1} f(p_1)$. By Assumption 1, there exists $s'_1 \in S_1(p'_1)$ such that $af(s_1) - bf(s'_1) \leq \|ap_1 - bp'_1\|_1$. Thus, $a \text{cav}_{p_1} f(p_1) - b \text{cav}_{p_1} f(p'_1) \leq af(s_1) - bf(s'_1) \leq \|ap_1 - bp'_1\|_1$. Applying this logic to $-f$ yields the result.

\Leftarrow Fix p_1, p'_1 in P_1 , $s_1 \in S_1(p_1)$, $a, b \geq 0$. Define $\gamma : S(p'_1) \times D_\theta^1 \rightarrow \mathbb{R}$ by

$$\gamma(s'_1, f) = af(s_1) - bf(s'_1) - \|ap_1 - bp'_1\|_1.$$

For each f , there exists $s'_1 \in S(p'_1)$ such that $f(s'_1) = \text{cav}_{p_1} f(p'_1)$. Then

$$\gamma(s'_1, f) \leq a \text{cav}_{p_1} f(p_1) - b \text{cav}_{p_1} f(p'_1) - \|ap_1 - bp'_1\|_1 \leq 0.$$

From Sion's minmax theorem, there exists $s'_1 \in S_1(p'_1), \forall f \in D_\theta^1, \gamma(s'_1, f) \leq 0$. Considering f and $-f$ gives the result. \square

Proof of Proposition 3 (1) $P_1 = \Delta(\Omega_1)$, all splittings admissible. Fix p_1, p'_1 in $\Delta(\Omega_1)$, $f \in D_\theta^1$, $a, b \geq 0$. Consider $s_1 = \sum_m \lambda_m \delta_{p_{1m}}$ a splitting of p_1 with finite support such $\text{cav}_{P_1} f(p_1) = f(s_1)$. This splitting is induced by the experiment which sends message m with probability $\lambda_m^{\omega_1} = \lambda_m \frac{p_{1m}(\omega_1)}{p_1(\omega_1)}$ in state ω_1 . Consider now the same experiment used at prior p'_1 , and denote by $s'_1 = \sum_m \lambda'_m \delta_{p'_{1m}}$ the splitting induced by that experiment at p'_1 . We have $p'_1 = \sum_m \lambda'_m \delta_{p'_{1m}}$, with

$$\lambda'_m = \sum_{\omega_1} \lambda_m p_{1m}(\omega_1) \frac{p'_1(\omega_1)}{p_1(\omega_1)} \text{ and } p'_{1m}(\omega_1) = \frac{p'_1(\omega_1) \lambda_m^{\omega_1}}{\lambda'_m}.$$

Then,

$$\begin{aligned} a \text{cav}_{P_1} f(p_1) - b \text{cav}_{P_1} f(p'_1) &\leq a f(s_1) - b f(s'_1) = \sum_m a \lambda_m f(p_{1m}) - b \lambda'_m f(p'_{1m}) \leq \sum_m \|a \lambda_m p_{1m} - b \lambda'_m p'_{1m}\|_1 \\ &= \sum_m \sum_{\omega_1} |a \lambda_m p_{1m}(\omega_1) - b \lambda'_m p'_{1m}(\omega_1)| = \sum_m \sum_{\omega_1} |a p_1(\omega_1) \lambda_m^{\omega_1} - b \lambda_m^{\omega_1} p'_1(\omega_1)| = \|a p_1 - b p'_1\|_1. \end{aligned}$$

(2) P_1 is finite. Assume that P_1 contains at least two points. Define $\theta \in (0, 1]$ by

$$\theta = \min\{\|\alpha p_1 - (1 - \alpha) p'_1\|_1, \alpha \in [0, 1], p_1, p'_1 \in P_1, p_1 \neq p'_1\}.$$

Fix p_1, p'_1 in P_1 , $f \in D_\theta^1$, $a, b \geq 0$. If $p_1 \neq p'_1$, we have for any $s_1 \in S_1(p_1)$, $s'_1 \in S_1(p'_1)$,

$$a f(s_1) - b f(s'_1) \leq (a + b) \theta \text{ and } \|a p_1 - b p'_1\|_1 \geq (a + b) \theta.$$

If $p_1 = p'_1$, given $s_1 \in S_1(p_1)$, we have $a f(s_1) - b f(s_1) \leq |a - b| \theta \leq |a - b| \|p_1\|_1$, since $\theta \leq 1$. \square

Proof of Proposition 4

Lemma 7. Let M be a non empty compact subset of an Euclidean space (set of messages), and let $x : \Omega_1 \rightarrow \Delta(M)$ be fixed. For p_1 in $\Delta(\Omega_1)$, let $\varphi(p_1)$ in $\Delta(\Delta(\Omega_1))$ be the splitting of p_1 induced by x , i.e.,

$$\varphi(p_1) = \int_{m \in M} \delta_{\nu(m)} d\zeta(m),$$

where ζ is the distribution of messages induced by p_1 and $\nu(m)$ is a conditional probability on Ω_1 given m . We have for any $\theta \in (0, 1]$:

$$\forall p_1, p'_1 \in \Delta(\Omega_1), \forall f \in D_\theta^1, \forall a, b \geq 0, |a f(\varphi(p_1)) - b f(\varphi(p'_1))| \leq \|a p_1 - b p'_1\|_1.$$

Proof. Define the probability λ on M by $\lambda(B) = (1/|\Omega_1|) \sum_{\omega_1 \in \Omega_1} x(B|\omega_1)$ for all Borel subsets B of M . By the Radon-Nikodym theorem, for each $\omega_1 \in \Omega_1$, there exists a measurable $g^{\omega_1} : M \rightarrow \mathbb{R}_+$ such that $dx(m|\omega_1) = g^{\omega_1}(m) d\lambda(m)$, i.e., for each B , $x(B|\omega_1) = \int_{m \in B} g^{\omega_1}(m) d\lambda(m)$.

Fix now p_1 and p'_1 . Let ζ be the distribution of m induced by p_1 : $d\zeta(m) = g(m)d\lambda(m)$, with $g(m) = \sum_{\omega_1 \in \Omega_1} p_1(\omega_1)g^{\omega_1}(m)$. A conditional distribution of ω_1 given m is given by the measurable map $\nu : M \rightarrow \Delta(\Omega_1)$ such that for all ω_1 in Ω_1 and Borel $B \subseteq M$:

$$p_1(\omega_1)x(B|\omega_1) = \int_{m \in B} \nu(\omega_1|m)d\zeta(m).$$

So for all (ω_1, m) , $p_1(\omega_1)g^{\omega_1}(m) = g(m)\nu(\omega_1|m)$. Similarly, let ζ' be the distribution of m induced by p'_1 : $d\zeta'(m) = g'(m)d\lambda(m)$, with $g'(m) = \sum_{\omega_1 \in \Omega_1} p'_1(\omega_1)g^{\omega_1}(m)$. Consider $\nu' : M \rightarrow \Delta(\Omega_1)$ such that for all ω_1 in Ω_1 and Borel $B \subseteq M$:

$$p'_1(\omega_1)x(B|\omega_1) = \int_{m \in B} \nu'(\omega_1|m)d\zeta'(m).$$

So for all (ω_1, m) , $p'_1(\omega_1)g^{\omega_1}(m) = g'(m)\nu'(\omega_1|m)$.

Fix f in D_θ^1 , $a, b \geq 0$. Since $\varphi(p_1) = \int_{m \in M} \delta_{\nu(m)}d\zeta(m)$,

$$f(\varphi(p_1)) = \int_{m \in M} f(\nu(m))d\zeta(m) = \int_{m \in M} f(\nu(m))g(m)d\lambda(m).$$

Therefore,

$$\begin{aligned} af(\varphi(p_1)) - bf(\varphi(p'_1)) &= \int_{m \in M} (af(\nu(m))g(m) - bf(\nu'(m))g'(m))d\lambda(m), \\ &\leq \int_{m \in M} \|ag(m)\nu(m) - b\nu'(m)g'(m)\|_1 d\lambda(m), \\ &= \int_{m \in M} \sum_{\omega_1 \in \Omega_1} |ap_1(\omega_1)g^{\omega_1}(m) - bp'_1(\omega_1)g^{\omega_1}(m)| d\lambda(m), \\ &= \sum_{\omega_1 \in \Omega_1} \int_{m \in M} g^{\omega_1}(m) |ap_1(\omega_1) - bp'_1(\omega_1)| d\lambda(m), \\ &= \|ap_1 - bp'_1\|_1. \end{aligned}$$

□

We now show that Assumption 1 is satisfied for S_1 . Fix p_1 and p'_1 in P_1 , $s_1 \in S_1(p_1)$ and a sequence T_n and σ_{1n} such that $(\mu_{T_n}(p, \sigma_{1n}))_n$ converges to s_1 . We have $\mu_{T_n}(p'_1, \sigma_{1n}) \in S_1(p'_1)$ for each n . Let s'_1 be a limit point of the sequence $(\mu_{T_n}(p'_1, \sigma_{1n}))_n$. Take $f \in D_\theta^1$, $a, b \geq 0$. The auxiliary strategy σ_{1n} defines an experiment $z : \Omega_1 \rightarrow \Delta((X \times M)^{T_n})$ with compact set of messages. From Lemma 7:

$$\forall n, af(\mu_{T_n}(p_1, \sigma_{1n})) - bf(\mu_{T_n}(p'_1, \sigma_{1n})) \leq \|ap_1 - bp'_1\|_1.$$

So

$$\begin{aligned} af(s_1) - bf(s'_1) &\leq af(s_1) - af(\mu_{T_n}(p_1, \sigma_{1n})) + (af(\mu_{T_n}(p_1, \sigma_{1n})) - bf(\mu_{T_n}(p'_1, \sigma_{1n}))) + bf(\mu_{T_n}(p'_1, \sigma_{1n})) - bf(s'_1) \\ af(s_1) - bf(s'_1) &\leq af(s_1) - af(\mu_{T_n}(p, \sigma_{1n})) + bf(\mu_{T_n}(p'_1, \sigma_{1n})) - bf(s'_1) + \|ap_1 - bp'_1\|_1 \end{aligned}$$

Passing to the limit as $n \rightarrow \infty$, we get $af(s_1) - bf(s'_1) \leq \|ap_1 - bp'_1\|_1$, concluding the proof. \square

Proof of Theorem 2 First note that uniqueness has been proved by Laraki and Renault (2020), Proposition 4, which implies:

Lemma 8. *Let $u : P_1 \times P_2 \rightarrow \mathbb{R}$ be a continuous function. There exists at most one continuous function $v : P_1 \times P_2 \rightarrow \mathbb{R}$ which is S_1 -concave S_2 -convex and satisfies (C1) and (C2) of Proposition 1 for u .*

The reader is referred to Laraki and Renault (2020) for the proof. We focus then on existence. From Assumption 1, there exists $\theta \in (0, 1]$ such that condition (1) is satisfied for both designers. Define D'_θ as the set of functions from $P_1 \times P_2$ to $[-\theta, \theta]$ such that $\forall p_1, p'_1 \in P_1, \forall p_2, p'_2 \in P_2, \forall a, b \geq 0$,

$$|af(p_1, p_2) - bf(p'_1, p_2)| \leq \|ap_1 - bp'_1\|_1 \text{ and } |af(p_1, p_2) - bf(p_1, p'_2)| \leq \|ap_2 - bp'_2\|_1.$$

Notice that any f in D'_θ is 1-Lipschitz in each variable, therefore D'_θ is a set of equicontinuous functions from $P_1 \times P_2$ to $[-\theta, \theta]$. The following proposition shows existence of $\text{MZ}(u)$ for $u \in D'_\theta$.

Proposition 7. *Assume $u \in D'_\theta$. There exists a function $v = \text{MZ}(u) \in D'_\theta$ which is S_1 -concave, S_2 -convex and satisfies (C1) and (C2) of Proposition 1. Moreover, for u, u' in D'_θ , $\|\text{MZ}(u) - \text{MZ}(u')\|_\infty \leq \|u - u'\|_\infty$.*

Proof. Proposition 7 follows from a series of several lemmas and propositions, where $u \in D'_\theta$ is assumed. For $\delta \in [0, 1)$ and $f \in D'_\theta$, define $\Psi(f) : P_1 \times P_2 \rightarrow \mathbb{R}$ by

$$\Psi(f)(p_1, p_2) = (1 - \delta)u(p_1, p_2) + \delta \text{Val}_{S_1(p_1) \times S_2(p_2)} f(s_1, s_2),$$

where $\text{Val}_{S_1(p_1) \times S_2(p_2)} f(s_1, s_2) = \max_{s_1 \in S_1(p_1)} \min_{s_2 \in S_2(p_2)} f(s_1, s_2) = \min_{s_2 \in S_2(p_2)} \max_{s_1 \in S_1(p_1)} f(s_1, s_2)$.

Lemma 9. *For any f in D'_θ , $\Psi(f)$ is well defined and belongs to D'_θ . Also, Ψ has a unique fixed point in D'_θ which we denote by v_δ .*

Proof. Fix δ and f in D'_θ . For every p_1, p_2 , $\text{Val}_{S_1(p_1) \times S_2(p_2)} f(s_1, s_2)$ is well defined by Sion's minmax Theorem. So the function $\Psi(f) : P_1 \times P_2 \rightarrow \mathbb{R}$ is well defined. We now show that $\Psi(f) \in D'_\theta$ for $f \in D'_\theta$. Fix $p_1, p'_1, p_2, a \geq 0, b \geq 0$, we want to prove that $a\Psi(f)(p_1, p_2) - b\Psi(f)(p'_1, p_2) \leq \|ap_1 - bp'_1\|_1$. Since $u \in D'_\theta$, we have to prove that $A \leq \|ap_1 - bp'_1\|_1$, where

$$A = \max_{s_1 \in S_1(p_1)} \min_{s_2 \in S_2(p_2)} af(s_1, s_2) - \max_{s'_1 \in S_1(p'_1)} \min_{s_2 \in S_2(p_2)} bf(s'_1, s_2).$$

Let $s_1 \in S_1(p_1)$ be optimal in $\max_{s_1 \in S_1(p_1)} \min_{s_2 \in S_2(p_2)} f(s_1, s_2)$. By Assumption 1, there exists $s'_1 \in S_1(p'_1)$ s.t. for all $s_2 \in \Delta(P_2)$, $bf(s'_1, s_2) \geq af(s_1, s_2) - \|ap_1 - bp'_1\|_1$. So $b\text{Val}_{S_1(p'_1) \times S_2(p_2)} f \geq a\min_{s_2 \in S_2(p_2)} f(s_1, s_2) - \|ap_1 - bp'_1\|_1$. Hence $A = \min_{s_2 \in S_2(p_2)} af(s_1, s_2) - b\text{Val}_{S_1(p'_1) \times S_2(p_2)} f \leq \|ap_1 - bp'_1\|_1$ and $\Psi(f) \in D'_\theta$. The operator Ψ is a δ -contraction, so by the contracting fixed point Theorem, Ψ has a unique fixed point v_δ . \square

The family of functions $(v_\delta)_\delta$ is equicontinuous since all belong to D'_θ . By Ascoli's Theorem, it admits a limit point $v \in D'_\theta$ for uniform convergence. Taking limit in the fixed point equation as $\delta \rightarrow 1$

implies that for all $(p_1, p_2) \in P_1 \times P_2$, $v(p_1, p_2) = \text{Val}_{S_1(p_1) \times S_2(p_2)} v(s_1, s_2)$. We now have the following lemma.

Lemma 10. *Any such limit point v is S_1 -concave, S_2 -convex and satisfies (C1) and (C2) of Proposition 1 for u . Thus, an MZ function $\text{MZ}(u)$ exists.*

Proof. The proofs of Propositions 2 and 3 in Laraki and Renault (2020) show that any limit point v of $(v_\delta)_\delta$ is S_1 -concave and S_2 -convex and has the following property: for all (p_1, p_2) in $P_1 \times P_2$, there exists $s_1 \in S_1(p_1)$ such that $v(p_1, p_2) = v(s_1, p_2) \leq u(s_1, p_2)$ and there exists $s_2 \in S_2(p_2)$ such that $v(p_1, p_2) = v(p_1, s_2) \geq u(p_1, s_2)$. We show now that v satisfies (C1), (C2) being obtained by a symmetric argument.

Fix (p_1^0, p_2^0) in $P_1 \times P_2$. Consider the correspondence S'_1 from P_1 to $\Delta(P_1)$ defined for p_1 in P_1 by,

$$S'_1(p_1) = \{s_1 \in S_1(p_1) : v(p_1, p_2^0) = v(s_1, p_2^0) \leq u(s_1, p_2^0)\}.$$

This correspondence admits measurable selections and any such selection f defines a strategy for Designer 1: at each p_1 , play $f(p_1)$. Starting from p_1^0 , the strategy f induces a martingale $(p_1^n)_n$ which converges a.s. to some random variable p_1^∞ with $\mathbb{E}(p_1^\infty) = p_1^0$ and $\mathbb{E}(v(p_1^\infty, p_2^0)) = v(p_1^0, p_2^0)$. Thus, the distribution of p_1^∞ denoted s_1^∞ belongs to $S_1(p_1^0)$ and satisfies $v(s_1^\infty, p_2^0) = v(p_1^0, p_2^0)$. The proof of Lemma 10 is concluded by the following claim.

Claim 1. $v(p_1^\infty, p_2^0) \leq u(p_1^\infty, p_2^0)$ almost surely.

Proof. Let $g : P_1 \rightarrow \mathbb{R}$ be the continuous function given by $g(p_1) = v(p_1, p_2^0) - u(p_1, p_2^0)$. We want to show that $g(p_1^\infty) \leq 0$ a.s. From the choice of f , for each p_1 in P_1 , $g(f(p_1)) = \mathbb{E}_{f(p_1)}(g) \leq 0$. Since g is continuous, $g(p_1^{n+1}) - g(p_1^n) \rightarrow_{n \rightarrow \infty} 0$ a.s., so by the dominated convergence theorem $\mathbb{E}((g(p_1^{n+1}) - g(p_1^n))^2) \rightarrow_{n \rightarrow \infty} 0$. Define the random variable:

$$Y^n = \mathbb{E}[(g(p_1^{n+1}) - g(p_1^n))^2 | p_1^n].$$

We have then $Y^n \geq 0$, Y^n is bounded and $\mathbb{E}(Y^n) \rightarrow_{n \rightarrow \infty} 0$, which imply that there exists a subsequence $(Y^{\varphi(n)})_n$ which converges to 0 almost surely. Moreover, since $\mathbb{E}[g(p_1^{n+1}) - g(f(p_1^n)) | p_1^n] = 0$, we have,

$$\begin{aligned} Y^n &= \mathbb{E}[(g(p_1^{n+1}) - g(f(p_1^n)))^2 | p_1^n] + \mathbb{E}[(g(f(p_1^n)) - g(p_1^n))^2 | p_1^n] \\ &= \mathbb{E}[(g(p_1^{n+1}) - g(f(p_1^n)))^2 | p_1^n] + (g(f(p_1^n)) - g(p_1^n))^2. \end{aligned}$$

Thus $Y^{\varphi(n)}$ tends to 0 and is the sum of two positive terms, so each term tends to 0. Thus

$$(g(f(p_1^{\varphi(n)})) - g(p_1^{\varphi(n)}))^2 \rightarrow_{n \rightarrow \infty} 0 \text{ a.s.}$$

Since $g(f(p_1^{\varphi(n)})) \leq 0$ for all n and $g(p_1^{\varphi(n)}) \rightarrow_{n \rightarrow \infty} g(p_1^\infty)$, we get $g(p_1^\infty) \leq 0$ a.s. □

This proves Lemma 10. □

We know then that any limit point of v_δ is the unique MZ function (Lemma 8). This implies that v_δ converges to $v = \text{MZ}(u)$ as $\delta \rightarrow 1$. Since $\|v_\delta(u) - v_\delta(u')\|_\infty \leq \|u - u'\|_\infty$ for all δ , we get $\|\text{MZ}(u) - \text{MZ}(u')\|_\infty \leq \|u - u'\|_\infty$. This ends the proof of Proposition 7. □

To complete the proof of Theorem 2, we must extend the existence result to all continuous functions. The argument is that functions in D'_θ (or their multiples) are dense within continuous functions.

Lemma 11.

1. The set $\mathcal{L} = \{lf, l \geq 0, f \in D_\theta^1\}$ is dense (for the uniform norm) in the set of continuous functions from P_1 to \mathbb{R} .
2. The set $\mathcal{L}' = \{lf, l \geq 0, f \in D'_\theta\}$ is dense (for the uniform norm) in the set of continuous functions from $P_1 \times P_2$ to \mathbb{R} .

Proof.

1. First observe that \mathcal{L} is a lattice. Take f, g in \mathcal{L} and l, l' such that f/l and g/l' belong to D_θ^1 and let $l'' = \max\{l, l'\}$. Then $\frac{f}{l''}, \frac{g}{l''}$ belong to D_θ^1 and $\max(f, g)/l''$ and $\min(f, g)/l''$ as well. We get that $\max(f, g)$ and $\min(f, g)$ are in \mathcal{L} .

Second, for every $p_1^0 \neq p_1^1$ in P_1 and real numbers α, β , there exists $f \in \mathcal{L}$ such that $f(p_1^0) = \alpha$ and $f(p_1^1) = \beta$. We know that there exists a linear mapping φ on \mathbb{R}^{Ω_1} such that $\varphi(p_1^0) = \alpha$ and $\varphi(p_1^1) = \beta$. Because φ is linear, there exists $l \geq 0$ such that for every p_1, p_1' and $a, b \geq 0$,

$$|a\varphi(p_1) - b\varphi(p_1')| = |\varphi(ap_1 - bp_1')| \leq l\|ap_1 - bp_1'\|_1$$

Thus φ restricted to p_1^1 belongs to \mathcal{L} . By the (lattice) Stone-Weierstrass Theorem, \mathcal{L} is dense in continuous functions.

2. As in the previous point, \mathcal{L}' is a lattice. It is thus enough to prove that for (p_1^0, p_2^0) and (p_1^1, p_2^1) , with $(p_1^0, p_2^0) \neq (p_1^1, p_2^1)$, in $P_1 \times P_2$, $\alpha, \beta \in \mathbb{R}$, there exists f in \mathcal{L}' such that $f(p_1^0, p_2^0) = \alpha$ and $f(p_1^1, p_2^1) = \beta$. Assume $p_1^0 \neq p_1^1$. From the previous point, we know that there exists g in \mathcal{L} such that $g(p_1^0) = \alpha$ and $g(p_1^1) = \beta$. Define $f(p_1, p_2) = g(p_1)$ for all p_1, p_2 . To see that $f \in \mathcal{L}'$, consider p_1, p_1', p_2, p_2' and $a \geq 0, b \geq 0$. We have

$$|af(p_1, p_2) - bf(p_1', p_2)| = |ag(p_1) - bg(p_1')| \leq l\|ap_1 - bp_1'\|_1$$

since $g \in \mathcal{L}$. Then,

$$|af(p_1, p_2) - bf(p_1, p_2')| = |g(p_1)||a - b| = |g(p_1)| \cdot |\|ap_2\|_1 - \|bp_2'\|_1| \leq l'\|ap_2 - bp_2'\|_1$$

with $l' = \max_{p_1} |g(p_1)|$. □

We may now conclude the proof of Theorem 2. Consider a continuous payoff function $u : P_1 \times P_2 \rightarrow \mathbb{R}$. For each $n \geq 1$, there exists u_n in \mathcal{L}' such that $\|u - u_n\|_1 \leq \frac{1}{n}$ and $l_n > 0$ such that $u_n/l_n \in D'_\theta$. By Lemma 10, $\text{MZ}(u_n/l_n)$ exists and we let $v_n = l_n \text{MZ}(u_n/l_n)$. This function is S_1 -concave, S_2 -convex and satisfies (C1) and (C2) of Proposition 1 for u_n . So $v_n = \text{MZ}(u_n)$. Fix now n, m and set $l := \max\{l_n, l_m\}$. We have u_n/l and u_m/l are in D'_θ , $v_n = l \text{MZ}(u_n/l)$ and $v_m = l \text{MZ}(u_m/l)$. Also we know that $\|v_n - v_m\|_1 \leq l\|u_n/l - u_m/l\|_1 = \|u_n - u_m\|_1$. Since $(u_n)_n$ is a Cauchy sequence, so is $(v_n)_n$. Hence $(v_n)_n$ converges to a continuous v which is S_1 -concave and S_2 -convex. By taking limits, v satisfies (C1) and (C2) for u . Therefore $v = \text{MZ}(u)$ exists. This ends the proof of Theorem 2. □

Proof of Proposition 6 Consider the stationary strategy σ_1 of Designer 1 who plays non-revealingly if $v(p_1, p_2) \leq u(p_1, p_2)$, and s_1 is given by (C1) otherwise. Consider any strategy σ_2 of Designer 2. The definition of σ_1 implies that for each n ,

$$\mathbb{E}[v(p_1^{n+1}, p_2^n) | p_1^n, p_2^n] = v(p_1^n, p_2^n).$$

Since v is S_2 -convex,

$$\mathbb{E}[v(p_1^{n+1}, p_2^{n+1}) | p_1^{n+1}, p_2^n] \geq v(p_1^{n+1}, p_2^n).$$

Taking expectation, for each n , $\mathbb{E}[v(p_1^{n+1}, p_2^{n+1})] \geq \mathbb{E}[v(p_1^{n+1}, p_2^n)] \geq \mathbb{E}[v(p_1^n, p_2^n)] \geq v(p_1^0, p_2^0)$ by induction. Denote $X = \{(p_1, p_2) \in P_1 \times P_2 : v(p_1, p_2) \leq u(p_1, p_2)\}$; by construction, $(p_1^{n+1}, p_2^n) \in X$ almost surely for each n . Since by assumption X is closed, $(p_1^{N^*}, p_2^{N^*}) \in X$ a.s., i.e., $u(p_1^{N^*}, p_2^{N^*}) \geq v(p_1^{N^*}, p_2^{N^*})$ a.s.

With an abuse of notation, denote by $(p_1^n, p_2^n)_n$ the martingale stopped at N^* , i.e., such that $(p_1^n, p_2^n) = (p_1^{N^*}, p_2^{N^*})$ for all $n > N^*$.

Claim 2. $\limsup_n v(p_1^n, p_2^n) \leq v(p_1^\infty, p_2^\infty)$ a.s.

Proof. Fix a realized play path and consider a converging subsequence of $(v(p_1^n, p_2^n))_n$ denoted by $(v(p_{1*}^n, p_{2*}^n))_n$. We show that $\lim_n v(p_{1*}^n, p_{2*}^n) \leq v(p_1^\infty, p_2^\infty)$. There are 2 cases.

1) Suppose that there exists n_0 such that for all $n \geq n_0$, $(p_{1*}^n, p_{2*}^n) \in X$. Then, for each $n \geq n_0$, $p_{1*}^{n+1} = p_{1*}^n = p_1^\infty$ and $\lim_n v(p_{1*}^n, p_{2*}^n) = \lim_n v(p_1^\infty, p_{2*}^n)$. Since v is u.s.c. in p_2 , $\lim_n v(p_1^\infty, p_{2*}^n) \leq v(p_1^\infty, p_2^\infty)$.

2) Otherwise, there exists a subsequence (p_{1**}^n, p_{2**}^n) of (p_{1*}^n, p_{2*}^n) with values in $\{(p_1, p_2) : v(p_1, p_2) > u(p_1, p_2)\}$. Since v is u.s.c. on the closure of this set, we obtain $v(p_1^\infty, p_2^\infty) \geq \lim_n v(p_{1**}^n, p_{2**}^n) = \lim_n v(p_{1*}^n, p_{2*}^n)$. \square

It follows that

$$v(p_1^0, p_2^0) \leq \limsup_n \mathbb{E}(v(p_1^n, p_2^n)) \leq \mathbb{E}(\limsup_n v(p_1^n, p_2^n)) \leq \mathbb{E}(v(p_1^\infty, p_2^\infty)) \leq \mathbb{E}(u(p_1^\infty, p_2^\infty)),$$

therefore Designer 1 guarantees $v(p_1^0, p_2^0)$. By symmetry this is the value of the game. \square

A.2 Example of a non-expansive splitting correspondence generated by a set of experiments

The following example is inspired by Example 3.13 in Renault and Venel (2017). Let $\Omega_1 = \{a, b, c, d\}$ and consider a single deterministic experiment $x : \Omega_1 \rightarrow \{m_0, m_1\}$ defined by:

$$x(a) = x(c) = m_0, x(b) = x(d) = m_1$$

That is, x reveals whether the state is in $\{a, c\}$ or $\{b, d\}$. Denote by $\varphi(p, x) \in \Delta(\Delta(\Omega_1))$ the splitting of prior p_1 induced by x .

- For $p'_1 = (1/4, 1/2, 1/4, 0)$, $\varphi(p, x) \in \Delta(\Delta(\Omega_1))$ puts probability $\frac{1}{2}$ on the posterior $(1/2, 0, 1/2, 0)$ and probability $\frac{1}{2}$ on the posterior $(0, 1, 0, 0)$.

- For $p_1'' = (1/4, 1/2, 0, 1/4)$, $\varphi(p_1', x) \in \Delta(\Delta(\Omega_1))$ puts probability $\frac{1}{4}$ on the posterior $(1, 0, 0, 0)$ and probability $\frac{3}{4}$ on the posterior $(0, 2/3, 0, 1/3)$.

Notice that $\|p_1' - p_1''\|_1 = 1/2$. Consider now $p_1^1 = (1, 0, 0, 0)$, $p_1^2 = (0, 2/3, 0, 1/3)$, $p_1^3 = (0, 1, 0, 0)$, $p_1^4 = (1/2, 0, 1/2, 0)$ and denote $S_1 = \{p_1^1, p_1^2, p_1^3, p_1^4\}$. Define a 1-Lipschitz function $f : S_1 \rightarrow \mathbb{R}$ by such that $f(p_1^1) = 2/3$, $f(p_1^2) = 1$, $f(p_1^3) = 1/3$ and $f(p_1^4) = -1/3$. The function f can be extended to a 1-Lipschitz function from $\Delta(\Omega_1)$ to \mathbb{R} by the formula:

$$\forall q \in \Delta(\Omega_1), f(q) = \sup_{\hat{p}_1 \in S_1} (f(\hat{p}_1) - \|q - \hat{p}_1\|_1).$$

We have $f(\varphi(p_1')) = \frac{1}{2} \cdot 1/3 + \frac{1}{2}(-1/3) = 0$ and $f(\varphi(p_1'')) = \frac{1}{4} \cdot 2/3 + \frac{3}{4} \cdot 1 = 11/12$, so

$$|f(\varphi(p_1')) - f(\varphi(p_1''))| = 11/12 > 1/2 = \|p_1' - p_1''\|_1.$$

The non expansivity condition of Laraki and Renault (2020) requires that

$$|f(\varphi(p_1')) - f(\varphi(p_1''))| \leq \|p_1' - p_1''\|_1, \forall f \text{ 1-Lipschitz}$$

and is therefore not satisfied here.

References

- Brian C. Albrecht. Political persuasion. *mimeo*, 2017.
- Pak Hung Au and Keiichi Kawai. Competitive disclosure of correlated information. *Economic Theory*, pages 1–33, 2019.
- Pak Hung Au and Keiichi Kawai. Competitive information disclosure by multiple senders. *Games and Economic Behavior*, 119:56–78, 2020.
- Robert J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.
- Robert J. Aumann and Sergiu Hart. Bi-convexity and bi-martingales. *Israel Journal of Mathematics*, 54(2):159–180, 1986.
- Robert J. Aumann and Sergiu Hart. Long cheap talk. *Econometrica*, 71(6):1619–1660, 2003.
- Robert J. Aumann and Michael Maschler. Game theoretic aspects of gradual disarmament. Report of the U.S. Arms Control and Disarmament Agency, ST-80, Chapter V, pp. 1–55, 1966.
- Robert J. Aumann and Michael Maschler. Repeated games with incomplete information: A survey of recent results. Report of the U.S. Arms Control and Disarmament Agency, ST-116, Chapter III, pp. 287–403, 1967.

- Robert J. Aumann, Michael Maschler, and Richard E. Stearns. *Repeated Games of Incomplete Information*. MIT Press, Cambridge, Massachusetts, 1995.
- Dirk Bergemann and Stephen Morris. Information design, bayesian persuasion, and bayes correlated equilibrium. *The American Economic Review*, 106(5):586–591, 2016a.
- Dirk Bergemann and Stephen Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11:487–522, 2016b.
- Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.
- Raphael Boleslavsky and Christopher Cotton. Grading standards and education quality. *American Economic Journal: Microeconomics*, 7(2):248–279, 2015.
- Raphael Boleslavsky and Christopher Cotton. Limited capacity in project selection: Competition through evidence production. *Economic Theory*, 65(2):385–421, 2018.
- Raphael Boleslavsky and Kyungmin Kim. Bayesian persuasion and moral hazard. *mimeo*, 2018.
- Raphael Boleslavsky, Christopher Cotton, and Haresh Gurnani. Demonstrations and price competition in new product release. *Management Science*, 63(6):2016–2026, 2016.
- Jesse Bull and Joel Watson. Hard evidence and mechanism design. *Games and Economic Behavior*, 58:75–93, 2007.
- Laura Doval and Jeffrey C. Ely. Sequential information design. *Econometrica*, 88(6):2575–2608, 2020.
- Ronald A Dye. Disclosure of nonproprietary information. *Journal of accounting research*, pages 123–145, 1985.
- Jeffrey C. Ely. Beeps. *American Economic Review*, 107(1):31–53, 2017.
- Françoise Forges. Note on nash equilibria in repeated games with incomplete information. *International Journal of Game Theory*, 13:179–187, 1984.
- Françoise Forges. An approach to communication equilibria. *Econometrica*, 54(6):1375–1385, 1986.
- Françoise Forges. Equilibria with communication in a job market example. *Quarterly Journal of Economics*, 105:375–398, 1990.
- Françoise Forges. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision*, 35:277–310, 1993.
- Françoise Forges. Games with incomplete information: from repetition to cheap talk and persuasion. *Annals of Economics and Statistics*, (137):3–30, 2020.
- Françoise Forges and Frédéric Koessler. Communication equilibria with partially verifiable types. *Journal of Mathematical Economics*, 41(7):793–811, 2005.

- Matthew Gentzkow and Emir Kamenica. Competition in persuasion. *The Review of Economic Studies*, 84(1):300–322, 2017.
- Jerry R. Green and Jean-Jacques Laffont. Partially verifiable information and mechanism design. *Review of Economic Studies*, 53(3):447–456, 1986.
- Martin Heuer. Asymptotically optimal strategies in repeated games with incomplete information. *International Journal of Game Theory*, 20:377–392, 1992.
- Emir Kamenica. Bayesian persuasion and information design. *Annual Review of Economics*, 11, 2018.
- Emir Kamenica and Matthew Gentzkow. Bayesian Persuasion. *American Economic Review*, 101(6): 2590–2615, 2011.
- Frédéric Koessler, Marie Laclau, and Tristan Tomala. Interactive information design. *Mathematics of Operations Research*, forthcoming, 2020.
- Frédéric Koessler, Marie Laclau, Jérôme Renault, and Tristan Tomala. Splitting games over finite sets. *mimeo*, 2021.
- Rida Laraki. The splitting game and applications. *International Journal of Game Theory*, 30(3): 359–376, 2001a.
- Rida Laraki. Variational inequalities, system of functional equations and incomplete information repeated games. *SIAM Journal on Control and Optimization*, 40:516–524, 2001b.
- Rida Laraki and Jérôme Renault. Acyclic gambling games. *Mathematics of Operation Research*, 45(4): 1237–1257, 2020.
- Mael Le Treust and Tristan Tomala. Persuasion with limited communication capacity. *Journal of Economic Theory*, 184, 2019.
- Fei Li and Peter Norman. Sequential persuasion. *Theoretical Economics*, 16:639–675, 2021.
- Miltiadis Makris and Ludovic Renou. Information design in multi-stage games. *mimeo*, 2018.
- Laurent Mathevet, Jacopo Perego, and Ina Taneva. On information design in games. *Journal of Political Economy*, 128(4):1370–1404, 2020.
- Ludmila Matyskova and Alfonso Montes. Bayesian persuasion with costly information acquisition. *mimeo*, 2021.
- Jean-François Mertens and Shmuel Zamir. The value of two-person zero-sum repeated games with lack of information on both sides. *International Journal of Game Theory*, 1(1):39–64, 1971.
- Jean-François Mertens and Shmuel Zamir. A duality theorem on a pair of simultaneous functional equations. *Journal of Mathematical Analysis and Applications*, 60(2):550–558, 1977.

- Jean-François Mertens, Sylvain Sorin, and Shmuel Zamir. *Repeated games*. Cambridge University Press, 2015.
- Paul Milgrom. Good news and bad news: Representation theorems and applications. *Bell Journal of Economics*, 12:380–391, 1981.
- Roger B. Myerson. Optimal coordination mechanisms in generalized principal-agent problems. *Journal of Mathematical Economics*, 10:67–81, 1982.
- Roger B. Myerson. Bayesian equilibrium and incentive-compatibility: An introduction. *Social goals and social organization: Essays in memory of Elisha Pazner*, pages 229–260, 1985.
- Roger B. Myerson. Multistage games with communication. *Econometrica*, 54:323–358, 1986.
- Miquel Oliu-Barton. The splitting game: Uniform value and optimal strategies. *Dynamic Games and Applications*, pages 1–23, 2017.
- Eduardo Perez-Richet. Interim bayesian persuasion: First steps. *The American Economic Review*, 104(5):469–474, 2014.
- Jean-Pierre Ponsard and Sylvain Sorin. Some results on zero-sum games with incomplete information: The dependent case. *International Journal of Game Theory*, 9(4):233–245, 1980.
- Jérôme Renault and Xavier Venel. Long-term values in markov decision processes and repeated games, and a new distance for probability spaces. *Mathematics of Operations Research*, 42(2):349–376, 2017.
- Jérôme Renault, Eilon Solan, and Nicolas Vieille. Optimal dynamic information provision. *Games and Economic Behavior*, 104:329–349, 2017.
- Régis Renault. Advertising in markets. In *Handbook of Media Economics*, volume 1, pages 121–204. Elsevier, 2015.
- Andrés Salamanca. The value of mediated communication. *Journal of Economic Theory*, 192:105191, 2021.
- Sylvain Sorin. On a pair of simultaneous functional equations. *Journal of mathematical analysis and applications*, 98(1):296–303, 1984.
- Sylvain Sorin. *A First Course on Zero-Sum Repeated Games*. Springer, 2002.
- Ina Taneva. Information design. *American Economic Journal: Microeconomics*, 11(4):151–85, 2019.
- Mark Whitmeyer. Dynamic competitive persuasion. *mimeo*, 2020.
- Jiemai Wu. Non-competing persuaders. *European Economic Review*, 127:103454, 2020.