

October 2025

“Information Design in Smooth Games”

Alex Smolin and Takuro Yamashita

Information Design in Smooth Games

Alex Smolin, Takuro Yamashita

July 9, 2025

Abstract

We study information design in games where players choose from a continuum of actions and have continuously differentiable payoffs. We show that an information structure is optimal when the equilibrium it induces can also be implemented in a principal-agent contracting problem. Building on this result, we characterize optimal information structures in symmetric linear-quadratic games. With common values, targeted disclosure is robustly optimal across all priors. With interdependent and normally distributed values, linear disclosure is uniquely optimal. We illustrate our findings with applications in venture capital, Bayesian polarization, and price competition.

Keywords: Bayesian persuasion, information design, dual certification, first-order approach, linear-quadratic games, targeted disclosure, Gaussian coupling, linear disclosure.

Smolin: Toulouse School of Economics, University of Toulouse Capitole and CEPR, alexey.v.smolin@gmail.com. Yamashita: Osaka University, tytakuroy@gmail.com. For valuable suggestions and comments, we would like to thank the Coeditor Rakesh Vohra, two anonymous referees, Dirk Bergemann, Deniz Dizdar, Laura Doval, Philippe Jehiel, Jiangtao Li, Xiao Lin, Elliot Lipnowski, Stephen Morris, Alessandro Pavan, Antonio Penta, Jacopo Perego, Fedor Sandomirskiy, Ludvig Sinander, Takashi Ui, Xavier Vives, and Alexander Wolitzky, as well as seminar participants at Toulouse School of Economics, Western University, University of Toronto, University of Surrey, University of Oxford, Hitotsubashi University, Singapore Management University, CMU/Pittsburgh, Pennsylvania State University, MIT/Harvard, Bonn Winter Theory Workshop 2022, SWET 2022, CMid2022, EC 2022, EEA-ESEM 2022, EARIE 2022, and Venice Winter Theory Workshop 2023. An extended abstract of a previous version of this paper appeared as “Information Design in Concave Games” in the proceedings to EC’22. Smolin gratefully acknowledges funding from the French National Research Agency (ANR) under the Investments for the Future program (grant ANR-17-EURE-0010) and through the AI Interdisciplinary Institute ANITI (grant ANR-23-IACL-0002). Yamashita gratefully acknowledges funding from the French National Research Agency (ANR) under the Investments for the Future program (grant ANR-17-EURE-0010), funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant 714693), and JSPS KAKENHI (grant 23K01311).

1 Introduction

As advances in IT infrastructure and artificial intelligence expand our capacity to collect, process, and generate data, an increasing number of firms, regulators, and individuals must decide which information to supply to strategically interacting agents. Which information should investors receive to allocate capital most efficiently? Which demand signals should be disclosed to improve market outcomes? What are the limits of Bayesian polarization? We show that these and related questions can be tractably addressed within a single framework, and that the optimal policies are simple, intuitive, and robust.

Our main analysis focuses on concave games of incomplete information, i.e., games in which each player’s action lies in a convex set and each player’s payoff is strictly concave in that action, but our methods extend to any smooth game. Concave games are a staple of applied economic modeling with fixed information structures, because equilibria can be tractably characterized by first-order conditions. We show that the same considerations render tractable the task of designing the information structure itself.

To solve the information-design problem, we adopt a duality-based approach. The dual can be viewed as an adversarial contracting problem between a principal and an omniscient agent who both observes the state and controls all players’ actions. The dual-certification theorem ([Theorem 1](#)) states that if a state-action distribution emerges in equilibrium under both an information structure (in the information-design problem) and a contract (in the adversarial-contracting problem), then that information structure and that contract are optimal for their respective problems. The contract thus serves as an optimality certificate and, under mild regularity conditions, is guaranteed to exist ([Theorem 2](#)). Furthermore, it is capable of certifying any optimal information structure ([Proposition 1](#)), which enables us to establish whether the optimal information structure is unique or, when several are optimal, to identify their common features.

We show that optimal information structures are prior-robust: once optimal under one prior, they remain optimal under any other prior as long as the implemented allocation rule’s support is contained within the original support in every state ([Proposition 2](#)). This has two consequences. First, if an optimal information structure is fully informative about the state, it remains optimal under all priors. Second, unless the problem is trivial, an optimal information structure cannot induce a full-support action distribution in every state, which argues against

adding full-support, extraneous, independent noise to individual signals.

We apply the solution method in two broad, symmetric settings with linear-quadratic payoffs for both the players and the designer. In each case, the certifying contract is symmetric across players and affine in actions, yet it yields qualitatively different optimal information policies.

In the first setting, the state is one-dimensional (Section 5). We derive the parameters of the certifying contract and use them to certify the optimality of *targeted disclosure*: an information structure that fully reveals the state to a subset of players while leaving the rest completely uninformed. The structure is asymmetric and, remarkably, distributionally robust: it remains optimal under any prior. The optimum, however, need not be unique. When the state is normally distributed, *Gaussian coupling*, which adds idiosyncratic normal noise to each player’s signal that cancels out in aggregate, is optimal as well.

We apply this result to two concrete applications. In Section 5.1, we analyze optimal capital fundraising when several investors independently choose how much to invest in a project of uncertain quality. We show that the information structure that maximizes the project’s expected return is exclusive disclosure, a targeted disclosure in which only one investor is informed. Relative to full or no disclosure, exclusive disclosure prevents the dissipation of returns as the investor pool grows, offering a possible rationale for a common venture-capital practice. In Section 5.2, we examine the limits of Bayesian polarization when each player forms a forecast from private information. We show that the information structure that maximizes a natural polarization index is a targeted disclosure that informs exactly half of the players. Hence, maximal polarization occurs when the population splits into two cohorts that are internally uniform but sharply divergent from each other. This finding underscores how media segregation can intensify societal polarization.

In the second setting, the state is multidimensional with jointly normal components (Section 6). We show that a certifying contract can be found within the class of symmetric affine contracts by solving a single-variable minimization problem. This contract certifies the optimality of *linear disclosure*: an information structure that gives each player some linear statistic of the state. The resulting information structure is noise-free and symmetric ex ante. Moreover, it is uniquely optimal: richer uncertainty imposes tighter constraints on design, fully determining the optimal information.

In [Section 6.1](#), we apply this result to optimal price recommendations in a differentiated-product duopoly with linear demand and stochastic demand shocks. We characterize the information structures that maximize any weighted average of consumer and producer surplus. The resulting recommended prices are linear functions of the demand shocks and display two distinct regimes: when the weight on consumer surplus is low, prices are positively correlated; when it is high, they become negatively correlated. The transition between regimes is discontinuous, underscoring the risk that the algorithms generating price recommendations—able to shift objectives and adapt policies far faster than human decision-makers—can destabilize markets.

Related Literature Our paper contributes to the recent and flourishing literature on information design. Much of this literature focuses on information design with a single player, called Bayesian persuasion ([Rayo and Segal \(2010\)](#), [Kamenica and Gentzkow \(2011\)](#)). Popular solution methods are belief based, i.e., they operate within the space of the receiver’s belief distributions.¹ A natural continuation of this research agenda is a study of information design in multiplayer games ([Bergemann and Morris \(2016\)](#), [Taneva \(2019\)](#)). In games, players’ beliefs constitute infinite hierarchies, which renders the belief-based approach less tractable ([Mathevet, Perego, and Taneva \(2020\)](#)). Instead, in games, an action-based approach rooted in the revelation principle is promising, as it frames the design problem as a linear program and enables the use of duality machinery.²³

[Galperti and Perego \(2018\)](#) and [Galperti, Levkun, and Perego \(2024\)](#) employ this action-based approach to study information design in games with finitely many actions, imposing minimal structure on payoffs. They develop an economic interpretation of the Lagrange multipliers associated with Bayes’ plausibility as the value of data records and propose the idea of pooling externalities across records; all these observations apply to our setting as we discuss

¹See, for example, [Dworczak and Martini \(2019\)](#), [Dizdar and Kováč \(2020\)](#), and [Dworczak and Kolotilin \(2024\)](#). The single-player case also covers scenarios in which there are many players but the information is required to be public.

²The duality methodology is routinely used in many disciplines to solve optimization problems. In mechanism design, duality methods have been recently used to study optimal delegation ([Amador, Werning, and Angeletos \(2006\)](#); [Amador and Bagwell \(2013\)](#)), matching ([Chiappori, Salanié, and Weiss \(2017\)](#); [Galichon and Salanié \(2022\)](#)), robust selling mechanisms ([Carroll \(2017\)](#); [Du \(2018\)](#); [Brooks and Du \(2021\)](#)), mediation ([Salamanca \(2021\)](#), [Ortner, Sugaya, and Wolitzky \(2024\)](#)), and limited commitment ([Lin and Liu \(2024\)](#)) among others. For a unified treatment, see [Vohra \(2011\)](#).

³Alternatively, one can develop original arguments tailored to the studied problem; see, for example, [Arieli and Babichenko \(2019\)](#), [Chan, Gupta, Li, and Wang \(2019\)](#), [Elliott, Galeotti, Koh, and Li \(2022\)](#), [Arieli, Babichenko, and Sandomirskiy \(2023\)](#), and [Candogan and Strack \(2023\)](#).

in Section 3.1.

In contrast, we study games with infinitely many actions and impose a concavity structure on payoffs, thus enabling us to rely on a first-order approach for incentives (Holmström (1979), Mirrlees (1999)) that leads to more succinct and tractable primal and dual problems. This approach was introduced in a Bayesian persuasion setting by Kolotilin (2012, 2018) and further refined by Kolotilin, Corrao, and Wolitzky (2025). These papers study an information-design problem with a single player and a one-dimensional state, identifying when censorship or, respectively, assortative disclosures are optimal. We deepen and extend this approach, adapting it to settings with multiple players and a multidimensional state.

Most of our current understanding of optimal information in games is drawn from the study of Gaussian signals in games with quadratic payoffs and a normally distributed state (e.g., Angeletos and Pavan (2007, 2009), Bergemann and Morris (2013), Bergemann, Heumann, and Morris (2015a, 2021), Ui (2020)). In that literature, as well as in a vast body of work in macroeconomics and finance, the Gaussian form of players' signals is imposed *ad hoc* for analytical convenience.

Recent arguments suggest that Gaussian signals are not only convenient but often optimal among all information structures. For example, Tamura (2012, 2018) established the optimality of a Gaussian signal in a setting with a single player by building on the statistical properties of a covariance matrix of posterior expectations. Bergemann, Heumann, and Morris (2017) and Miyashita and Ui (2023) extend this argument to games.⁴ Our results align with these findings and show that in many scenarios the optimal Gaussian structures are symmetric and noise-free; moreover, the certification method allows direct identification of the optimal informational parameters. Furthermore, we provide tools that establish when Gaussian information structures are uniquely optimal while also showing that, in many settings, other structures, such as targeted disclosure, can likewise be optimal, and robustly so.

2 Design Problem

We study a standard information design problem as presented by Bergemann and Morris (2016), extended to accommodate a continuum of players' actions.

⁴In their Section 4.4, Bergemann et al. (2017) argue that Gaussian structures, possibly asymmetric and with extraneous noise, span all implementable covariance matrices of equilibrium actions.

Payoffs There are N players indexed by i , $1 \leq N < \infty$, and an information designer. Each player chooses an action $a_i \in A_i = \mathbb{R}$. We denote an action profile by $a \in A = \times_i A_i$ and write (a_i, a_{-i}) when highlighting player i 's action.

A state ω is distributed over a Polish set Ω , according to a full-support prior $\mu_0 \in \Delta(\Omega)$. The action profile and the state jointly determine payoffs through

$$u_i : A \times \Omega \rightarrow \mathbb{R}, \quad (1)$$

$$v : A \times \Omega \rightarrow \mathbb{R}, \quad (2)$$

for each player i and for the designer, respectively. The tuple $((A_i, u_i)_{i=1}^N, \mu_0)$ constitutes the *basic game*.

Information The players and the designer start with a commonly known prior belief about the state ω that coincides with the prior μ_0 . The designer can provide additional information to players by choosing an information structure $\mathcal{I} = (S, \pi)$ that consists of a measurable signal set $S = \times_i S_i$ and a likelihood function $\pi \in \Delta(S \times \Omega)$ that has μ_0 as its state marginal distribution. This information structure determines the sets of private signals the players can observe and, through the likelihood function, their informational content.

First, the designer chooses an information structure \mathcal{I} . Second, the state ω and the signal profile s are realized according to \mathcal{I} . Finally, each player privately observes his signal s_i and chooses an action a_i .

We will regularly refer to two benchmark information structures: *full disclosure*, with $S_i = \Omega$ and $s_i \equiv \omega$ for all i ; and *no disclosure*, with $|S_i| = 1$ for all i .

Equilibrium The basic game together with the information structure chosen by the designer determine a Bayesian game of incomplete information. In that game, each player's behavior is described by a strategy that maps any received signal to a possibly random action, $\sigma_i : S_i \rightarrow \Delta(A_i)$, and we consider as an equilibrium concept a Bayes Nash equilibrium:

Definition 1. (Bayes Nash Equilibrium) *For a given information structure \mathcal{I} , a strategy profile $\sigma = (\sigma_1, \dots, \sigma_N)$ constitutes a Bayes Nash equilibrium if*

$$\mathbb{E}_{\mathcal{I}, \sigma_i, \sigma_{-i}}[u_i(a_i, a_{-i}, \omega)] \geq \mathbb{E}_{\mathcal{I}, \sigma'_i, \sigma_{-i}}[u_i(a'_i, a_{-i}, \omega)] \quad (3)$$

for all i and $\sigma'_i : S_i \rightarrow \Delta(A_i)$.⁵

An information structure and a strategy profile determine a distribution over the action profiles in each state $\alpha : \Omega \rightarrow \Delta(A)$, which we call an *allocation rule*, and the corresponding designer's expected payoff $\mathbb{E}_{\mathcal{I}, \sigma}[v(a, \omega)]$. The *value* of an information structure is defined as the maximal designer's expected payoff that can arise in equilibrium of the induced game: if the game has multiple equilibria, the designer can choose the one she prefers, whereas if no equilibrium exists, the value is undefined. An information-design problem consists of finding an information structure with a maximal value without placing any additional restrictions on the sets of signals or the likelihood function, apart from a mild "admissibility" condition. (This condition, and other omitted formal details, are deferred to [Appendix A](#).)

Definition 2. (Optimal Information Structure) *An information structure is optimal if there does not exist an information structure with a strictly higher value.*

The search for an optimal information structure is complicated by the scale of the basic game: multiple players, after receiving private signals, choose actions from a continuum while anticipating one another's behavior. Our central simplifying assumption is:

Assumption 1. (Concave Payoffs) *For all $i = 1, \dots, N$, $\omega \in \Omega$, and $a_{-i} \in A_{-i}$, $u_i(a_i, a_{-i}, \omega)$ is continuously differentiable in a_i , weakly concave in a_i , and obtains its maximum at some finite value.*

We call a basic game in which [Assumption 1](#) is satisfied a *concave game*.⁶ In a concave game, for any player i and equilibrium belief $\mu \in \Delta(A_{-i} \times \Omega)$ about the others' actions and the state, a best response a_i^* exists and satisfies the first-order condition:

$$\left. \frac{\partial}{\partial a_i} \mathbb{E}_\mu [u_i(a_i, a_{-i}, \omega)] \right|_{a_i=a_i^*} = \mathbb{E}_\mu \left[\left. \frac{\partial}{\partial a_i} u_i(a_i, a_{-i}, \omega) \right] \right|_{a_i=a_i^*} \triangleq \mathbb{E}_\mu [\dot{u}_i(a_i^*, a_{-i}, \omega)] = 0, \quad (4)$$

where we denote the player's marginal payoff function by

$$\dot{u}_i(a, \omega) \triangleq \frac{\partial u_i(a, \omega)}{\partial a_i}.$$

⁵We write $\mathbb{E}_{\mathcal{I}, \sigma_i, \sigma_{-i}}[\cdot]$ for the expectation under information structure \mathcal{I} and strategy profile (σ_i, σ_{-i}) . Throughout the paper an integral is left undefined whenever the integrand is not integrable with respect to the relevant measure.

⁶This notion of a concave game is related to but distinct from the notion of a concave game of [Rosen \(1965\)](#). In particular, it requires neither differentiability nor continuity of a player's payoff in the other players' actions or the state.

An equilibrium under any given information structure is characterized by a system of such conditions, one for each player's signal.

The information-design problem can be simplified by appealing to the revelation principle (Myerson (1982)): it is without loss of generality to focus on *direct information structures* that inform each player about a recommended action $S = A$ and are such that all players are obedient, i.e., are willing to follow the recommendations. Each direct information structure corresponds to a measure $\pi \in \Delta(A \times \Omega)$, and the information-design problem can be formulated as a constrained maximization over these measures:

$$V^P \triangleq \sup_{\pi \in \Delta(A \times \Omega)} \int_{A \times \Omega} v(a, \omega) d\pi \quad (5)$$

$$\text{s.t. } \int_{A'_i \times A_{-i} \times \Omega} \dot{u}_i(a, \omega) d\pi = 0 \quad \forall i = 1, \dots, N, \text{ measurable } A'_i \subseteq A_i, \quad (6)$$

$$\int_{A \times \Omega'} d\pi = \int_{\Omega'} d\mu_0 \quad \forall \text{ measurable } \Omega' \subseteq \Omega. \quad (7)$$

Constraints (6) are a proper formulation of first-order conditions (4) in light of a continuum of recommended actions. These constraints capture players' obedience and effectively require that for each player i , the linear projection of π on A_i weighted by the marginal utilities equals zero measure. Constraints (7) capture Bayes' plausibility and, likewise, require that the linear projection of π on Ω equals the prior μ_0 .

3 Solution Method

Problem (5–7) is linear in π , yet unwieldy to solve directly. In the spirit of linear programming, we view it as a *primal problem* and call any $\pi \in \Delta(A \times \Omega)$ a *primal measure*. If a primal measure satisfies the constraints of the primal problem, then we call that measure *implementable by information*. We can construct a dual problem as follows (e.g., Anderson and Nash (1987)):

$$V^D \triangleq \inf_{\lambda \in \times_i \mathcal{M}(A_i), \gamma \in \mathcal{M}(\Omega)} \int_{\Omega} \gamma(\omega) d\mu_0 \quad (8)$$

$$\text{s.t. } \sum_{i=1}^N \lambda_i(a_i) \dot{u}_i(a, \omega) + \gamma(\omega) \geq v(a, \omega) \quad \forall a \in A, \omega \in \Omega,$$

where $\mathcal{M}(X)$ denotes the space of measurable real-valued functions on X . The minimization arguments, the dual variables (λ, γ) , represent the Lagrange multipliers associated with the

primal incentive constraints (6) and the feasibility constraints (7), respectively.

The problem (8) is a generalization of the dual problem of Kolotilin (2018) and Kolotilin et al. (2025) to multiple players. By the arguments analogous to those of Galperti et al. (2024), the optimal $\gamma^*(\omega)$ measures the marginal benefit of increasing the frequency of state ω . Importantly, the optimal $\gamma^*(\omega)$ can be solved away: the objective in (8) is additively separable in $\gamma(\omega)$, and the constraints at different states ω are linked only through the variables λ . Hence, for any λ and ω , the optimal choice of $\gamma(\omega)$ is the smallest value consistent with the dual constraints:

$$\gamma^*(\omega; \lambda) = \sup_{a \in A} w(a, \omega, \lambda),$$

where the *dual payoff function* w is defined as

$$w(a, \omega, \lambda) \triangleq v(a, \omega) - \sum_{i=1}^N \lambda_i(a_i) \dot{u}_i(a, \omega). \quad (9)$$

As a result, the problem (8) can be restated as follows:

$$V^D = \inf_{\lambda \in \times_i \mathcal{M}(A_i)} \mathbb{E}[\sup_{a \in A} w(a, \omega, \lambda)]. \quad (10)$$

Problem (10) admits a simple economic interpretation as a *contracting problem* between a *dual principal* and a *dual agent*. First, the principal chooses an incentive contract λ that consists of N functions $\lambda_i(a_i)$ and determines the agent's payoff according to (9). Second, the state ω is realized. Finally, the agent perfectly observes the state and chooses the whole action profile $a \in A$ to maximize his payoff. The contracting is *adversarial* in that the principal aims to minimize the agent's expected payoff.

If the best responses exist at all states and induce the joint action-state measure $\pi(a, \omega)$, then we say that λ *implements* π *by incentives* and that π *is implementable by incentives*, by contract λ .

Theorem 1. (Weak Duality. Dual Certification) *If $\pi \in \Delta(A \times \Omega)$ is implementable by information and is implementable by incentives by contract λ , then (i) π solves the information-design problem, (ii) λ solves the adversarial-contracting problem, and (iii) $V^P = V^D$.*

Theorem 1 offers a solution method based on optimality certification. When the conditions

of the theorem hold, we say that λ is a (dual) *certificate* of π , that λ *certifies* the optimality of π , and that π is *certifiably optimal*. Similarly, if an information structure induces a certifiably optimal measure, then we call that information structure *certifiably optimal*.

We highlight two general properties of certificates. First, every certificate λ^* has a clear economic meaning: it represents Lagrange multipliers associated with the obedience constraints (6). Accordingly, $\lambda_i^*(a_i)$ measures the marginal value for the information designer from perturbing the obedience constraint of action a_i . Second, each certificate is “universal”:

Proposition 1. (Universality) *If λ certifies the optimality of measure π , and π' is another optimal measure, then λ also certifies the optimality of measure π' .*

By Proposition 1, a certificate does more than certify a single information structure—it constrains the entire set of optima. Every optimal measure must be certified by the same certificate; equivalently, it must be the dual agent’s best response to the same contract in the dual problem. We exploit this property in Section 5, where multiple optima exist, to isolate their common features, and in Section 6 to show that the optimal measure is unique.

The difference $V^D - V^P$ between the optimal values of primal and dual problems is non-negative and constitutes a *duality gap*. The solution to the information-design problem can be certified if and only if (i) solutions to both primal and dual problems exist and (ii) the duality gap is equal to zero, $V^D = V^P$. Thus, either all optimal information structures can be certified or none of them can.

Any constructed certificate, by its very existence, proves that the duality gap is zero and that both primal and dual solutions exist. However, it may be useful to know *a priori* whether one can expect to find such a certificate in a given problem.

Condition 1. (Regularity) *(i) v and each u_i are continuous; (ii) the full-disclosure value is finite; (iii) there exists a sequence of compact pairs $\{(\Omega_k, A_k)\}_{k \in \mathbb{N}}$ such that $\Omega_k \subseteq \Omega_{k+1}$ and $\mu_0(\Omega_k) \rightarrow 1$ as $k \rightarrow \infty$, and for every k the value of the information-design problem with the prior μ_0 restricted to Ω_k is unchanged when the action space is reduced from A to A_k .*

Theorem 2. (Strong Duality) *If Condition 1 holds, then $V^D = V^P$.*

To prove Theorem 2, we first apply Fenchel-Rockafellar duality and establish a zero duality gap when the state and action spaces are compact, the payoff functions are continuous, and the

optimal value is finite. For non-compact spaces, we approximate the problem with a sequence of compact subproblems and show that the primal and dual values are the limits of their compact counterparts. Part (iii) of [Condition 1](#) is important for this step, because it ensures that if the state space is bounded, the optimal actions can be bounded as well.

4 Robustness of Optimal Information Structures

By [Theorem 1](#), an allocation rule induced by a certifiably optimal information structure must be chosen freely by a fully informed dual agent. This observation has two consequences. First, the prior is irrelevant for the implementability of an allocation rule by incentives since the prescribed action profiles must be optimal state-by-state. Second, if the optimal allocation rule randomizes over several action profiles at a given state, the dual agent must be indifferent among those profiles and could therefore randomize over them with any probabilities. Hence, only the support of the action profiles matters for implementability by incentives:

Proposition 2. (Robustness) *Let two concave information-design problems differ only in their priors, if at all. Let $\pi_1 \in \Delta(A \times \Omega)$ be certifiably optimal in the first problem. If $\pi_2 \in \Delta(A \times \Omega)$ is implementable by information in the second problem and $\text{supp } \pi_2(\cdot \mid \omega) \subseteq \text{supp } \pi_1(\cdot \mid \omega)$ for all $\omega \in \Omega$, then π_2 is certifiably optimal in the second problem.*

[Proposition 2](#) shows that certifiably optimal information structures are, to some extent, *prior-robust*: once optimal under one prior, an information structure remains optimal under any other prior, provided it still implements an allocation rule whose support is no larger in every state. It is specific to our setting, which features a continuum of actions and thus more flexible players' best responses, and does not hold in generic games with finitely many actions (cf. [Kamenica and Gentzkow \(2011\)](#)).

Generally, the larger the support is, the easier it is to construct multiple information structures that implement allocation rules within that support. In the extreme case, if the action support covers the whole action space, then the support condition of [Proposition 2](#) has no bite, and any information structure can be certified to be optimal.

Corollary 1. (Full-Support Noise) *If measure $\pi \in \Delta(A \times \Omega)$ is certifiably optimal and $\text{supp } \pi(\cdot \mid \omega) = A$ for all $\omega \in \Omega$, then any information structure is certifiably optimal.*

Corollary 1 presents a case against using extraneous noises that induce full-support action profiles in concave games with finitely many players. The information structures that employ such noises can never be certifiably optimal, except in trivial cases in which the designer’s expected payoff is invariant to the information provided. This finding resonates with the analysis of Taneva (2019), who studied a two-player binary setting and showed that sending conditionally independent signals is never strictly optimal, as well as with the analysis of Candogan and Strack (2023), who established optimality of partitional signals in a class of games. However, such extraneous independent noises may optimally appear in the limit information structure as the number of players grows to infinity, as we discuss in Section 7.

Furthermore, Proposition 2 enables us to assess the optimality of full state transparency. We call an information structure *fully informative about the state* if each player can deduce the state with certainty from her private signal. There could be many such structures, each differing in the state-by-state coordination of players’ actions, including full disclosure. However, any such structure can implement the same allocation rule under all priors.

Corollary 2. (Full State Information) *If a certifiably optimal information structure is fully informative about the state, then it is certifiably optimal under all priors.*

Corollary 2 shows that, once full disclosure is optimal, it is extremely robust: it remains optimal under any prior (cf. Jehiel (2015)).

5 Linear-Quadratic Games: Common Value

The certification approach developed in the previous section applies to any concave game (see Online Appendix C) and, more generally, to any smooth game (Section 7). In this section, we demonstrate the approach in symmetric linear-quadratic games with a common state. We first develop general results and then illustrate them through two applications: capital fundraising and expectation polarization.

We define a *linear quadratic symmetric game with a common state* as a game with $\Omega \subseteq \mathbb{R}$, and the following payoff structure:

$$u_i(a, \omega) = (\omega + b_1)a_i + \frac{1}{2}(q^o - \frac{q^c}{N})a_i^2 + q^c a_i \bar{a} + f_1(a_{-i}, \omega), \quad (11)$$

$$v(a, \omega) = (h\omega + b_2)\bar{a} + p^o \check{a} + p^c \bar{a}^2 + f_2(\omega), \quad (12)$$

where $\bar{a} \triangleq \sum_j a_j/N$, $\check{a} \triangleq \sum_j a_j^2/N$, f_1 and f_2 are arbitrary functions of the indicated arguments, $h \geq 0$, and q^o , q^c , p^o , p^c satisfy the concavity and “genericity” conditions:

$$q^o < 0, \quad q^o + q^c < 0, \quad (13)$$

$$p^o + p^c \neq h(q^o + q^c), \quad q^o p^c \neq p^o q^c. \quad (14)$$

Condition (13) ensures that u_i is strictly concave in a_i for all $N \geq 1$, whereas condition (14) avoids knife-edge cases (e.g., division by zero). This game is strategically equivalent to, and thus admits the same solution as, a simpler *normalized* game, in which $\mathbb{E}[\omega] = 0$,

$$\begin{aligned} u_i(a, \omega) &= \omega a_i + \frac{1}{2}(q^o - \frac{q^c}{N})a_i^2 + q^c a_i \bar{a}, \\ v(a, \omega) &= (h\omega + b)\bar{a} + p^o \check{a} + p^c \bar{a}^2, \end{aligned}$$

where h , q^o , q^c , p^o , p^c are as in the original game and $b = h\mathbb{E}[\omega] + b_2 - 2(\mathbb{E}[\omega] + b_1)(p^o + p^c)/(q^o + q^c)$. Therefore, we can focus on an analysis of the normalized game.

Claim 1. *In a linear quadratic symmetric game with a common state, Condition 1 is satisfied and thus $V^D = V^P$.*

We prove Claim 1 by showing that when the state lies in a bounded interval, every action profile implementable by any information structure, optimal or otherwise, is bounded.

Consequently, optimality certificates exist. To find such certificates, observe that the players’ marginal payoffs are linear:

$$\dot{u}_i(a, \omega) = \omega + q^o a_i + q^c \bar{a}, \quad (15)$$

and the dual payoff function takes the form

$$w(a, \omega, \lambda) = (h\omega + b)\bar{a} + p^o \check{a} + p^c \bar{a}^2 - \sum_{i=1}^N \lambda_i(a_i)(\omega + q^o a_i + q^c \bar{a}). \quad (16)$$

Given the linear-quadratic nature of the environment, we conjecture that an optimal allocation rule is linear in the state. Because the allocation rule must be a best response given the

dual payoff (16), it is natural to posit a certificate that is affine and symmetric across players:

$$\lambda_i(a_i) = -\frac{1}{N}(xa_i + x_0). \quad (17)$$

With this certificate, the dual payoff function becomes

$$w(a, \omega, x, x_0) = x_0\omega + ((h+x)\omega + b + x_0(q^o + q^c))\bar{a} + (p^c + xq^c)\bar{a}^2 + (p^o + q^ox)\check{a}.$$

If the coefficient in front of \check{a} is nonzero, $p^o + q^ox \neq 0$, then the dual best-response is uniquely defined, linear in state, and symmetric across players; thus, such x could certify only the optimality of full disclosure or no disclosure. To certify the optimality of partial disclosure, we must have

$$x^* = -\frac{p^o}{q^o},$$

$$w(a, \omega, x^*, x_0) = x_0\omega + \left(\left(h - \frac{p^o}{q^o}\right)\omega + b + x_0(q^o + q^c)\right)\bar{a} + \left(p^c - \frac{p^o q^c}{q^o}\right)\bar{a}^2.$$

With this choice of x , the dual payoff function is a function solely of the average action \bar{a} . Because of (15) and $\mathbb{E}[\omega] = 0$, any incentive compatible expected average action must equal to 0; therefore we must have

$$x_0^* = -\frac{b}{q^o + q^c},$$

$$\bar{a}^*(\omega) = \frac{p^o - hq^o}{2(p^c q^o - p^o q^c)}\omega. \quad (18)$$

Any allocation rule that results in this average action is implementable by incentives by the affine contract (38). By Theorem 1, if the same allocation rule can be implemented by information, then that allocation rule is optimal.⁷ A natural way to implement such a linear average action is to fully inform a subset of players:

Definition 3. (Targeted Disclosure) *For $k \in \{0, 1, \dots, N\}$, a k -targeted disclosure is an information structure that fully reveals the state to k players while providing no information to others.*

⁷This observation resonates with the research on robust information design in regime change games, such as Inostroza and Pavan (2025) and Morris, Oyama, and Takahashi (2024), which is fundamentally concerned with multiple informational implementations of the same aggregate behavior.

A 0-targeted disclosure is no disclosure and N -targeted disclosure is full disclosure. Under a k -targeted disclosure, there exists an equilibrium where each uninformed player plays $a_i^{ND} \equiv 0$, whereas each informed player plays

$$a_i^I(\omega) = -\frac{N}{q^c k + q^o N} \omega,$$

resulting in the average action

$$\bar{a}(\omega) = \frac{k}{N} a_i^I(\omega) = -\frac{k}{q^c k + q^o N} \omega. \quad (19)$$

Theorem 3. (Optimality of Targeted Disclosure) *In a normalized linear-quadratic symmetric game with a common state, if*

$$k^* \triangleq \frac{q^o(hq^o - p^o)}{q^o(2p^c - hq^c) - p^o q^c} N \quad (20)$$

is in $\{0, 1, \dots, N\}$, then k^ -targeted disclosure is optimal. If $k^* \notin [0, N]$, then either (i) $q^o p^c < p^o q^c$ and $p^o + p^c < h(q^o + q^c)$, or $q^o p^c > p^o q^c$ and $p^o < hq^o$, in which case no disclosure is optimal, or (ii) $q^o p^c < p^o q^c$ and $p^o + p^c > q^o + q^c$, or $q^o p^c > p^o q^c$ and $(q^o - p^o)(hq^o + q^c) > 2(p^c q^o - p^o q^c)$, in which case full disclosure is optimal.*

Proof. If k is in $\{0, 1, \dots, N\}$, then by [Theorem 1](#), comparing (18) and (19), a k -targeted disclosure is optimal if

$$\begin{aligned} -\frac{k}{q^c k + q^o N} &= \frac{p^o - hq^o}{2(p^c q^o - p^o q^c)} \\ \Leftrightarrow k &= \frac{q^o(hq^o - p^o)}{q^o(2p^c - hq^c) - p^o q^c} N. \end{aligned}$$

If $k \notin [0, N]$, then given the conditions (13)–(14), either (i) or (ii) holds.⁸ In case of (i), the optimality of no disclosure can be certified by an affine contract with $(x, x_0) = (-h, -b/(q^o + q^c))$. In case of (ii), the optimality of full disclosure can be certified by an affine contract with $(x, x_0) = (h - 2(p^o + p^c)/(q^o + q^c), -b/(q^o + q^c))$. (See [Appendix A.3](#) for details.) \square

[Theorem 3](#) highlights the importance of targeted disclosure in linear-quadratic games with

⁸Note that conditions (i) and (ii) are not equivalent to $k^* < 0$ and $k^* > N$. Examples exist in which full disclosure is optimal even when $k^* < 0$, and others in which no disclosure is optimal even when $k^* > N$.

a common state. First, if $k^* \notin (0, N)$, then either 0-targeted disclosure or N -targeted disclosure is optimal. Second, if k^* is in $\{0, 1, \dots, N\}$, then k^* -targeted disclosure is exactly optimal, and the parameter space that yields it spans a wide range of economically relevant settings. Third, if $k^* \in (0, N)$ but is not an integer, a $\lceil k^* \rceil$ -targeted disclosure is approximately optimal, with the optimality loss vanishing as N grows large.⁹ Indeed, by weak duality, the payoff from any allocation rule that achieves average action (18) is an upper bound on the designer's objective. As N grows, the average action (19) generated by $\lceil k^* \rceil$ -targeted disclosure, converges to that bound. Because payoffs are continuous in the allocation rule, the designer's payoff under $\lceil k^* \rceil$ -targeted disclosure converges to the optimum.

The optimality of targeted disclosure is notable for two reasons. First, when $k^* \in [1, N - 1]$ the policy is asymmetric, even though the underlying environment is symmetric. Second, because the argument never relies on the prior, the same targeted-disclosure policy is optimal for *all* priors—underscoring its exceptional robustness in applications (cf. Proposition 2).

At the same time, optimal information structures need not be unique. When targeted disclosure is optimal, the designer is free to choose which players receive information. That choice leaves aggregate equilibrium outcomes unchanged but shifts surplus among players. For example, an ex-ante symmetric variant of k^* -targeted disclosure first publicly draws a random subset of k^* players and then applies k^* -targeted disclosure to that subset. Furthermore, qualitatively distinct information structures can also be optimal, specifically when the state is normally distributed.

Definition 4. (Gaussian Coupling) *For $\beta \in \mathbb{R}$ and $\sigma^2 > 0$, a (β, σ^2) -Gaussian coupling is an information structure such that for all i and ω ,*

$$s_i = \beta\omega + \varepsilon_i - \frac{1}{N-1} \sum_{j \neq i} \varepsilon_j,$$

where each noise term $\varepsilon_i \sim N(0, \sigma^2)$ is independent from ω and ε_j for $j \neq i$.

A Gaussian coupling distorts the state by adding Gaussian noise to each player's signal, with the individual noises coupled in such a way that they vanish upon aggregation—so that the sum of all signals is a deterministic function of the state ω . When $\beta \neq 0$, knowledge of all signals perfectly reveals the state. However, observing only s_i provides imperfect information

⁹ $\lceil k^* \rceil$ is the lowest integer larger than k^* .

about the state; in this sense, a Gaussian coupling splits the complete state information across players. When the state is normally distributed, a Gaussian coupling corresponds to a Gaussian information structure—that is, the signals and the state are jointly normally distributed.

Theorem 4. (Optimality of Gaussian Coupling) *In a normalized linear-quadratic symmetric game with a common state, if $\omega \sim N(0, \sigma_\omega^2)$ and k^* as defined in (20) lies in $(0, N)$,¹⁰ then a (β, σ^2) –Gaussian coupling is optimal, where*

$$\beta = \frac{p^o - hq^o}{2(p^c q^o - p^o q^c)}, \quad (21)$$

$$\sigma^2 = \frac{N-1}{N} \frac{\beta(1 + (q^c + q^o)\beta)}{-q^o} \sigma_\omega^2. \quad (22)$$

In the optimal equilibrium, $a_i \equiv s_i$.

Proof. Consider (β, σ^2) –Gaussian coupling with (β, σ^2) as defined by (21) and (22). If each player plays $a_i \equiv s_i$, then $\bar{a}(\omega) = \beta\omega = \bar{a}^*(\omega)$, as defined in (18). By Theorem 1, if $a_i \equiv s_i$ is incentive-compatible and $\sigma^2 > 0$, then the (β, σ^2) –Gaussian coupling is optimal.

To show incentive compatibility, observe that $\omega + q^o a_i + q^c \bar{a}$ and a_i are jointly normally distributed. Therefore, the incentive compatibility is equivalent to

$$\begin{aligned} \mathbb{E}[(\omega + q^o a_i + q^c \bar{a})a_i] &= 0, \\ \mathbb{E}[(\omega + q^o a_i + q^c \beta\omega)a_i] &= 0, \\ (1 + q^c \beta)\beta\sigma_\omega^2 + q^o \left(\beta^2\sigma_\omega^2 + \sigma^2 + \frac{1}{N-1}\sigma^2 \right) &= 0, \end{aligned}$$

which, after rearrangement, is equivalent to (22).

Finally, it is straightforward to verify that if $k^* \in (0, N)$, then $\sigma^2 > 0$; see Appendix A.3 for details. \square

Theorem 4 highlights an alternative way of providing information optimally—by carefully coupling the players’ individual noises. In the case of a normally distributed state, this information structure comes closest to the existing Gaussian information design literature (e.g., Angeletos and Pavan (2007, 2009), Bergemann and Morris (2013), Bergemann, Heumann, and Morris (2015a, 2021), Ui (2020)).

¹⁰By Theorem 3, if $k^* \notin (0, N)$, then either full disclosure or no disclosure is optimal.

The optimality of a Gaussian coupling depends crucially on the prior’s Gaussian form, making this solution less robust than targeted disclosure. Therefore, in the applications that are presented in the next two sections, we focus on optimal targeted disclosure; however, we encourage the reader to bear in mind the fact that if the state is normally distributed in those applications, then providing information via a Gaussian coupling is also optimal.

5.1 Persuading Investors

An important question in securities regulation is how much disclosure should be required from firms seeking capital (e.g., Carvajal, Rostek, and Sublet (2018)). Although many factors matter, in this section we focus on two. On the benefit side, greater transparency helps investors reduce risk and make better decisions; on the cost side, it intensifies cream-skimming, as investors herd into the same projects and crowd out investment gains.¹¹ We characterize the optimal balance between these forces within our framework and show that the resulting solution aligns with key venture-capital practices.

Formally, we consider an investment game in the spirit of Angeletos and Pavan (2007) and Bergemann and Morris (2013). There are $N \geq 2$ investors, who simultaneously decide how much to invest in a project, $A_i = \mathbb{R}$. The profitability of the project is uncertain; it depends on the unknown project quality $\theta \in \mathbb{R}$ and on the total amount of investment $\check{a} \triangleq \sum_{i=1}^N a_i$. The ex post payoff of player i is

$$u_i(a, \theta) = (\theta - r\check{a})a_i - ca_i, \quad (23)$$

where $r > 0$ is the congestion parameter and $c > 0$ is the opportunity cost of investment. As $r > 0$, the project features decreasing returns to scale, i.e., its average profitability decreases in the total investment. We can conveniently rewrite payoff (23) as

$$u_i(a, \omega) = r(\omega - \check{a})a_i, \quad (24)$$

where the state $\omega \in \Omega$ is a normalized project quality defined as $\omega \triangleq (\theta - c)/r$.

Given (24), for any belief $\mu \in \Delta(A_{-i} \times \Omega)$, the player i ’s best response can be found via

¹¹This phenomenon parallels cream-skimming among competing contractors on online platforms, which can likewise be mitigated through information design (Romanyuk and Smolin (2019)).

the first-order condition to equal

$$a_i^*(\mu) = \mathbb{E}_\mu \left[\frac{\omega - \check{a}_{-i}}{2} \right], \quad (25)$$

where $\check{a}_{-i} \triangleq \sum_{j \neq i} a_j$, so that the best response linearly increases in the normalized quality expectation and linearly decreases in the expected amount of total investment made by other players. The players' actions are thus strategic substitutes.

The information designer has full control over the information regarding the project quality and can privately convey it to each player, thereby persuading that player to invest more or less. The designer aims to maximize the total profits generated by the project with her ex post payoff being¹²

$$v(a, \omega) = \sum_{i=1}^N (\omega - \check{a}) a_i = (\omega - \check{a}) \check{a} = \omega \check{a} - \check{a}^2. \quad (26)$$

Investment Control If the designer could control the individual investment directly, then she would set the total investment to respond to the state as $\check{a}^{FB}(\omega) = \omega/2$. As we later show, this first-best allocation rule cannot be implemented via pure information control; however, it can be approximated.

No Disclosure and Full Disclosure Before proceeding with the characterization of the optimal information structure, we consider two extreme information structures that present natural benchmarks: no disclosure and full disclosure. Under no disclosure, all players base their investment decisions only on the prior estimate of project quality. By condition (27), the unique equilibrium is one in which each player invests

$$a_i^{ND} = \frac{1}{N+1} \mathbb{E}[\omega].$$

The investment is uniform across projects with different qualities, resulting in total investment $\check{a}^{ND} = N\mathbb{E}[\omega]/(N+1)$ and in the designer's payoff

$$v^{ND} = \mathbb{E} \left[(\omega - \check{a}^{ND}) \check{a}^{ND} \right] = \frac{N}{(N+1)^2} \mathbb{E}^2[\omega].$$

¹²Because the expected investment and the expected investment costs are invariant to information, this payoff also captures the objective of maximizing the total investor welfare.

In contrast, under full disclosure, each player has complete information about the project quality and incorporates it into his investment decision. For any commonly known state, the ensuing game admits a strict potential and has a unique equilibrium (Neyman (1997)). This equilibrium is symmetric and each player invests proportionally to the normalized quality as

$$a_i^{FD}(\omega) = \frac{1}{N+1}\omega,$$

resulting in total investment $\check{a}^{FD}(\omega) = N\omega/(N+1)$. The designer's payoff is

$$v^{FD} = \mathbb{E} \left[(\omega - \check{a}^{FD}(\omega)) \check{a}^{FD}(\omega) \right] = \frac{N}{(N+1)^2} \mathbb{E}[\omega^2] = \frac{N}{(N+1)^2} (\mathbb{E}^2[\omega] + \mathbb{V}[\omega]).$$

Comparing the designer's payoffs across the two benchmarks shows that full disclosure strictly dominates no disclosure. The advantage of full disclosure grows with the variance of the normalized project quality ω ; consequently, it increases with the project-quality variance $\mathbb{V}[\theta]$ and decreases with the congestion parameter r .

However, both information structures suffer from a scaling problem: as the number of players goes to infinity, the designer's payoff and thus the total project profit converges to zero. In the limit, the individual rent as well as the total profit are dissipated. We now show that this problem can be mitigated by adopting an optimal information structure.

Optimal Information For any information structure, take the ex-ante expectation of both sides of (25) and apply the law of iterated expectations (cf. Bergemann et al. (2017)). This yields each player's expected investment:

$$\mathbb{E}[a_i] = \frac{1}{N+1} \mathbb{E}[\omega]. \tag{27}$$

The expected total investment stays constant at

$$\mathbb{E}[\check{a}] = \frac{N}{N+1} \mathbb{E}[\omega] = \frac{N}{N+1} \frac{\mathbb{E}[\theta] - c}{r} > \frac{1}{2} \mathbb{E}[\omega] = \mathbb{E}[\check{a}^{FB}(\omega)],$$

and, consequently, information control can never achieve a full control benchmark. The total investment increases in the number of players, converging to the expected normalized quality as this number goes to infinity. Intuitively, as the number of players grows, congestion is

exacerbated, each individual investment decreases, and each player internalizes less of the congestion effect induced by his investment.

While the designer cannot affect the total amount of investment, she can direct the investment toward more productive projects. In doing so, in accordance with (26), the designer must balance two conflicting objectives. On the one hand, the designer wants investment to track project quality to maximize $\mathbb{E}[\omega\check{a}]$. On the other hand, the designer wants to reduce investment volatility—and thus congestion—to maximize $\mathbb{E}[-\check{a}^2]$.

Proposition 3. (Persuading Investment by Exclusivity) *For any number of players, a 1-targeted disclosure is optimal.*

Proof. The setting is an instance of the general setting (11-12) with $h = 1$, $q^o = -1$, $q^c = -N$, $p^o = 0$, and $p^c = -N$. The formula (20) for an optimal k^* becomes:

$$k^* = \frac{q^o(hq^o - p^o)}{q^o(2p^c - hq^c) - p^oq^c}N = \frac{-1(-1 - 0)}{-1(-2N + N) - 0}N = \frac{1}{N}N = 1, \quad (28)$$

and thus by Theorem 3, 1-targeted disclosure is optimal. \square

Proposition 3 shows that an optimal information structure takes a very simple form: the designer designates a single player and provides him with exclusive and full access to information. This player makes fully informed decisions. All other players invest the same amount regardless of the project quality. The same information structure is optimal for any number of investors, congestion and cost parameters, and project quality distribution.

Interestingly, this asymmetric treatment of investors mirrors common venture-capital practice: firms cultivate close relationships with a small set of venture capitalists who learn more about the firm than the broader pool of potential investors (e.g., Bernstein, Giroud, and Townsend (2016)). Although several explanations could account for this pattern, our results highlight its informational benefit: by creating a sharp information asymmetry among investors, it curbs congestion and raises both investment profits and investor welfare.

Under the exclusive disclosure, the optimal total investment is

$$\check{a}^*(\omega) = \sum_{i=1}^N a_i^*(\omega) = \frac{1}{2} \left(\omega + \frac{N-1}{N+1} \mathbb{E}[\omega] \right) \quad \forall \omega \in \Omega.$$

Thus, the optimal responsiveness of total investment to project quality coincides with that

under direct investment control. In a sense, while not being able to avoid congestion on average across different projects, optimal information design ensures that the total investment responds to the project quality linearly with a first-best responsiveness coefficient.

Importantly, optimal information design avoids rent dissipation as the number of players grows to infinity. Indeed, the optimal designer's payoff equals

$$v^* = \mathbb{E} \left[\left(\omega - \frac{\omega}{2} - \frac{N-1}{N+1} \mathbb{E}[\omega] \right) \left(\frac{\omega}{2} + \frac{N-1}{N+1} \mathbb{E}[\omega] \right) \right] = \frac{N}{(N+1)^2} \mathbb{E}^2[\omega] + \frac{1}{4} \mathbb{V}[\omega].$$

As such, the project's total profit always stays above $\mathbb{V}[\omega]/4$ and converges to this level as $N \rightarrow \infty$. The limiting payoff increases with the variance of project quality: under the optimal information structure, just as under full disclosure, riskier projects generate higher realized profits even when their expected quality is the same.¹³

5.2 Polarizing Predictions

Numerous studies document societal polarization, with individuals reaching sharply divergent views of the same issues (e.g., [Alesina, Miano, and Stantcheva \(2020\)](#)). Conventional explanations stress behavioral biases or naïveté, yet recent evidence shows that most people can reliably distinguish fake news from accurate reporting ([Angelucci and Prat \(2024\)](#)). This tension prompts a sharper question: To what extent can polarization arise among fully rational agents solely because they receive different information, and which information structure generates the greatest polarization?

We address this question in a prediction game, in which multiple players try to predict a common underlying state. The state is one-dimensional $\omega \in \mathbb{R}$ and distributed according to the prior μ_0 . There are $N \geq 2$ players, each choosing a prediction $a_i \in A_i = \mathbb{R}$. The ex post payoff to player i is

$$u_i(a, \omega) = -(a_i - \omega)^2.$$

¹³The setting in this section can also be viewed as Cournot competition with linear production costs and uncertain linear demand. In that interpretation, the rent dissipation under both the no-disclosure and full-disclosure benchmarks parallels the zero-profit outcome of large competitive markets. Targeted disclosure then amounts to informing only a single firm about the demand state, and this policy maximizes total producer surplus. The structure bears some resemblance to collusion that designates one firm as a monopolist, but information control is weaker than direct production control, because every firm still chooses a positive output even when uninformed. This perspective helps illuminate the role of trade associations, which cannot engage in illegal collusion yet can manage information flows (see, e.g., [Kirby \(1988\)](#); [Vives \(1990\)](#)).

Because each player's payoff depends only on his own prediction and the state, for any given belief $\mu \in \Delta(A_{-i} \times \Omega)$, player i 's best prediction is simply the posterior expectation of the state:

$$a_i^*(\mu) = \mathbb{E}_\mu[\omega].$$

We focus on the question of inducing maximal polarization of the players' predictions as measured by the pairwise squared sum:

$$v(a, \omega) = \sum_{i,j} (a_i - a_j)^2.$$

Given the payoff, the designer benefits from sending private signals: any public information structure, including full disclosure or no disclosure, leads the players to make the same predictions and consequently minimizes the designer's objective. An optimal information structure, on the one hand, must provide some state information to move players' predictions and, on the other hand, should heterogeneously obfuscate the information to counteract truth drifting.

Proposition 4. (Polarizing by Segregation) *Let the number of players N be even. Then, an $N/2$ -targeted disclosure is optimal.*

Proof. The setting is an instance of the general setting (11-12) with $h = 0$, $q^o = -2$, $q^c = 0$, $p^o = 2N^2$, and $p^c = -2N^2$. The formula (20) for an optimal k^* becomes:

$$k^* = \frac{q^o(hq^o - p^o)}{q^o(2p^c - hq^c) - p^oq^c} N = \frac{-2(-2N^2)}{-2(-4N^2) - 0} N = \frac{N}{2}, \quad (29)$$

and thus by Theorem 3, $N/2$ -targeted disclosure is optimal. \square

Proposition 4 shows that, irrespective of the prior, expectation polarization can be achieved by a strikingly simple information structure that informs only half of the population.

This finding aligns with earlier work in information design. In a two-player, binary-state model, Arieli et al. (2021) characterized the feasible joint belief distributions and reached a parallel conclusion: to maximize polarization of posterior expectations, it is optimal to inform one of the two players. Proposition 4 extends this insight to any number of players and states. Arieli and Babichenko (2024) continue the study of polarization of posterior beliefs. They show

that for any even number of players, informing half of them is optimal, because it achieves appropriate statistical bounds. Interestingly, even though belief polarization and expectation polarization differ once the state space is larger than binary, and the underlying arguments are not interchangeable, both problems share the same optimal information structure.

These results contribute to the debate over the sources of societal polarization. Recent studies show that political news consumption is sharply segregated, either because it flows through endogenous peer networks (Bowen et al. (2023)) or because it is curated by social-media algorithms (Braghieri et al. (2024)). As a result, different groups may learn about different topics. This kind of segregation—learning about different topics, rather than learning different facts about the same topic—can maximally polarize a society.

6 Linear-Quadratic Games: Interdependent Values

In this section, we apply the certification approach to strategic environments with interdependent values and thus with a multidimensional state. For tractability, we focus on jointly normally distributed state components. As in the previous section, we first present general results and then illustrate them with an economic application: informing competitive pricing.

Formally, we study symmetric games with $\omega \sim N(\mu \mathbf{1}, \sigma^2 M(1, \rho))$, where $M(x, y)$ denotes the $N \times N$ matrix whose diagonal entries equal x and off-diagonal entries equal y , and $\rho \in (-1/(N-1), 1)$. The players' payoffs are

$$u_i(a, \omega) = (\omega_i p^o + \check{\omega}_{-i} p^c + b_1) a_i + \frac{q^o}{2} a_i^2 + q^c a_i \check{a}_{-i} + f_1(a_{-i}, \omega), \quad (30)$$

$$v(a, \omega) = \sum_{i=1}^N (\omega_i \tilde{p}^o + b_2) a_i + \left(\sum_{i=1}^N \sum_{j \neq i} \omega_i a_j \right) \tilde{p}^c + \left(\sum_{i=1}^N a_i^2 \right) \tilde{q}^o + \left(\sum_{i=1}^N \sum_{j \neq i} a_i a_j \right) \tilde{q}^c + f_2(\omega), \quad (31)$$

where $\check{\omega}_{-i} \triangleq \sum_{j \neq i} \omega_j$, $\check{a}_{-i} \triangleq \sum_{j \neq i} a_j$, f_1 and f_2 are arbitrary functions of the indicated arguments, and q^o , q^c , p^o , p^c , \tilde{q}^o , \tilde{q}^c , \tilde{p}^o , \tilde{p}^c are commonly known parameters that satisfy the

concavity and “genericity” conditions:

$$q^o - q^c < 0, \quad q^o + (N - 1)q^c < 0, \quad p^o > |p^c|, \quad (32)$$

$$q^o \tilde{q}^c \neq q^c \tilde{q}^o, \quad (33)$$

$$(q^o - q^c)(\tilde{p}^o - \tilde{p}^c) \neq (\tilde{q}^o - \tilde{q}^c)(p^o - p^c), \quad (34)$$

$$(q^o + (N - 1)q^c)(\tilde{p}^o + (N - 1)\tilde{p}^c) \neq (\tilde{q}^o + (N - 1)\tilde{q}^c)(p^o + (N - 1)p^c). \quad (35)$$

Condition (32) ensures that u_i is strictly concave in a_i for all $N \geq 1$, whereas other conditions avoid knife-edge cases (e.g., division by zero).

This game is strategically equivalent to—and thus admits the same solution as—a simpler *normalized* game, in which $\omega \sim N(0, M(1, \rho))$,

$$\begin{aligned} u_i(a, \omega) &= (\omega_i p^o + \check{\omega}_{-i} p^c) a_i + \frac{q^o}{2} a_i^2 + q^c a_i \check{a}_{-i}, \\ v(a, \omega) &= \sum_{i=1}^N (\omega_i \tilde{p}^o + b) a_i + \left(\sum_{i=1}^N \sum_{j \neq i} \omega_i a_j \right) \tilde{p}^c + \left(\sum_{i=1}^N a_i^2 \right) \tilde{q}^o + \left(\sum_{i=1}^N \sum_{j \neq i} a_i a_j \right) \tilde{q}^c, \end{aligned}$$

and where $q^o, q^c, p^o, p^c, \tilde{q}^o, \tilde{q}^c, \tilde{p}^o, \tilde{p}^c$ are the same as in the original game. = Therefore, we can focus on the analysis of the normalized game.

By arguments analogous to those in the previous section, we obtain:

Claim 2. *In a linear quadratic symmetric game with a multidimensional state, [Condition 1](#) is satisfied and thus $V^D = V^P$.*

Consequently, optimality certificates exist. Given the greater complexity of the present setting, we proceed slightly differently from the previous section. We still rely on affine contracts for certification; however, instead of calculating the optimal contract parameters in closed form, we invoke an envelope theorem. We show that an optimal symmetric affine contract implements an obedient allocation rule, thereby certifying both its global optimality in the dual problem and the optimality of the corresponding measure in the primal problem.

Formally, observe that the players’ marginal payoffs are linear:

$$\dot{u}_i(a, \omega) = \omega_i p^o + \check{\omega}_{-i} p^c + q^o a_i + q^c \check{a}_{-i}, \quad (36)$$

and the dual payoff function is

$$w(a, \omega, \lambda) = v(a, \omega) - \sum_{i=1}^N \lambda_i(a_i)(\omega_i p^o + \check{\omega}_{-i} p^c + q^o a_i + q^c \check{a}_{-i}). \quad (37)$$

Given the linear-quadratic nature of the environment, we conjecture that an optimal allocation rule is linear in the state. As this allocation rule must be a best-response given the dual payoff (37), we posit a certificate that is affine and symmetric across players:

$$\lambda_i(a_i) = -(x a_i + x_0). \quad (38)$$

With this certificate, the dual payoff function can be written as

$$\begin{aligned} w(a, \omega, x, x_0) &= \omega^T M(\tilde{p}^o + x p^o, \tilde{p}^c + x p^c) a + a^T M(\tilde{q}^o + x q^o, \tilde{q}^c + x q^c) a \\ &\quad + (b + x_0(q^o + (N-1)q^c))\check{a}. \end{aligned}$$

Because of (36) and $\mathbb{E}[\omega] = 0$, any incentive-compatible expected individual action must equal zero, and therefore we conjecture that the coefficient in front of the last sum equals zero and thus

$$\begin{aligned} x_0^* &= -\frac{b}{q^o + (N-1)q^c}, \\ w(a, \omega, x, x_0^*) &= \omega^T M(\tilde{p}^o + x p^o, \tilde{p}^c + x p^c) a + a^T M(\tilde{q}^o + x q^o, \tilde{q}^c + x q^c) a \\ &= \omega^T M(\bar{p}^o(x), \bar{p}^c(x)) a + a^T M(\bar{q}^o(x), \bar{q}^c(x)) a. \end{aligned} \quad (39)$$

where we defined

$$\bar{q}^o(x) \triangleq \tilde{q}^o + x q^o, \quad \bar{q}^c(x) \triangleq \tilde{q}^c + x q^c, \quad (40)$$

$$\bar{p}^o(x) \triangleq \tilde{p}^o + x p^o, \quad \bar{p}^c(x) \triangleq \tilde{p}^c + x p^c. \quad (41)$$

For the optimal allocation rule $a^*(\omega)$ to maximize (39), $M(\bar{q}^o(x), \bar{q}^c(x))$ must be negative semi-definite; otherwise, the value of the dual payoff diverges to $+\infty$.

Lemma 1. $M(\bar{q}^o(x), \bar{q}^c(x))$ is negative semi-definite if and only if

$$x \geq \underline{x} \triangleq \max \left\{ -\frac{\tilde{q}^o - \tilde{q}^c}{q^o - q^c}, -\frac{\tilde{q}^o + (N-1)\tilde{q}^c}{q^o + (N-1)q^c} \right\}.$$

$M(\bar{q}^o(x), \bar{q}^c(x))$ is negative definite if and only if $x > \underline{x}$.

If $x > \underline{x}$, $M(\bar{q}^o(x), \bar{q}^c(x))$ is negative definite, and the optimal best-response is

$$a^*(\omega, x) = -\frac{1}{2}M^{-1}(\bar{q}^o(x), \bar{q}^c(x))M(\bar{p}^o(x), \bar{p}^c(x))\omega. \quad (42)$$

The resulting expected value of the dual payoff as a function of x is

$$W(x) \triangleq \mathbb{E}[w(a^*(\omega, x), \omega, x, x_0^*)] \quad (43)$$

$$= -\frac{1}{4}\mathbb{E} \left[\omega^T M(\bar{p}^o(x), \bar{p}^c(x))M^{-1}(\bar{q}^o(x), \bar{q}^c(x))M(\bar{p}^o(x), \bar{p}^c(x))\omega \right]. \quad (44)$$

Lemma 2. $\lim_{x \downarrow \underline{x}} W(x) = \lim_{x \uparrow +\infty} W(x) = +\infty$, and $\min_{x \geq \underline{x}} W(x)$ exists.

By Lemma 2, the dual payoff admits a minimum at some $x^* > \underline{x}$. We show that this x^* certifies an optimal information structure.

Definition 5. (Linear Disclosure) For $R = (r_1, \dots, r_N) \in \mathbb{R}^{N \times N}$, an R -linear disclosure is an information structure such that for all i and ω ,

$$s_i = r_i^T \omega = \sum_{j=1}^N r_{i,j} \omega_j.$$

A linear disclosure informs player i about a linear statistic of the state with the responsiveness weights r_i . When $N = 1$, a linear disclosure corresponds to no disclosure if $r = 0$, and to full disclosure otherwise. When $N \geq 2$ and R has full rank, each signal is imperfectly informative, yet the full signal profile perfectly reveals the state. In that case, like a Gaussian coupling, a linear disclosure spreads complete state information across players. When the state is normally distributed, a linear disclosure is a Gaussian information structure.

Theorem 5. (Optimality of Linear Disclosure) *In a normalized symmetric linear-quadratic Gaussian game with a multidimensional state, an R^* -linear disclosure is optimal, where*

$$R^* = -\frac{1}{2}M^{-1}(\bar{q}^o(x^*), \bar{q}^c(x^*))M(\bar{p}^o(x^*), \bar{p}^c(x^*)), \quad (45)$$

and x^* is any point in $\arg \min_{x > \underline{x}} W(x)$. In the optimal equilibrium, $a_i \equiv s_i$. The induced measure $\pi \in \Delta(A \times \Omega)$ is uniquely optimal.

Proof. By Lemma 2, the dual payoff admits a minimum at some $x^* > \underline{x}$. Consider any such x^* . By (42), the dual best-response is $a^*(\omega, x^*) = R^*\omega$; hence, this allocation rule is implementable by incentives, by the contract $\lambda_i(a_i) = x^*a_i + x_0^*$. The allocation rule is symmetric in i and linear in ω . Recall that

$$W(x) = \mathbb{E}[v(a^*(\omega, x), \omega) + \sum_{i=1}^N (xa_i^*(\omega, x) + x_0^*)\dot{u}_i(a^*(\omega, x), \omega)],$$

Since $a^*(\omega, x)$ is an interior maximizer and x is an interior minimizer of the dual payoff, it follows by the Envelope Theorem that

$$\left. \frac{dW}{dx} \right|_{x=x^*} = \mathbb{E} \left[\sum_{i=1}^N a_i^*(\omega, x^*) \dot{u}_i(a^*(\omega, x^*), \omega) \right] = 0.$$

By symmetry in i , it follows that for all i ,

$$\mathbb{E} [a_i^*(\omega, x^*) \dot{u}_i(a^*(\omega, x^*), \omega)] = 0.$$

Because $a_i^*(\omega, x^*)$ is linear in ω and the components of ω are jointly normally distributed, $a^*(\omega, x^*)$ and $\dot{u}_i(a^*(\omega, x^*), \omega)$ are also jointly normally distributed. Furthermore, $\mathbb{E} [\dot{u}_i(a^*(\omega, x^*), \omega)] = 0$, and thus $a^*(\omega, x^*)$ and $\dot{u}_i(a^*(\omega, x^*), \omega)$ are independent, and for all i and ω ,

$$\mathbb{E} [\dot{u}_i(a^*(\omega, x^*), \omega) \mid a_i^*(\omega, x^*)] = 0. \quad (46)$$

Therefore, the allocation rule $a^*(\omega, x^*)$ is implementable by information, by R^* -linear disclosure. By Theorem 1, the optimality follows.

Finally, as the dual agent's best response to $\lambda(x^*, x_0^*)$ is unique, it follows from Proposition 1

that the corresponding measure is uniquely optimal. \square

Theorem 5 shows that an optimal information structure can be found within a simple class of symmetric noise-free Gaussian information structures (cf. Bergemann et al. (2015b)). Moreover, its parameters can be identified by solving a one-dimensional minimization of $W(x)$.

In contrast with the common-state setting of **Section 5**, the optimal measure is unique here. Formally, in that setting the optimal certificate leaves the dual agent indifferent among many action profiles, whereas in this interdependent-value environment it yields a single best response. Conceptually, having one state component per player introduces richer uncertainty, which in turn fully determines the optimal recommendations.

6.1 Informing Competitive Pricing

The information design machinery is useful to understand and guide the design of digital platforms and, more generally, of algorithmic information processing (e.g., Bergemann and Bonatti (2024), Ichihashi and Smolin (2025)). In this section, we view the designer as a platform that knows demand conditions better than firms do, thanks to a larger, more recent sales dataset and superior analytics. The platform can privately convey this information to each firm by granting it access to personalized data analysis or direct price recommendations, thereby persuading the firm regarding its pricing decisions. The platform aims to maximize a weighted average of consumer and producer surplus. We characterize the information structure such a designer would optimally design and the resulting price behavior.

Formally, we consider a differentiated duopoly game in the spirit of Vives (1984) and Gal-Or (1985).¹⁴ There are two firms that operate in the market. Each firm sells a single product and competes in price with its opponent, so action a_i is the price set by firm i . Demand is ex ante symmetric across firms and is linear in prices and demand shocks:

$$q_i(a, \omega) = \omega_i - a_i + \eta a_{-i}, \quad (47)$$

where ω_i captures an individual demand shock for firm i 's product and η is a *cross-price*

¹⁴This formalization complements that adopted by Elliott et al. (2022), who study information design in unit-demand competition. Naturally, the details of their optimal information structures differ from ours; however, their interpretation of information provision as market segmentation and their broader discussion of how information shapes competitive outcomes can be applied in our setting.

sensitivity that satisfies $|\eta| < 1$, so that $-1 < \eta < 0$ corresponds to the case of complementary products, whereas $0 < \eta < 1$ corresponds to the case of substitute products.

The state is two-dimensional $\omega = (\omega_1, \omega_2) \in \mathbb{R}^2$ and comprises the individual demand shocks. The shocks are independently and identically distributed according to a normal distribution, $\omega_i \sim N(\omega_0, \sigma^2)$.¹⁵ The firms do not know the state but can be informed about it by the designer. The firms have quadratic costs of production so that their profits are

$$u_i(a, \omega) = a_i q_i(a, \omega) - c q_i(a, \omega)^2. \quad (48)$$

The resulting ex post values of consumer surplus and producer surplus are

$$CS(a, \omega) = \frac{a_1^2}{2} + \frac{a_2^2}{2} - \eta a_1 a_2 - a_1 \omega_1 - a_2 \omega_2, \quad (49)$$

$$PS(a, \omega) = u_1(a, \omega) + u_2(a, \omega). \quad (50)$$

The designer's payoff is a convex combination of consumer and producer surpluses, with $\delta \in [0, 1]$ representing the weight assigned to the consumer surplus:

$$v(a, \omega) = \delta CS(a, \omega) + (1 - \delta) PS(a, \omega). \quad (51)$$

The optimal designer's choices in the extreme cases $\delta = 1$ and $\delta = 0$ correspond to consumer-optimal and producer-optimal information structures, respectively, whereas the choice in the case $\delta = 1/2$ corresponds to the socially efficient information structure. As the welfare weight δ spans the interval $[0, 1]$, the corresponding solutions span the Pareto frontier in the space of consumer and producer surpluses.

Price Control We begin the analysis by studying a hypothetical scenario in which the designer can directly control the prices set by the firms. This scenario constitutes a first-best benchmark; it provides an upper bound on the designer's payoff and illustrates the designer's preferred pricing.

This problem admits a solution only if δ is not excessively high: there is a threshold value $\bar{\delta} \in (0, 1)$, such that if $\delta > \bar{\delta}$, then the designer can arbitrarily increase her payoff by setting

¹⁵As standard, this specification allows the prices and quantities to be negative.

arbitrarily large negative prices, because the monetary transfer to consumers outweighs any allocation inefficiency. In contrast, if $\delta < \bar{\delta}$, then the designer's problem is well-behaved; it is concave and admits a unique solution that can be found by first-order conditions to (51):

$$a_i^{FB}(\omega) = r_o^{FB}\omega_i + r_c^{FB}\omega_{-i},$$

where r_o^{FB} and r_c^{FB} measure the optimal responsiveness of each firm's price to its own demand shock and to the shock of its competitor (explicit formulas appear in [Section A.6](#)).

No Information and Full Disclosure The designer does not control prices directly but rather indirectly through demand information she supplies to firms. Before deriving the generally optimal policy, it is instructive to analyze two information benchmarks: no disclosure and full disclosure.

Under no disclosure, the firms' beliefs stay at the prior, and the equilibrium prices satisfy the first-order conditions derived from (48). In equilibrium, each firm sets a price

$$a_i^{ND} = \frac{1 + 2c}{2(1 + c) - \eta(1 + 2c)}\omega_0.$$

Lacking demand information, firms fix their prices at a level proportional to the expected demand.

In contrast, under full disclosure, the demand shocks are always commonly known. In equilibrium, each firm responds linearly to the shocks perfectly anticipating the price of its opponent:

$$a_i^{FD}(\omega) = r_o^{FD}\omega_i + r_c^{FD}\omega_{-i} = \frac{2(1 + c)(1 + 2c)}{4(1 + c)^2 - \eta^2(1 + 2c)^2}\omega_i + \frac{\eta(1 + 2c)^2}{4(1 + c)^2 - \eta^2(1 + 2c)^2}\omega_{-i}.$$

In a sense, this behavior enriches price-setting under no disclosure. If $\omega_1 = \omega_2 = \omega_0$, then prices are the same as those under no disclosure, $a_i^{FD}(\omega_0, \omega_0) = a_i^{ND}$. If $\omega_1 \neq \omega_2$, then demand is asymmetric across firms, and prices are adjusted to reflect competitive advantages. However, the prices average to the prices under no disclosure, $\mathbb{E}[a_i^{FD}(\omega)] = a_i^{ND}$.

Optimal Information The choice of any of the extreme information structures has drawbacks. Providing no disclosure misses the opportunity to strengthen the link between demand

and allocation and thus potentially limits efficiency. Providing full disclosure may exacerbate competition and dissipate firm profits. Instead, and as a direct consequence of [Theorem 5](#), the optimal information structure is partially informative and takes a form of a linear disclosure.

Proposition 5. (Managing Competition by Linear Statistics) *If $\delta \neq \hat{\delta}(\eta, c)$, an optimal direct information structure exists, is unique, and recommends for some coefficients a_0^* , r_o^* and r_c^* prices*

$$a_i(\omega) = a_0^* + r_o^* \omega_i + r_c^* \omega_{-i}. \quad (52)$$

Proof. The setting is an instance of the general framework [\(30-31\)](#) with $N = 2$, $\rho = 0$, $b_1 = b_2 = 0$ and $p_o = 1 + 2c$, $p_c = 0$, $q_o = -2(1 + c)$, $q_c = \eta(1 + 2c)$, $\tilde{p}_o = \delta(-1) + (1 - \delta)(1 + 2c)$, $\tilde{p}_c = (1 - \delta)(-2c\eta)$, $\tilde{q}_o = \delta/2 + (1 - \delta)(-1 - c(1 + \eta^2))$, $\tilde{q}_c = -\delta\eta/2 + (1 - \delta)\eta(1 + 2c)$.

The non-genericity conditions [\(33-35\)](#) reduce to

$$\delta \neq \hat{\delta}(\eta, c) \triangleq \frac{2 + 6c + 4c^2 - 8c|\eta| - 8c^2|\eta| + 2c\eta^2 + 4c^2\eta^2}{5 + 8c + 4c^2 - |\eta| - 10c|\eta| - 8c^2|\eta| + 2c\eta^2 + 4c^2\eta^2}. \quad (53)$$

If $\eta = 0$, condition [\(33\)](#) is violated. If $\eta < 0$, at $\delta = \hat{\delta}$ condition [\(34\)](#) is violated. If $\eta > 0$, at $\delta = \hat{\delta}$ condition [\(35\)](#) is violated. Otherwise, the conditions hold.

The resulting lower bound on x is

$$\underline{x} = \max \left\{ -\frac{\tilde{q}_o - \tilde{q}_c}{q_o - q_c}, -\frac{\tilde{q}_o + \tilde{q}_c}{q_o + q_c} \right\}. \quad (54)$$

The expected dual payoff reduces to

$$W(x) = -\frac{\sigma^2(\bar{q}_o(x)\bar{p}_o(x)^2 + \bar{q}_o(x)\bar{p}_c(x)^2 - 2\bar{q}_c(x)\bar{p}_o(x)\bar{p}_c(x))}{2(\bar{q}_o(x)^2 - \bar{q}_c(x)^2)}. \quad (55)$$

By [Theorem 5](#) and normalization [\(A.5\)](#), whenever $\delta \neq \hat{\delta}$, the optimal direct information structure is unique and is given by [\(52\)](#) with

$$\begin{aligned} r_o^* &= -\frac{\bar{q}_o(x^*)\bar{p}_o(x^*) - \bar{q}_c(x^*)\bar{p}_c(x^*)}{2(\bar{q}_o(x^*)^2 - \bar{q}_c(x^*)^2)}, \\ r_c^* &= -\frac{\bar{q}_o(x^*)\bar{p}_c(x^*) - \bar{q}_c(x^*)\bar{p}_o(x^*)}{2(\bar{q}_o(x^*)^2 - \bar{q}_c(x^*)^2)}, \end{aligned}$$

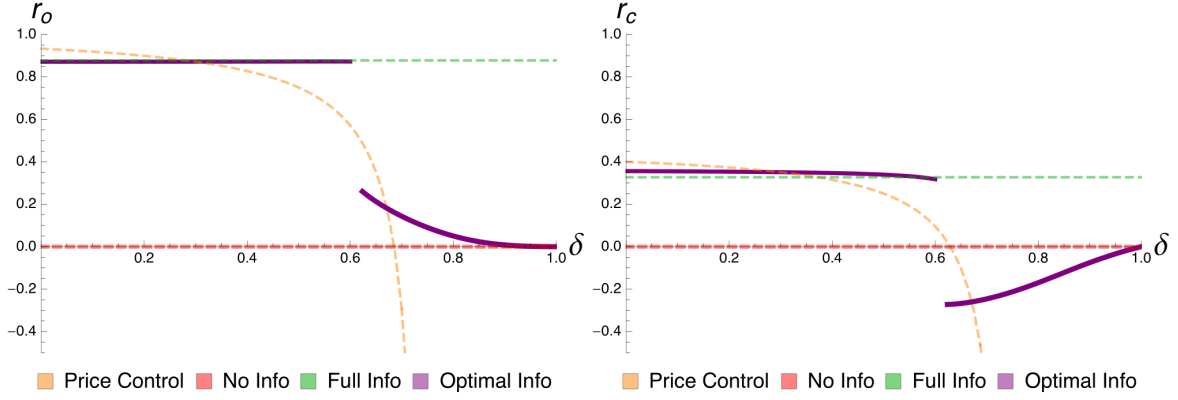


Figure 1: Price responsiveness to own demand shock (left) and opponent's demand shock (right) under price control, no disclosure, full disclosure, and optimal information structure. Calculated at $c = 1$, $\omega_0 = 3$, $\sigma^2 = 1$, and $\eta = 1/2$ for different values of consumer weight δ .

and x^* being any minimizer of (55) over $x^* > \underline{x}$. The constant term a_0^* can be found via mean action invariance: $a_0^* = a_i^{ND} - (r_o^* + r_c^*)\omega_0$. \square

The uniquely optimal direct information structure treats the firms symmetrically. Under it, each firm can only infer a linear combination of the individual demand shocks, and thus not able to infer whether the recommendation of a given price stems from its own demand conditions or the conditions of its competitor. Generically, $r_c^* \neq r_c^{FD}$ and $r_o^* \neq r_o^{FD}$, so providing full disclosure is suboptimal; instead, each firm receives personalized information that differs from its competitor's. This finding suggests that restricting disclosure to be public, as often done in the oligopoly literature on information sharing (Vives (1990, 1999)), though natural in some contexts, may entail losses.

Numerical Example To illustrate the optimal information structure and comparative statistics with respect to consumer weight, we fix the baseline parameters of the problem at $c = 1$, $\omega_0 = 3$, $\sigma^2 = 1$, and $\eta = 1/2$, so that the products are imperfect substitutes. The firms' equilibrium strategies under direct price control, no disclosure, full disclosure, and optimal information are linear in the state, and their responsiveness coefficients are depicted in Figure 1 as functions of consumer weight δ .

Figure 1 reveals the presence of two distinct optimal information regimes: the coordination regime that benefits firms and the anti-coordination regime that benefits consumers. For




Figure 2: Optimal price volatility (left) and price correlation (right). Calculated at $c = 1$, $\omega_0 = 3$, $\sigma^2 = 1$, and $\eta = 1/2$ for different values of consumer weight δ .

$\delta < \hat{\delta} = 11/18$, the optimally induced prices resemble those under full disclosure, yet full disclosure itself is not optimal. At $\delta = 0$, to benefit firms, the designer induces larger responses to the opponent's demand. As δ increases, the price responsiveness decreases. Around $\delta = \hat{\delta}$, the optimal information structure undertakes a discontinuous structural change: r_o plummets in absolute value, whereas r_c changes its sign, so firms begin responding oppositely to the same demand shock. As δ further increases in the region $\delta > \hat{\delta}$, the responsiveness parameters gradually decrease in their absolute values. At $\delta = 1$, both parameters equal zero: a designer who wishes to maximize solely consumer surplus should not inform firms at all. Why does the optimal information structure exhibit discontinuous change with respect to the consumer weight? Formally, the optimal certificate changes continuously for all $\delta \in (0, 1)$. However, at $\delta = \hat{\delta}$, the certificate permits a range of the dual agent's best responses that connect one branch of best responses to another, so that the implemented allocation rule displays a jump. Intuitively, the pricing induced under the optimal information structure aims to approximate its first-best counterpart under direct price control. However, without the ability to enforce prices directly, the designer opts for one of the two information regimes that better approximates it.

The two optimal information regimes translate into distinct patterns of price volatility and correlation, which can be measured by the standard deviation and the Pearson correlation

coefficient, respectively:

$$\sigma_i = \text{var}(a_i)^{1/2} = \sigma(r_o^2 + r_c^2)^{1/2},$$

$$\rho = \frac{\text{cov}(a_i, a_{-i})}{\sigma_i \sigma_{-i}} = \frac{2r_o r_c}{r_o^2 + r_c^2}.$$

For lower consumer weights, the price volatility is high, and the prices are positively correlated (Figure 2). This region is marked by coordination and high volatility of prices. In contrast, for higher consumer weights, the price volatility is low, and the product prices are negatively correlated. This region is associated with anticoordination and low volatility of prices.

This comparative statics yields two main insights. First, patterns of price volatility and correlations can serve as effective diagnostics for discerning which side of the market—consumers or producers—the information structure is designed to favor. Second, even a slight shift in the design objective can trigger a dramatic change in the information provision and in resulting market behavior. This, in turn, can serve as a cautionary tale, underscoring the potential for market instability introduced by algorithms, which may pursue shifting objectives and adapt their policies more rapidly than human decision makers.

7 Extensions

Bounded Action Spaces We have assumed that action spaces are unbounded, to avoid corner solutions and simplify the exposition. However, the certification approach can be easily extended to games with upper and lower bounds on each player’s action space, as demonstrated in Appendix B. The primal problem formulation remains the same, except that the boundary first-order conditions must be inequalities rather than equalities. The dual problem formulation is identical, with the added sign constraints on the contract functions evaluated at the boundary actions.

Infinite Economies Our analysis considers games with a finite number of players. It does not directly cover games with a continuum of players, as considered in the literature on infinite economies, e.g., by Angeletos and Pavan (2007) and Bergemann and Morris (2013), even though we anticipate that an analogous certification approach, suitably extended, would work in those settings as well.

Nonetheless, our analysis enables determining an optimal information structure in a game with a fixed number of players and then examining its limit as the number of players approaches infinity. Our results suggest that in many quadratic economies with a normally distributed state, Gaussian information structures are certifiably optimal and an optimal aggregate behavior is a deterministic function of a state. With a finite number of players, the latter condition necessitates the designer to either not introduce extraneous noise at all or to ensure that the noises are carefully coupled. With an infinite number of players, this condition can be met by adding independent noises and relying on the law of large numbers; the optimal aggregate behavior can be ensured by the large population size rather than by noise correlation. However, our analysis suggests that even in infinite economies, alternative information structures such as targeted disclosure may also be optimal, and robustly so.

General Smooth Games The tractability of the certification approach in concave games hinges on the capacity of first-order conditions to succinctly represent players’ incentives. In general smooth games, i.e., games in which each player’s payoff is continuously differentiable in their own action, these first-order conditions may be insufficient, as they can select a suboptimal local maximum or even a minimum in a player’s payoff. However, they remain necessary, which enables a straightforward extension of our analysis: in a smooth but not necessarily concave information-design problem, one constructs a relaxed primal problem that features only first-order conditions. This problem can be solved using the certification approach. In the final step, one verifies whether players obey the recommendation of the information structure found. If so, this information structure solves the original information-design problem.

8 Conclusion

In this paper, we introduced a certification approach to solving information-design problems in games and demonstrated its effectiveness and tractability in symmetric linear-quadratic settings. Our findings shed light on disclosure practices that guide investment, on the socially efficient control of information in markets, and on the limits of Bayesian polarization. Along the way, we provided theoretical justification for the use of targeted disclosure, Gaussian coupling, and linear disclosure.

Our analysis lays the groundwork and offers tools for studying information design in general

smooth games. We see at least two promising avenues for further research. First, our theoretical framework could be applied to other important economic settings, such as contests, public-good provision, and labor or financial markets. Second, the framework could be expanded to incorporate elements like information elicitation, information spillovers across players, and dynamic interaction. We expect all of these extensions to be feasible within the approach we have outlined.

References

- ALESINA, A., A. MIANO, AND S. STANTCHEVA (2020): “The Polarization of Reality,” in *AEA Papers and Proceedings*, vol. 110, 324–328.
- AMADOR, M. AND K. BAGWELL (2013): “The Theory of Optimal Delegation with an Application to Tariff Caps,” *Econometrica*, 81, 1541–1599.
- AMADOR, M., I. WERNING, AND G.-M. ANGELETOS (2006): “Commitment vs. Flexibility,” *Econometrica*, 74, 365–396.
- ANDERSON, E. AND P. NASH (1987): *Linear Programming in Infinite-dimensional Spaces: Theory and Applications*, Wiley.
- ANGELETOS, G.-M. AND A. PAVAN (2007): “Efficient Use of Information and Social Value of Information,” *Econometrica*, 75, 1103–1142.
- (2009): “Policy with Dispersed Information,” *Journal of the European Economic Association*, 7, 11–60.
- ANGELUCCI, C. AND A. PRAT (2024): “Is Journalistic Truth Dead? Measuring How Informed Voters Are about Political News,” *American Economic Review*, 114, 887–925.
- ARIELI, I. AND Y. BABICHENKO (2019): “Private Bayesian Persuasion,” *Journal of Economic Theory*, 182, 185–217.
- (2024): “A Population’s Feasible Posterior Beliefs,” *Journal of Economic Theory*, 215, 105764.

- ARIELI, I., Y. BABICHENKO, AND F. SANDOMIRSKIY (2023): “Persuasion as Transportation,” *Working paper*.
- ARIELI, I., Y. BABICHENKO, F. SANDOMIRSKIY, AND O. TAMUZ (2021): “Feasible Joint Posterior Beliefs,” *Journal of Political Economy*, 129, 2546–2594.
- BERGEMANN, D. AND A. BONATTI (2024): “Data, Competition, and Digital Platforms,” *American Economic Review*, 114, 2553–2595.
- BERGEMANN, D., T. HEUMANN, AND S. MORRIS (2015a): “Information and Volatility,” *Journal of Economic Theory*, 158, 427–465.
- (2015b): “Information and Volatility,” *Journal of Economic Theory*, 158, 427–465.
- (2017): “Information and Interaction,” *Working paper*.
- (2021): “Information, Market Power, and Price Volatility,” *RAND Journal of Economics*, 52, 125–150.
- BERGEMANN, D. AND S. MORRIS (2013): “Robust Predictions in Games with Incomplete Information,” *Econometrica*, 81, 1251–1308.
- (2016): “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games,” *Theoretical Economics*, 11, 487–522.
- BERNSTEIN, S., X. GIROUD, AND R. R. TOWNSEND (2016): “The Impact of Venture Capital Monitoring,” *Journal of Finance*, 71, 1591–1622.
- BOWEN, T. R., D. DMITRIEV, AND S. GALPERTI (2023): “Learning from Shared News: When Abundant Information Leads to Belief Polarization,” *Quarterly Journal of Economics*, 138, 955–1000.
- BRAGHERI, L., S. EICHMEYER, R. LEVY, M. M. MOBIUS, J. STEINHARDT, AND R. ZHONG (2024): “Level Slant and Polarization of News Consumption on Social Media,” *Working paper*.
- BROOKS, B. AND S. DU (2021): “Optimal Auction Design with Common Values: An Informationally Robust Approach,” *Econometrica*, 89, 1313–1360.

- CANDOGAN, O. AND P. STRACK (2023): “Optimal Disclosure of Information to Privately Informed Agents,” *Theoretical Economics*, 18, 1225–1269.
- CARROLL, G. (2017): “Robustness and Separation in Multidimensional Screening,” *Econometrica*, 85, 453–488.
- CARVAJAL, A., M. ROSTEK, AND G. SUBLET (2018): “Information Design and Capital Formation,” *Journal of Economic Theory*, 176, 255–292.
- CHAN, J., S. GUPTA, F. LI, AND Y. WANG (2019): “Pivotal Persuasion,” *Journal of Economic theory*, 180, 178–202.
- CHIAPPORI, P.-A., B. SALANIÉ, AND Y. WEISS (2017): “Partner Choice, Investment in Children, and the Marital College Premium,” *American Economic Review*, 107, 2109–2167.
- DIZDAR, D. AND E. KOVÁČ (2020): “A Simple Proof of Strong Duality in the Linear Persuasion Problem,” *Games and Economic Behavior*, 122, 407–412.
- DU, S. (2018): “Robust Mechanisms under Common Valuation,” *Econometrica*, 86, 1569–1588.
- DWORCZAK, P. AND A. KOLOTILIN (2024): “The Persuasion Duality,” *Theoretical Economics*, 19, 1701–1755.
- DWORCZAK, P. AND G. MARTINI (2019): “The Simple Economics of Optimal Persuasion,” *Journal of Political Economy*, 127, 1993–2048.
- ELLIOTT, M., A. GALEOTTI, A. KOH, AND W. LI (2022): “Market Segmentation through Information,” *Working paper*.
- GAL-OR, E. (1985): “Information Sharing in Oligopoly,” *Econometrica*, 329–343.
- GALICHON, A. AND B. SALANIÉ (2022): “Cupid’s Invisible hand: Social Surplus and Identification in Matching Models,” *Review of Economic Studies*, 89, 2600–2629.
- GALPERTI, S., A. LEVKUN, AND J. PEREGO (2024): “The Value of Data Records,” *Review of Economic Studies*, 91, 1007–1038.
- GALPERTI, S. AND J. PEREGO (2018): “A Dual Perspective on Information Design,” *Working Paper*.

- HOLMSTRÖM, B. (1979): “Moral Hazard and Observability,” *Bell Journal of Economics*, 74–91.
- ICHIHASHI, S. AND A. SMOLIN (2025): “Buyer-Optimal Algorithmic Recommendations,” *Working paper*.
- INOSTROZA, N. AND A. PAVAN (2025): “Adversarial Coordination and Public Information Design,” *Theoretical Economics*, 20, 763–813.
- JEHIEL, P. (2015): “On Transparency in Organizations,” *The Review of Economic Studies*, 82, 736–761.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KIRBY, A. J. (1988): “Trade Associations as Information Exchange Mechanisms,” *The RAND Journal of Economics*, 138–146.
- KOLOTILIN, A. (2012): “Optimal Information Disclosure: Quantity vs. Quality,” *Working paper*.
- (2018): “Optimal Information Disclosure: A Linear Programming Approach,” *Theoretical Economics*, 13, 607–635.
- KOLOTILIN, A., R. CORRAO, AND A. WOLITZKY (2025): “Persuasion and Matching: Optimal Productive Transport,” *Journal of Political Economy*, 133, 1334–1381.
- LIN, X. AND C. LIU (2024): “Credible Persuasion,” *Journal of Political Economy*, 132, 2228–2273.
- MATHEVET, L., J. PEREGO, AND I. TANEVA (2020): “On Information Design in Games,” *Journal of Political Economy*, 128, 1370–1404.
- MIRPLEES, J. A. (1999): “The Theory of Moral Hazard and Unobservable Behaviour: Part I,” *Review of Economic Studies*, 66, 3–21.
- MIYASHITA, M. AND T. UI (2023): “LQG Information Design,” *Working paper*.

- MORRIS, S., D. OYAMA, AND S. TAKAHASHI (2024): “Implementation via Information Design in Binary-Action Supermodular Games,” *Econometrica*, 92, 775–813.
- MYERSON, R. (1982): “Optimal Coordination Mechanism in Generalized Principal-Agent Problems,” *Journal of Mathematical Economics*, 10, 67–81.
- NEYMAN, A. (1997): “Correlated Equilibrium and Potential Games,” *International Journal of Game Theory*, 26, 223–227.
- ORTNER, J., T. SUGAYA, AND A. WOLITZKY (2024): “Mediated Collusion,” *Journal of Political Economy*, 132, 1247–1289.
- RAYO, L. AND I. SEGAL (2010): “Optimal Information Disclosure,” *Journal of Political Economy*, 118, 949–987.
- ROMANYUK, G. AND A. SMOLIN (2019): “Cream Skimming and Information Design in Matching Markets,” *American Economic Journal: Microeconomics*, 11, 250–276.
- ROSEN, J. (1965): “Existence and Uniqueness of Equilibrium Points for Concave N-Person Games,” *Econometrica*, 33, 520–534.
- RUDIN, W. (1976): *Principles of Mathematical Analysis*, New York: McGraw-Hill.
- SALAMANCA, A. (2021): “The Value of Mediated Communication,” *Journal of Economic Theory*, 192, 105191.
- TAMURA, W. (2012): “A Theory of Multidimensional Information Disclosure,” *Working paper*.
- (2018): “Bayesian Persuasion with Quadratic Preferences,” *Working paper*.
- TANEVA, I. (2019): “Information Design,” *American Economic Journal: Microeconomics*, 11, 151–85.
- UI, T. (2020): “LQG Information Design,” *Working paper*.
- VIVES, X. (1984): “Duopoly Information Equilibrium: Cournot and Bertrand,” *Journal of Economic Theory*, 34, 71–94.
- (1990): “Trade Association Disclosure Rules, Incentives to Share Information, and Welfare,” *RAND Journal of Economics*, 409–430.

——— (1999): *Oligopoly Pricing: Old Ideas and New Tools*, MIT press.

VOHRA, R. (2011): *Mechanism Design: a Linear Programming Approach*, vol. 47, Cambridge University Press.

A Appendix

A.1 Formalism Omitted in Section 2

Derivation of First-Order Condition (4) For $\varepsilon > 0$, define

$$\begin{aligned}\phi_{i\varepsilon}^+(a_i, a_{-i}, \omega) &\triangleq \frac{u_i(a_i + \varepsilon, a_{-i}, \omega) - u_i(a_i, a_{-i}, \omega)}{\varepsilon}, \\ \phi_{i\varepsilon}^-(a_i, a_{-i}, \omega) &\triangleq \frac{u_i(a_i, a_{-i}, \omega) - u_i(a_i - \varepsilon, a_{-i}, \omega)}{\varepsilon}.\end{aligned}$$

Definition 6. (Admissible Equilibria) *An equilibrium is admissible, if for any player i and equilibrium belief $\mu_i \in \Delta(A_{-i} \times \Omega)$, there exists a best response $a_i^* \in A_i$, $\Delta > 0$, and $\psi : A_{-i} \times \Omega \rightarrow \mathbb{R}$ such that $\int_{A_{-i} \times \Omega} \psi(a_{-i}, \omega) d\mu_i < +\infty$ and for all $\varepsilon \in (0, \Delta)$, $(a_{-i}, \omega) \in A_{-i} \times \Omega$, $|\phi_{i\varepsilon}^+(a_i^*, a_{-i}, \omega)| \leq \psi(a_{-i}, \omega)$ and $|\phi_{i\varepsilon}^-(a_i^*, a_{-i}, \omega)| \leq \psi(a_{-i}, \omega)$.*

This admissibility effectively requires players' payoff differences to be locally well behaved at the best responses. It ensures the interchangeability of operators in (4) and is satisfied in all equilibria that we characterize. (If A and Ω were compact, then admissibility would be trivially satisfied whenever each $u_i(a, \omega)$ was continuously differentiable in a_i .)

Lemma 3. (First-Order Approach) *Under Assumption 1, for any player i and admissible equilibrium belief $\mu_i \in \Delta(A_{-i} \times \Omega)$, a_i is player i 's best response if and only if condition (4) is satisfied.*

Proof. For any i and admissible equilibrium belief μ_i , a_i^* satisfies the following condition:

$$\int_{A_{-i} \times \Omega} u_i(a_i^*, a_{-i}, \omega) d\mu_i \geq \int_{A_{-i} \times \Omega} u_i(a_i, a_{-i}, \omega) d\mu_i, \quad \forall a_i \in A_i. \quad (56)$$

Under Assumption 1, condition (56) is satisfied if and only if for any $\varepsilon > 0$, the following holds:

$$\int_{A_{-i} \times \Omega} \phi_{i\varepsilon}^+(a_i^*, a_{-i}, \omega) d\mu_i \leq 0, \quad (57)$$

$$\int_{A_{-i} \times \Omega} \phi_{i\varepsilon}^-(a_i^*, a_{-i}, \omega) d\mu_i \geq 0. \quad (58)$$

Moreover, $\phi_{i\varepsilon}^+$ is decreasing in ε and $\phi_{i\varepsilon}^-$ is increasing in ε ; thus, conditions (57), (58) are equivalent to

$$\begin{aligned} 0 &\geq \lim_{\varepsilon \rightarrow 0} \int_{A_{-i} \times \Omega} \phi_{i\varepsilon}^+(a_i^*, a_{-i}, \omega) d\mu_i = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \int_{A_{-i} \times \Omega} u_i(a_i^* + \varepsilon, a_{-i}, \omega) - u_i(a_i^*, a_{-i}, \omega) d\mu_i \\ &= \frac{\partial^+}{\partial a_i} \int_{A_{-i} \times \Omega} u_i(a_i, a_{-i}, \omega) d\mu_i \Big|_{a_i=a_i^*}, \\ 0 &\leq \lim_{\varepsilon \rightarrow 0} \int_{A_{-i} \times \Omega} \phi_{i\varepsilon}^-(a_i^*, a_{-i}, \omega) d\mu_i = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \int_{A_{-i} \times \Omega} u_i(a_i^*, a_{-i}, \omega) - u_i(a_i^* - \varepsilon, a_{-i}, \omega) d\mu_i, \\ &= \frac{\partial^-}{\partial a_i} \int_{A_{-i} \times \Omega} u_i(a_i, a_{-i}, \omega) d\mu_i \Big|_{a_i=a_i^*}. \end{aligned}$$

By Lebesgue's dominated convergence theorem (Theorem 11.32, Rudin (1976)), because belief μ_i arises in an admissible equilibrium, the following holds:

$$\begin{aligned} \lim_{\varepsilon \rightarrow 0} \int_{A_{-i} \times \Omega} \phi_{i\varepsilon}^+(a_i^*, a_{-i}, \omega) d\mu_i &= \int_{A_{-i} \times \Omega} \lim_{\varepsilon \rightarrow 0} \phi_{i\varepsilon}^+(a_i^*, a_{-i}, \omega) d\mu_i = \int_{A_{-i} \times \Omega} \dot{u}_i(a_i^*, a_{-i}, \omega) d\mu_i, \\ \lim_{\varepsilon \rightarrow 0} \int_{A_{-i} \times \Omega} \phi_{i\varepsilon}^-(a_i^*, a_{-i}, \omega) d\mu_i &= \int_{A_{-i} \times \Omega} \lim_{\varepsilon \rightarrow 0} \phi_{i\varepsilon}^-(a_i^*, a_{-i}, \omega) d\mu_i = \int_{A_{-i} \times \Omega} \dot{u}_i(a_i^*, a_{-i}, \omega) d\mu_i, \end{aligned}$$

and thus

$$\frac{\partial}{\partial a_i} \int_{A_{-i} \times \Omega} u_i(a_i, a_{-i}, \omega) d\mu_i \Big|_{a_i=a_i^*} = \int_{A_{-i} \times \Omega} \dot{u}_i(a_i^*, a_{-i}, \omega) d\mu_i = 0.$$

□

A.2 Formalism Omitted in Section 3

Proof of Theorem 1 The first step is

Lemma 4. (Weak Duality) $V^P \leq V^D$.

Proof. Take any dual variables (λ, γ) that satisfy the constraints of the dual problem (8). Take any measure π that satisfies the constraints of the primal problem (5). Integrating both sides of the dual constraints over $a \in A$ and $\omega \in \Omega$ against measure π yields:

$$\int_{A \times \Omega} v(a, \omega) d\pi \leq \int_{A \times \Omega} \sum_{i=1}^N \lambda_i(a_i) \dot{u}_i(a, \omega) d\pi + \int_{A \times \Omega} \gamma(\omega) d\pi = \int_{\Omega} \gamma(\omega) d\mu_0, \quad (59)$$

where the equality follows because π satisfies the primal constraints. The left-hand side of (59) is the value of the primal problem given measure π . The right-hand side of (59) is the value of the dual problem given dual variables (λ, γ) . As the inequality (59) holds for any allowed values of primal measure and dual variables, it also holds at the respective maximization and minimization limits. \square

Continuing with the proof of Theorem 1, take any primal measure π implementable by information, i.e., that satisfies the constraints of the primal problem (5). If it is implementable by incentives, then there exist dual variables λ that implement this measure in the dual problem (10), and

$$\begin{aligned} V^D &= \inf_{\lambda' \in \times_i \mathcal{M}(A_i)} \mathbb{E}_{\mu_0} \left[\sup_{a \in A} w(a, \omega, \lambda') \right] \\ &\leq \mathbb{E}_{\pi} [w(a, \omega, \lambda)] \\ &= \int_{A \times \Omega} v(a, \omega) d\pi - \int_{A \times \Omega} \sum_{i=1}^N \lambda_i(a_i) \dot{u}_i(a, \omega) d\pi \\ &= \int_{A \times \Omega} v(a, \omega) d\pi \leq V^P, \end{aligned}$$

where the first inequality follows from the implementability of π in the dual problem and the last three steps follow from the feasibility of π in the primal problem.

Furthermore, by Lemma 4, $V^D \geq V^P$. Combining the two inequalities, we obtain

$$V^D = \int_{A \times \Omega} v(a, \omega) d\pi = V^P,$$

which proves the optimality of measure π .

Proof of Proposition 1 For any optimal measure π , the following holds:

$$\begin{aligned} \mathbb{E}_{\pi} [v(a, \omega)] &= V^P, \\ \int_{A'_i \times A_{-i} \times \Omega} \dot{u}_i(a, \omega) d\pi &= 0 \quad \forall i = 1, \dots, N, \text{ measurable } A'_i \subseteq A_i. \end{aligned}$$

As such, if the dual agent is offered contract λ and plays according to an optimal measure then his expected payoff is

$$\begin{aligned}\mathbb{E}_\pi[w(a, \omega, \lambda)] &= \mathbb{E}_\pi[v(a, \omega) - \sum_{i=1}^N \lambda_i(a_i) \dot{u}_i(a, \omega)] \\ &= \mathbb{E}_\pi[v(a, \omega)] - \mathbb{E}_\pi[\sum_{i=1}^N \lambda_i(a_i) \dot{u}_i(a, \omega)] = V^P.\end{aligned}$$

That is, the dual agent obtains the same expected payoff of V^P by playing according to an optimal measure, whether π or π' , irrespective of a contract.

Moreover, if λ certifies the optimality of measure π , then playing according to π is a best response of the dual agent to contract λ . However, since, as showed above, π' delivers the same expected dual payoff as π , it follows that π' is also a best response of the dual agent to contract λ .

Finally, as an optimal measure, π' satisfies the constraints of the primal problem. Hence, π' is implementable by information and by incentives; by definition, λ certifies the optimality of π' .

Proof of Theorem 2 First, consider the case in which Ω and all A_i are compact (separable, metric) spaces. By weak duality, it suffices to establish that $V^P = V^D$. For notational convenience, denote $X_0 = \Omega$, $X_i = A_i$, and $X = \Omega \times A$.

Lemma 5. (Strong Duality. Compact Spaces) *If X is compact, then $V^P = V^D$.*

Proof. Let $E = C(X)$ be the set of continuous functions on X , associated with the sup norm. Because X is compact, $C(X)$ coincides with the set of bounded continuous functions on X . Its dual is $E^* = M(X)$, the set of all Radon measures on X .

For each $f \in E$, define

$$\theta(f) = \begin{cases} 0, & \text{if } f(x) \geq v(x), \forall x, \\ +\infty, & \text{otherwise.} \end{cases}$$

Let $F \subseteq E$ be such that $f \in F$ if and only if there exists γ with $f(x) = \gamma(x_0) + \sum_{i \neq 0} \lambda_i(x_i) \dot{u}_i(x)$ for all x . Define

$$\xi(f) = \inf_{\gamma} \left\{ \int_{X_0} \gamma(x_0) d\mu_0 \mid f(x) = \gamma(x_0) + \sum_{i \neq 0} \lambda_i(x_i) \dot{u}_i(x), \forall x \right\},$$

if $f \in F$, and let $\xi(f) = +\infty$ if $f \in E \setminus F$. Thus, if $f \in F$, then for any $\varepsilon > 0$, there exists γ such that $f(x) = \gamma(x_0) + \sum_{i \neq 0} \lambda_i(x_i) \dot{u}_i(x)$ and $\xi(f) > \int_{X_0} \gamma(x_0) d\mu_0 - \varepsilon$.

Observe that both θ, ξ are convex functions. Convexity of θ is immediate. For ξ , take $f, f' \in E$ such that $\xi(f), \xi(f') < +\infty$. In particular, for any $\varepsilon > 0$, there exists $\gamma(\varepsilon)$ and $\gamma'(\varepsilon)$, such that

$$\begin{aligned} f(x) &= \gamma(x_0; \varepsilon) + \sum_{i \neq 0} \lambda_i(x_i; \varepsilon) \dot{u}_i(x), \\ f'(x) &= \gamma'(x_0; \varepsilon) + \sum_{i \neq 0} \lambda'_i(x_i; \varepsilon) \dot{u}_i(x), \end{aligned}$$

and

$$\begin{aligned} \int_{X_0} \gamma(x_0; \varepsilon) d\mu_0 &\leq \xi(f) + \varepsilon, \\ \int_{X_0} \gamma'(x_0; \varepsilon) d\mu_0 &\leq \xi(f') + \varepsilon. \end{aligned}$$

Now, let

$$f'' = \alpha f + (1 - \alpha) f'$$

for some $\alpha \in (0, 1)$. Then,

$$f''(x) = \alpha \gamma(x_0; \varepsilon) + (1 - \alpha) \gamma'(x_0; \varepsilon) + \sum_{i \neq 0} (\alpha \lambda_i(x_i; \varepsilon) + (1 - \alpha) \lambda'_i(x_i; \varepsilon)) \dot{u}_i(x)$$

and thus,

$$\xi(f'') \leq \int_{X_0} (\alpha \gamma(x_0; \varepsilon) + (1 - \alpha) \gamma'(x_0; \varepsilon)) d\mu_0 \leq \alpha \xi(f) + (1 - \alpha) \xi(f') + \varepsilon,$$

implying $\xi(f'') \leq \alpha \xi(f) + (1 - \alpha) \xi(f')$, and hence ξ is convex.

Furthermore, if we let $\bar{v} = \max_x v(x)$ and $f(x) = \bar{v}$ for all x , then we have $\theta(\bar{f}), \xi(\bar{f}) < \infty$ and θ is continuous at \bar{f} . Thus, we can apply Fenchel-Rockafeller duality:

$$\inf_{f \in E} (\theta(f) + \xi(f)) = \max_{\pi \in E^*} (-\theta^*(-\pi) - \xi^*(\pi)), \quad (60)$$

where

$$\begin{aligned} -\theta^*(-\pi) &= -\sup_{f \in E} \left(\int f d(-\pi) - \theta(f) \right), \\ -\xi^*(\pi) &= -\sup_{f \in E} \left(\int f d\pi - \xi(f) \right). \end{aligned}$$

The left-hand side of (60) is weakly higher than V^D . Thus, the proof completes if the right-hand side equals V^P .

We have the following:

$$\begin{aligned} -\theta^*(-\pi) &= -\sup_{f \in E} \left(\int f d(-\pi) - \theta(f) \right), \\ &= \inf_{f \in E, f \geq v} \int f d\pi, \\ &= \begin{cases} \int v d\pi & \text{if } \pi \in M_+(X), \\ -\infty & \text{if } \pi \notin M_+(X). \end{cases} \end{aligned}$$

For $-\xi^*(\pi)$, we have

$$-\xi^*(\pi) = -\sup_{f \in E} \left(\int f d\pi - \xi(f) \right) = \inf_{f \in F} \left(-\int_X f d\pi + \xi(f) \right).$$

We first examine $-\int_X f d\pi + \xi(f)$ for any fixed $f \in F$. Recall that for any $\varepsilon > 0$, there exists γ such that $f(x) = \gamma(x_0) + \sum_{i \neq 0} \lambda_i(x_i) \dot{u}_i(x)$ and $\xi(f) > \int_{X_0} \gamma(x_0) d\mu_0 - \varepsilon$. Thus,

$$\begin{aligned} -\int_X f d\pi + \xi(f) &> -\int_X (\gamma + \sum_{i \neq 0} \lambda_i \dot{u}_i) d\pi + \int_{X_0} \gamma d\mu_0 - \varepsilon \\ &= \int_{X_0} \gamma d(\mu_0 - \pi_0) - \sum_{i \neq 0} \int_X \lambda_i \dot{u}_i d\pi - \varepsilon. \end{aligned}$$

Therefore, if π is such that $\pi_0 = \mu_0$ and $\pi_{-0} \equiv 0$, we have $-\int_X f d\pi + \xi(f) > -\varepsilon$ for any $\varepsilon > 0$, and hence $-\int_X f d\pi + \xi(f) = 0$. On the other hand, if π does not satisfy either $\pi_0 = \mu_0$ or $\pi_{-0} \equiv 0$, then for any $M \in \mathbb{R}$, it is always possible to find some $f \in F$ with $-\int_X f d\pi + \xi(f) > M$. In conclusion,

$$-\xi^*(\pi) = \begin{cases} 0, & \text{if } \pi_0 = \mu_0, \pi_{-0} \equiv 0, \\ -\infty, & \text{otherwise.} \end{cases}$$

Therefore the right-hand side of (60) is

$$\max_{\pi \in E^*} (-\theta^*(-\pi) - \xi^*(\pi)) = \max_{\pi \in M_+(X): \pi_0 = \mu_0, \pi_{-0} \equiv 0} \int v d\pi = V^P.$$

□

We use Lemma 5 to establish the duality in the non-compact case.

Note: For the rest of the proof we assume that Condition 1 holds and normalize the designer-preferred equilibrium payoff under full disclosure to 0 in all states.

The k -th primal problem is

$$\begin{aligned} & \max_{\pi \in \Delta(X)} \int_X v(x) d\pi \\ & \text{s.t.} \quad \int_{X'_0 \times X_{-0}} d\pi = \mu_0(X'_0 \cap X_0^k), \quad \forall X'_0 \subseteq X_0, \\ & \quad \int_{X'_i \times X_{-i}} \dot{u}_i(x) d\pi = 0, \quad \forall i \neq 0, \quad \forall X'_i \subseteq X_i. \end{aligned}$$

Denote by P_k and D_k the primal and dual values in the k -th problem, respectively. By Lemma 5, $P_k = D_k$.

Lemma 6. (Primal Value Monotonicity) *For all k , $P_k \leq P_{k+1} \leq P^*$.*

Proof. In the $(k+1)$ -th problem, one of the feasible information structures for the designer is to fully disclose x_0 if $x_0 \in X_0^{k+1} \setminus X_0^k$ and to provide an information structure optimal in k -th problem if $x_0 \in X_0^k$. This strategy delivers precisely P_k to the designer: whenever $x_0 \in X_0^{k+1} \setminus X_0^k$ is realized, this is fully observed by the agents and they play a full-disclosure Nash equilibrium, yielding payoff 0 for the designer (by normalization). Thus, $P_k \leq P_{k+1}$.

The fact that $P_k \leq P^*$ can be established analogously. □

Now, consider the dual version of the k -th problem. Fix any $\varepsilon > 0$ and k . In the k -th problem, there exists $\lambda^k \in C(X_{-0})$ such that

$$\left| D_k - \int_{X_0^k} \max_{x_{-0}} \{v(x) - \sum_{i \neq 0} \lambda_i^k(x_i) \dot{u}_i(x)\} d\mu_0 \right| < \varepsilon.$$

Fix any $\delta > 0$. By [Lemma 5](#) and [Lemma 6](#), there exists $k(\delta)$ such that, for any $k > k(\delta)$,

$$\left| \int_{X_0} \max_{x_{-0}} \{v(x) - \sum_{i \neq 0} \lambda_i^k(x_i) \dot{u}_i(x)\} d\mu_0 - \int_{X_0^k} \max_{x_{-0}} \{v(x) - \sum_{i \neq 0} \lambda_i^k(x_i) \dot{u}_i(x)\} d\mu_0 \right| < \delta,$$

because μ on X_0^k converges to μ on X_0 in a weak-* topology, and the term inside the integral is continuous in x_0 .

In the original dual problem, we have

$$D^* \leq \int_{X_0} \max_{x_{-0}} \{v(x) - \sum_{i \neq 0} \lambda_i^k(x_i) \dot{u}_i(x)\} d\mu_0.$$

Therefore,

$$P^* \leq D^* \leq \int_{X_0} \max_{x_{-0}} \{v(x) - \sum_{i \neq 0} \lambda_i^k(x_i) \dot{u}_i(x)\} d\mu_0 < D_k + \varepsilon + \delta = P_k + \varepsilon + \delta \leq P^* + \varepsilon + \delta.$$

Because $\varepsilon > 0$ and $\delta > 0$ are arbitrary, we conclude that $P^* = D^*$.

A.3 Formalism Omitted in [Section 5](#)

Game Normalization The game normalization is achieved by defining $\omega' = \omega - \mathbb{E}[\omega]$, $a'_i = a_i + \frac{\mathbb{E}[\omega] + b_1}{q^o + q^c}$, $b = h\mathbb{E}[\omega] + b_2 - 2(p^o + p^c) \frac{\mathbb{E}[\omega] + b_1}{q^o + q^c}$, and ignoring the strategically irrelevant additive terms, i.e., those in u_i that do not depend on a_i and those in v that do not depend on a .

Proof of [Claim 1](#) The result is an immediate corollary of [Claim 2](#), obtained by setting $x_{0i} \equiv x_{0j}$ for all i, j .

Proof of [Claim 2](#) Clearly, u_i and v are continuous, and the full-disclosure value is bounded. To establish the third property of [Condition 1](#), we show that if the state is bounded, then the players' actions can be bounded too. We can rewrite player i 's obedience condition as

$$\begin{aligned} \mathbb{E}[x_{0i} + X_{-i} - 2x_i \mid x_i] &= 0, \\ \mathbb{E}[x_{0i} + X_{-i} \mid x_i] &= 2x_i, \end{aligned} \tag{61}$$

where, for notational convenience, we denote $X_{0i} = \Omega_i$, $X_i = A_i$, $X = \Omega \times A$, and $X_{-i} = \beta \sum_{j \neq i} x_j$, with $\beta < \frac{2}{N-1}$.

Assume x_{0i} is supported on a compact interval, say $[\underline{X}_0, \overline{X}_0]$, for all i . We show there exists $Z > 0$ such that any feasible (i.e., obedient) π does not support any $|x_i| > Z$ for any i . Let π_i denote π 's marginal for x_i , and π_{-0} denote π 's marginal for x_{-0} .

Let us take the expectation of (61) for each i , with respect to $x_i > Z$. Regarding the left-hand side,

$$\int_{x_i > Z} \mathbb{E}[x_{0i} + X_{-i} \mid x_i] d\pi_i(x_i) = \int_{x_i > Z} \mathbb{E}[x_{0i} \mid x_i] d\pi_i(x_i) + \int_{x_i > Z} \int_{X_{-i} \in \mathbb{R}} X_{-i} d\pi_{-0}(x_{-0}),$$

and hence, summing this expression across all i , we obtain

$$\sum_{i=1}^N \int_{x_i > Z} \mathbb{E}[x_{0i} \mid x_i] d\pi_i(x_i) + \sum_{i=1}^N \int_{x_i > Z} \int_{X_{-i} \in \mathbb{R}} X_{-i} d\pi_{-0}(x_{-0}).$$

Regarding the right-hand side,

$$2 \int_{x_i > Z} x_i d\pi_i(x_i) \geq (2 - \alpha) \int_{x_i > Z} Z d\pi_i(x_i) + \alpha \int_{x_i > Z} \int_{X_{-i} \in \mathbb{R}} x_i d\pi_{-0}(x_{-0}),$$

for any $\alpha \in (0, 2)$. Summing this expression across all i , we obtain

$$(2 - \alpha) \sum_{i=1}^N \int_{x_i > Z} Z d\pi_i(x_i) + \alpha \sum_{i=1}^N \int_{x_i > Z} \int_{X_{-i} \in \mathbb{R}} x_i d\pi_{-0}(x_{-0}).$$

Lemma 7. *For some $\alpha \in (0, 2)$,*

$$\int_{x_{-0}} \sum_{i=1}^N (\alpha x_i - X_{-i}) 1_{\{x_i > Z\}} d\pi_{-0}(x_{-0}) \geq 0.$$

Proof. Let us examine each sub-domain of the integration. Fix any nonempty subset $K \subseteq \{1, \dots, N\}$

with $k = |K|$. On the region of x_{-0} where $x_i > Z$ if $i \in K$ (and $x_i \leq Z$ otherwise), we have

$$\begin{aligned} \sum_{i=1}^N (\alpha x_i - X_{-i}) 1_{\{x_i > Z\}} &= \sum_{i \in K} (\alpha x_i - X_{-i}) = \sum_{i \in K} \alpha x_i - \sum_{i \in K} \sum_{j \neq i} \beta x_j \\ &= \alpha \sum_{i \in K} x_i - \beta \left(\sum_{i \in K} \sum_{\substack{j \neq i \\ j \in K}} x_j + \sum_{i \in K} \sum_{j \notin K} x_j \right) \geq \alpha k Z - \beta [k(k-1)Z + k(N-k)Z] = kZ(\alpha - \beta(N-1)), \end{aligned}$$

which can be made non-negative by choosing $\alpha \in (\beta(N-1), 2)$. This is possible since $\beta(N-1) < 2$. \square

The lemma implies that

$$\sum_{i=1}^N \int_{x_i > Z} (\mathbb{E}[x_{0i}|x_i] - (2 - \alpha)Z) d\pi_i(x_i) > 0.$$

However, given that $x_{0i} \leq \bar{X}_0$, by choosing $Z > \frac{\bar{X}_0}{2-\alpha}$, this inequality would be non-positive, which is a contradiction. Therefore, no feasible π can support $x_i > Z$.

Proof of Theorem 3, continued. If $k^* \notin [0, N]$, then either (i) $q^o p^c < p^o q^c$ and $p^o + p^c < h(q^o + q^c)$, or $q^o p^c > p^o q^c$ and $p^o < h q^o$, in which case no disclosure is optimal, or (ii) $q^o p^c < p^o q^c$ and $p^o + p^c > h(q^o + q^c)$, or $q^o p^c > p^o q^c$ and $(h q^o - p^o)(q^o + q^c) > 2(p^c q^o - p^o q^c)$.

Indeed, if neither (i) or (ii) holds then it must be the case that $q^o p^c > p^o q^c$, $p^o \geq h q^o$, and $(h q^o - p^o)(q^o + q^c) \leq 2(p^c q^o - p^o q^c)$. But in this case, because of our assumptions $q^o < 0$, $q^o(h q^o - p^o) \geq 0$ and therefore,

$$\frac{k^*}{N} = \frac{q^o(h q^o - p^o)}{(2(p^c q^o - p^o q^c) - (h q^o - p^o)(q^o + q^c)) + q^o(h q^o - p^o)} \in [0, 1].$$

Consequently, if $k^* \notin [0, N]$, then either (i) or (ii) holds.

Case (i). If (i) holds, then no disclosure is optimal, which can be certified by an affine contract with $(x, x_0) = (-h, -\frac{b}{q^o + q^c})$. Under this contract, the dual payoff function is

$$w(a, \omega) = x_0 \omega + (p^o - h q^o) \check{a} + (p^c - h q^c) \bar{a}^2 = x_0 \omega + \frac{1}{N^2} a^T M ((p^o - h q^o)N + p^c - h q^c, p^c - h q^c) a,$$

where $M(x, y)$ denotes the $N \times N$ matrix whose diagonal entries equal x and off-diagonal entries equal y . By Theorem 1, no disclosure is optimal if $a = (0, \dots, 0)$ maximizes $w(a, \omega)$ for all ω . This happens

if $M((p^o - hq^o)N + p^c - hq^c, p^c - hq^c)$ is negative semidefinite (NSD). Note that matrix $M(x, y)$ has two distinct eigenvalues: $x + (N - 1)y$ and $x - y$. Therefore, $M(x, y)$ is NSD if $x + (N - 1)y \leq 0$ and $x - y \leq 0$ which for $M((p^o - hq^o)N + p^c - hq^c, p^c - hq^c)$ corresponds to:

$$\begin{aligned} p^o &\leq hq^o, \\ p^o + p^c &\leq h(q^o + q^c). \end{aligned}$$

Subcase 1: $hq^c \geq p^c$. In this case, NSD is equivalent to $p^o - hq^o \leq 0$. In case $q^o p^c > p^o q^c$ and $p^o < hq^o$, this is satisfied. In case $q^o p^c < p^o q^c$ and $p^o + p^c < h(q^o + q^c)$, note that, by the initial assumption, we have $p^c = -p^o - \eta$ and $q^c = q^o - \varepsilon$ for some η and $\varepsilon > 0$, where $p^o + p^c < h(q^o + q^c)$ implies $h\varepsilon < \eta$ and $q^o p^c < p^o q^c$ implies $\eta/\varepsilon < p^o/q^o$, which jointly imply $p^o/(hq^o) > 1$; equivalently, $p^o < hq^o$.

Subcase 2: $hq^c < p^c$. In this case, NSD is equivalent to $p^o - hq^o + p^c - hq^c \leq 0$. In case $q^o p^c < p^o q^c$ and $p^o + p^c < h(q^o + q^c)$, this is satisfied. In case $q^o p^c > p^o q^c$ and $p^o < hq^o$, again note that $p^c = -p^o - \eta$ and $q^c = q^o - \varepsilon$ for some η and $\varepsilon > 0$. Because $q^o p^c > p^o q^c$, we have $-q^o \eta > -p^o \varepsilon$, which implies $\eta/\varepsilon > p^o/q^o > h$ because $p^o < hq^o < 0$. Thus, $\eta > h\varepsilon$; equivalently, $p^o + p^c < h(q^o + q^c)$.

Case (ii). If (ii) holds, then full disclosure is optimal, which can be certified by an affine contract with $(x, x_0) = (h - \frac{2(p^o + p^c)}{q^o + q^c}, -\frac{b}{q^o + q^c})$. Under this contract, the dual payoff function is

$$\begin{aligned} w(a, \omega) &= x_0 \omega + \left(2h - \frac{2(p^o + p^c)}{q^o + q^c}\right) \omega \bar{a} + \left(p^o + q(h - \frac{2(p^o + p^c)}{q^o + q^c})\right) \check{a} + \left(p^c + q^c(h - \frac{2(p^o + p^c)}{q^o + q^c})\right) \bar{a}^2 \\ &= x_0 \omega + \frac{1}{N^2} \left(\left(2h - \frac{2(p^o + p^c)}{q^o + q^c}\right) N \omega 1_N^T a + a^T M(x, y) a \right), \end{aligned}$$

where $1_N = (1, \dots, 1)^T$ is an N -dimensional vector of ones, $x = N(p^o + q^o(2 - h - \frac{2(p^o + p^c)}{q^o + q^c})) + p^c + q^c(h - \frac{2(p^o + p^c)}{q^o + q^c})$ and $y = p^c + q^c(h - \frac{2(p^o + p^c)}{q^o + q^c})$.

By [Theorem 1](#), full disclosure is optimal if $a = (-\frac{\omega}{q^o + q^c}, \dots, -\frac{\omega}{q^o + q^c})$ maximizes $w(a, \omega)$ for all ω . If $M(x, y)$ is negative definite (ND), then the maximizing a can be found via F.O.C.:

$$\left(2h - \frac{2(p^o + p^c)}{q^o + q^c}\right) N \omega 1_N + 2M(x, y)a = 0,$$

implying $a = (-\frac{\omega}{q^o + q^c}, \dots, -\frac{\omega}{q^o + q^c})$. Therefore, it remains to establish that $M(x, y)$ is ND.

$M(x, y)$ is ND if $x + (N - 1)y \leq 0$ and $x - y \leq 0$ which corresponds to

$$\begin{aligned} p^o + q^o \left(h - \frac{2(p^o + p^c)}{q^o + q^c} \right) &< 0, \\ p^o + q^o \left(h - \frac{2(p^o + p^c)}{q^o + q^c} \right) + p^c + q^c \left(h - \frac{2(p^o + p^c)}{q^o + q^c} \right) &< 0. \end{aligned}$$

Subcase 1: $p^c + q^c(h - \frac{2(p^o + p^c)}{q^o + q^c}) < 0$. It suffices to show $p^o + q^o(h - \frac{2(p^o + p^c)}{q^o + q^c}) < 0$; equivalently, $(hq^o - p^o)(q^o + q^c) > 2(p^c q^o - p^o q^c)$. In case $p^c q^o > p^o q^c$ and $(hq^o - p^o)(q^o + q^c) > 2(p^c q^o - p^o q^c)$, this is clearly satisfied. In case $q^o p^c < p^o q^c$ and $p^o + p^c > h(q^o + q^c)$, noting that $p^c = -p^o - \eta$ and $q^c = -q^o - \varepsilon$ for some η and $\varepsilon > 0$, we have $-q^o \eta < -p^o \varepsilon$ and $h\varepsilon > \eta$. Thus,

$$\begin{aligned} p^o + q^o \left(h - \frac{2(p^o + p^c)}{q^o + q^c} \right) &= p^o + q^o \left(h - \frac{2\eta}{\varepsilon} \right) \\ &= \frac{1}{\varepsilon}(p^o \varepsilon + q^o(h\varepsilon - 2\eta)) \\ &< \frac{q^o}{\varepsilon}(h\varepsilon - \eta) < 0. \end{aligned}$$

Subcase 2: $p^c + q^c(h - \frac{2(p^o + p^c)}{q^o + q^c}) > 0$. It suffices to show $p^o + q^o(h - \frac{2(p^o + p^c)}{q^o + q^c}) + p^c + q^c(h - \frac{2(p^o + p^c)}{q^o + q^c}) < 0$; equivalently, $h(q^o + q^c) - p^o - p^c < 0$. In case $q^o p^c - p^o q^c < 0$ and $p^o + p^c - h(q^o + q^c) > 0$, this is clearly satisfied. In case $q^o p^c - p^o q^c > 0$ and $(hq^o - p^o)(q^o + q^c) > 2(p^c q^o - p^o q^c)$, noting that $p^c = -p^o - \eta$ and $q^c = -q^o - \varepsilon$ for some η and $\varepsilon > 0$, we have $-q^o \eta > -p^o \varepsilon$ and $q^o(2\eta - h\varepsilon) > p^o \varepsilon$, implying (by summing them side-by-side) $q^o(\eta - h\varepsilon) > 0$; equivalently, $\eta - h\varepsilon < 0$. Thus,

$$h(q^o + q^c) - p^o - p^c = \eta - h\varepsilon < 0.$$

Proof of Theorem 4, continued. By concavity conditions $-q^o > 0$, and thus $\sigma^2 > 0$ if and only if

$$\begin{aligned} \beta(1 + q^c \beta + q^o \beta) &> 0, \\ \frac{(p^o - hq^o)}{2(p^c q^o - p^o q^c)} \left(1 + \frac{(q^c + q^o)(p^o - hq^o)}{2(p^c q^o - p^o q^c)} \right) &> 0, \\ \frac{(p^o - hq^o)(2(p^c q^o - p^o q^c) + (q^c + q^o)(p^o - hq^o))}{4(p^c q^o - p^o q^c)^2} &> 0. \end{aligned}$$

Because $q^o < 0$, setting $c_1 \triangleq q^o(hq^o - p^o)$ and $c_2 \triangleq q^o(2p^c - hq^c) - p^o q^c$, the inequality holds if and only if

$$q^o(hq^o - p^o)(2(p^c q^o - p^o q^c) + (q^c + q^o)(p^o - hq^o)) > 0,$$

$$c_1(c_2 - c_1) > 0.$$

Moreover,

$$\frac{k^*}{N} = \frac{q^o(hq^o - p^o)}{q^o(2p^c - hq^c) - p^o q^c} = \frac{c_1}{c_2}.$$

The result follows, because for any $c_1, c_2 \in \mathbb{R}$, $c_1/c_2 \in (0, 1)$ if and only if $c_1(c_2 - c_1) > 0$.

A.4 Formalism Omitted in Section 5.1

No Disclosure Under no disclosure, each player's action cannot depend on the state and is thus uniquely determined by condition (27).

Full Disclosure For any ω , the ensuing game admits a strictly concave potential $\Psi(a, \omega) = (\omega - \frac{\check{a}}{2})\check{a} - \sum_{i=1}^N a_i^2$ and thus has a unique equilibrium. Parameterize a symmetric linear strategy profile:

$$a_i(\omega) = k_0 + k_1 \omega.$$

The best-response condition (25) can be rewritten as

$$a_i(\omega) = -\frac{N-1}{2}k_0 + \frac{1-k_1(N-1)}{2}\omega,$$

which determines the equilibrium parameters at $k_0 = 0$ and $k_1 = \frac{1}{N+1}$.

A.5 Formalism Omitted in Section 6

Game Normalization Under any information structure, the expected marginal payoffs must equal to zero. Thus,

$$\mathbb{E}[\omega_i p^o + \check{\omega}_{-i} p^c + b_1 + q^o a_i + q^c \check{a}_{-i}] = \mathbb{E}[\omega](p^o + (N-1)p^c) + b_1 + \mathbb{E}[q^o a_i + q^c \check{a}_{-i}] = 0.$$

By symmetry, it follows that for all i ,

$$\mathbb{E}[a_i] = a_0 \triangleq -\frac{\mathbb{E}[\omega](p^o + (N-1)p^c) + b_1}{q^o + (N-1)q^c}.$$

In the normalized game, the expected actions equal zero. The normalization is achieved by defining $\omega' = (\omega - \mathbb{E}[\omega])/\sigma$, $a'_i = (a_i - a_0)/\sigma$, $b = (b_2 + (\tilde{p}^o + (N-1)\tilde{p}^c)\mathbb{E}[\omega] + 2(\tilde{q}^o + (N-1)\tilde{q}^c)a_0)/\sigma$, rescaling the payoffs by σ^2 , and ignoring the strategically irrelevant additive terms, i.e., those in u_i that do not depend on a_i and those in v that do not depend on a .

Proof of Lemma 1 Matrix $M(x, y)$ has two distinct eigenvalues: $x + (N-1)y$ and $x - y$. Therefore, it is negative semidefinite if and only if $x + (N-1)y \leq 0$ and $x - y \leq 0$; it is negative definite if and only if $x + (N-1)y < 0$ and $x - y < 0$. Therefore, $M(\bar{q}^o, \bar{q}^c) = M(\tilde{q}^o + xq^o, \tilde{q}^c + xq^c)$ is negative semidefinite if and only if

$$\begin{aligned} x(q^o + (N-1)q^c) &\leq -\tilde{q}^o - (N-1)\tilde{q}^c, \\ x(q^o - q^c) &\leq \tilde{q}^o - \tilde{q}^c. \end{aligned}$$

By concavity assumptions, $q^o + (N-1)q^c < 0$ and $q^o - q^c < 0$. The result follows.

Proof of Lemma 2 Note that $M(x, y)M(z, w) = M(xz + (N-1)yw, xw + yz + (N-2)yw)$ and $M(x, y)^{-1} = M(\frac{x+(N-2)y}{(x-y)(x+(N-1)y)}, -\frac{y}{(x-y)(x+(N-1)y)})$, if invertible. Thus,

$$\begin{aligned} a^*(\omega, x) &= -\frac{1}{2}M^{-1}(\bar{q}^o(x), \bar{q}^c(x))M(\bar{p}^o(x), \bar{p}^c(x))\omega \\ &= -\frac{1}{2} \frac{M(\bar{q}^o(x)\bar{p}^o(x) + (N-2)\bar{q}^c(x)\bar{p}^o(x) - (N-1)\bar{q}^c(x)\bar{p}^c(x), \bar{q}^o(x)\bar{p}^c(x) - \bar{q}^c(x)\bar{p}^o(x))}{(\bar{q}^o(x) - \bar{q}^c(x))(\bar{q}^o(x) + (N-1)\bar{q}^c(x))} \omega, \end{aligned}$$

and

$$\begin{aligned} W(x) &= -\frac{1}{4}\mathbb{E} \left[\omega^T M(\bar{p}^o(x), \bar{p}^c(x))M^{-1}(\bar{q}^o(x), \bar{q}^c(x))M(\bar{p}^o(x), \bar{p}^c(x))\omega \right] \\ &= -\frac{N}{4} \frac{Z(x)}{(\bar{q}^o(x) - \bar{q}^c(x))(\bar{q}^o(x) + (N-1)\bar{q}^c(x))}, \end{aligned}$$

where

$$Z(x) \triangleq \bar{q}^o(x) \left(\bar{p}^o(x)^2 + \bar{p}^c(x)^2 (N-1)(1 + (N-2)\rho) + 2(N-1)\rho \bar{p}^o(x) \bar{p}^c(x) \right) + \\ + \bar{q}^c(x) \left(\bar{p}^o(x)^2 (N-2 - \rho(N-1)) - \rho(N-1)^2 \bar{p}^c(x)^2 - 2(N-1)\bar{p}^o(x) \bar{p}^c(x) \right).$$

First, consider the case of $x \downarrow \underline{x}$. In this case, the term dividing $Z(x)$ is going to zero (from above). Thus, if $Z(\underline{x}) \neq 0$, then as $x \downarrow \underline{x}$, $W(x)$ diverges to $+\infty$.

(i) $\underline{x} = -\frac{\tilde{q}^o - \tilde{q}^c}{q^o - q^c}$. As $-\frac{\tilde{q}^o - \tilde{q}^c}{q^o - q^c} \leq -\frac{\tilde{q}^o - \tilde{q}^c}{q^o + (N-1)q^c}$, we have $\tilde{q}^o q^c \leq q^o \tilde{q}^c$, and by the genericity assumption, $\tilde{q}^o q^c < q^o \tilde{q}^c$. This implies that $\bar{q}^o = \bar{q}^c < 0$. Thus,

$$Z(\underline{x}) = \bar{q}^c (N-1)(1 - \rho)(\bar{p}^o - \bar{p}^c)^2 < 0.$$

(ii) $\underline{x} = -\frac{\tilde{q}^o + (N-1)\tilde{q}^c}{q^o + (N-1)q^c}$. As $-\frac{\tilde{q}^o - \tilde{q}^c}{q^o - q^c} \geq -\frac{\tilde{q}^o - \tilde{q}^c}{q^o + (N-1)q^c}$, we have $\tilde{q}^o q^c \geq q^o \tilde{q}^c$, and by the genericity assumption, $\tilde{q}^o q^c > q^o \tilde{q}^c$. This implies that $\bar{q}^o + (N-1)\bar{q}^c = 0$, and $\bar{q}^c > 0$. Thus,

$$Z(\underline{x}) = -\bar{q}^c (1 + \rho(N-1))(\bar{p}^o + (N-1)\bar{p}^c)^2 < 0.$$

Now, consider the case of $x \uparrow +\infty$. In this case, for the denominator of the dual payoff, we have

$$\frac{(\bar{q}^o - \bar{q}^c)(\bar{q}^o + (N-1)\bar{q}^c)}{x^2} \rightarrow (q^o - q^c)(q^o + (N-1)q^c) > 0,$$

while for the numerator we have

$$\frac{Z(x)}{x^3} \rightarrow q^o \left((p^o)^2 + (p^c)^2 (N-1)(1 + (N-2)\rho) + 2(N-1)\rho p^o p^c \right) \\ + q^c \left((p^o)^2 (N-2 - \rho(N-1)) - \rho(N-1)^2 (p^c)^2 - 2(N-1)p^o p^c \right). \quad (62)$$

It suffices to show that the right-hand side of (62) is strictly negative, because then the dual payoff is x times a strictly positive constant, which diverges to infinity as x grows.

First, note that

$$\begin{aligned}
& (p^o)^2 + (p^c)^2(N-1)(1+(N-2)\rho) + 2(N-1)\rho p^o p^c \\
&= (p^o - \rho(N-1)p^c)^2 + (p^c)^2((N-1)(1+(N-2)\rho) - \rho^2(N-1)^2) \\
&> (p^c)^2(N-1)(1+(N-2)\rho - \rho^2(N-1)) \\
&= (p^c)^2(N-1)(1+(N-1)\rho)(1-\rho) \geq 0.
\end{aligned}$$

(i) $q^c < 0$. Because $q^o < q^c$, the right-hand side of (62) is lower than

$$\begin{aligned}
& q^c \left((p^o)^2 + (p^c)^2(N-1)(1+(N-2)\rho) + 2(N-1)\rho p^o p^c \right) \\
&+ q^c \left((p^o)^2(N-2-\rho(N-1)) - \rho(N-1)^2(p^c)^2 - 2(N-1)p^o p^c \right) \\
&= q^c(p^o - p^c)^2(N-1)(1-\rho) < 0.
\end{aligned}$$

(ii) $q^c > 0$. Because $q^o + (N-1)q^c < 0$, the right-hand side of (62) is lower than

$$\begin{aligned}
& - (N-1)q^c \left((p^o)^2 + (p^c)^2(N-1)(1+(N-2)\rho) + 2(N-1)\rho p^o p^c \right) \\
&+ q^c \left((p^o)^2(N-2-\rho(N-1)) - \rho(N-1)^2(p^c)^2 - 2(N-1)p^o p^c \right) \\
&= -q^c(p^o + (N-1)p^c)^2(1+\rho(N-1)) < 0.
\end{aligned}$$

The result follows.

A.6 Formalism Omitted in Section 6.1

We can microfound the linear demand as being generated by a continuum of consumers that differ in their tastes. Each consumer has a type $\omega = (\omega_1, \omega_2) \in \mathbb{R}^2$ and decides how much of the firms' products to consume, $q = (q_1, q_2)$. The payoff of a type- ω consumer who consumes quantities q at prices a is

$$-a^T q + \frac{1}{2}(\omega - q)^T H^{-1}(\omega - q),$$

where H is an $N \times N$ negative semidefinite matrix:

$$H \triangleq \begin{pmatrix} -1 & \eta \\ \eta & -1 \end{pmatrix}.$$

For any price vector $a \in A$, the quantities demanded by a type- ω consumer are $q(a, \omega) = \omega + Ha$.

Consumer and producer surpluses can be written as

$$\begin{aligned} CS(a, \omega) &= a^T B_{CS} \omega - \frac{1}{2} a^T C_{CS} a, \\ PS(a, \omega) &= a^T (\omega + Ha) - c(\omega + Ha)^T (\omega + Ha) = -c\omega^T \omega + a^T B_{PS} \omega - \frac{1}{2} a^T C_{PS} a, \end{aligned}$$

where the payoff coefficient matrices are

$$\begin{aligned} B_{CS} &\triangleq -I, \quad C_{CS} \triangleq H, \\ B_{PS} &\triangleq I - 2cH, \quad C_{PS} \triangleq -2H + 2cH^2, \\ B &\triangleq \delta B_{CS} + (1 - \delta) B_{PS}, \quad C \triangleq \delta C_{CS} + (1 - \delta) C_{PS}. \end{aligned}$$

Price Control The designer's first-order condition is

$$B\omega - Ca = 0,$$

which results in the first-best responsiveness matrix $R^{FB} = C^{-1}B$, and thus,

$$\begin{aligned} r_o^{FB} &= \frac{2 + \delta(6\delta - 7) + 4c^2(1 - \delta)^2(1 - \eta^2) + 2c(1 - \delta)(3 - 5\delta - (1 - \delta)\eta^2)}{(1 - \eta^2)(2 - 3\delta + 2c(1 - \delta)(1 - \eta))(2 - 3\delta + 2c(1 - \delta)(1 + \eta))}, \\ r_c^{FB} &= \frac{\eta(2 + \delta(6\delta - 7) + 4c^2(1 - \delta)^2(1 - \eta^2) + 4c(1 - \delta)(1 - 2\delta))}{(1 - \eta^2)(2 - 3\delta + 2c(1 - \delta)(1 - \eta))(2 - 3\delta + 2c(1 - \delta)(1 + \eta))}. \end{aligned}$$

The threshold value $\bar{\delta}$ is the one that equalizes the determinant of C to zero:

$$\bar{\delta} = \frac{2 + 2c(1 - |\eta|)}{3 + 2c(1 - |\eta|)}.$$

No Disclosure and Full Disclosure Equilibrium pricing behavior is derived from the system of first-order conditions:

$$\mathbb{E}_\mu[q_i(a_i, a_{-i}, \omega) + \frac{\partial q_i(a_i, a_{-i}, \omega)}{\partial a_i}(a_i - 2ca_i)] = 0, \quad i = 1, 2.$$

Numerical Example At $c = 1$, $\eta = 1/2$, by condition (53) $\hat{\delta} = 11/18$, and by condition (54),

$$\underline{x} = \max \left\{ \frac{18\delta - 15}{22}, \frac{4\delta - 3}{10} \right\}.$$

Given that $\omega_0 = 3$ and $\sigma^2 = 1$, the expected dual payoff, as given in (55) can be expressed as

$$W(x) = \frac{2}{25(45 + 72\delta^2 + 216x + 220x^2 - 2\delta(57 + 134x))} \left(-12015 + 31717\delta^3 - 57902x - 57015x^2 + 3600x^3 - \delta^2(69463 + 119098x) + \delta(50278 + 167872x + 95745x^2) \right).$$

The value of x that minimizes $W(x)$ satisfies the following equation:

$$(2996\delta^4 - 7880\delta^3 + 7490\delta^2 - 3000\delta + 414) + (6728\delta^3 - 20948\delta^2 + 20234\delta - 6210)x + (-52780\delta^2 + 88084\delta - 36368)x^2 + (77184\delta - 62208)x^3 - 31680x^4 = 0. \quad (63)$$

Equation (63) can be solved in radicals for any $\delta \in [0, 1]$. A solution x^* greater than \underline{x} corresponds to an optimal linear certificate. Numerical calculations show that such an x is unique for all $\delta \neq \hat{\delta}$ and its value is continuous in δ . The numerical values derived from these calculations were used to create the figures presented in the main text.

B Bounded Action Spaces

Consider the concave information-design problem as in the main text, but let the action space of each player be $A_i = [\underline{a}_i, \bar{a}_i]$, $-\infty < \underline{a}_i < \bar{a}_i < +\infty$.¹⁶ For any given $\mu \in \Delta(A_{-i} \times \Omega)$, the player's best-response action $a_i^*(\mu)$, if interior, must be unimprovable by local deviations to lower and higher

¹⁶The extension to half-bounded spaces is straightforward.

actions and hence satisfies the first-order condition

$$\mathbb{E}_\mu [\dot{u}_i(a_i^*, a_{-i}, \omega)] = 0.$$

In contrast, the optimal boundary actions must only be unimprovable by one-sided local deviations. As such, the player's best response if located on the boundary must satisfy the following:

$$\mathbb{E}_\mu [\dot{u}_i(\underline{a}_i^*, a_{-i}, \omega)] \leq 0, \quad (64)$$

$$\mathbb{E}_\mu [\dot{u}_i(\bar{a}_i^*, a_{-i}, \omega)] \geq 0. \quad (65)$$

We can write the resulting primal information-design problem as follows:

$$V_B^P \triangleq \sup_{\pi \in \Delta(A \times \Omega)} \int_{A \times \Omega} v(a, \omega) d\pi \quad (66)$$

$$\text{s.t. } \int_{A'_i \times A_{-i} \times \Omega} \dot{u}_i(a, \omega) d\pi = 0 \quad \forall i = 1, \dots, N, \text{ measurable } A'_i \subseteq (\underline{a}_i, \bar{a}_i), \quad (67)$$

$$\int_{A'_i \times A_{-i} \times \Omega} \dot{u}_i(a, \omega) d\pi \leq 0 \quad \forall i = 1, \dots, N, \text{ measurable } A'_i \subseteq [\underline{a}_i, \bar{a}_i), \quad (68)$$

$$\int_{A'_i \times A_{-i} \times \Omega} \dot{u}_i(a, \omega) d\pi \geq 0 \quad \forall i = 1, \dots, N, \text{ measurable } A'_i \subseteq (\underline{a}_i, \bar{a}_i], \quad (69)$$

$$\int_{A \times \Omega'} d\pi = \int_{\Omega'} d\mu_0 \quad \forall \text{ measurable } \Omega' \subseteq \Omega. \quad (70)$$

This primal problem entails the dual problem

$$V_B^D \triangleq \inf_{\lambda \in \times_i \mathcal{M}(A_i), \gamma \in \mathcal{M}(\Omega)} \int_{\Omega} \gamma(\omega) d\mu_0 \quad (71)$$

$$\text{s.t. } \sum_{i=1}^N \lambda_i(a_i) \dot{u}_i(a, \omega) + \gamma(\omega) \geq v(a, \omega) \quad \forall a \in A, \omega \in \Omega, \quad (72)$$

$$\lambda_i(\underline{a}_i) \leq 0, \lambda_i(\bar{a}_i) \geq 0 \quad \forall i = 1, \dots, N. \quad (73)$$

The presence of additional obedience constraints (68), (69) in the primal problem translates into the sign constraints on the Lagrange multipliers (73) in the dual problem. As in the case of unbounded

actions, the dual problem (71) can be simplified and rewritten as

$$V_B^D = \inf_{\lambda \in \times_i \mathcal{M}(A_i)} \mathbb{E}_{\mu_0} [\sup_{a \in A} w(a, \omega, \lambda)] \quad (74)$$

$$\text{s.t. } \lambda_i(\underline{a}_i) \leq 0, \lambda_i(\bar{a}_i) \geq 0 \quad \forall i = 1, \dots, N. \quad (75)$$

The adversarial-contracting interpretation remains intact, but the space of allowed contracts is limited at the boundary actions by the presence of sign constraints.

Lemma 8. (Weak Duality with Bounded Action Spaces) $V_B^P \leq V_B^D$.

Proof. Take any dual variables (λ, γ) that satisfy the constraints of the dual problem (71). Take any measure π that satisfies the constraints of primal problem (66). Integrating both sides of the dual constraints over $a \in A$ and $\omega \in \Omega$ against measure π yields

$$\int_{A \times \Omega} v(a, \omega) d\pi \leq \int_{A \times \Omega} \sum_{i=1}^N \lambda_i(a_i) \dot{u}_i(a, \omega) d\pi + \int_{A \times \Omega} \gamma(\omega) d\pi \leq \int_{\Omega} \gamma(\omega) d\mu_0, \quad (76)$$

where the second inequality follows because π satisfies the primal constraints and the Lagrange multipliers satisfy the dual constraints. (This inequality holds as equality in the case of unbounded actions.) The left-hand side of (76) is the value of the primal problem given measure π , whereas the right-hand side of (76) is the value of the dual problem given dual variables (λ, γ) . As inequality (76) holds for any allowed values of primal measure and dual variables, it also holds at the respective maximization and minimization limits. \square

In the case of bounded actions, we call a measure $\pi \in \Delta(A \times \Omega)$ *implementable by information* if it satisfies the constraints of the primal problem (66) and *implementable by incentives* if there exists a feasible contract in the dual problem (74) that induces this measure as a best response.

Theorem 6. (Optimality Certification with Bounded Action Spaces) *In the case of bounded actions, if $\pi \in \Delta(A \times \Omega)$ is implementable by information and is implementable by incentives by contract λ , then (i) π solves the information-design problem, (ii) λ solves the adversarial-contracting problem, and (iii) $V^P = V^D$.*

Proof. Take any primal measure π implementable by information, i.e., one that satisfies the constraints of the primal problem (66). If it is implementable by incentives, then there exist dual variables λ that

implement this measure in the dual problem (74), and

$$\begin{aligned}
V_B^D &= \inf_{\lambda' \in \times_i \mathcal{M}(A_i)} \mathbb{E}_{\mu_0} \left[\sup_{a \in A} w(a, \omega, \lambda') \right] \\
&\leq \mathbb{E}_{\pi} [w(a, \omega, \lambda)] \\
&= \int_{A \times \Omega} v(a, \omega) d\pi - \int_{A \times \Omega} \sum_{i=1}^N \lambda_i(a_i) u_i(a, \omega) d\pi \\
&\leq \int_{A \times \Omega} v(a, \omega) d\pi \leq V_B^P,
\end{aligned}$$

where the first inequality follows from the implementability of π in the dual problem, whereas the last three steps follow from the feasibility of π in the primal problem and the sign constraints on dual variables. (The next-to-last step holds as equality in the case of unbounded actions.)

Furthermore, by Lemma 8, $V_B^D \geq V_B^P$. Hence,

$$V_B^D = \int_{A \times \Omega} v(a, \omega) d\pi = V_B^P,$$

which proves the optimality of measure π . □

C Online Appendix

In this appendix, we apply the certification approach to a game with non-quadratic payoffs.

Imagine a designer who can inform a representative retailer and a representative consumer about a new two-sided application. The application is either high quality ($\omega = 1$) or low quality ($\omega = 0$).

The retailer chooses a search intensity $a_1 \in [0, 1]$ in the application, and the consumer chooses $a_2 \in [0, 1]$. Each player i incurs the cost

$$c(a_i) = \begin{cases} \hat{a}^3 a_i, & a_i \in [0, \hat{a}), \\ \frac{a_i^4}{4} + \frac{3\hat{a}^4}{4}, & a_i \in [\hat{a}, 1]. \end{cases}$$

where $\hat{a} = 3/4$. Thus c is smooth and convex, and \hat{a} minimizes the average-cost function.

The players' payoffs are

$$u_i(a, \omega) = \omega a_i a_{-i} - c(a_i).$$

A higher search intensity a_i raises player i 's chance of finding a match, and this benefit is greater when the state is high and when the other player searches more intensively. Hence the game exhibits strategic complementarities.

The designer aims to raise overall search intensity to boost engagement and advertising revenue. Her payoff is

$$v(a, \omega) = a_1 + a_2.$$

Public Information First, suppose the platform sends a public signal, identical for both players. The optimal signal can then be derived via the standard concavification procedure. Let μ denote the common belief that $\omega = 1$. Given this belief, the players play either (i) $a_1 = a_2 = a^*(\mu) \geq \hat{a}$, where

$$\mu a^*(\mu) = (a^*(\mu))^3,$$

so $a^*(\mu) = \sqrt{\mu} \geq \hat{a}$; or (ii) $a_1 = a_2 = 0$. Consequently, the designer's indirect utility equals $\sqrt{\mu}$ whenever $\sqrt{\mu} \geq \hat{a}$ and 0 otherwise.

Hence, if the prior satisfies $\mu_0 \geq \hat{a}^2$, no disclosure is optimal. If $\mu_0 < \hat{a}^2$, the optimal public signal

splits posteriors between 0 and \hat{a}^2 ; concretely, (i) when $\omega = 0$, the posterior is either 0 or \hat{a}^2 , leading to actions $(a_1, a_2) = (0, 0)$ or (\hat{a}, \hat{a}) ; (ii) when $\omega = 1$, the posterior is always \hat{a}^2 , so $(a_1, a_2) = (\hat{a}, \hat{a})$.

Optimal Information Now, we show that the above public information is optimal even within the class of all private information structures.

The dual payoff is

$$w(a, \omega) = a_1 + a_2 - \lambda_1(a_1)(\omega a_2 - c'(a_1)) - \lambda_2(a_2)(\omega a_1 - c'(a_2)).$$

Posit the following certificate:

$$\lambda_i(a_i) = \begin{cases} -\frac{a_i}{\hat{a}^3}, & \text{if } a_i < \hat{a}, \\ -\frac{1}{a_i^2}, & \text{if } a_i \geq \hat{a}, \end{cases}$$

so that

$$w(a, \omega) = \begin{cases} \omega\left(\frac{a_2}{a_1^2} + \frac{a_1}{a_2^2}\right), & \text{if } a_1, a_2 \geq \hat{a}, \\ \omega\left(\frac{a_2}{a_1^2} + \frac{a_1 a_2}{\hat{a}^3}\right), & \text{if } a_1 \geq \hat{a} > a_2, \\ \omega\left(\frac{a_1}{a_2^2} + \frac{a_1 a_2}{\hat{a}^3}\right), & \text{if } a_2 \geq \hat{a} > a_1, \\ \omega\left(\frac{2a_1 a_2}{\hat{a}^3}\right), & \text{if } a_1, a_2 < \hat{a}. \end{cases}$$

From the perspective of the dual agent, if $\omega = 0$, then any a is optimal, including $a = (0, 0)$ and (\hat{a}, \hat{a}) . If $\omega = 1$, then $w(a, \omega)$ is strictly increasing in a_i whenever $a_i < \hat{a}$. Thus, we can focus on the case $a_1, a_2 \geq \hat{a}$. Because $w(a, \omega)$ is convex on this domain, its maximum must occur at one of the corner points $a = (\hat{a}, \hat{a})$, $(\hat{a}, 1)$, $(1, \hat{a})$, or $(1, 1)$. A straightforward calculation shows that (\hat{a}, \hat{a}) is optimal whenever $1 - \hat{a} - \hat{a}^2 \leq 0$, which indeed holds for $\hat{a} = 3/4$. The result follows.