J. Econ. Theory 238 (2026) 106272



Contents lists available at ScienceDirect

Journal of Economic Theory

journal homepage: www.elsevier.com/locate/jet



Strategic ignorance and information design ⋆

Ina Taneva [a,*], Thomas Wiseman [b]

[a] University of Edinburgh and CEPR, EH8 9JT, Edinburgh, UK
[b] University of Texas at Austin, 78712, Austin, Texas, USA

ARTICLE INFO

JEL classification:
D82
D83

Keywords:
Information design
Strategic ignorance
Correlated equilibrium

ABSTRACT

We study information design in strategic settings when agents can publicly refuse to view their private signals. Ignoring the constraints that agents must be willing to view their signals may lead to substantial divergence between the designer's intent and actual outcomes, even in the case where the designer seeks to maximize the agents' payoffs. We introduce the appropriate equilibrium concept — ignorance-permissive Bayes correlated equilibrium — and characterize implementable distributions over states and actions. The designer's optimal response to strategic ignorance generates qualitative properties that standard information design cannot: the designer may provide redundant or even counterproductive information, asymmetric information structures may be strictly optimal in symmetric environments, providing information conditional on players' viewing choices rather than all at once may hurt the designer, and communication between players may help her. Optimality sometimes requires that players ignore their signals with positive probability.

1. Introduction

We argue that in modeling information design, it can be important to incorporate the incentives of agents to accept information as well as the designer's incentive to provide it. In the standard setting of information design (e.g., Bergemann and Morris, 2019, Taneva, 2019), a designer commits to disclosing information about an uncertain payoff relevant state to a group of interacting agents. Through the release of information, the designer incentivizes the agents to take actions that will benefit her. An implicit assumption is that players will agree to get informed according to the information structure chosen by the designer, which comprises joint distributions of agents' private messages conditional on each possible realization of the state. Crucially, that setting does not permit players to refuse to observe the signals and to credibly demonstrate this choice to the other players. In many strategic environments, however, an agent may benefit from publicly remaining uninformed. Therefore, if we augment the standard information design framework with a pre-play stage where players publicly choose whether or not to observe the signal sent by the designer, then in many settings it is unreasonable to assume that players can be induced to play under the designer-chosen information structure. In such cases, the intended information structure provided by the designer gets transformed through the strategic choices of the agents into a very different informational environment. In this paper, we study when and how a designer modifies her choice of information structure in response.

⋆ We are grateful to the editor Pierpaolo Battigalli, an associate editor, and two referees for their suggestions, which significantly improved the paper. We thank Alesandro Arcuri, Andreas Blume, Ben Brooks, Hector Chade, Inga Deimen, Francesc Dilmé, Laura Doval, Françoise Forges, Amanda Friedenberg, Willie Fuchs, Johannes Hörner, Elliot Lipnowski, Erik Madsen, Arina Nikandrova, Jacopo Perego, Eddie Schlee, Denis Shishkin, Vasiliki Skreta, Mark Whitmeyer, Leeat Yariv, Gabriel Ziegler, and seminar audiences for helpful comments and discussions. Siyu Chen provided excellent research assistance.

* Corresponding author.
E-mail addresses: ina.taneva@ed.ac.uk (I. Taneva), wiseman@austin.utexas.edu (T. Wiseman).

https://doi.org/10.1016/j.jet.2026.106272
Received 9 March 2025; Received in revised form 24 October 2026; Accepted 25 October 2026
Available online 29 October 2026
0022-0531/© 2026 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

	H	L
E	3, 0	1, 1
M	2, 2	0, 0
F	0, 0	−2, 1

$\omega = e$

	H	L
E	0, 0	−2, 1
M	2, 2	0, 0
F	3, 0	1, 1

$\omega = f$

Fig. 1. Procurement auction example.

	HH	HL	LH	LL
E	1.5	0.5	0.5	−0.5
M	2	1	1	0
F	1.5	0.5	0.5	−0.5

$\Pr(\omega = e) = \frac{1}{2}$

Fig. 2. Small company's expected payoffs at the prior.

Most of the literature on information design following Kamenica and Gentzkow (2011) focuses on the case of a single agent, where information always has weakly positive value. In that case the issue of robustness to strategic ignorance does not arise. The gain from ignoring information comes when other players change their behavior in response: that indirect, strategic benefit may outweigh the agent's reduced ability to tailor his own action to the state and to the other players' actions. Suppose, for example, that the designer is a government agency that wants to find a supplier of internet connectivity through a procurement auction. The agency does not have the technical expertise to determine its own connectivity needs, but it can provide a report on its operations, work protocols, etc., which will let the bidders identify its needs and the corresponding best solution. There are two bidders: a large company, with many clients, and a small company, which would serve only this agency. We model their interaction in the payoff matrices in Fig. 1, one for each equally likely state of the world, $\omega \in \{e, f\}$, corresponding to whether the agency needs solution E or solution F. The row player is the small company, which has three possible actions. Action E represents a choice to invest, ahead of the auction, in technology that will let it provide solution E at a low cost and hence a low bid; action F is the equivalent choice for solution F; action M corresponds to no investment and a high bid to reflect the high costs of delivery without the preliminary investment. The column player is the large company, which serves many other clients and will not find it profitable to invest in a bespoke solution. Its choices are to bid high (H) or low (L) in the auction.

The agency wants both companies to submit low bids, and it would like to have the right bespoke solution if the small company wins the auction. Specifically, it gets a payoff of 1 if (E, L) is played in state e or (F, L) is played in state f, and 0 otherwise. The agency can achieve its goals through information design by providing a detailed report that both companies inspect: when the realized state is common knowledge, then the small company has a dominant strategy to match the state. The large company's best response is the low bid L, so the agency gets payoff 1. If, however, the small company can *credibly* signal to the big company that it has not read the report, for example by preparing and submitting its bid before or immediately after the report is made available, then as long as the big company *observes* that choice, the small company would choose not to get informed about the agency's needs. Under the prior distribution over states, no investment (M) strictly dominates (against any state-contingent strategy of the large company) blindly investing in either solution, as shown in Fig. 2. The large company's best response is high bid H in both states, so by ignoring the report the small company increases its payoff from 1 to 2. The agency, though, gets payoff 0.

There are many other economic settings where demonstrating ignorance is valuable, as we discuss below, and we will show that incorporating the designer's incentive to provide information broadens the range of such settings. The requirement that agents must be incentivized to view their signals, then, imposes new constraints on the designer's choice of information structure. Those constraints are conceptually analogous to the participation constraints in mechanism design. Our goal in this paper is to understand the impact of those "Look constraints" on the set of implementable outcomes. Formally, we augment the baseline environment (that is, where agents must view their signals) with a simultaneous-move pre-play stage where the players publicly choose whether to "Look" at their private signals or "Ignore" them. We find that in two applications prominent in the literature on information design in games, currency attacks and a binary investment game, if the designer provides the information structure that would be optimal in the baseline environment, then there is no equilibrium where all players choose to Look at their signals. As a consequence, the outcome is not what the designer intended. Instead, for each application we derive the designer's optimal information structure when accounting for the presence of strategic ignorance.

As is standard in the literature on information design (e.g., Bergemann and Morris, 2016; Taneva, 2019), we assume that the designer costlessly commits to an information structure without observing the state and that the agents cannot communicate with

each other, and we restrict attention to the best equilibrium for the designer.[1] Our other key assumption is that each agent's choice of whether to Look at or Ignore his signal is both observed by the other agents and irrevocable. That is, agents publicly commit to their choices of whether or not to become informed. Otherwise, that choice would not influence other players' subsequent actions, and the choice to Look would be weakly dominant, just as in the single-agent case. By giving agents this commitment power, our paper can be viewed as an attempt to level the playing field in terms of the commitment assumption that is liberally granted in the information design literature, but only to the designer/sender. In this sense, we argue that in settings where agents can exercise public commitment to ignorance, the designer needs to account for the presence of these strategic motives and potentially adjust the information structure to ensure robustness of the desired outcome.

We view our model as the most immediate departure from the standard setting that allows us to explore robustness to strategic ignorance in the context of information design. There are other reasonable formulations, such as allowing the designer to also choose the extensive form of the game, to reveal information incrementally in multiple rounds, or to commit to different information structures contingent on the Look-Ignore choices of the agents, and we believe these other variations to be interesting avenues for future research on how a designer can respond to strategic ignorance.

1.1. Preview of results

As a high-level overview, the nature of our results is as follows. We first provide a general characterization of the set of implementable outcome distributions for *any* finite game. A designer with *any* objective function can find the optimal outcome distribution by maximizing her expected payoff, which is linear in the outcome probabilities, over this (semi-algebraic) set. We then present a number of results concerning properties of optimal information design under strategic ignorance and reversals of well-established results from standard information design. We establish these latter results through examples and demonstrate the underlying intuition.

A given joint distribution over actions and states is *implementable* if it is the outcome of a perfect Bayesian equilibrium with a "no-signaling-what-you-don't-know" refinement (PBE*, Fudenberg and Tirole (1991)) of the two-stage game for some information structure.[2] In Theorem 1, we characterize the implementable outcome distributions over actions and states under strategic ignorance in general finite environments. To this end, we define the concept of *ignorance-permissive Bayes correlated equilibrium* (IPBCE). Theorem 1 demonstrates that the set of IPBCE outcome distributions is equivalent to the set of PBE* outcome distributions across all information structures.[3] For our environment with strategic ignorance, this is a counterpart of the celebrated equivalence result of Bergemann and Morris (2016), which established that the set of Bayes correlated equilibrium (BCE) outcome distributions is equivalent to the union across all information structures of the set of BNE outcome distributions. IPBCE outcomes are exactly those BCE outcomes that the designer can still achieve under the additional constraints of strategic ignorance.

A corollary of the theorem (Corollary 1) states that it is without loss of generality to restrict the designer to direct *contingent* information structures, where messages correspond to (pure) action recommendations for each possible choice of the other players in the pre-play Look-Ignore stage. What changes relative to the baseline environment is that the direct information structures with *single* action recommendations are no longer enough.

Our first set of results (Theorems 2–4) outlines the properties of optimal robust design. Theorem 2 identifies a condition from the literature that ensures that an outcome is robust to strategic ignorance. Theorem 3 shows that in some cases the designer's optimal outcome is implementable only in an equilibrium where some players Ignore their signals with positive probability.[4] Importantly, that outcome distribution cannot be replicated by an information structure that sometimes sends an uninformative message to some players. Theorem 4 demonstrates that giving the players the option to Ignore messages does not necessarily make them better off. Even if the designer's goal is to maximize players' payoffs, all players may get lower payoffs than under the baseline where messages are automatically observed.[5]

[1] That is, we assume, first, that after a player deviates at the Look-Ignore stage, the worst continuation BNE of the resulting belief system for the deviator is played. Second, we assume that on path agents play the designer's preferred BNE among those that satisfy the Look-Ignore constraints. We note, as a subtlety, that there may be other BNEs at the second stage, given the on-path Look-Ignore choices, that give the designer a higher payoff but that would not make the specified Look-Ignore choices optimal at the first stage.

[2] By specifying the designer's chosen information structure as part of the two-stage game, we sidestep the issue of players' updating their beliefs about the state based on the designer's choice. That is, our game begins with the information structure already chosen, and the "no-signaling-what-you-don't-know" refinement applies to the updating after surprise moves by the players at the Look-Ignore stage. Watson (2025) provides a condition called plain consistency that would apply to the designer's choice from the infinite set of possible information structures as well.

[3] Strategic ignorance can also restrict the set of implementable correlated equilibria (CE) in games of complete information. In the companion working paper (Taneva and Wiseman, 2026), we provide a two-player example in which the worst Nash equilibrium gives a better payoff than the worst CE. Interestingly, that situation can never arise in any two-player binary-action complete information game.

[4] We thank Elliot Lipnowski for pushing us to investigate this question. This result is related to corresponding mechanism design results of rejecting a mechanism in equilibrium (e.g., Celik and Peters, 2011; Balzer and Schneider, 2019, and Correia-da-Silva, 2020). There are differences on many dimensions, but the main conceptual departure from these papers is that in our model the refusal to get informed does not convey information about the private types of agents, as there are no such types. Instead, agents endogenously influence their outside option by deciding whether or not to accept information.

[5] A related implication is that a collusive agreement (corresponding to a designer who seeks to maximize players' payoffs) on what types of information to obtain and observe may not be sustainable. Bergemann et al. (2017), for example, study the information structures over bidders' values that would minimize the distribution of winning bids in a first price auction.

Our second set of findings (Propositions 1–5) showcases the reversal of well-known results from standard information design. A consequence of the adjustments the designer needs to make in the presence of strategic ignorance is that she may end up providing what would be considered redundant (Proposition 1) or counterproductive (Proposition 2) information from the perspective of standard design. Also surprisingly, we find that the designer may need to use an asymmetric information structure even in a completely symmetric environment (Proposition 3). An additional substantive difference is that while the set of implementable outcomes in the standard information design environment is always decreasing in the amount of exogenous information that players have (Bergemann and Morris, 2016), that monotonicity may fail in the presence of strategic ignorance.

The next two results are based on modifications of our main environment. In Proposition 4, we show that providing a recommendation on how to punish a player only after that player has deviated by Ignoring the original signal may give the designer a worse outcome than providing all contingent recommendations simultaneously at the beginning. The reason is that incentives must be provided to view each separate signal. In contrast, in the standard information design setting, providing information only when it is needed cannot give the designer a worse outcome (e.g., Makris and Renou, 2023).

Another salient qualitative reversal from the standard information design setting is that allowing the players to communicate with each other after receiving their private signals may improve outcomes for the designer, as we show in Proposition 5. In contrast, in the standard environment, a designer can be only hurt when players communicate, because any information she wants them to have she can impart directly.

Finally, it is important to emphasize that the issue of robustness to strategic ignorance is distinct from the question of equilibrium selection – that is, of whether agents will play the designer's preferred equilibrium when multiple equilibria exist. We maintain the assumption that the designer's preferred equilibrium is played (advantageous selection) throughout, and so we consider an outcome robust if it can be achieved in *some* equilibrium of the dynamic game.

1.2. Applications and relation to literature

A setting that broadly matches our model is the U.S. Forest Service auctions for timber harvest contracts, studied by Athey and Levin (2001). The composition (in terms of timber species) of the tract being auctioned determines its value to bidders. The Forest Service can provide bidders with its own estimates, and it can also allow bidders to "cruise" designated portions of the tract and gather private information.[6] By choosing which portions of the tract to make available to different bidders, the Forest Service can control the precision and correlation of those private signals. Bidders have the option to abstain from cruising the tract before the auction and the identities of bidders who undertake a cruise are public record. Another matching setting is online meetings using Teams or similar software. Both public messages and private messages to subsets of attendees are feasible, and as long as invitees can observe who joined the meeting, then a decision not to join corresponds to how we model strategic ignorance. As another example, members of networking and social media platforms can choose to unfollow certain other members or unsubscribe from particular services to send a publicly observable signal that will change the perceptions and actions of their followers.

Previous research has identified many settings where incentives for strategic ignorance of a payoff-relevant state arise. In all of those examples, as in our setting, the decision to remain ignorant has strategic value only if it is observed by other agents. In the context of relationship-specific investments that may create a hold-up problem, a public commitment by the party with the bargaining power to not obtain the private information available to the vulnerable party may incentivize the latter to make an optimal investment in the relationship (Tirole, 1986; Rogerson, 1992; Gul, 2001). Committing to ignorance can prevent a situation of asymmetric information and the resulting adverse selection problems of Akerlof (1970) or preserve incentives for efficient risk-sharing as in Hirshleifer (1971), Rothschild and Stiglitz (1976), and Schlee (2001). Strategic ignorance about demand can be utilized by a less risk-averse firm to create risk and thus induce a more risk-averse opponent in a Cournot duopoly game to scale back its production, resulting in a higher price, as in Palfrey (1982), or by a Stackelberg leader to maintain his first-mover advantage as in Gal-Or (1987). Similarly, a public commitment to information avoidance can be used to convincingly strengthen one's bargaining position (Schelling, 1956).[7] Other papers have pointed out benefits of strategic ignorance in the context of procurement costs (Kessler, 1998), private-values in second-price auctions (McAdams, 2012), buyer valuations in bilateral trade (Roesler and Szentes, 2017), and sellers' marginal production costs in consumer search (Atayev, 2022). In the context of strategic communication, Deimen and Szalay (2019) show that an expert optimally chooses to remain partially ignorant about his own preferred choice in order to credibly influence a decision-maker with his advice.

Many of the papers mentioned above consider strategic ignorance as a choice between becoming perfectly informed about the state or remaining fully uninformed. We find that endogenizing the information provided by the designer may broaden the class of settings where players' strategic incentives to ignore information are a relevant concern. In the investment game in Section 4.2, players faced with a choice between learning the state perfectly or learning nothing would want to learn the state. We will see, however, that if the designer provides the information structure that would maximize her objective in the baseline case where players must observe their messages, then strategic ignorance becomes important: the players will choose to Ignore.

Motivated by a similar question, Arcuri (2021) considers a weaker form of robustness to strategic ignorance: an information structure S satisfies the "hear-no-evil" condition if for each player i, there is some BNE at the action stage under S that player i

[6] Large forest product companies retain in-house cruisers, so the marginal cost of cruising a tract for a given auction is zero. (See Athey and Levin, 2001, p.381.)

[7] Golman et al. (2017) provide a detailed overview of the different motives behind the avoidance of free and payoff-relevant information along with many examples from the theoretical and experimental literature.

prefers to the worst BNE for him under the information structure that results if he unilaterally Ignores his message. That definition allows for the possibility that a player prefers his worst BNE after deviating to Ignore over his equilibrium outcome. In contrast, we require a stronger form of robustness: the actually played BNE outcome distribution needs to give each player a higher payoff than what he could obtain by deviating.

Because of the pre-play Look-Ignore stage, our paper is related to the literature on sequential information design and information design in multi-stage games (Doval and Ely, 2020; de Oliveira and Lamba, 2019, and Makris and Renou, 2023), and more broadly to the literature on generalizations of correlated equilibrium to multi-stage games (e.g., communication equilibrium[8] (Forges, 1986; Myerson, 1986) and extensive form correlated equilibrium (von Stengel and Forges, 2008)). The main difference relative to these papers is that in our model there is a single information structure provided to the players at the beginning of the interaction and the extensive form is fixed.

2. Model & characterization result

There is a set $\mathcal{I}$ of $N > 1$ expected-utility maximizing agents who will play a simultaneous-move stage game. Each player i has a finite set of actions A_i with $a_i \in A_i$ denoting a generic action for player i; $A \equiv A_1 \times \ldots \times A_N$ is the set of action profiles with $a \in A$ denoting a generic action profile.[9] There is a finite set of states of the world Ω, with generic element ω. Agents' payoffs are given by $u : A \times \Omega \to \mathbb{R}^N$, where agent i's payoff function $u_i : A \times \Omega \to \mathbb{R}$ depends on the action profile and the (ex ante unknown) state. The designer has a utility function $u^D : A \times \Omega \to \mathbb{R}$, so that her payoff also depends on the agents' actions and the state. The agents and the designer share a common full-support prior μ over Ω. Let $G = ((A, u), \mu)$ be the *basic game*.

An *information structure* (T, P) consists of 1) a finite set of possible signal realizations T_i for each agent i, with $T \equiv T_1 \times \ldots \times T_N$; and 2) conditional signal distributions $P : \Omega \to \Delta(T)$, one for each state.

Given a basic game G, the designer publicly commits to an information structure (T, P). Play then proceeds as follows: the state $\omega \in \Omega$ is realized according to μ. Then the vector of signals $t \in T$ is drawn according to $P(\cdot|\omega)$, and the designer sends each agent i his private signal t_i.

At the *Look-Ignore stage*, each agent makes a choice $s_i \in S_i \equiv \{\ell, g\}$: whether to *Look* ($\ell$) at his signal and learn the realization of t_i, or to *Ignore* (g) it and remain uninformed. The Look-Ignore choices are public[10] and simultaneous. Given a profile $s \in S \equiv \{\ell, g\}^N$ of realized choices from the Look-Ignore stage, let $\mathcal{L}(s) := \{i : s_i = \ell\}$ denote the set of players who chose Look, and let $\mathcal{G}(s) := \mathcal{I} \setminus \mathcal{L}(s)$. Given an information structure (T, P) and a subset of players $\mathcal{L} \subseteq \mathcal{I}$, denote by $\left(T_{\mathcal{L}}, P_{\mathcal{L}}\right)$ the informational environment where it is common certainty that all $i \in \mathcal{L}$ have been informed according to (T, P) while all $i \in \mathcal{G} := \mathcal{I} \setminus \mathcal{L}$ do not observe any signal realization. That is, $\left(T_{\mathcal{L}}, P_{\mathcal{L}}\right)$ is the information structure induced by (T, P), and the (publicly observed) choices of Look by the agents in $\mathcal{L}$ and of Ignore by the agents in $\mathcal{G}$. Upon choosing Look and observing t_i and s, agent i updates his beliefs about the state and the signals observed by other agents by applying Bayes' rule to his own signal realization t_i, $\left(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)}\right)$ and the prior μ. An agent who chooses to Ignore his signal uses $\left(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)}\right)$ and μ to form beliefs about t_{-i} and does not update his beliefs about the state.

Given (T, P) and s, define the *action stage* by the Bayesian game $G\left(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)}\right)$. At this stage, each agent i chooses an action $a_i \in A_i$, and payoffs are realized. For a given information structure (T, P), we will refer to the basic game augmented by the Look-Ignore and the action stage as the *dynamic game*, denoted by $G^*(T, P)$. An *outcome* $v \in \Delta(A \times \Omega)$ is a joint distribution over action profiles and states. A strategy for player i in dynamic game G^* is a tuple $(\gamma_i, (\tilde{\beta}_i^s)_s)$ with $\gamma_i \in \Delta\{\ell, g\}$, $\tilde{\beta}_i^s : T_i \to \Delta A_i$ if $i \in \mathcal{L}(s)$, and $\tilde{\beta}_i^s \in \Delta A_i$ if $i \in \mathcal{G}(s)$. Let $\gamma := (\gamma_i)_{i \in \mathcal{I}}$ and $\tilde{\beta}^s := (\tilde{\beta}_i^s)_{i \in \mathcal{I}}$.

Our solution concept for a dynamic game G^* is perfect Bayesian equilibrium with a "no-signaling-what-you-don't-know" refinement. In particular, continuation play in the action stage $G\left(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)}\right)$ after a subset of agents $\mathcal{L}(s)$ choose Look must constitute a BNE of that game (Definition 1). In the Look-Ignore stage, each agent optimally chooses in order to maximize his expected continuation payoffs (Definition 2). Given a realized profile s of choices from the Look-Ignore stage, the information structure $\left(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)}\right)$ is common knowledge. Agent i who has chosen Look and observed t_i updates his beliefs about ω and t by using Bayes' rule. (The agent also observes s, but "no-signaling-what-you-don't-know" implies that the Look-Ignore choices are uninformative about ω and t.) Similarly, an agent who has chosen Ignore observes only s, so he does not update his beliefs about ω and t.

Definition 1. Given (T, P) and $s \in S$, $\tilde{\beta}^s$ is a BNE of $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$ if:
for each $i \in \mathcal{L}(s)$, $t_i \in T_i$, and $a_i \in A_i$ with $\tilde{\beta}_i^s(a_i|t_i) > 0$, we have

$$\sum_{a_{-i}, t_{\mathcal{L}(s)\setminus i}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(t_i, t_{\mathcal{L}(s)\setminus i}|\omega) \left(\prod_{j \in \mathcal{L}(s)\setminus i} \tilde{\beta}_j^s(a_j|t_j) \prod_{k \in \mathcal{G}(s)} \tilde{\beta}_k^s(a_k) \right) u_i(a_i, a_{-i}, \omega)$$
$$\geq \sum_{a_{-i}, t_{\mathcal{L}(s)\setminus i}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(t_i, t_{\mathcal{L}(s)\setminus i}|\omega) \left(\prod_{j \in \mathcal{L}(s)\setminus i} \tilde{\beta}_j^s(a_j|t_j) \prod_{k \in \mathcal{G}(s)} \tilde{\beta}_k^s(a_k) \right) u_i(a_i', a_{-i}, \omega), \quad (1)$$

[8] Communication equilibria involve eliciting the information that agents have before sending action recommendations, while in BCE agents do not have any information that needs to be elicited.

[9] We also use the notation $A_{-i} = \times_{j \neq i} A_j$ with generic element $a_{-i} \in A_{-i}$ and $A_{-ik} = \times_{j \neq i,k} A_j$ with generic element $a_{-ik} \in A_{-ik}$. We apply this convention to other relevant sets as well.

[10] Our results, with the exception of Theorem 3, generalize to the case where only the agents who choose to Look observe the Look-Ignore choices of all agents, while the agents who choose to Ignore do not observe those choices.

for all $a_i' \in A_i$;
and for each $i \in \mathcal{G}(s)$ and $a_i \in A_i$ with $\tilde{\beta}_i^s(a_i) > 0$, we have

$$\sum_{a_{-i}, t_{\mathcal{L}(s)}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)}|\omega) \left(\prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j|t_j) \prod_{k \in \mathcal{G}(s) \setminus i} \tilde{\beta}_k^s(a_k) \right) u_i(a_i, a_{-i}, \omega)$$
$$\geq \sum_{a_{-i}, t_{\mathcal{L}(s)}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)}|\omega) \left(\prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j|t_j) \prod_{k \in \mathcal{G}(s) \setminus i} \tilde{\beta}_k^s(a_k) \right) u_i(a_i', a_{-i}, \omega), \tag{2}$$

for all $a_i' \in A_i$.
Then $v(\tilde{\beta}^s) \in \Delta(A \times \Omega)$ defined as

$$v(\tilde{\beta}^s)(a, \omega) := \sum_{t_{\mathcal{L}(s)}} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)}|\omega) \left(\prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j|t_j) \prod_{i \in \mathcal{G}(s)} \tilde{\beta}_i^s(a_i) \right) \tag{3}$$

for all $a \in A$ and $\omega \in \Omega$ is a BNE outcome (distribution)[11] of $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$.

Definition 2. A strategy profile $\left(\gamma, \left(\tilde{\beta}^s\right)_s\right)$ is a perfect Bayesian equilibrium satisfying the no-signaling-what-you-don't-know refinement (PBE*) of $G^*(T, P)$ if for each $s \in S$, $\tilde{\beta}^s$ is a BNE of $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$, and for each $i \in \mathcal{I}$ and $s_i \in \{\ell, g\}$ with $\gamma_i(s_i) > 0$,

$$\sum_{s_{-i}, a, \omega} \prod_{j \neq i} \gamma_j(s_j) v(\tilde{\beta}^{s_i, s_{-i}})(a, \omega) u_i(a, \omega) \geq \sum_{s_{-i}, a, \omega} \prod_{j \neq i} \gamma_j(s_j) v(\tilde{\beta}^{s_i', s_{-i}})(a, \omega) u_i(a, \omega), \tag{4}$$

for all $s_i' \in \{\ell, g\}$.
Then $v \in \Delta(A \times \Omega)$ defined as

$$v(a, \omega) := \sum_{s \in S} \prod_{i \in \mathcal{I}} \gamma_i(s_i) v(\tilde{\beta}^s)(a, \omega)$$

for all $a \in A$ and $\omega \in \Omega$ is a PBE* outcome of $G^*(T, P)$.

Definition 3. Let $PBE^*(G^*(T, P))$ denote the set of PBE* outcomes of $G^*(T, P)$.

We say that an outcome v is *robust to strategic ignorance* if $v \in PBE^*(G^*(T, P))$ for some (T, P). That is, when accounting for strategic ignorance, the designer chooses an information structure (T, P) to maximize her expected payoff across the set of all PBE* outcomes $\cup_{(T,P)} PBE^*(G^*(T, P))$. Hence, the designer's problem can be formalized as:

$$\max_{(T,P)} \sum_{a, \omega} u^D(a, \omega) v(a, \omega) \quad \text{s.t.} \quad v \in PBE^*\left(G^*(T, P)\right).$$

2.1. Characterization

The designer maximizes over a very large space, the set of all information structures (T, P). In the standard information design environment, the set of BNE outcomes across all possible information structures equals the set of BCE outcomes. The latter set is much easier to work with (as in Bergemann and Morris, 2016 and Taneva, 2019), because it circumvents the need to specify the information structures explicitly. In our environment with strategic ignorance, the analogous result – that the set of PBE* outcomes across all possible information structures equals the set of BCE outcomes of the two-stage game in our setting – is too strong. For any choice of information structure, the designer's messages can provide correlation of strategies (with the state or with the strategies of other players) only at the action stage and not at the Look-Ignore stage, and only for those players who choose Look. We call a BCE that incorporates those constraints on the correlation structure an ignorance-permissive Bayes correlated equilibrium (IPBCE). Hence, an IPBCE is an element

$$(\gamma, \beta^g, \pi) \in \times_i \left(\Delta\{\ell, g\} \times \left(\times_{s_{-i}} \Delta A_i \right) \right) \times \Delta(\mathcal{A} \times \Omega),$$

where γ denotes the distribution of Look-Ignore recommendations, β^g denotes the distributions of post-Ignore recommendations, $\mathcal{A}_i \equiv A_i^{|S_{-i}|}$ denotes the set of agent i's (pure) mappings from possible Look-Ignore choice profiles of the opponents S_{-i} to own pure actions A_i, and π denotes the joint distribution of post-Look recommendations and the state, which is consistent with the prior over Ω. We will denote a generic element of $\mathcal{A}_i$ by m_i, for message, and denote the action recommended after combination s_{-i} of other agents' Look-Ignore choices by $m_i(s_{-i}) \in A_i$. Let $\mathcal{A} \equiv \mathcal{A}_1 \times \ldots \times \mathcal{A}_N$.

Given (γ, β^g, π), the timing of IPBCE is as follows:

1. Look-Ignore recommendations $s \in S$ are drawn from γ, post-Ignore recommendations from β^g, and (m, ω) from π.

[11] Throughout the paper, we use the terms "outcome" and "outcome distribution" interchangeably.

2. The realization s_i is privately recommended to each agent i.
3. Each agent i chooses $\tilde{s}_i$, which is publicly observed.
4. The realized draws from Step 1 corresponding to the choices in Step 3 are privately recommended: $m_i \in \mathscr{A}_i$ to each i such that $\tilde{s}_i = \ell$, and $a_i(\tilde{s}_{-i}) \in A_i$ to each i such that $\tilde{s}_i = g$.
5. Each agent i makes an action choice $\tilde{a}_i$.

An IPBCE is a triple (γ, β^g, π) such that for all $\tilde{s} \in S$, the action recommendations sent in Step 4 are obedient, and the Look-Ignore recommendations $s \in S$ sent in Step 2 are obedient. We next provide a formal definition.

For each $s \in S$, let $\pi(m_{\mathcal{L}(s)}, \omega) := \sum_{m_{\mathcal{G}(s)}} \pi(m_{\mathcal{L}(s)}, m_{\mathcal{G}(s)}, \omega)$, where $m_{\mathcal{L}} := (m_i)_{i\in\mathcal{L}}$ and $m_{\mathcal{G}} := (m_i)_{i\in\mathcal{G}}$. Similarly, let $a_{\mathcal{G}} := (a_i)_{i\in\mathcal{G}}$.

Definition 4. (γ, β^g, π) is an ignorance-permissive Bayes correlated equilibrium (IPBCE) of G^* if

1. (**Consistency with the prior**) $\pi(\mathscr{A} \times \{\omega\}) = \mu(\omega)$ for all $\omega \in \Omega$;
2. (**Obedience for agent i who chooses Look**) for every $s \in S$, $i \in \mathcal{L}(s)$, $m_i \in \mathscr{A}_i$, and $a_i' \in A_i$

$$\sum_{m_{\mathcal{L}(s)\setminus i}, a_{\mathcal{G}(s)}, \omega} \pi(m_i, m_{\mathcal{L}(s)\setminus i}, \omega) \prod_{k\in\mathcal{G}(s)} \beta_k^g(a_k|s_{-k}) u_i(m_i(s_{-i}), (m_j(s_{-j}))_{j\in\mathcal{L}(s)\setminus i}, a_{\mathcal{G}(s)}, \omega)$$
$$\geq \sum_{m_{\mathcal{L}(s)\setminus i}, a_{\mathcal{G}(s)}, \omega} \pi(m_i, m_{\mathcal{L}(s)\setminus i}, \omega) \prod_{k\in\mathcal{G}(s)} \beta_k^g(a_k|s_{-k}) u_i(a_i', (m_j(s_{-j}))_{j\in\mathcal{L}(s)\setminus i}, a_{\mathcal{G}(s)}, \omega) \tag{5}$$

3. (**Obedience for agent i who chooses Ignore**) for every $s \in S$, $i \in \mathcal{G}(s)$, and $a_i, a_i' \in A_i$ such that $\beta_i^g(a_i|s_{-i}) > 0$

$$\sum_{m_{\mathcal{L}(s)}, a_{\mathcal{G}(s)\setminus i}, \omega} \pi(m_{\mathcal{L}(s)}, \omega) \prod_{k\in\mathcal{G}(s)\setminus i} \beta_k^g(a_k|s_{-k}) u_i(a_i, (m_j(s_{-j}))_{j\in\mathcal{L}(s)}, a_{\mathcal{G}(s)\setminus i}, \omega)$$
$$\geq \sum_{m_{\mathcal{L}(s)}, a_{\mathcal{G}(s)\setminus i}, \omega} \pi(m_{\mathcal{L}(s)}, \omega) \prod_{k\in\mathcal{G}(s)\setminus i} \beta_k^g(a_k|s_{-k}) u_i(a_i', (m_j(s_{-j}))_{j\in\mathcal{L}(s)}, a_{\mathcal{G}(s)\setminus i}, \omega) \tag{6}$$

4. (**Obedience for agent i at the Look-Ignore stage**) for every $i \in I$, s_i such that $\gamma_i(s_i) > 0$, and $s_i' \in S_i$

$$\sum_{s_{-i}, m_{\mathcal{L}(s)}, a_{\mathcal{G}(s)}, \omega} \prod_{j\neq i} \gamma_j(s_j) \pi(m_{\mathcal{L}(s)}, \omega) \prod_{k\in\mathcal{G}(s)} \beta_k^g(a_k|s_{-k}) u_i((m_j(s_{-j}))_{j\in\mathcal{L}(s)}, a_{\mathcal{G}(s)}, \omega)$$
$$\geq \sum_{s'_{-i}, m_{\mathcal{L}(s')}, a_{\mathcal{G}(s')}, \omega} \prod_{j\neq i} \gamma_j(s_j') \pi(m_{\mathcal{L}(s')}, \omega) \prod_{k\in\mathcal{G}(s')} \beta_k^g(a_k|s'_{-k}) u_i((m_j(s'_{-j}))_{j\in\mathcal{L}(s')}, a_{\mathcal{G}(s')}, \omega) \tag{7}$$

where $s \equiv (s_i, s_{-i})$ and $s' \equiv (s_i', s_{-i}')$.

A few features of this equilibrium concept are worth emphasizing. First, notice that the concept requires obedience of the action recommendations even if an agent chooses to disobey his Look-Ignore recommendation ((5) and (6) above). This reflects the requirement that after any Look-Ignore choice behavior be sequentially rational. This in turn implies that the harshest punishment that can be inflicted upon an agent who is recommended to Look but chooses to Ignore (or vice versa) is the payoff from the worst BNE for that agent in the subsequent Bayesian game. Second, notice that the distribution π includes a whole vector of action recommendations for each agent, one for each possible outcome of the other agents' Look-Ignore choices, and only one of those elements of the recommendation vector is relevant for the outcome of the Look-Ignore stage that has actually realized. Finally, because π captures the recommendations that agents who choose to Look will receive, it gives (upon conditioning on the state) the canonical information structure that the designer can choose in order to implement the IPBCE outcome in a PBE*. We further elaborate on this below.

Definition 5. Given an IPBCE (γ, β^g, π), let $v(\gamma, \beta^g, \pi) \in \Delta(A \times \Omega)$ defined as

$$v(\gamma, \beta^g, \pi)(a, \omega) := \sum_{s\in S} \prod_{i\in I} \gamma_i(s_i) \left(\sum_{m_{\mathcal{L}(s)} : (m_j(s_{-j}))_{j\in\mathcal{L}(s)} = a_{\mathcal{L}(s)}} \pi\left(m_{\mathcal{L}(s)}, \omega\right) \right) \prod_{k\in\mathcal{G}(s)} \beta_k^g(a_k|s_{-k})$$

for all $a \in A$ and $\omega \in \Omega$, denote the resulting IPBCE outcome. Let $IPBCE(G^*)$ denote the set of IPBCE outcomes for a game G^*.

That definition highlights a difference between BCEs and IPBCEs. In our setting, where agents have no prior information about the state, a BCE can be defined directly in terms of the outcome distribution it induces, while an IPBCE induces an outcome distribution but also contains additional information (about possibly off-path behavior, for example). Two different IPBCEs may induce the same outcome.

With these definitions in hand, we now state our characterization result.

Theorem 1. $\cup_{(T,P)} PBE^*(G^*(T,P)) = IPBCE(G^*)$.

The equivalence of the two sets implies that the designer can maximize the expected value of her payoff function over the set of IPBCE outcome distributions. Once she finds the optimal IPBCE outcome and a corresponding IPBCE, she can implement the same outcome by providing the information structure obtained by conditioning the π from the IPBCE on ω; then there exists a PBE* where agents make the same Look-Ignore and action choices as in the IPBCE.

Indeed, an implication of that equivalence result is that without loss of generality we can restrict the designer to selecting a *direct contingent* information structure.[12] In a direct contingent information structure, each signal realization for agent i corresponds to a list of recommended actions, one for each combination of Look-Ignore choices by the other $N-1$ agents. That is, in a direct contingent information structure $T_i = \mathscr{A}_i$ for each agent i. Say that an outcome $v \in \Delta(A \times \Omega)$ is *implementable with direct contingent recommendations* if there exists a conditional message distribution $P : \Omega \to \Delta(\mathscr{A})$ such that v is a PBE* outcome of $G^*(\mathscr{A}, P)$.

Corollary 1. *An outcome v is a PBE* outcome if and only if it is implementable with direct contingent recommendations.*

In the rest of the paper, we exploit Theorem 1 and its corollary to characterize the solution to the designer's problem: we derive the optimal information structure $(\mathscr{A}, P^*)$ directly from the optimal IPBCE. This is conceptually equivalent to the designer optimizing over direct contingent information structures and choosing $(\mathscr{A}, P^*)$, and then nature optimizing on the designer's behalf over the set of PBE* of $G^*(\mathscr{A}, P^*)$. Hence, the advantageous equilibrium selection over outcomes under the PBE* solution concept allows for direct optimization over outcomes under the IPBCE solution concept.

2.2. Properties of the IPBCE set

Non-Convexity. Unlike the BCE set, the IPBCE set is not necessarily convex. The reason is the designer's inability to correlate the actions of players who Ignore their signals with the actions of other players. That inability means that obedience constraints cannot be pooled across two IPBCEs, while this is possible across two BCEs. The non-convexity extends to the set of IPBCE outcome distributions as well and drives our finding that the designer's optimal information structure in a symmetric environment may be asymmetric.

Non-Monotonicity. In the standard information design environment, the set of implementable outcomes – the BCE set – is decreasing in the amount of exogenous information about the state that the players start out with (Bergemann and Morris, 2016). The intuition is simple: the designer has the option to provide additional information if it would be useful, but too much information may interfere with obedience. Under strategic ignorance, however, this monotonicity no longer holds. Giving the players exogenous information removes the Look constraint on information that the designer might have wanted to provide anyway.

Our analysis of the investment game of Section 4.2 clearly demonstrates both of these properties of the IPBCE set, and we will further elaborate on them there.

3. Properties of robust information design

The characterization above would be much simpler if we focused only on equilibria where all agents choose to Look at their private messages with probability one, that is $\gamma_i(\ell) = 1$ for all $i \in \mathcal{I}$. Surprisingly, that restriction turns out not to be innocuous, as we show in Theorem 3 below.[13] First, though, we present a condition from the literature that ensures that the canonical information structure implementing the designer's optimal outcome in the standard information design setting will be robust to strategic ignorance and implementable in an equilibrium where all agents Look.

A key tension for the designer is whether or not, if an agent i chooses to Ignore his message, the other agent(s) are still willing to follow their original recommendations even though their expectations of agent i's action have changed. If so, then agent i cannot gain from the deviation. It follows that if a BCE is "1-fault tolerant" (Eliaz, 2002), or equivalently, if it has an "index of resilience" of 1 or higher in the language of Kalai and Kalai (2024), then it is robust to strategic ignorance. Such BCEs have the property that for any deviation by a single player to an action that does not condition on his recommendation, the recommendations for the other $N-1$ players are still obedient. (Formally those authors define their robustness concepts only for Nash equilibria, but the definitions extend straightforwardly to BCEs, as specified below.)

Definition 6. An outcome $v \in \Delta(A \times \Omega)$ is a 1-fault tolerant BCE of G if

1. (**Consistency with the prior**) $v(A \times \{\omega\}) = \mu(\omega)$ for all $\omega \in \Omega$,
2. (**Obedience**) for every $i \in \mathcal{I}$
$$\sum_{a_{-i},\omega} v(a_i, a_{-i}, \omega) u_i(a_i, a_{-i}, \omega) \geq \sum_{a_{-i},\omega} v(a_i, a_{-i}, \omega) u_i(a_i', a_{-i}, \omega) \tag{8}$$
for all $a_i, a_i' \in A_i$, and
3. (**1-fault obedience**) for every $i, k \in \mathcal{I}$ with $i \neq k$
$$\sum_{a_{-i},\omega} v(a_i, a_{-i}, \omega) u_i(a_i, a_k', a_{-ik}, \omega) \geq \sum_{a_{-i},\omega} v(a_i, a_{-i}, \omega) u_i(a_i', a_k', a_{-ik}, \omega) \tag{9}$$
for all $a_i, a_i' \in A_i$ and $a_k' \in A_k$.

Theorem 2. *Any 1-fault tolerant BCE of G is robust to strategic ignorance and can be implemented through simple direct action recommendations in a PBE* of G^* where all agents choose to Look with probability one.*

[12] Proposition 1 in Section 4.1 demonstrates that, in contrast to the standard information design environment, single action recommendations are not sufficient to obtain all implementable outcomes.

[13] Relatedly, we can without loss of generality disregard equilibria in which any agent plays Ignore with certainty, as this is simply equivalent to the designer choosing a completely uninformative message for that agent and the agent choosing to Look with certainty.

3.1. The necessity of ignorance

When the baseline solution is not robust to strategic ignorance, restricting attention to equilibria where all agents Look is with loss of generality. In particular, when some agent i needs information in order to implement a desirable (for the designer) outcome, but the only way to deter another agent j from Ignoring his signal is through a continuation equilibrium that exists only when there is common knowledge that player i is *un*informed, then that incompatibility yields randomization by agent i over whether or not to Look as the optimal solution.

Theorem 3. *The designer's optimal PBE* outcome υ may be implementable only if $\gamma_i(g) > 0$ for some $i \in \mathcal{I}$.*

A two-agent example (in Appendix C) formalizes that intuition. There, the optimal solution is a compromise. Sometimes Player 2 Looks at his signal and plays the designer's desired action, while Player 1 is incentivized to Look by the possibility that Player 2 may Ignore his signal and then be willing to punish Player 1 harshly for deviating. Hence, the designer does strictly better by relying on an equilibrium where one agent (Player 2) randomizes between Look and Ignore instead of a pure Look equilibrium. The essential point is that common knowledge of Player 2 being uninformed cannot be replicated through an information structure which leaves Player 2 uninformed with the right probability. Upon choosing to Ignore this information structure, Player 1 will have no way of knowing whether Player 2 has received the uninformative signal or not, and hence cannot be effectively punished.

3.2. The harm of ignorance

There are many examples from game theory where, in equilibrium, flexibility harms a player. Our next result establishes that this negative effect can arise from the ability of agents to Ignore their messages.

Theorem 4. *The payoffs to all players from any designer-optimal IPBCE may be strictly lower than their payoffs from any designer-optimal BCE, even when the designer's objective is to maximize the sum of players' payoffs.*

We know that the designer is always (weakly) worse off under strategic ignorance due to the added incentive constraints. In the case where the preferences of the designer and the agents are completely aligned, the players collectively may be strictly worse off when one of them has a temptation to Ignore his signal under the baseline optimal information structure. Appendix D gives a two-player, symmetric example where not only is the sum of players' payoffs lower under strategic ignorance, but in fact both players get strictly lower payoffs in any optimal IPBCE. In other examples, the potential for strategic ignorance harms the players only indirectly, by leading the designer to adjust the information structure in a way that benefits her but is detrimental to the players. That is, the result that strategic ignorance may be harmful does not rely on the presence or absence of a designer with a particular objective.

4. Non-robust properties of standard information design

In this section we demonstrate that some standard information design results are not robust to the presence of strategic ignorance. Propositions 1 and 2 demonstrate that in order to achieve robustness to strategic ignorance, the designer may need to provide information that is unnecessary or even counterproductive for the purpose of her objective maximization, from the point of view of standard information design. It is also possible that the designer may need to provide *less* information than under the baseline, and the direction of the optimal adjustment (toward more or less information) can depend on fine details of the environment. The third result, Proposition 3, shows that asymmetric information structures may be uniquely optimal in completely symmetric environments, which is never the case in the standard information design setting.

Next, we consider two modifications of our environment. In the first one, the designer can give only simple (non-contingent) action recommendations, but she can do that in multiple rounds, once before the Look-Ignore choices have been made and then conditional on the specific Look-Ignore profile that has realized. The second modification allows for strategic communication between the players.

4.1. Redundant and counterproductive information provision

When the canonical information structure implementing the designer's optimal outcome in the standard information design is not robust to strategic ignorance, the designer has a variety of ways to adjust the information structure in response. And as we demonstrate next, the optimal response may be to provide more information about the state to the players. In some instances, this extra information is not used on path and would be considered redundant from the perspective of standard information design (Proposition 1), while in other instances the extra information is used on path and would be considered counterproductive in the absence of strategic ignorance (Proposition 2).

Proposition 1. *Direct simple action recommendations are not sufficient for optimal design under strategic ignorance: all optimal information structures under strategic ignorance may require information about the state that a player does not use on path.*

The designer will provide redundant information whenever the following combination of circumstances holds: punishing a player for deviating to Ignore requires more information than is used on-path, and that additional information does not interfere with the on-path obedience constraints. The proof (in Appendix E) is in the form of a two-player, two-state example.

In this example, there is no tension between giving the players the information that they need on path, and giving them the information that they need to punish a deviation. Fully revealing the state works for both situations, even though Player 2 does not

	$K_W = N-1$, $K_S > 0$	$K_W < N-1$, $K_S = 0$	$K_W = N-1$, $K_S = 0$	$K_W < N-1$, $K_S > 0$
k	$\frac{x+2}{2}$	$\frac{1}{2}$	$\frac{x-1}{2}$	2
n	1	0	$\frac{1}{2}$	$\frac{1}{2}$

$$\Pr(\omega = W) = \Pr(\omega = S) = \frac{1}{2}$$

Fig. 3. k is dominant for an uninformed player.

need any information on path. However, if the designer does not reveal the state to Player 2, Player 1's Look constraint cannot be satisfied as he cannot be effectively punished for deviating to Ignore: the unique BNE when neither player knows the state gives Player 1 a high payoff. Therefore, sending a simple action recommendation of how to play on-path does not give Player 2 the necessary information about the state because he plays the same action with probability one on path in both states.

Next, we show that the designer may need to optimally provide an information structure that is more informative about the state for all players than would be optimal in the absence of strategic ignorance, and this information is used on path by the players. In the baseline environment without strategic ignorance, this would be considered counterproductive information provision – it results in a lower payoff for the designer. The designer will provide counterproductive information when punishing a player for deviating to Ignore requires more information than is used on-path to achieve any desired outcome, and that additional information *does* interfere with the on-path obedience constraints.

Proposition 2. *All optimal information structures under strategic ignorance may give strictly more information about the state to all $i \in \mathcal{I}$ than the optimal information structure in the absence of strategic ignorance, and all $i \in \mathcal{I}$ may use that additional information on path to the designer's detriment.*

Proof. This result is demonstrated in the following model of currency attacks. There are $N \geq 2$ symmetric players deciding whether or not to attack a currency ($(attac)k$ or $n(ot)$). The currency may be either weak or strong with equal probability. If the currency is weak, then one player is enough for a successful attack, and so attacking is strictly dominant. If the currency is strong, then the attack succeeds if and only if at least two players attack. We capture that setting with the following payoff function, where player i's payoff depends on the state $\omega \in \{W(eak), S(trong)\}$, his own action, and the number K of other players who play k:

$$u_i(k, K; W) = \begin{cases} 2 & \text{if } K < N-1 \\ x & \text{if } K = N-1 \end{cases}, u_i(k, K; S) = \begin{cases} -1 & \text{if } K = 0 \\ 2 & \text{if } K > 0 \end{cases},$$

and

$$u_i(n, K; W) = \begin{cases} 0 & \text{if } K < N-1 \\ 1 & \text{if } K = N-1 \end{cases}, u_i(n, K; S) = \begin{cases} 0 & \text{if } K = 0 \\ 1 & \text{if } K > 0 \end{cases}$$

We assume $x \geq 2$, so that when the currency is weak the payoff from attacking when all other players attack is at least as high as when some of them do not.

The designer wants to prevent a successful attack: she gets a payoff of 1 if $(n, \ldots, n)$ is played in state W, or if at least $N-1$ players play n in state S, and she gets 0 otherwise.

Baseline. At the prior, k is dominant, so the designer must provide the players some information in order to get a positive payoff. The best that she can do is to publicly recommend n for sure in state S, and to publicly recommend n in state W as often as possible subject to the players' attaching a high enough probability to $\omega = S$ after recommendation n for n to be obedient. Formally, the designer's optimal information structure with *single* action recommendations, $(A, \tilde{P})$, is

$$\tilde{P}((k, \ldots, k)|\omega = W) = \tilde{P}((n, \ldots, n)|\omega = W) = \frac{1}{2}, \tilde{P}((n, \ldots, n)|\omega = S) = 1.$$

The obedience constraint binds for a player who gets recommendation n: the updated probability of state S is $0.5/(0.5 + 0.25) = 2/3$, so both k and n yield expected payoff 0 given that the other players will choose n.

The designer's payoff is $\frac{3}{4}$, and the players' payoff is $\frac{1}{4}x$.

With strategic ignorance. If the players can publicly Ignore their signals, then under that baseline information structure $(A, \tilde{P})$ there is no equilibrium in which all players Look at their recommendations and follow them.

First consider the case that $x > 2$, and suppose that Player i deviates to Ignore. When Player i is uninformed, then k is dominant at the action stage: as shown in Fig. 3, it gives a strictly higher payoff than n does against any strategy profile mapping states to actions for the other $N-1$ players. We can summarize a strategy profile for the other players as $\left(K_W, K_S\right)$ denoting the number who play k in each state.

In either state, the unique best response for any other player when Player i chooses k is k. The outcome is thus $(k, \ldots, k)$ regardless of the designer's recommendations, and Player i's resulting payoff is $\frac{1}{2} \cdot x + \frac{1}{2} \cdot 2 > \frac{1}{4}x$. It follows that deviating to Ignore is profitable.

In fact, when $x > 2$ the designer cannot achieve any outcome other than $(k, \dots, k)$ regardless of the realized state, by the same reasoning. That action profile gives the players their maximum possible payoff in either state, and under any information structure they can achieve it in a BNE by deviating to Ignore. Requiring robustness to strategic ignorance completely undoes the designer's ability to use information design to her advantage.[14]

If $x = 2$, then the situation changes. From Fig. 3, we see that now an uninformed Player i's expected payoff from playing k against a strategy of (k in state W, n in state S) by each other player (that is, $K_W = N - 1, K_S = 0$) is $\frac{1}{2}$; both k and n are best responses.

The information structure $(A, \tilde{P})$ still does not work: a player's message gives him only partial information about the state, and so he cannot play strategy (k in state W, n in state S). An uninformed Player i's unique best response to anything other than the strategy profile of (k in state W, n in state S) for all opponents is k, and the rest of the argument is the same as in the $x > 2$ case.

In contrast to the $x > 2$ case, though, now the designer can achieve a positive payoff. In particular, if a player's message perfectly reveals the state, then (k in state W, n in state S) becomes a feasible strategy. Hence, consider the direct contingent information structure $(\mathscr{A}, P^*)$, which recommends action k to every player after every profile of others' Look-Ignore choices with probability 1 in state W, and action n to every player after every profile of others' Look-Ignore choices with probability 1 in state S.

Under $(\mathscr{A}, P^*)$, it is an equilibrium for all players to Look at and follow their recommendations, yielding payoff $\frac{x}{2} = 1$. If Player i deviates to Ignore, then there is a BNE where he plays n and all other players follow their recommendation by playing (k in state W, n in state S). That BNE gives Player i a payoff of $\frac{1}{2} < 1$, so the deviation to Ignore is not profitable. The designer's payoff is $\frac{1}{2}$.

In contrast, any information structure that does not fully reveal the state to all players gives the designer a payoff of 0. The outcome must be $(k, \dots, k)$ because any player whose opponents are not fully informed can achieve that outcome by choosing to Ignore. Once the state is perfectly revealed to all players, they use that information on path by always playing the unique equilibrium $(k, \dots, k)$ in state W. □

That analysis shows that the pattern of the designer's response to the possibility of strategic ignorance is very sensitive to the details of the environment. At the baseline, under $(A, \tilde{P})$ the designer optimally provides partial information about the state. With strategic ignorance, for any value of the payoff x that is strictly greater than 2, the designer optimally provides no information, while at $x = 2$ she reveals the state perfectly.

Broadly speaking, the logic of Proposition 2 shows that under strategic ignorance the designer tends to provide more information than in the standard information design environment when the players need additional information in order to punish deviations to Ignore. Conversely, as the analysis behind Proposition 3 in the next section demonstrates, the designer may optimally provide less information. This occurs when, under the baseline optimal information structure, a player's deviation to Ignore changes how another player uses the information contained in his message. The other player's resulting best response may then raise the deviator's expected payoff. By coarsening the information she provides, the designer eliminates the gain from such a deviation.

4.2. Strict optimality of asymmetric information structures

In the standard information design environment without strategic ignorance, if the basic game and the designer's objective are player- and state-symmetric, then there is always a player- and state-symmetric optimal information structure. That result is implied by the convexity of the BCE set: if there is an optimal asymmetric BCE outcome distribution, then the mirror image of that distribution is also an optimal BCE distribution, and so is the equally weighted convex combination of these two outcome distributions. That convex combination corresponds to a player- and state-symmetric direct information structure. However, when players can exercise strategic ignorance, asymmetric information structures can be *strictly* optimal even in completely symmetric environments due to the non-convexity of the IPBCE set. The proof of this result is presented in a version of the parameterized basic game from Taneva (2019), which also showcases the non-convexity of the IPBCE set.

Proposition 3. *Asymmetric information structures can be strictly optimal in completely symmetric environments when players can exercise strategic ignorance.*

Proof. There are two symmetric firms seeking to coordinate on one of two possible projects. Which project has the potential to succeed depends on a binary unknown state of the world and we assume each state is equally likely. The profitability of a successful project increases with the total investment, so choosing the right project yields a higher payoff if the other firm invests in it as well. We capture that setting in the payoff matrices in Fig. 4.

The designer wants the project to fail. In particular, she gets a payoff of 1 if (F, F) is played in state e or (E, E) is played in state f, and 0 otherwise.

Baseline. In the baseline information design environment, where agents automatically observe their private signals from the designer, we can rely on the analysis in Taneva (2019). Without loss of generality, we maximize over the set of symmetric[15] BCE outcome

[14] This example illustrates the distinction between equilibrium selection and strategic ignorance. "Always play $(k, \dots, k)$" is a BNE under the baseline information structure $(A, \tilde{P})$, but under advantageous selection we assume that instead the agents play the designer's preferred BNE. In contrast, if a player deviates to Ignore, then the *unique* BNE under the resulting information structure is "always play $(k, \dots, k)$". Hence, there is no equilibrium where anything other than "always play $(k, \dots, k)$" is played on path. We are still selecting the designer's preferred equilibrium of the dynamic game, but there is only one outcome to select from.

[15] This is due to the convexity of the BCE set, as argued previously, and due to the linearity of the designer's expected payoff in the probabilities. Additionally, note that Fig. 5 represents distributions over action profiles conditional on each state; the unconditional outcome distributions are obtained by multiplying these with the prior.

	E	F
E	$2,2$	$1,0$
F	$0,1$	$0,0$

$\omega = e$

	E	F
E	$0,0$	$0,1$
F	$1,0$	$2,2$

$\omega = f$

Fig. 4. Investment game.

	E	F
E	r	$q-r$
F	$q-r$	$1-2q+r$

$\omega = e$

	E	F
E	$1-2q+r$	$q-r$
F	$q-r$	r

$\omega = f$

Fig. 5. Parameterized symmetric BCE outcome distributions.

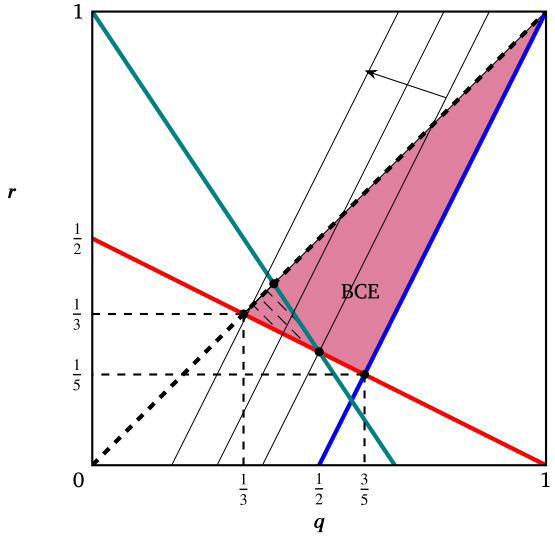


Fig. 6. BCE set and designer's expected payoff

Notes: Red line is the single obedience constraint. Blue line is a feasibility constraint on the probability distribution ($1-2q+r \geq 0$). Dashed 45-degree line is a feasibility constraint on the probability distribution ($q \geq r$). BCE set is the triangle in solid purple. Green line represents the Look obedience constraint for direct symmetric information structures. Hatched triangle is the set of direct symmetric information structures that are Ignored. Solid gray parallel lines are the level lines for the designer's expected payoff.

distributions represented in Fig. 5, where q captures the probability that each player receives the action recommendation that corresponds to the state (i.e., the total probability of being recommended action E in state e or action F in state f for each player) and r captures the probability that both players together receive the action recommendation that corresponds to the state (i.e., the probability of (E,E) in state e or (F,F) in state f). Naturally, r counts toward q, and so we must have $q \geq r$.

The set of BCE outcome distributions is then the triangle in solid purple in Fig. 6. The red line represents the obedience constraint[16], while the 45-degree line and the blue line represent the constraints on the parameters that ensure the outcome is a proper probability distribution (namely, $q \geq r$ and $1-2q+r \geq 0$). The level lines for the designer's expected payoff are given by the solid gray parallel lines, with increasing levels as they shift to the left and up.

The optimum BCE outcome distribution is at the leftmost point of the BCE set $\tilde{r} = \tilde{q} = \frac{1}{3}$. This point corresponds to the optimal direct information structure of the designer $(A, \tilde{P})$ given by:

$$\tilde{P}(E,E|\omega = e) = \tilde{P}(F,F|\omega = f) = \tilde{r} = \frac{1}{3},$$

$$\tilde{P}(F,F|\omega = e) = \tilde{P}(E,E|\omega = f) = 1 - 2\tilde{q} + \tilde{r} = \frac{2}{3}.$$

[16] There are two obedience constraints – one for action E and one for action F – but they reduce to the same inequality due to the symmetry.

$P^*(\cdot\|e)$	EF	FE	FF
EF	$\frac{1-\alpha}{2} \approx 0.2115$	0	0
FE	0	$\alpha \approx 0.577$	0
FF	0	0	$\frac{1-\alpha}{2} \approx 0.2115$

$P^*(\cdot\|f)$	EF	FE	FF
EF	$1-\alpha \approx 0.423$	0	0
FE	0	$1-\alpha \approx 0.423$	0
FF	0	0	$2\alpha - 1 \approx 0.154$

Fig. 7. Optimal direct contingent information structure.

The designer's payoff is $\frac{2}{3}$, and each firm's payoff is $2 \cdot \frac{1}{3} = \frac{2}{3}$.

Under $(A, \tilde{P})$, the designer sends a public signal. She exploits the firms' desire to coordinate their investment by recommending the "correct" project with probability $\frac{1}{3} < 0.5$. Each firm is just willing to obey the recommendation given that the other firm will. Switching to the other project means matching the state with higher probability but mismatching the other firm: obedience yields 2 with probability $\frac{1}{3}$, and switching yields 1 with probability $\frac{2}{3}$.[17]

With strategic ignorance. If the firms can publicly Ignore their signals, then that baseline information structure $(A, \tilde{P})$ will not lead to the designer's desired outcome. There is no equilibrium in which both firms Look at their signals and then follow their recommendations. To see why not, suppose that Firm 1 chooses to Ignore his signal while Firm 2 Looks at his. The worst BNE for Firm 1 under the resulting information structure involves Firm 1 randomizing uniformly between E and F. Firm 2's best response is to choose the opposite of the project that the designer recommended: now that Firm 1 cannot coordinate by following the designer's recommendation, Firm 2 just wants to pick the project that is more likely to succeed. Under $(A, \tilde{P})$, the project that the designer recommends is more likely to be the wrong one, so Firm 2 will pick the other project.

In that BNE, Firm 1 gets an expected payoff of $\frac{1}{2} \cdot \left(\frac{2}{3} \cdot 2 + \frac{1}{3} \cdot 1\right) = \frac{5}{6}$, which is strictly greater than the payoff $\frac{2}{3}$ from playing the designer's preferred BNE under $(A, \tilde{P})$. Thus, Firm 1 gained by choosing to Ignore his signal. By deviating to Ignore, Firm 1 forgoes the chance to coordinate perfectly with Firm 2. But because Firm 2 will now choose the correct project more frequently, Firm 1 has increased the probability of choosing correctly *conditional* on matching the other firm. The complementarity in payoffs means that at $(A, \tilde{P})$ that tradeoff is beneficial.

The green line in Fig. 6 represents the Look obedience constraint for any direct symmetric information structure. Above this constraint the deviation to Ignore is no longer beneficial. The hatched triangle thus depicts the set of direct symmetric information structures that will be Ignored.[18]

In order to satisfy the constraint that firms be willing to Look at their signals[19], the designer's optimal adjustment involves reducing the probability that Firm 2 will choose the correct project if Firm 1 deviates to Ignore. One component is to lower the frequency of recommending the wrong project from $\frac{2}{3}$. The second component is to introduce asymmetry between the states: the designer is less likely to recommend the wrong action in state f than in state e. That change creates a post-Ignore BNE worse than the one where Firm 1 randomizes uniformly and Firm 2 chooses the project matching the more likely state. Instead, Firm 1 puts higher probability on F, and in order to try to coordinate with him Firm 2 is willing to choose F even after the designer recommends F on path (meaning that state e is more likely). Overall, the reduction in the probability of coordinating with Firm 2 *conditional* on matching the state leaves Firm 1 worse off after deviating to Ignore. That effect, combined with the fact that the designer recommends the correct action more frequently on path, makes Firm 1 willing to Look.

Specifically, we calculate that an optimal IPBCE (γ, β^g, π) corresponds to the asymmetric direct contingent information structure $(\mathscr{A}, P^*)$ in Fig. 7, where $\alpha \equiv \frac{1}{\sqrt{3}} \approx 0.577$, and no symmetric information structure can achieve the same payoff for the designer. At $(\mathscr{A}, P^*)$ both firms Look with probability 1 (that is, $\gamma_i(\ell) = 1$ for all i), and a firm that deviates to Ignore randomizes with probability $\beta_i^g(F|s_{-i} = \ell) = \sqrt{3} - 1 \approx 0.732$ on project F at the action stage. The designer's payoff equals the probability that she recommends the wrong action,

$$\mathbb{E}(u^D) = \frac{1}{2}\left(\frac{1-\alpha}{2} + \alpha\right) + \frac{1}{2}(1-\alpha) \approx 0.606,$$

and the firms' payoff is 2 times the probability of a correct recommendation:

$$\mathbb{E}(u) = \frac{1}{2}2\left(\frac{1-\alpha}{2}\right) + \frac{1}{2}2(\alpha) \approx 0.789.$$

Deviating to Ignore would yield the same payoff, so the Look constraint is satisfied with equality.

□

Non-Convexity. We note that the state-wise mirror image of the optimal IPBCE above, denoted by $(\gamma, \bar{\beta}^g, \bar{\pi})$, is also an optimal IPBCE. An equally weighted convex combination of those two, however, is not. Indeed, while the BCE set is always convex, the IPBCE set may not be, and in this example it is not. In particular, after Firm 1 deviates to Ignore, then the strategy $\frac{1}{2}\beta^g + \frac{1}{2}\bar{\beta}^g$ calls for him to

[17] We note the role of advantageous equilibrium selection here. Under the baseline optimal information structure $(A, \tilde{P})$, there is also a BNE where the firms do the opposite of their recommendations, and that BNE gives them higher payoffs. In fact, the designer's preferred BNE gives them payoffs below those of the worst BNE (randomizing uniformly between projects) in the basic game without a designer.

[18] The derivations are in Appendix F.

[19] Numerical estimation showed that the optimal IPBCE does not require mixing between Look and Ignore.

randomize uniformly between E and F. But then Firm 2's best response is to pick the project that is more likely to succeed rather than to follow the designer's recommendation, so the obedience constraint fails. Hence, this convex combination of the two optimal IPBCEs is not an IPBCE itself. This non-convexity extends to the outcome sets: the equally weighted convex combination of the two IPBCE outcome distributions is not an IPBCE outcome distribution.

Non-Monotonicity with Respect to Exogenous Information. In the standard information design environment, the set of implementable outcomes decreases in the amount of exogenous information about the state that the players start out with (Bergemann and Morris, 2016). Our analysis of the investment game, however, demonstrates that this monotonicity may fail under strategic ignorance. In particular, consider the designer's optimal baseline outcome distribution, implemented by $(A, \tilde{P})$. When the players have no exogenous information, this outcome is not implementable under strategic ignorance. However, if they start out with exactly the exogenous information structure implied by $(A, \tilde{P})$ (a public signal of the state that is "correct" with probability $\frac{1}{3}$), then the designer could simply reveal nothing further and there will be a BNE that implements her baseline optimal outcome distribution. Hence, by increasing the exogenous information of the players in this way, an outcome that was previously not implementable in the presence of strategic ignorance becomes implementable. If we further increase the exogenous information of the agents to fully reveal the state, then the baseline optimal distribution becomes again not implementable. Thus, the set of implementable outcome distributions under strategic ignorance is non-monotone in the players' exogenous information.

4.3. Recommendations contingent on Look-Ignore choices

We know that in general the designer is hurt by the fact that players see their recommendations for all possible combinations of Look-Ignore choices at once, because the information about the state contained in off-path recommendations may interfere with the obedience constraint for the on-path recommendation. This effect suggests that the designer would benefit from the ability to instead provide only on-path recommendations initially, and then send additional recommendations only if some player deviates at the Look-Ignore stage. However, in the example from the proof of Proposition 1 we can show that the designer does *strictly* better by giving both recommendations (one for when the other player chooses Look and one for when he chooses Ignore) at once. More specifically, providing the designer with the option to give Player 2 more information after Player 1 chooses Ignore does not help.

Proposition 4. *Giving recommendations in multiple rounds rather than all at once may be strictly worse for the designer when players can exercise strategic ignorance.*

Proof. Consider the example from the proof of Proposition 1 given by the state-contingent payoff matrices in Fig. E.1. In that example, at the baseline without strategic ignorance, the designer's optimal direct information structure $(A, \tilde{P})$ is

$$\tilde{P}(E, Y|\omega = e) = \tilde{P}(F, Y|\omega = f) = 1.$$

Recall from the proof of Proposition 1 that the designer's optimal outcome is robust to strategic ignorance if she can provide the direct contingent recommendations all at once. Specifically, the outcome can be implemented by providing the following direct contingent information structure $(\mathscr{A}, P^*)$:

$$P^*(EE, YY|\omega = e) = P^*(FF, YX|\omega = f) = 1.$$

Now both players' messages reveal the state.

Giving punishment recommendations only after a deviation to Ignore. Here, we have in mind the following modification of our main setup: the designer first gives on-path recommendations only. If both players choose Look, then the designer sends no further messages. If Player i deviates to Ignore, then the designer sends a second message to Player j, with a recommendation for what to play now. Player j then chooses whether or not to Look at that second message.

In this setting, the designer cannot implement the desired outcome distribution by first providing the on-path recommendations given by the information structure $(A, \tilde{P})$ and providing the second message only after the other player has chosen to Ignore. Suppose that Player 1 deviates to Ignore. Now Player 2 can decide whether to Look at the subsequent recommendation (Y in state e and X in state f) or Ignore it. If Player 2 chooses to Look at the second recommendation, then the expected payoffs in the continuation BNE, where (E, Y) is played in state e and (E, X) is played in state f, are $\left(1, \frac{1}{2}\right)$. If Player 2 chooses to Ignore the second recommendation, then both players' beliefs equal the prior, and in that case action F is strictly dominant for Player 1, to which Player 2's unique best response is X. Thus, the outcome will be (F, X) with expected payoffs $\left(\frac{3}{2}, \frac{3}{2}\right)$. Therefore, conditional on Player 1 choosing to Ignore the initial message, it is the unique best response for Player 2 to Ignore the second message given the continuation BNEs. Consequently, Player 1 will optimally Ignore his first message and get a payoff of $\frac{3}{2}$, instead of choosing to Look at it and get a payoff of 1. Hence, (F, X) is played in both states, which results in a payoff of 0 for the designer. We conclude that the designer cannot do as well here as she did in the previous section with direct *contingent* recommendation, where she got a payoff of 1. □

Interpretation. If the designer gives both the on-path and the punishment recommendations at once, then Player 2's Look constraint is satisfied. He expects that Player 1 will choose Look, and so Player 2 is indifferent between Look and Ignore. If there were any positive probability that Player 1 might choose Ignore, then Player 2 would strictly prefer Ignore. But because that probability is zero, the designer effectively gets the "punishment Look constraint" of Player 2 for free.

On the other hand, giving just on-path recommendations to start does not reveal the state to Player 2, but he needs to know it in order to punish Player 1: (F, X) played in both states is the unique BNE when neither player knows the state. Once Player 1 has

deviated to Ignore, we can no longer satisfy Player 2's second Look constraint to get him to learn the state and punish Player 1. Hence, in this example, giving punishment information to a player only after his opponent has deviated to Ignore means having to satisfy a second Look constraint, and that effect makes the designer worse off.

4.4. Communication between players

Finally, we introduce the possibility of players strategically communicating with each other. In the standard information design environment without strategic ignorance, the possibility of players communicating their private signals with each other can never be strictly beneficial for the designer: once the designer has provided the optimal information structure, any change resulting from communication between the players must weakly lower the designer's expected payoff. However, in the presence of strategic ignorance, the designer can leverage the players' incentives for strategic information sharing to her own benefit, in order to relax some of the Look constraints.

Proposition 5. *Strategic communication between players can be strictly beneficial for the designer when players can exercise strategic ignorance.*

The formal proof is in Appendix G. The intuition is as follows: suppose that at the designer's optimal baseline information structure, both players choosing Look is not an equilibrium at the Look-Ignore stage because Player 1 would prefer to deviate to Ignore. However, Player 1 prefers the outcome from the continuation equilibrium after both players have chosen Look to the outcome from the continuation equilibrium after both players have chosen Ignore. Therefore, Player 1 can be incentivized to choose Look by sending an informative signal to him only, after which he would want to pass it on to Player 2. As long as Player 2 also prefers the outcome after both players have chosen Look to the outcome after he chooses Ignore while Player 1 has chosen Look, he would agree to observe the information that Player 1 wants to pass on to him. Essentially, by sending information to Player 1 only, the designer is able to rule out Player 1's deviation to the outcome where Player 1 chooses Ignore while Player 2 chooses Look.

Coded messages. We can build on that reasoning to argue that when players can communicate, the designer may do better than using direct contingent action recommendations by sending coded messages that are only informative when combined. For example, each player gets a binary signal whose marginal distribution is uniform and independent of the state. The signals are perfectly correlated in state 2 and perfectly negatively correlated in state 1. Thus, seeing one signal gives no information, but knowing whether or not they match perfectly identifies the state. In that way, Player 1 can pass on a signal without knowing the meaning that Player 2 will assign to it. That message structure would be useful in a setting where Player 2 is willing to punish Player 1 effectively only when Player 2 does not know one component of a multidimensional state, but in order to play the designer's desired actions on path Player 2 must know that component and Player 1 must not.

5. Discussion and conclusion

We have shown that the ability of agents to publicly refuse information has important consequences for information design in strategic settings. Requiring robustness to strategic ignorance significantly alters optimal information structures and the ensuing outcomes in leading economic applications. Moreover, it generates new qualitative predictions and undoes standard results from the information design literature. Our findings are also relevant in settings where agents seek to coordinate on what pre-play information to gather: the agreement that maximizes expected payoffs ex ante may not be sustainable.

In future work, we believe that it will be productive to expand our analysis from static (that is, one shot, simultaneous move) games to extensive form games. More specifically, it would be interesting to consider different possible extensive forms of the Look-Ignore stage, either as a choice made by the designer or, alternatively, by the agents. Another relevant extension would be to allow the agents to choose arbitrary garblings of their signals instead of the two extremes of either perfectly observing their signal or remaining completely uninformed. A particularly interesting related topic is the optimal design of monitoring structures in repeated games where players can publicly ignore their signals of each other's actions.

CRediT authorship contribution statement

Ina Taneva: Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization; **Thomas Wiseman:** Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the authors used Claude (Anthropic) to check the proofs and examples, to improve the clarity of parts of the text, and to fix typos and grammar. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Appendix A. Proof of Theorem 1

Proof. First we prove that $IPBCE(G^*) \subseteq \cup_{(T,P)} PBE^*(G^*(T,P))$. Take any $\upsilon(\gamma, \beta^g, \pi) \in IPBCE(G^*)$. Consider the information structure $(\mathscr{A}, P)$ with $P(m|\omega) := \pi(m,\omega)/\mu(\omega)$ for all $m \in \mathscr{A}, \omega \in \Omega$.

Given profile $s \in S$ of Look-Ignore choices, let $\mathscr{A}_{\mathcal{L}(s)} \equiv \times_{i\in\mathcal{L}(s)}\mathscr{A}_i$ and $P_{\mathcal{L}(s)}(m_{\mathcal{L}(s)}|\omega) := \pi(m_{\mathcal{L}(s)},\omega)/\mu(\omega)$. In $G(\mathscr{A}_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$ consider the following strategy for each player $i \in \mathcal{L}(s)$:

$$\tilde{\beta}_i^s(a_i|m_i) = \begin{cases} 1, & \text{if } a_i = m_i(s_{-i}) \\ 0, & \text{if } a_i \neq m_i(s_{-i}), \end{cases}$$

for all $m_i \in \mathscr{A}_i$, and for each player $i \in \mathcal{G}(s)$, consider $\tilde{\beta}_i^s(a_i) = \beta_i^g(a_i|s_{-i})$.

Given any $s \in S$, the interim payoff to agent $i \in \mathcal{L}(s)$ observing message $m_i \in \mathscr{A}_i$ and choosing action $a_i \in A_i$ when his opponents play according to $\tilde{\beta}_{-i}^s$ is given by

$$\begin{aligned} &\sum_{a_{-i}, m_{\mathcal{L}(s)\setminus i}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(m_i, m_{\mathcal{L}(s)\setminus i}|\omega) \prod_{j\in\mathcal{L}(s)\setminus i} \tilde{\beta}_j^s(a_j|m_j) \prod_{k\in\mathcal{G}(s)} \tilde{\beta}_k^s(a_k) u_i(a_i, a_{-i}, \omega) \\ &= \sum_{m_{\mathcal{L}(s)\setminus i}, a_{\mathcal{G}(s)}, \omega} \pi(m_i, m_{\mathcal{L}(s)\setminus i}, \omega) \prod_{k\in\mathcal{G}(s)} \beta_k^g(a_k|s_{-k}) u_i(a_i, (m_j(s_{-j}))_{j\in\mathcal{L}(s)\setminus i}, a_{\mathcal{G}(s)}, \omega). \end{aligned} \tag{A.1}$$

Hence, by (5) we obtain

$$\begin{aligned} &\sum_{a_{-i}, m_{\mathcal{L}(s)\setminus i}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(m_i, m_{\mathcal{L}(s)\setminus i}|\omega) \prod_{j\in\mathcal{L}(s)\setminus i} \tilde{\beta}_j^s(a_j|m_j) \prod_{k\in\mathcal{G}(s)} \tilde{\beta}_k^s(a_k) u_i(m_i(s_{-i}), a_{-i}, \omega) \\ &\geq \sum_{a_{-i}, m_{\mathcal{L}(s)\setminus i}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(m_i, m_{\mathcal{L}(s)\setminus i}|\omega) \prod_{j\in\mathcal{L}(s)\setminus i} \tilde{\beta}_j^s(a_j|m_j) \prod_{k\in\mathcal{G}(s)} \tilde{\beta}_k^s(a_k) u_i(a_i', a_{-i}, \omega) \end{aligned} \tag{A.2}$$

for each $i \in \mathcal{L}(s)$, $m_i \in \mathscr{A}_i$, and $a_i' \in A_i$. This establishes the BNE interim incentive compatibility constraint (1) for each $i \in \mathcal{L}(s)$, $m_i \in \mathscr{A}_i$, and $a_i \in A_i$ such that $\tilde{\beta}_i^s(a_i|m_i) > 0$.

Given any $s \in S$, the interim payoff to agent $i \in \mathcal{G}(s)$ choosing action $a_i \in A_i$ when his opponents play according to $\tilde{\beta}_{-i}^s$ is given by

$$\begin{aligned} &\sum_{a_{-i}, m_{\mathcal{L}(s)}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(m_{\mathcal{L}(s)}|\omega) \prod_{j\in\mathcal{L}(s)} \tilde{\beta}_j^s(a_j|m_j) \prod_{k\in\mathcal{G}(s)\setminus i} \tilde{\beta}_k^s(a_k) u_i(a_i, a_{-i}, \omega) \\ &= \sum_{m_{\mathcal{L}(s)}, a_{\mathcal{G}(s)\setminus i}, \omega} \pi(m_{\mathcal{L}(s)}, \omega) \prod_{k\in\mathcal{G}(s)\setminus i} \beta_k^g(a_k|s_{-k}) u_i(a_i, (m_j(s_{-j}))_{j\in\mathcal{L}(s)}, a_{\mathcal{G}(s)\setminus i}, \omega). \end{aligned} \tag{A.3}$$

Hence, by (6) we obtain

$$\begin{aligned} &\sum_{a_{-i}, m_{\mathcal{L}(s)}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(m_{\mathcal{L}(s)}|\omega) \prod_{j\in\mathcal{L}(s)} \tilde{\beta}_j^s(a_j|m_j) \prod_{k\in\mathcal{G}(s)\setminus i} \tilde{\beta}_k^s(a_k) u_i(a_i, a_{-i}, \omega) \\ &\geq \sum_{a_{-i}, m_{\mathcal{L}(s)}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(m_{\mathcal{L}(s)}|\omega) \prod_{j\in\mathcal{L}(s)} \tilde{\beta}_j^s(a_j|m_j) \prod_{k\in\mathcal{G}(s)\setminus i} \tilde{\beta}_k^s(a_k) u_i(a_i', a_{-i}, \omega) \end{aligned} \tag{A.4}$$

for each $i \in \mathcal{G}(s)$, a_i such that $\beta_i^g(a_i|s_{-i}) > 0$, and $a_i' \in A_i$. This establishes the BNE interim incentive compatibility constraint (2) for each $i \in \mathcal{G}(s)$ and $a_i \in A_i$ with $\tilde{\beta}_i^s(a_i) > 0$.

By Definition 1 we conclude that for all $s \in S$, $\tilde{\beta}^s = (\tilde{\beta}_i^s)_i$ is a BNE of $G(\mathscr{A}_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$. Then $\upsilon(\tilde{\beta}^s)$ defined as

$$\begin{aligned} \upsilon(\tilde{\beta}^s)(a,\omega) &:= \sum_{m_{\mathcal{L}(s)}} \mu(\omega) P_{\mathcal{L}(s)}(m_{\mathcal{L}(s)}|\omega) \left(\prod_{j\in\mathcal{L}(s)} \tilde{\beta}_j^s(a_j|m_j) \prod_{k\in\mathcal{G}(s)} \tilde{\beta}_k^s(a_k) \right) \\ &= \sum_{m_{\mathcal{L}(s)}} \pi(m_{\mathcal{L}(s)}, \omega) \left(\prod_{j\in\mathcal{L}(s)} \tilde{\beta}_j^s(a_j|m_j) \prod_{k\in\mathcal{G}(s)} \tilde{\beta}_k^s(a_k) \right) \end{aligned} \tag{A.5}$$

for all $a \in A$ and $\omega \in \Omega$ is a BNE outcome of $G(\mathscr{A}_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$.

Notice that for each $i \in \mathcal{I}$ and $s_i, s_i' \in S_i$ such that $\gamma_i(s_i) > 0$, (7) can be equivalently written as

$$\begin{aligned} &\sum_{s_{-i}, a, \omega} \prod_{j\neq i} \gamma_j(s_j) \left(\sum_{m_{\mathcal{L}(s)}} \pi(m_{\mathcal{L}(s)}, \omega) \left(\prod_{j\in\mathcal{L}(s)} \tilde{\beta}_j^s(a_j|m_j) \prod_{k\in\mathcal{G}(s)} \tilde{\beta}_k^s(a_k) \right) \right) u_i(a_i, a_{-i}, \omega) \\ &= \sum_{s_{-i}, a, \omega} \prod_{j\neq i} \gamma_j(s_j) \upsilon(\tilde{\beta}^s)(a,\omega) u_i(a_i, a_{-i}, \omega) \\ &\geq \sum_{s_{-i}', a, \omega} \prod_{j\neq i} \gamma_j(s_j') \upsilon(\tilde{\beta}^{s'})(a,\omega) u_i(a_i, a_{-i}, \omega) \end{aligned}$$

$$= \sum_{s'_{-i},a,\omega} \prod_{j\neq i} \gamma_j(s'_j)\left(\sum_{m_{\mathcal{L}(s')}} \pi(m_{\mathcal{L}(s')},\omega)\left(\prod_{j\in\mathcal{L}(s')} \tilde{\beta}_j^{s'}(a_j|m_j) \prod_{k\in\mathcal{G}(s')} \tilde{\beta}_k^{s'}(a_k)\right)\right) u_i(a_i,a_{-i},\omega), \tag{A.6}$$

where $s \equiv (s_i, s_{-i})$ and $s' \equiv (s'_i, s'_{-i})$, which establishes (4).

Hence, $(\gamma, (\tilde{\beta}^s)_s)$ is a PBE* of $G^*(\mathscr{A}, P)$. Then $\hat{v} \in \Delta(A\times\Omega)$ defined as

$$\hat{v}(a,\omega) := \sum_{s\in S} \prod_{i\in\mathcal{I}} \gamma_i(s_i) v(\tilde{\beta}^s)(a,\omega)$$

for all $a \in A$ and $\omega \in \Omega$ is a PBE* outcome of $G^*(\mathscr{A}, P)$, that is $\hat{v} \in PBE^*(G^*(\mathscr{A}, P))$. Notice that for all $a \in A$ and $\omega \in \Omega$

$$\hat{v}(a,\omega) = \sum_{s\in S} \prod_{i\in\mathcal{I}} \gamma_i(s_i) v(\tilde{\beta}^s)(a,\omega) \tag{A.7}$$

$$= \sum_{s\in S} \prod_{i\in\mathcal{I}} \gamma_i(s_i)\left(\sum_{m_{\mathcal{L}(s)}:(m_j(s_{-j}))_{j\in\mathcal{L}(s)}=a_{\mathcal{L}(s)}} \pi\left(m_{\mathcal{L}(s)},\omega\right)\right) \prod_{k\in\mathcal{G}(s)} \beta_k^g(a_k|s_{-k}) = v(\gamma,\beta^g,\pi)(a,\omega). \tag{A.8}$$

Thus, $v(\gamma,\beta^g,\pi) \in PBE^*(G^*(\mathscr{A}, P))$.

Next, we prove that $IPBCE(G^*) \supseteq \cup_{(T,P)} PBE^*(G^*(T,P))$. Take any $\bar{v} \in \cup_{(T,P)} PBE^*(G^*(T,P))$. Hence, there exists an information structure (T,P) and a PBE* strategy profile $(\gamma,(\tilde{\beta}^s)_s)$ of $G^*(T,P)$ such that

$$\bar{v}(a,\omega) = \sum_{s\in S} \prod_{i\in\mathcal{I}} \gamma_i(s_i) \sum_{t_{\mathcal{L}(s)}} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)}|\omega)\left(\prod_{j\in\mathcal{L}(s)} \tilde{\beta}_j^s(a_j|t_j) \prod_{k\in\mathcal{G}(s)} \tilde{\beta}_k^s(a_k)\right)$$

for all $a \in A$ and $\omega \in \Omega$.

For each $i \in \mathcal{I}$ define $\beta_i^g : S_{-i} \to \Delta A_i$ in the following way: for each $s \in S$ such that $s_i = g$, $\beta_i^g(a_i|s_{-i}) = \tilde{\beta}_i^s(a_i)$ for all $a_i \in A_i$. Let $\beta^g = \times_i \beta_i^g$. Symmetrically, for each $i \in \mathcal{I}$, $m_i \in \mathscr{A}_i$, and $t_i \in T_i$, let

$$Q_i(m_i \mid t_i) := \prod_{s_{-i}\in S_{-i}} \tilde{\beta}_i^{(\ell,s_{-i})}\left(m_i(s_{-i}) \,\middle|\, t_i\right), \tag{A.9}$$

where $Q_i(\cdot \mid t_i) \in \Delta(\mathscr{A}_i)$, and for each $m \in \mathscr{A}$ and $\omega \in \Omega$, define $\pi \in \Delta(\mathscr{A}\times\Omega)$ by

$$\pi(m,\omega) := \mu(\omega) \sum_{t\in T} P(t \mid \omega) \prod_{i\in\mathcal{I}} Q_i(m_i \mid t_i). \tag{A.10}$$

That is, (ω, t) is drawn as in $G^*(T,P)$, and then, independently across agents and across the profiles $s_{-i} \in S_{-i}$, the coordinate $m_i(s_{-i})$ of agent i's recommendation vector is drawn from the action distribution that i actually plays under the PBE* strategy profile, given t_i, in the continuation game in which he Looks and his opponents choose s_{-i}.[20]

Since each $Q_i(\cdot \mid t_i)$ is a product of probability measures over the coordinates of m_i, we have $\sum_{m_i\in\mathscr{A}_i} Q_i(m_i \mid t_i) = 1$; hence π is a probability distribution and $\pi(\mathscr{A}\times\{\omega\}) = \mu(\omega)$ for all $\omega \in \Omega$, which is Condition 1 of Definition 4.

The distribution π defined in (A.10) satisfies, for every $s \in S$, $a_{\mathcal{L}(s)} \in \times_{i\in\mathcal{L}(s)} A_i$, and $\omega \in \Omega$,

$$\sum_{m_{\mathcal{L}(s)}:(m_j(s_{-j}))_{j\in\mathcal{L}(s)}=a_{\mathcal{L}(s)}} \pi(m_{\mathcal{L}(s)},\omega) = \sum_{t_{\mathcal{L}(s)}} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)} \mid \omega) \prod_{i\in\mathcal{L}(s)} \tilde{\beta}_i^s(a_i \mid t_i). \tag{A.11}$$

In addition, note the following two properties of π. First, conditional on t_i, the vector m_i is independent of (ω, t_{-i}, m_{-i}). Second, by (A.9), $Q_i(m_i \mid t_i) > 0$ implies $\tilde{\beta}_i^{(\ell,s_{-i})}(m_i(s_{-i}) \mid t_i) > 0$ for every $s_{-i} \in S_{-i}$.

Now fix $s \in S$, $i \in \mathcal{L}(s)$, and $m_i \in \mathscr{A}_i$, and multiply both sides of (1) – taken at the action $a_i = m_i(s_{-i})$ – by $Q_i(m_i \mid t_i)$ and sum across t_i. Since m_i is independent of (ω, t_{-i}, m_{-i}) conditional on t_i, we obtain

$$\begin{aligned}
&\sum_{t_i} Q_i(m_i \mid t_i) \sum_{a_{-i},t_{\mathcal{L}(s)\setminus i},\omega} \mu(\omega) P_{\mathcal{L}(s)}(t_i, t_{\mathcal{L}(s)\setminus i} \mid \omega) \prod_{j\in\mathcal{L}(s)\setminus i} \tilde{\beta}_j^s(a_j \mid t_j) \prod_{k\in\mathcal{G}(s)} \tilde{\beta}_k^s(a_k)\, u_i\left(m_i(s_{-i}), a_{-i}, \omega\right) \\
&= \sum_{m_{\mathcal{L}(s)\setminus i},\, a_{\mathcal{G}(s)},\, \omega} \pi(m_i, m_{\mathcal{L}(s)\setminus i}, \omega) \prod_{k\in\mathcal{G}(s)} \beta_k^g(a_k \mid s_{-k})\, u_i\left(m_i(s_{-i}), (m_j(s_{-j}))_{j\in\mathcal{L}(s)\setminus i}, a_{\mathcal{G}(s)}, \omega\right) \\
&\geq \sum_{m_{\mathcal{L}(s)\setminus i},\, a_{\mathcal{G}(s)},\, \omega} \pi(m_i, m_{\mathcal{L}(s)\setminus i}, \omega) \prod_{k\in\mathcal{G}(s)} \beta_k^g(a_k \mid s_{-k})\, u_i\left(a'_i, (m_j(s_{-j}))_{j\in\mathcal{L}(s)\setminus i}, a_{\mathcal{G}(s)}, \omega\right) \\
&= \sum_{t_i} Q_i(m_i \mid t_i) \sum_{a_{-i},t_{\mathcal{L}(s)\setminus i},\omega} \mu(\omega) P_{\mathcal{L}(s)}(t_i, t_{\mathcal{L}(s)\setminus i} \mid \omega) \prod_{j\in\mathcal{L}(s)\setminus i} \tilde{\beta}_j^s(a_j \mid t_j) \prod_{k\in\mathcal{G}(s)} \tilde{\beta}_k^s(a_k)\, u_i\left(a'_i, a_{-i}, \omega\right)
\end{aligned} \tag{A.12}$$

for each $a'_i \in A_i$, where the two equalities use (A.10) and the marginalization steps behind (A.11). Because every t_i such that $Q_i(m_i \mid t_i) > 0$ satisfies $\tilde{\beta}_i^s(m_i(s_{-i}) \mid t_i) > 0$, (1) applies term by term and the inequality is preserved. This establishes (5) for each $s \in S$, $i \in \mathcal{L}(s)$, $m_i \in \mathscr{A}_i$, and $a'_i \in A_i$.

[20] Independence across the coordinates of m_i is convenient but not essential: any coupling of the marginals $\{\tilde{\beta}_i^{(\ell,s_{-i})}(\cdot \mid t_i)\}_{s_{-i}}$ would serve, since a coupling's support is contained in the product of the supports. What *is* essential is that the coupling be a function of t_i alone, so that m_i is independent of (ω, t_{-i}, m_{-i}) conditional on t_i.

For all $s \in S$, $i \in \mathcal{G}(s)$ and $a_i \in A_i$ with $\tilde{\beta}_i^s(a_i) > 0$, (2) can be equivalently written as

$$\begin{aligned}
&\sum_{a_{-i},\omega}\left(\sum_{t_{\mathcal{L}(s)}}\mu(\omega)P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)}|\omega)\prod_{j\in\mathcal{L}(s)}\tilde{\beta}_j^s(a_j|t_j)\right)\prod_{k\in\mathcal{G}(s)\setminus i}\tilde{\beta}_k^s(a_k)u_i(a_i,a_{-i},\omega)\\
&=\sum_{m_{\mathcal{L}(s)},a_{\mathcal{G}(s)\setminus i},\omega}\pi(m_{\mathcal{L}(s)},\omega)\prod_{k\in\mathcal{G}(s)\setminus i}\beta_k^g(a_k|s_{-k})u_i(a_i,(m_j(s_{-j}))_{j\in\mathcal{L}(s)},a_{\mathcal{G}(s)\setminus i},\omega)\\
&\geq\sum_{m_{\mathcal{L}(s)},a_{\mathcal{G}(s)\setminus i},\omega}\pi(m_{\mathcal{L}(s)},\omega)\prod_{k\in\mathcal{G}(s)\setminus i}\beta_k^g(a_k|s_{-k})u_i(a_i',(m_j(s_{-j}))_{j\in\mathcal{L}(s)},a_{\mathcal{G}(s)\setminus i},\omega)\\
&=\sum_{a_{-i},\omega}\left(\sum_{t_{\mathcal{L}(s)}}\mu(\omega)P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)}|\omega)\prod_{j\in\mathcal{L}(s)}\tilde{\beta}_j^s(a_j|t_j)\right)\prod_{k\in\mathcal{G}(s)\setminus i}\tilde{\beta}_k^s(a_k)u_i(a_i',a_{-i},\omega),
\end{aligned}\tag{A.13}$$

for all $a_i, a_i' \in A_i$ such that $\beta_i^g(a_i|s_{-i}) > 0$, which establishes (6).

For all $i \in \mathcal{I}$ and $s_i \in \{\ell, g\}$ with $\gamma_i(s_i) > 0$, (4) can be written as

$$\begin{aligned}
&\sum_{s_{-i},a,\omega}\prod_{j\neq i}\gamma_j(s_j)\left(\sum_{t_{\mathcal{L}(s)}}\mu(\omega)P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)}|\omega)\prod_{j\in\mathcal{L}(s)}\tilde{\beta}_j^s(a_j|t_j)\right)\prod_{k\in\mathcal{G}(s)}\tilde{\beta}_k^s(a_k)u_i(a_i,a_{-i},\omega)\\
&=\sum_{s_{-i},m_{\mathcal{L}(s)},a_{\mathcal{G}(s)},\omega}\prod_{j\neq i}\gamma_j(s_j)\pi(m_{\mathcal{L}(s)},\omega)\prod_{k\in\mathcal{G}(s)}\beta_k^g(a_k|s_{-k})u_i((m_j(s_{-j}))_{j\in\mathcal{L}(s)},a_{\mathcal{G}(s)},\omega)\\
&\geq\sum_{s'_{-i},m_{\mathcal{L}(s')},a_{\mathcal{G}(s')},\omega}\prod_{j\neq i}\gamma_j(s'_j)\pi(m_{\mathcal{L}(s')},\omega)\prod_{k\in\mathcal{G}(s')}\beta_k^g(a_k|s'_{-k})u_i((m_j(s'_{-j}))_{j\in\mathcal{L}(s')},a_{\mathcal{G}(s')},\omega)\\
&=\sum_{s'_{-i},a,\omega}\prod_{j\neq i}\gamma_j(s'_j)\left(\sum_{t_{\mathcal{L}(s')}}\mu(\omega)P_{\mathcal{L}(s')}(t_{\mathcal{L}(s')}|\omega)\prod_{j\in\mathcal{L}(s')}\tilde{\beta}_j^{s'}(a_j|t_j)\right)\prod_{k\in\mathcal{G}(s')}\tilde{\beta}_k^{s'}(a_k)u_i(a_i,a_{-i},\omega)
\end{aligned}\tag{A.14}$$

for all $s_i' \in \{\ell, g\}$, where $s \equiv (s_i, s_{-i})$ and $s' \equiv (s_i', s_{-i}')$, which establishes (7).

Hence, (γ, β^g, π) is an IPBCE of G^*. Then, $v(\gamma, \beta^g, \pi) \in \Delta(A \times \Omega)$ is an IPBCE outcome of G^*, that is, $v(\gamma, \beta^g, \pi) \in IPBCE(G^*)$. Notice that

$$\begin{aligned}
&v(\gamma,\beta^g,\pi)(a,\omega)\\
&=\sum_{s\in S}\prod_{i\in\mathcal{I}}\gamma_i(s_i)\left(\sum_{m_{\mathcal{L}(s)}:(m_j(s_{-j}))_{j\in\mathcal{L}(s)}=a_{\mathcal{L}(s)}}\pi\left(m_{\mathcal{L}(s)},\omega\right)\right)\prod_{k\in\mathcal{G}(s)}\beta_k^g(a_k|s_{-k})\\
&=\sum_{s\in S}\prod_{i\in\mathcal{I}}\gamma_i(s_i)\sum_{t_{\mathcal{L}(s)}}\mu(\omega)P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)}|\omega)\left(\prod_{j\in\mathcal{L}(s)}\tilde{\beta}_j^s(a_j|t_j)\prod_{k\in\mathcal{G}(s)}\tilde{\beta}_k^s(a_k)\right)\\
&=\bar{v}(a,\omega)
\end{aligned}\tag{A.15}$$

for all $a \in A$ and $\omega \in \Omega$. Thus, $\bar{v} \in IPBCE(G^*)$.

□

Appendix B. Proof of Theorem 2

Proof. Let $v \in \Delta(A \times \Omega)$ be a 1-fault tolerant BCE of G. Intuitively, we construct the desired PBE* in this way: the designer sends simple direct action recommendations according to v. All players Look and follow their recommendations. If a single player k deviates to Ignore, then the other players still follow their recommendations, and player k plays a best response; by the definition of 1-fault tolerance, that strategy is a BNE. It follows immediately that player k cannot benefit from a deviation to Ignore, because he faces the same behavior from the other players but has weakly less information to use in best responding. To complete the construction, after multiple deviations to Ignore we specify an arbitrary BNE of the resulting continuation game.

Formally, we construct the information structure (T, P) and the PBE* strategy profile $\left(\gamma, \left(\tilde{\beta}^s\right)_s\right)$ of the dynamic game $G^*(T, P)$ as follows.

- Let $T = A$, and for all $a \in A, \omega \in \Omega$ let $P(a|\omega) = v(a, \omega)/\mu(\omega)$. Hence, (T, P) is an information structure with simple direct action recommendations.
- Take $\gamma_i(\ell) = 1$ for all $i \in \mathcal{I}$. That is, all players choose Look with probability 1.
- Let $s^L = (\ell, \ldots, \ell)$ be the Look-Ignore profile where all players choose Look, and let the set $S^1 := \{s \in S : |\mathcal{L}(s)| = N - 1\}$ denote all the profiles at which a single player chooses Ignore.
  - For all $s \in S^1 \cup \{s^L\}$ (that is, after at most one player chooses Ignore), set

$$\tilde{\beta}_i^s(a_i'|a_i) = \begin{cases} 1, & \text{if } a_i' = a_i \\ 0, & \text{if } a_i' \neq a_i \end{cases}\tag{B.1}$$

for each player $i \in \mathcal{L}(s)$, and for player $k \in \mathcal{G}(s)$ (if any) let $\tilde{\beta}_k^s$ be a best response to $\left(\tilde{\beta}_i^s\right)_{i\neq k}$.

– For all $s \notin S^1 \cup \left\{s^L\right\}$, let $\tilde{\beta}^s$ be any BNE of $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$, which must exist by finiteness.

On path (that is, after all players choose Look), the specified strategy is a BNE of the continuation game $G(T_{\mathcal{L}(s^L)}, P_{\mathcal{L}(s^L)})$ because it corresponds to the BCE υ. Thus, the outcome of $\left(\gamma, \left(\tilde{\beta}^s\right)_s\right)$ is υ, as desired. After multiple deviations to Ignore (that is, at any $s \notin S^1 \cup \left\{s^L\right\}$), $\tilde{\beta}^s$ is a BNE of $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$ by construction. It remains only to show that 1) $\tilde{\beta}^s$ is a BNE of $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$ for $s \in S^1$; and 2) choosing Look is a best response for all players at the Look-Ignore stage (Condition (4)).

To show that $\tilde{\beta}^s$ is a BNE of $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$ for $s \in S^1$, we use the definition of 1-fault tolerance. For any $s \in S^1$ with $k \in \mathcal{G}(s)$, summing over a_k' on both sides of Condition (9) according to $\tilde{\beta}_k^s$ yields

$$\begin{aligned}&\sum_{a_k', a_{-ik}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(a_i, a_{-ik}|\omega)\left(\tilde{\beta}_k^s(a_k')\right) u_i(a_i, a_k', a_{-ik}, \omega)\\ \geq &\sum_{a_k', a_{-ik}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(a_i, a_{-ik}|\omega)\left(\tilde{\beta}_k^s(a_k')\right) u_i(a_i', a_k', a_{-ik}, \omega)\end{aligned} \tag{B.2}$$

for each $i \in \mathcal{L}(s)$ and all $a_i, a_i' \in A_i$; we are using the fact that

$$P_{\mathcal{L}(s)}(a_i, a_{-ik}|\omega) = \sum_{a_k} P(a_i, a_k, a_{-ik}|\omega)$$

for each $a \in A$ and $\omega \in \Omega$. The strategies defined in (B.1) imply that

$$\prod_{j \in \mathcal{L}(s)\setminus i} \tilde{\beta}_j^s(a_j|a_j) = 1,$$

so Condition (B.2) is equivalent to

$$\begin{aligned}&\sum_{a_{-i}', a_{-ik}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(a_i, a_{-ik}|\omega)\left(\tilde{\beta}_k^s(a_k') \prod_{j \in \mathcal{L}(s)\setminus i} \tilde{\beta}_j^s(a_j'|a_j)\right) u_i(a_i, a_k', a_{-ik}', \omega)\\ \geq &\sum_{a_{-i}', a_{-ik}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(a_i, a_{-ik}|\omega)\left(\tilde{\beta}_k^s(a_k') \prod_{j \in \mathcal{L}(s)\setminus i} \tilde{\beta}_j^s(a_j'|a_j)\right) u_i(a_i', a_k', a_{-ik}', \omega).\end{aligned} \tag{B.3}$$

Thus, the specified strategies are best responses for all $i \in \mathcal{L}(s)$ (Condition (1)). For player $k \in \mathcal{G}(s)$, $\tilde{\beta}_k^s$ is a best response to $\left(\tilde{\beta}_i^s\right)_{i\neq k}$ by definition. Hence, by Definition 1, $\tilde{\beta}^s$ is a BNE of $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$ for all $s \in S^1$.

Finally, we show that choosing Look is a best response for each player i at the Look-Ignore stage, given that all other players are choosing Look. Because under the strategies defined in (B.1) other players do not change their behavior in response to player i's deviation to Ignore, that result follows immediately from the obedience constraint (Condition (8)) for the BCE υ: choosing Ignore is equivalent to disregarding the recommendation under υ. Condition (8) implies that deviating from the recommendation to any pure action cannot be profitable, so deviating to a mixed action cannot be profitable either, and Condition (4) is satisfied. □

Appendix C. Proof of Theorem 3

Proof. Consider the following game, where $\Omega = \{e, f\}$ and each state is equally likely. There are two players with action sets $A_1 = \{E, M, F, M'\}$ and $A_2 = \{L, R_e, R_f, P_e, P_f, Q\}$. The players' state contingent payoffs are given in Fig. C.1.

The designer gets a payoff of 1 if (E, R_e) is played in state e, or if (F, R_f) is played in state f, and payoff 0 otherwise. The actions R_e and R_f are duplicates from the players' point of view. Their role in the example is to make it so that Player 2 needs to know the state in order to play the designer's desired action.

Suppose that the state is common knowledge. In state e, E is dominant for Player 1, and R_e is a best response for Player 2. In state f, F is dominant for Player 1, and R_f is a best response for Player 2. The expected payoff vector for the players is $(1, 1)$, and the designer gets an expected payoff of 1.

At the prior, expected payoffs are given in Fig. C.2.

Suppose it is common knowledge that Player 1 knows the state and that Player 2's beliefs equal the prior. In state e, E is dominant for Player 1, and in state f, F is dominant. In both cases, either R_e or R_f is a best response for Player 2. In both cases, irrespective of which best response Player 2 plays, the expected payoff vector for the players is $(1, 1)$. However, the designer only gets a payoff of 1 if Player 2 plays R_e in state $\omega = e$ and R_f in state $\omega = f$.

Crucially, in this setting, Player 1 can be punished effectively for choosing Ignore only if it is common knowledge that Player 2's belief equals the prior. The reason is the following. Suppose first that it is common knowledge that both players' beliefs equal the prior. Then M strictly dominates E and F, and M weakly dominates M' for Player 1: M' is a weak best response for Player 1 if and only if Player 2 plays Q with probability 1. Q is a best response to M'. So (M', Q) is an equilibrium with payoff $(0, 2)$. Next, suppose Player 2 assigns belief $p > \frac{1}{2}$ to state ω. Then Q is not a best response to M': Q gives payoff 2, while P_ω gives expected payoff $3p + (1 - p) = 2p + 1 > 2$.

	L	R_e	R_f	P_e	P_f	Q
E	$3,0$	$1,1$	$1,1$	$3,-1$	$3,-1$	$1,-1$
M	$2,2$	$0,0$	$0,0$	$2,0$	$2,0$	$0,0$
F	$0,0$	$-2,1$	$-2,1$	$-2,-1$	$-2,-1$	$-2,-1$
M'	$0,2$	$-1,0$	$-1,0$	$-1,3$	$-1,1$	$0,2$

$\omega = e$

	L	R_e	R_f	P_e	P_f	Q
E	$0,0$	$-2,1$	$-2,1$	$-2,-1$	$-2,-1$	$-2,-1$
M	$2,2$	$0,0$	$0,0$	$2,0$	$2,0$	$0,0$
F	$3,0$	$1,1$	$1,1$	$3,-1$	$3,-1$	$1,-1$
M'	$0,2$	$-1,0$	$-1,0$	$-1,1$	$-1,3$	$0,2$

$\omega = f$

Fig. C.1. State-contingent payoffs.

	L	R_e	R_f	P_e	P_f	Q
E	$\frac{3}{2},0$	$-\frac{1}{2},1$	$-\frac{1}{2},1$	$\frac{1}{2},-1$	$\frac{1}{2},-1$	$-\frac{1}{2},-1$
M	$2,2$	$0,0$	$0,0$	$2,0$	$2,0$	$0,0$
F	$\frac{3}{2},0$	$-\frac{1}{2},1$	$-\frac{1}{2},1$	$\frac{1}{2},-1$	$\frac{1}{2},-1$	$-\frac{1}{2},-1$
M'	$0,2$	$-1,0$	$-1,0$	$-1,2$	$-1,2$	$0,2$

Fig. C.2. Expected payoffs at the prior.

Therefore, if it is common knowledge that Player 1's belief equals the prior, and that there is ex ante strictly positive probability that Player 2 has some information (i.e., assigns belief $p > \frac{1}{2}$ to one state or the other), then (M', Q) is not an equilibrium. Instead, M is dominant against state-contingent strategies of Player 2 and the unique equilibrium is (M, L), giving payoff $(2, 2)$.

A mixed Look-Ignore outcome: Suppose the designer's information structure is given by $(\mathscr{A}, P)$ with

$$P\big(EE, R_eL|\omega = e\big) = P\big(FF, R_fL|\omega = f\big) = 1$$

which perfectly informs both players of the state. The first term in each player's message is the action recommendation to follow after the other player has chosen Look (ℓ), while the second term is the action recommendation to follow after the other player has chosen Ignore (g).

Given this information structure, the following is an equilibrium of the Look-Ignore stage: Player 1 plays ℓ, i.e. $\gamma_1(\ell) = 1$, and Player 2 randomizes with equal probability over ℓ and g, that is $\gamma_2(\ell) = \gamma_2(g) = \frac{1}{2}$. On path, the payoff for the players is $(1, 1)$, regardless of Player 2's Look-Ignore choice, and in expectation the designer gets a payoff of $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{1}{2} = \frac{3}{4}$.

Next, we argue that following the action recommendations of the direct information structure specified above is incentive compatible for some post-Ignore contingent strategies, i.e. it is an equilibrium of the action stage:

- After (ℓ, ℓ): Player 1's recommendation specifies his dominant action for the revealed state (E or F), and Player 2's recommendation is a best response. The payoff vector is $(1, 1)$.
- After (ℓ, g): Player 1's recommendation specifies his dominant action (E or F). Player 2's post-Ignore strategy is $\beta_2^g(R_e|\ell) = \beta_2^g(R_f|\ell) = \frac{1}{2}$, where he randomizes between R_e and R_f, both of which are best responses. The payoff vector is $(1, 1)$.
- After (g, ℓ): Player 1's post-Ignore strategy is $\beta_1^g(M|\ell) = 1$; M is a best response to Player 2's recommendation L. For Player 2, L is the strict best response to M. The payoff vector is $(2, 2)$.
- After (g, g): Consider the post-Ignore strategies $\beta_1^g(M'|g) = 1$ and $\beta_2^g(Q|g) = 1$. At the prior, M' is a best response to Q, and Q is a best response to M'. The payoff vector is $(0, 2)$.

At the Look-Ignore stage:

- Given that Player 1 plays ℓ, Player 2 is indifferent between ℓ and g, as he gets a payoff of 1 either way. Hence, Player 2 is willing to mix, as required.
- Given that Player 2 chooses ℓ with probability $\frac{1}{2}$, Player 1's payoff from ℓ is $\frac{1}{2}\cdot 1+\frac{1}{2}\cdot 1=1$. Deviating to g gives Player 1 a payoff of $\frac{1}{2}\cdot 2+\frac{1}{2}\cdot 0=1$. Thus, ℓ is a best response for Player 1, as required.

Trying to replicate in a pure Look-Look equilibrium: For the designer to get a payoff $p>\frac{1}{2}$, Player 2 must match the state with probability at least p, so with strictly positive probability his recommendation must give him some information about the state.

Consequently, if Player 1 deviates to g at the Look-Ignore stage, then the continuation play after (g,ℓ) must be (M,L), giving a payoff vector $(2,2)$. Thus, Player 1 must get a payoff of at least 2 after (ℓ,ℓ) in order to satisfy his Look constraint. It follows that the designer's preferred action profiles (which give Player 1 a payoff of 1) can be played with probability no higher than $\frac{1}{2}$: Player 1's highest possible payoff is 3, and $1x+3(1-x)\geq 2$ implies that $x\leq\frac{1}{2}$. We conclude that the mixed Look-Ignore outcome above cannot be duplicated in a pure Look-Look equilibrium. □

Appendix D. Proof of Theorem 4

In the game constructed below, an information structure that reveals the state perfectly gives rise to a unique BNE that maximizes the players' expected payoffs. However, if players have the ability to exercise strategic ignorance, then, given the continuation equilibria, it is a dominant strategy to Ignore that baseline information structure. The game has the flavor of a prisoners' dilemma at the Look-Ignore stage, where Look corresponds to Cooperate, and Ignore corresponds to Defect. Roughly, an informed Player 2's best response to an uninformed Player 1's optimal action is much better for Player 1 than the best response to an informed Player 1's optimal action would be. That benefit from ignorance outweighs Player 1's loss from not being able to tailor his own action to the state. Against an uninformed opponent, a player also benefits from being uninformed. Thus, Ignore is strictly dominant at the baseline information structure and, as a consequence, the highest achievable sum of players' payoffs is lower than if messages were automatically observed. In fact, the payoff to each individual player is lower in any optimal IPBCE (that is, one that maximizes the sum of players' payoffs) than in the optimal BCE (the baseline information structure).

Proof.

Consider the following symmetric game, where each state $\omega\in\{1,2\}$ is equally likely. The players' state contingent payoffs are given in Fig. D.1. At the prior, expected payoffs are given in Fig. D.2, so that Y is strictly dominant for each player.

	X	Y	E_1	F_1	E_2	F_2
X	0, 0	1, 0.1	8, 1	11, 2	−10, −2	−10, −2
Y	0.1, 1	3, 3	6, 7	6, 2	6, −2	6, −2
E_1	1, 8	7, 6	10, 10	11.5, 6	10, 0	11.5, 0
F_1	2, 11	2, 6	6, 11.5	9, 9	6, 0	9, 0
E_2	−2, −10	−2, 6	0, 10	0, 6	0, 0	0, 0
F_2	−2, −10	−2, 6	0, 11.5	0, 9	0, 0	0, 0

$\omega=1$

	X	Y	E_1	F_1	E_2	F_2
X	0, 0	1, 0.1	−10, −2	−10, −2	8, 1	11, 2
Y	0.1, 1	3, 3	6, −2	6, −2	6, 7	6, 2
E_1	−2, −10	−2, 6	0, 0	0, 0	0, 10	0, 6
F_1	−2, −10	−2, 6	0, 0	0, 0	0, 11.5	0, 9
E_2	1, 8	7, 6	10, 0	11.5, 0	10, 10	11.5, 6
F_2	2, 11	2, 6	6, 0	9, 0	6, 11.5	9, 9

$\omega=2$

Fig. D.1. State-contingent payoffs.

	X	Y	E_1	F_1	E_2	F_2
X	0, 0	1, 0.1	−1, −0.5	0.5, 0	−1, −0.5	0.5, 0
Y	0.1, 1	3, 3	6, 2.5	6, 0	6, 2.5	6, 0
E_1	−0.5, −1	2.5, 6	5, 5	5.75, 3	5, 5	5.75, 3
F_1	0, 0.5	0, 6	3, 5.75	4.5, 4.5	3, 5.75	4.5, 4.5
E_2	−0.5, −1	2.5, 6	5, 5	5.75, 3	5, 5	5.75, 3
F_2	0, 0.5	0, 6	3, 5.75	4.5, 4.5	3, 5.75	4.5, 4.5

$$\Pr(\omega = 1) = \tfrac{1}{2}$$

Fig. D.2. Expected payoffs at the prior.

	E_1E_2	E_1F_2	F_1E_2	F_1F_2
X	8	9.5	9.5	11
Y	6	6	6	6
E_1	5	5	5.75	5.75
F_1	3	3	4.5	4.5
E_2	5	5.75	5	5.75
F_2	3	4.5	3	4.5

$$\Pr(\omega = 1) = \tfrac{1}{2}$$

Fig. D.3. Player 1's expected payoffs after (g, ℓ).

Playing action profile (E_ω, E_ω) in state ω is the unique state contingent action profile that maximizes the sum of the players' expected utilities, yielding a total payoff of 20. Achieving that maximum thus requires both players to be perfectly informed of the state, in which case it is a BNE to play this action profile in each state. If players automatically observe their private signals, and the designer uses an information structure that recommends action E_ω in state ω to each player, then the players are willing to obey their recommendations, as described below. However, with the possibility of strategic ignorance, we need to consider the following cases:

- After (ℓ, ℓ): In this case the state is common knowledge. In state ω, action E_ω strictly dominates every action except F_ω. The unique best response to any mixing between E_ω and F_ω is E_ω. Thus, the unique BNE is $\left(E_\omega, E_\omega\right)$ in each $\omega \in \{1, 2\}$, and the payoffs are $u(\ell, \ell) = (10, 10)$.
- After (g, ℓ): In this case it is common knowledge that Player 2 knows the state and that Player 1's beliefs are given by the prior. As above, in state ω, action E_ω strictly dominates every action except F_ω for Player 2. Thus, Player 2 has four undominated state-contingent strategies: E_1E_2, E_1F_2, F_1E_2, and F_1F_2, where the first element denotes the action in state 1 and the second element denotes the action in state 2. Player 1's expected payoffs against those strategies are given in Fig. D.3. Hence, for uninformed Player 1, X is the unique best response against any state-contingent undominated strategy of Player 2. Player 2's best response to X is F_1F_2. Thus, the unique BNE is $\left(X, F_1F_2\right)$, and the payoffs are $u(g, \ell) = (11, 2)$.
- After (ℓ, g): This case is symmetric to the preceding one.
- After (g, g): In this case it is common knowledge that both players' beliefs are given by the prior distribution, and, hence, Y is strictly dominant. Thus, the unique BNE is (Y, Y), and the payoffs are $u(g, g) = (3, 3)$.

Equilibrium at the Look-Ignore Stage: After each combination of Look-Ignore choices, we have shown that there is a unique BNE. Using these as the continuation payoffs, we can write the payoff matrix at the Look-Ignore stage as in Fig. D.4.

Ignore is strictly dominant, so the outcome is that both players choose Ignore and wind up with payoff 3. Thus, it is not possible to achieve total payoffs 20 when strategic ignorance is possible. As we show next, the total-payoff-maximizing outcome under the optimal information structure adjusted for the possibility of strategic ignorance is symmetric, which implies, because the total payoff must be less than 20, that both players get a lower payoff than if messages were automatically observed.

Total-Payoff-Maximizing Outcome with Strategic Ignorance: In the presence of strategic ignorance, computations show that the unique IPBCE payoff vector that maximizes the sum of players' expected payoffs is $\left(\frac{100}{11}, \frac{100}{11}\right)$. This outcome can be implemented by

	ℓ	g
ℓ	10, 10	2, 11
g	11, 2	3, 3

Fig. D.4. Payoffs at the Look-Ignore stage.

providing the following direct contingent information structure $(\mathcal{A}, P^*)$, where $q = \frac{1}{11}$ and for each $\omega \in \{1, 2\}$ we have:

$$P^*\left(E_\omega F_\omega, E_\omega F_\omega | \omega\right) = 1 - q, \quad P^*\left(E_{-\omega} F_{-\omega}, E_{-\omega} F_{-\omega} | \omega\right) = q.$$

The first term in each player's message is the action recommendation to follow after the other player Looks, and the second is the action recommendation for after the other player Ignores. Hence, upon observing the message $E_\omega F_\omega$, a player attaches a posterior belief of $1 - q = \frac{10}{11}$ to the state ω.

In words, the information structure publicly gives both players the "correct" recommendation with probability $1 - q$, and tells them the wrong state with probability q. The posterior belief $1 - q$ is high enough that the action recommendations are obedient, as in the baseline information structure. Now both players choosing Look is an equilibrium, because the Look constraint is satisfied with equality. Looking and following the recommendations, given that the other player will do so, yields expected payoff $(1 - q) \cdot 10 + q \cdot 0 = \frac{100}{11}$. Deviating to Ignore (again assuming that the other player will Look and follow his recommendation) yields $(1 - q) \cdot 11 + q \cdot (-10) = \frac{100}{11}$: as in the baseline case, X is the best response for a player who has unilaterally deviated to Ignore, but now with probability q the other player will play $F_{-\omega}$ rather than F_ω, and the deviator will get -10 instead of 11. □

Appendix E. Proof of Proposition 1

Proof. Consider the following state-contingent payoff matrices in Fig. E.1. The state space is $\Omega = \{e, f\}$ and each of the states is equally likely. The designer gets a payoff of 1 if (E, Y) is played in state e or (F, Y) in state f, and 0 otherwise.

	X	Y
E	0, 0	1, 1
F	2, 2	1, 1

$\omega = e$

	X	Y
E	1, 0	0, 0
F	1, 1	1, 1

$\omega = f$

Fig. E.1. State-contingent payoffs.

At the prior, the players' expected payoffs are given in Fig. E.2.

	X	Y
E	$\frac{1}{2}, 0$	$\frac{1}{2}, \frac{1}{2}$
F	$\frac{3}{2}, \frac{3}{2}$	1, 1

$\Pr(\omega = e) = \frac{1}{2}$

Fig. E.2. Expected payoffs at the prior.

Baseline. If players are forced to see the messages, then the designer's optimal direct information structure $(A, \tilde{P})$ is

$$\tilde{P}(E, Y | \omega = e) = \tilde{P}(F, Y | \omega = f) = 1.$$

The designer's payoff is 1, and the players' payoffs are $(1, 1)$. Player 1's message reveals the state. Against Y, Player 1 is indifferent between his two actions in state e, and action F is the unique best response in state f. Player 2's message reveals nothing, and his recommended action Y is a strict best response at the prior to Player 1's state-dependent strategy.

With strategic ignorance. The baseline information structure with single action recommendations does not implement the desired outcome if players can choose publicly whether to Look at their signals, because Player 1 would deviate to Ignore. Given that Player 2 does not have any information about the state, he cannot play a state-contingent strategy. As shown in Fig. E.2, F is strictly dominant for Player 1 at the prior, and Player 2's best response is to play X. Therefore, by deviating to Ignore, Player 1 gets a payoff of $\frac{3}{2}$, which is higher than the payoff of 1 that he gets if he chooses to Look.

Nevertheless, the designer's optimal outcome is robust to strategic ignorance, because she can provide an information structure with direct contingent recommendations. Such an information structure also specifies which action to play off-path, if the other player deviates to Ignore. More precisely, the outcome can be implemented by providing the following direct contingent information structure $(\mathscr{A}, P^*)$:

$$P^*(EE, YY|\omega = e) = P^*(FF, YX|\omega = f) = 1,$$

where the first term in each player's message is the action recommendation to follow after the other player Looks, and the second is the action recommendation for after the other player Ignores. Now both players' messages reveal the state. Note that $(\mathscr{A}, P^*)$ is the original baseline direct information structure augmented with post-Ignore punishment recommendations. The details are as follows.

If Player 1 plays E after choosing to Ignore, while Player 2 plays Y after choosing to Ignore, then both players choosing Look is an equilibrium at the Look-Ignore stage. In particular, Player 1's recommended actions at the action stage are the same whether or not Player 2 decides to Ignore, so it is immediate that Look is a best response for Player 2. After Player 1 deviates to Ignore, then the specified continuation ((E, Y) in state e and (E, X) in state f) gives him a payoff of 1, which also makes him indifferent between Look and Ignore. Thus, both Look constraints hold with equality.

After $s = (g, \ell)$, that is, after Player 1 deviates to Ignore, Player 2's recommended action Y in state e is a strict best response to E, and his recommended action X in state f is a weak best response to E. At the prior, both E and F are best responses for Player 1 to Player 2's recommended state-dependent actions. Thus, obedience is satisfied. After $s = (\ell, g)$, that is, after Player 2 deviates to Ignore, the recommendations are the same as under $(A, \tilde{P})$, so obedience is satisfied. The same holds for the recommendations after $s = (\ell, \ell)$.

Thus, the information structure $(\mathscr{A}, P^*)$ implements the designer's optimal outcome. In fact, any information structure that implements the optimal outcome must fully reveal the state to Player 2, even though he does not use that knowledge on path. The reason is that after a deviation to Ignore by Player 1, Player 2 needs that information in order to follow the strategy of (Y in state e, X in state f). An uninformed Player 1's strict best response to any other state-contingent strategy is F, and Player 2's best response to F is X unless he assigns probability 1 to state f. Thus, unless Player 2's message fully reveals the state, Player 1 can get a payoff of $\frac{3}{2}$ from deviating to Ignore, and the designer's desired outcome will not be achieved. □

Appendix F. Investment game: derivations

Consider the parameterized symmetric direct information structures of Fig. 5. We would like to determine which outcome distributions can be implemented with these information structures in pure Look equilibria, i.e., in equilibria where both players choose Look with probability one. The payoff to a player from choosing to Look and following the action recommendation while the other player is also choosing to Look is given by $r + q$. If a player chooses Ignore while the other player chooses Look, in the continuation BNE that is worst for the Ignorer, the player who Looked plays the action corresponding to the most likely state given his signal (his best response to an Ignorer who randomizes uniformly), and the Ignorer's payoff does not depend on his own mixing probability. Next, we characterize the biggest set of the parameterized direct symmetric information structures that will be "Looked" at by both players.

Case 1: If $q \geq 1/2$, the expected payoff to the agent who chooses Ignore is $\frac{1}{2}(1 + q)$. This is greater than the payoff from following the action recommendations if $\frac{1}{2}(1 + q) > r + q \Leftrightarrow r < \frac{1}{2} - \frac{1}{2}q$ which directly contradicts the red obedience constraint in Fig. 6 given by $r \geq \frac{1}{2} - \frac{1}{2}q$. Hence, for all direct symmetric information structures with $q \geq 1/2$, strategic ignorance is not an issue, as Look is a best response to Look. Basically, in this case, the agent who chooses Look continues to play the same strategy and follow the recommendations, irrespective of whether the other agent chooses Look or Ignore, so choosing Ignore is never strictly profitable.

Case 2: If $q \leq 1/2$, the expected payoff to the agent who chooses Ignore is $\frac{1}{2}(2 - q)$. This is greater than the payoff from following the action recommendations if $\frac{1}{2}(2 - q) > r + q \Leftrightarrow r < 1 - \frac{3}{2}q$. This constraint is represented by the line in green in Fig. 6 and the area to the left of it. Hence, the direct symmetric information structures that are affected by the agents' ability to exercise strategic ignorance are represented by the hatched triangle below the green constraint.

Appendix G. Proof of Proposition 5

Proof. The example that proves this result builds upon the complete information game in Fig. G.1. X is strictly dominant for Player 1 and F is a strict best response to that, so the unique equilibrium of the game is (X, F), giving the vector of payoffs $(2, 2)$.

Next we add two states of nature, which give rise to the payoff matrices in Fig. G.2. The idea is that for Player 1 to want to play anything other than Y he needs to know the state. If he plays X_1 in state 2, or vice versa, then he gets a bad payoff. At the prior, Y is strictly dominant for Player 1 against any state-contingent strategy of Player 2.

Similarly, for Player 2 to want to play E or F he needs to know the state. If he plays E_1 or F_1 in state 2, or vice versa, then he gets a bad payoff. At the prior, the "safe" action G is strictly dominant for Player 2 against any state-contingent strategy of Player 1. Notice that G is bad for Player 1, as it gives him weakly lower payoffs than any other action of Player 2.

	E	F
X	$4,1$	$2,2$
Y	$3,2$	$0,0$

Fig. G.1. Complete information game.

	E_1	E_2	G	F_1	F_2
X_1	$4,1$	$4,-100$	$1,\frac{1}{2}$	$2,2$	$2,-100$
X_2	$-100,1$	$-100,-100$	$-100,\frac{1}{2}$	$-100,2$	$-100,-100$
Y	$3,2$	$3,-100$	$0,\frac{1}{2}$	$0,0$	$0,-100$

$\omega = 1$

	E_1	E_2	G	F_1	F_2
X_1	$-100,-100$	$-100,1$	$-100,\frac{1}{2}$	$-100,-100$	$-100,2$
X_2	$4,-100$	$4,1$	$1,\frac{1}{2}$	$2,-100$	$2,2$
Y	$3,-100$	$3,2$	$0,\frac{1}{2}$	$0,-100$	$0,0$

$\omega = 2$

Fig. G.2. State-contingent payoffs.

	E_1	E_2	G	F_1	F_2
X_1	$-48,-49.5$	$-48,-49.5$	$-49.5,\frac{1}{2}$	$-49,-49$	$-49,-49$
X_2	$-48,-49.5$	$-48,-49.5$	$-49.5,\frac{1}{2}$	$-49,-49$	$-49,-49$
Y	$3,-49$	$3,-49$	$0,\frac{1}{2}$	$0,-50$	$0,-50$

$\Pr(\omega = 1) = \frac{1}{2}$

Fig. G.3. Expected payoffs at the prior.

At the prior, when both players are uninformed, their expected payoffs are given in Fig. G.3. The designer gets a payoff of 1 if (X_ω, F_ω) is played in state ω, and 0 otherwise.

Baseline. If players are forced to see the messages, then the designer's optimal information structure sends messages that reveal the state perfectly and the players play the corresponding equilibrium. Formally, the optimal direct information structure $(A, \tilde{P})$ is

$$\tilde{P}(X_1, F_1|\omega = 1) = \tilde{P}(X_2, F_2|\omega = 2) = 1.$$

The designer gets a payoff of 1. The players get $(2,2)$.

With strategic ignorance. Note that a designer payoff of 1 requires that (X_ω, F_ω) be played with probability one when the state is $\omega \in \{1,2\}$. Hence, both players must know the state on path under *any* information structure that achieves this payoff. Any such information structure, which reveals the state perfectly to both players, does not work if players can choose publicly whether to Look at their signals, because Player 1 would deviate to Ignore. At the prior, Y is strictly dominant for Player 1. Given that Player 2 knows the state, his best response is E_ω, so Player 1 gets a payoff 3, which is higher than the payoff of 2 he gets if he chooses to Look.

Allowing communication. With communication between players, the designer can restore the outcome from the optimal baseline information structure $(A, \tilde{P})$. She achieves that by revealing the state to Player 1 only, and having Player 1 subsequently reveal the state to Player 2.[21] The designer achieves her maximal payoff of 1, and the players' payoffs are $(2,2)$.

Those strategies are an equilibrium. If Player 1 deviates to Ignore, then both players are uninformed. At the prior, Y is strictly dominant for Player 1, and G is strictly dominant for Player 2. The outcome is (Y, G), giving Player 1 a payoff of $0 < 2$. If Player 1 chooses to Look but deviates and does not reveal the state to Player 2, then the outcome is (X_ω, G), giving Player 1 a payoff of $1 < 2$. Player 1 also cannot gain by disclosing the state only sometimes, whether in one state only or with some probability. Finally,

[21] We assume communication between the players is in the form of verifiable all-or-nothing disclosure, where the receiver can refuse to be informed.

if Player 2 deviates and refuses to Look at what Player 1 tells him, then the outcome is again (X_ω, G), giving Player 2 a payoff of $\frac{1}{2} < 2$. Therefore, there are no profitable deviations for either player. □

References

Akerlof, G.A., 1970. The market for 'lemons': quality uncertainty and the market mechanism. Q. J. Econ. 84 (3), 488–500.
Arcuri, A., 2021. Hear-no-evil equilibrium. Working Paper. Cornell University.
Atayev, A., 2022. Information asymmetry and search intensity. Working Paper. ZEW – Leibniz Centre for European Economic Research in Mannheim.
Athey, S., Levin, J., 2001. Information and competition in U.S. Forest Service timber auctions. J. Polit. Econ. 109 (2), 375–417.
Balzer, B., Schneider, J., 2019. Belief management and optimal arbitration. Working Paper.
Bergemann, D., Brooks, B., Morris, S., 2017. First-price auctions with general information structures: implications for bidding and revenue. Econometrica 85 (1), 107–143.
Bergemann, D., Morris, S., 2016. Bayes correlated equilibrium and the comparison of information structures in games. Theor. Econ. 11, 487–522.
Bergemann, D., Morris, S., 2019. Information design: a unified perspective. J. Econ. Liter. 57 (1), 44–95.
Celik, G., Peters, M., 2011. Equilibrium rejection of a mechanism. Games Econ. Behav. 73, 375–387.
Correia-da-Silva, J., 2020. Self-rejecting mechanisms. Games Econ. Behav. 120, 434–457.
Deimen, I., Szalay, D., 2019. Delegated expertise, authority, and communication. Am. Econ. Rev. 109 (4), 1349–1374.
Doval, L., Ely, J., 2020. Sequential information design. Econometrica 88 (6), 2575–2608.
Eliaz, K., 2002. Fault tolerant implementation. Rev. Econ. Stud. 69, 589–610.
Forges, F., 1986. An approach to communication equilibria. Econometrica 54 (6), 1375–1385.
Fudenberg, D., Tirole, J., 1991. Perfect Bayesian equilibrium and sequential equilibrium. J. Econ. Theory 53, 236–260.
Gal-Or, E., 1987. First mover disadvantages with private information. Rev. Econ. Stud. 54 (2), 279–292.
Golman, R., Hagmann, D., Loewenstein, G., 2017. Information avoidance. J. Econ. Liter. 55 (1), 96–135.
Gul, F., 2001. Unobservable investment and the hold-up problem. Econometrica 69 (2), 343–376.
Hirshleifer, J., 1971. The private and social value of information and the reward to inventive activity. Am. Econ. Rev. 61 (4), 561–574.
Kalai, A.T., Kalai, E., 2024. Beyond dominance and Nash: ranking equilibria by critical mass. Games Econ. Behav. 144, 378–394.
Kamenica, E., Gentzkow, M., 2011. Bayesian persuasion. Am. Econ. Rev. 101 (6), 2590–2615.
Kessler, A.S., 1998. The value of ignorance. RAND J. Econ. 29 (2), 339–354.
Makris, M., Renou, L., 2023. Information design in multi-stage games. Theor. Econ. 18, 1475–1509.
McAdams, D., 2012. The value of ignorance. Econ. Lett. 114, 83–85.
Myerson, R., 1986. Multistage games with communication. Econometrica 54 (2), 323–358.
de Oliveira, H., Lamba, R., 2019. Rationalizing dynamic choices. Working Paper. Penn State University.
Palfrey, T.R., 1982. Risk advantages and information acquisition. Bell J. Econ. 13 (1), 219–224.
Roesler, A.-K., Szentes, B., 2017. Buyer-optimal learning and monopoly pricing. Am. Econ. Rev. 107 (7), 2072–2080.
Rogerson, W.P., 1992. Contractual solutions to the hold-up problem. Rev. Econ. Stud. 59 (4), 777–793.
Rothschild, M., Stiglitz, J.E., 1976. Equilibrium in competitive insurance markets: an essay on the economics of imperfect information. Q. J. Econ. 90 (4), 629–649.
Schelling, T.C., 1956. An essay on bargaining. Am. Econ. Rev. 46 (3), 281–306.
Schlee, E.E., 2001. The value of information in efficient risk sharing arrangements. Am. Econ. Rev. 91, 509–524.
von Stengel, B., Forges, F., 2008. Extensive-form correlated equilibrium: definition and computational complexity. Math. Oper. Res. 33 (4), 1002–1022.
Taneva, I., 2019. Information design. Am. Econ. J.: Microecon. 11 (4), 151–185.
Taneva, I., Wiseman, T., 2026. Ignoring harmful correlations. Working Paper. University of Edinburgh.
Tirole, J., 1986. Procurement and renegotiation. J. Polit. Econ. 94 (2), 235–259.
Watson, J., 2025. Perfect Bayesian equilibrium: consistency conditions for practical definitions. Working Paper.