

Strategic Ignorance and Information Design

Ina Taneva

University of Edinburgh and CEPR
ina.taneva@ed.ac.uk

Thomas Wiseman*

University of Texas at Austin
wiseman@austin.utexas.edu

This version: June 2023
First Version: December 2021

Abstract

We study information design in strategic settings when agents can publicly commit to not 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 — robust correlated equilibrium — and characterize implementable distributions over states and actions. Requiring robustness to strategic ignorance can explain qualitative properties that standard information design cannot: the designer may provide redundant or even counter-productive information, asymmetric information structures may be strictly optimal in symmetric environments, providing information conditional on players’ 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 modelling information design, it is important to incorporate incentives 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

*We thank Alesandro Arcuri, Andreas Blume, Ben Brooks, Hector Chade, Inga Deimen, Francesc Dilme, Laura Doval, Françoise Forges, Amanda Friedenber, Willie Fuchs, Johannes Hörner, Elliot Lipnowski, Erik Madsen, Arina Nikandrova, Jacopo Perego, Eddie Schlee, Denis Shishkin, Vasiliki Skreta, Mark Whitmeyer, Leeat Yariv, and seminar audiences for helpful comments and discussions. Siyu Chen provided excellent research assistance.

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.

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. 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 Figure 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.

	<i>H</i>	<i>L</i>	
<i>E</i>	3, 0	1, 1	
<i>M</i>	2, 2	0, 0	
<i>F</i>	0, 0	-2, 1	

 $\omega = e$

	<i>H</i>	<i>L</i>
<i>E</i>	0, 0	-2, 1
<i>M</i>	2, 2	0, 0
<i>F</i>	3, 0	1, 1

 $\omega = f$

Figure 1: Procurement auction example

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 it 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 Figure 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.

	<i>HH</i>	<i>HL</i>	<i>LH</i>	<i>LL</i>
<i>E</i>	1.5	0.5	0.5	-0.5
<i>M</i>	2	1	1	0
<i>F</i>	1.5	0.5	0.5	-0.5

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

Figure 2: Small company's expected payoffs at the prior

There are many other economic settings where committing to 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.

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.

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*) of the two-stage game for some information structure. That is, given the information structure, 1) for each combination of Look-Ignore choices in the first stage, agents play a Bayes Nash equilibrium (BNE) of the corresponding incomplete information game in the second stage; and 2) the Look-Ignore choices in the first stage constitute an equilibrium given the continuation play specified in 1).

¹That 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.

1.1 Preview of Results

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 *robust correlated equilibrium* (RCE), which captures the appropriate restrictions on the correlation structure for the environment of interest. RCE allows the action recommendations of players who are recommended to Look at their signals to be correlated across players and with the state. However, it does not allow for any form of correlation of the players' Look-Ignore recommendations or of the action recommendations of agents who are recommended to Ignore. Theorem 1 demonstrates that the set of RCE outcome distributions is equivalent to the set of PBE* outcome distributions across all information structures.²

We next show (Theorem 2) 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. Here a player's message specifies a vector of actions, one for each combination of Look-Ignore choices by the other players.

Our characterization demonstrates a subtlety: some outcomes are implementable only if we allow randomization at the Look-Ignore stage (Theorem 3). Surprisingly, in some cases the designer's optimal outcome can be achieved only in an equilibrium where some players Ignore their signals with positive probability.³ Also surprisingly, we find that the designer may need to use an asymmetric information structure even in a symmetric environment (the investment game in Section 3.1).

When the designer's optimal information structure from the baseline environment fails to be robust to strategic ignorance – because some player's “on-path” payoff when everyone Looks at their signals is lower than his “post-deviation” payoff in the worst continuation BNE after he deviates unilaterally to Ignore – the designer has two methods of adjusting the

²Notice that strategic ignorance can also restrict the set of implementable correlated equilibria (CE) in games of complete information, as we show in the companion working paper [Taneva and Wiseman \(2023\)](#). 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, a result we also prove in the paper.

³We thank Elliot Lipnowski for pushing us to investigate this question.

information structure in order to satisfy the Look constraints. The first method is to raise the on-path payoff of the player(s) whose Look constraint is violated. The second method is to lower the post-deviation payoff. Those changes interact with each other. If raising the on-path payoff involves changing the information that players get, then that change also affects the set of BNEs after a deviation to Ignore: the players who Looked still have that different information. Analogously, giving players different information in order to lower the payoff from the worst post-deviation BNE changes the on-path information structure as well. As a consequence, giving the players the option to Ignore messages does not necessarily make them better off. In some cases, all players get lower payoffs when strategic ignorance is possible than under the baseline where messages are automatically observed.⁴

A substantive difference between the standard information design environment and the environment with strategic ignorance of this paper is that while the set of implementable outcomes is always decreasing in the amount of exogenous information that players have in the former (Bergemann and Morris (2016)), that monotonicity may fail in the presence of strategic ignorance. We demonstrate this point in the context of the investment game example in Section 3.1.

In order to further explore the impact of strategic ignorance, we consider a couple of extensions in Section 4. The first is to allow multistage communication by the designer. In our main analysis, we assume that the designer sends signals only once. That is, a player sees all of his recommendations before choosing an action, rather than just his recommendation for the realized Look-Ignore decisions. An implication, as discussed above, is that any information that the designer gives him to help punish a potential deviation to Ignore by another player is also available on path. That extra information may limit what behavior the designer can induce on path. For example, Player 2 may need information about the state in order to punish Player 1 effectively, but knowing the state may make him unwilling to play the designer's preferred action on path. We find, however, 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. The reason is that providing signals separately means

⁴A 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, Brooks, and Morris (2017), for example, study the information structures over bidders' values that would minimize the distribution of winning bids in a first price auction.

that players must be incentivized to view each separate signal. Instead of facing a single constraint that players must be willing to view the bundle of recommendations when they expect others to follow the equilibrium strategy, now the designer faces a new constraint after each potential deviation. The result that dynamic, sequential recommendations may be strictly worse for the designer contrasts with the baseline information design setting where agents must observe their signals (e.g., [Makris and Renou \(2022\)](#)).

Our second extension yields another qualitative difference: allowing the players to communicate with each other after receiving their private signals may improve outcomes for the designer. Suppose that Player 2 is willing to punish Player 1 effectively only when Player 2 does not know the state,⁵ but that Player 2 must be informed on-path in order to play the designer's state-contingent desired action. In that case, the designer cannot always achieve her desired outcome, because she cannot both deter a deviation to Ignore by Player 1 and give Player 2 the necessary information on path. She can, however, solve that problem if the players can communicate, by giving the information intended for Player 2 to Player 1. If Player 1 chooses Look, then he can pass on Player 2's information to Player 2 (assuming that he has the incentive to do so, and that Player 2 has the incentive to receive it). If Player 1 deviates to Ignore, then Player 2 remains uninformed and willing to punish, and so that deviation is deterred.

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 there are multiple equilibria, and specifically when there is one equilibrium where players coordinate their randomization by following their signals and another equilibrium where they disregard their signals and randomize independently of each other. 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 *any* equilibrium of the dynamic game (the Look-Ignore stage followed by the action choice stage). The distinction between equilibrium selection and robustness to strategic ignorance is especially clear when there is a unique BNE at the action stage after any of the possible outcomes of the Look-Ignore stage, yet given those continuation payoffs, Ignore is strictly dominant at the Look-Ignore stage. It follows that

⁵For example, because the punishment action is dominated by a different action in each state, but is undominated at the prior.

in the unique PBE* of the dynamic game all players remain uninformed.

1.2 Applications and Relation to Literature

Previous research has identified many settings where incentives for strategic ignorance of a payoff-relevant state arise. In the context of relationship-specific investments which 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)).⁶ 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.

In this paper, the designer offers information both about a common payoff-relevant state and about the information received by other players, as well as a component of pure correlation. Refusing information results in both remaining uninformed about the state and being unable to coordinate one's actions with those of other players.⁷

⁶Golman, Hagmann, and Loewenstein (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.

⁷Schelling (1960), van Damme (1989), and Ben-Porath and Dekel (1992), in contrast, study settings where committing to remain uninformed of the *previous* action choices of an opponent can be beneficial for reversing the opponent's first-mover advantage. Whitmeyer (2022) investigates how the receiver in a signaling game may learn more from the sender by publicly committing, ahead of the sender's choice, to observe only a garbled version of the sender's action. Similarly, companies may limit their ability to monitor

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 player's strategic incentives to ignore information are a relevant concern. In the investment game in Section 3.1, 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.

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.⁸ 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.

Other real-life situations that reflect some of the features of our analysis include pharmaceutical executives and regulators refusing to be informed of the detailed results of ongoing clinical trials in order to create plausible deniability when unsafe and inefficacious drugs enter the market (e.g., the licensing case of Ketek, an antibiotic drug manufactured by Sanofi-Aventis and linked to liver failure) and members of networking and social media platforms choosing to unfollow certain other members or unsubscribe from particular services to send a publicly observable signal which will change the perceptions and actions of their followers.

The most closely related work is [Arcuri \(2021\)](#), which we became aware of shortly before posting the first draft of our paper. 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 prefers to the worst BNE for him under the information structure

the specific test results of employees in order to incentivize the take-up of training and licensing courses.

⁸Large 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.)

that results if he unilaterally Ignores his message. Then an outcome σ mapping states to action distributions is a “hear-no-evil Bayes correlated equilibrium” if it corresponds to a BNE of some information structure S that satisfies the hear-no-evil condition. That definition allows for the possibility that a player i prefers his worst BNE after deviating to Ignore over his outcome under σ .

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 (2022)). The main conceptual difference with these papers is that we allow the designer to provide information only once. Specifically, the designer cannot send additional messages to players contingent on their Look-Ignore choices. Importantly, as outlined above, conditioning the information sent to players on their previous decisions, and thus providing dynamic, sequential recommendations, may be detrimental to the designer in our setting. In contrast, this is always beneficial in the baseline information design setting explored in the aforementioned papers. An additional difference relative to Doval and Ely (2020) is that the extensive form in our environment is fixed, with players taking actions simultaneously in both stages.

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 ; $A \equiv A_1 \times \dots \times A_N$ is the set of action profiles. There is a finite set of states of the world Ω , with generic element ω . Agents’ payoffs are given by $u : A \times \Omega \rightarrow \mathbb{R}^N$, where agent i ’s payoff function $u_i : A \times \Omega \rightarrow \mathbb{R}$ depends on the action profile and the (ex ante unknown) state. The designer has a utility function $u^D : A \times \Omega \rightarrow \mathbb{R}$, 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 $T \equiv T_1 \times \dots \times T_N$; and 2) conditional signal distributions $P : \Omega \rightarrow \Delta(T)$, 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* (ℓ) at his signal and learn the realization of t_i , or to *Ignore* (g) it and remain uninformed. The Look-Ignore choices are public 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) , denote by $(T_{\mathcal{L}}, P_{\mathcal{L}})$ 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, $(T_{\mathcal{L}}, P_{\mathcal{L}})$ 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 , $(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$ and the prior μ . An agent who chooses to Ignore his signal uses $(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$ 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 $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$. 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 mapping from states to distributions over action profiles. A strategy for player i in dynamic game G^* is a tuple $(\gamma_i, (\tilde{\beta}_i^s)_s)$ with $\gamma_i \in \Delta\{\ell, g\}$, $\tilde{\beta}_i^s : T_i \rightarrow \Delta A_i$ 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 $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$ after 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 $(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$ 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 do 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

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

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

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

for all $a'_i \in A_i$.

Then $v(\tilde{\beta}^s) \in \Delta(A \times \Omega)$ defined as

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

for all $a \in A$ and $\omega \in \Omega$ is a BNE outcome (distribution)⁹ of $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$.

Definition 2. A strategy profile $(\gamma, (\tilde{\beta}^s)_s)$ 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 $\gamma_i(s_i) > 0$,

$$\begin{aligned} & \sum_{s_{-i}, a, \omega} \prod_{j \neq i} \gamma_j(s_j) v(\tilde{\beta}^{s_i, s_{-i}})(a, \omega) u_i(a_i, a_{-i}, \omega) \\ & \geq \sum_{s_{-i}, a, \omega} \prod_{j \neq i} \gamma_j(s_j) v(\tilde{\beta}^{s'_i, s_{-i}})(a, \omega) u_i(a_i, a_{-i}, \omega), \quad (4) \end{aligned}$$

⁹Throughout the paper, we use the terms “outcome” and “outcome distribution” interchangeably.

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 $G^*(T, P)$.

The designer chooses an information structure (T, P) to maximize her expected payoff across the set of all PBE* outcomes $\cup_{(T,P)} PBE^*(G^*(T, P))$.

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 BNEs across all possible information structures equals the set of BCEs. 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* across all possible information structures equals the set of BCEs of the two-stage game in our setting – is too strong. The designer 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 satisfies those constraints on the correlation structure a robust correlated equilibrium (RCE), defined formally in Appendix A. Theorem 1 establishes that the outcome of any PBE* can be achieved in a RCE.

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

An implication of that equivalence result is that without loss of generality we can restrict the designer to selecting a *direct contingent* information structure. 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 $T_i = \mathcal{A}_i$ for each agent i , where $\mathcal{A}_i \equiv A_i^{|S_{-i}|}$ denotes the set of agent i 's (pure) mappings from S_{-i} to A_i . 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 $\mathcal{A} \equiv \mathcal{A}_1 \times \dots \times \mathcal{A}_N$. Say that an outcome $v \in \Delta(A \times \Omega)$ is *implementable with direct contingent recommendations* if there exists a conditional message distribution $P : \Omega \rightarrow \Delta(\mathcal{A})$ such that v is a PBE* outcome of $G^*(\mathcal{A}, P)$.

Theorem 2. *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 2 in order to characterize the solution to the designer's problem: we derive the optimal information structure (\mathcal{A}, P) directly from the optimal RCE. Since we assume advantageous equilibrium selection, conceptually we can think of the designer as optimizing over direct contingent information structures (\mathcal{A}, P) by choosing P^* , and then nature optimizing on the designer's behalf over the set of PBE* of $G^*(\mathcal{A}, P^*)$. In the investment game example of Section 3.1, we construct the optimal RCE and explicitly show how to derive the optimal direct contingent information structure from it. In the remaining examples, for the sake of brevity we present only the optimal direct contingent information structure.

2.2 The necessity of ignorance

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, though, that restriction turns out not to be innocuous.¹⁰

Theorem 3. *A PBE* outcome v may be implementable only if $\gamma_i(g) > 0$ for some $i \in \mathcal{I}$.*

In Appendix B we present a two-agent example where the designer does strictly better by relying on an equilibrium where one agent randomizes between Look and Ignore, which proves this result. In the example, the binding constraint for the designer is to incentivize Player 1 to Look at his signal. The structure of the basic game is such that there is no BNE that gives Player 1 a low enough payoff to deter his deviation to Ignore unless it is common knowledge that Player 2 also is completely uninformed. On path, however, the designer

¹⁰Relatedly, 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 Look with certainty.

must give Player 2 information so that he can play her state-dependent desired action. 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.

2.3 The harm of ignorance

Given the many examples from game theory where in equilibrium flexibility harms a player, it is not surprising that in some cases all agents are worse off when they have the option to Ignore their messages, relative to the baseline where messages are automatically observed. We provide such an example in the Online Appendix.

In that example, an information structure that reveals the state perfectly maximizes the players' payoffs. However, if players have the ability to exercise strategic ignorance, then it is a conditionally dominant strategy to Ignore that information structure. The example has the flavor of a prisoners' dilemma, 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. As a result, a designer who wants to maximize the total expected payoff of the players must provide a direct information structure that is less than perfectly informative, and the players get lower payoffs than they would if messages were automatically observed.

In this particular example, the ability to strategically ignore information is harmful to the players due to their own choices given a fixed information structure. It is also possible to construct examples where the potential for strategic ignorance harms the players 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.

3 Two economic examples

Here we examine the impact of strategic ignorance on information design in two economic settings, investment choice and currency attacks. Each setting illustrates two important

general findings. First, even when strategic ignorance would not benefit any agent in the underlying basic game, it can still impose restrictions on an information designer. Second, a key tension for the designer is whether or not, if an agent i deviates and Ignores his message, the other agent(s) are still willing to follow their original recommendations. If so, then agent i cannot gain from the deviation. If not – because their recommendations no longer provide information about player i 's action, although they are still informative about the state – then unless there is another BNE worse for player i than the original target outcome, the designer must adjust the information structure. The designer has a variety of ways to make that adjustment. In the investment game, her optimal response is to provide less information about the state of the world. For currency attacks, she provides more information.

3.1 Investment game

First we study a version of the parameterized basic game from [Taneva \(2019\)](#). 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 [Figure 3](#).

$$\begin{array}{c}
 \begin{array}{cc}
 & E & F \\
 E & \boxed{2, 2} & \boxed{1, 0} \\
 F & \boxed{0, 1} & \boxed{0, 0}
 \end{array}
 \quad , \quad
 \begin{array}{cc}
 & E & F \\
 E & \boxed{0, 0} & \boxed{0, 1} \\
 F & \boxed{1, 0} & \boxed{2, 2}
 \end{array} \\
 \omega = e & & \omega = f
 \end{array}$$

Figure 3: Investment game

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.

3.1.1 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 symmetric¹¹ BCE outcome distributions represented in Figure 4,

	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

$\omega = f$

Figure 4: Parameterized symmetric BCE outcome distributions

where q captures the probability that each player receives the action recommendation that corresponds to the state and r captures the probability that both players together receive the action recommendation that corresponds to the state.

The set of BCE outcome distributions is then the triangle in solid purple in Figure 5. The red line represents the obedience constraint¹², while the 45-degree line and the blue line represent the constraints on the parameters that ensure the outcome is a proper probability distribution. The level lines for the designer's expected payoff are given by the solid black 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:

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

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

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

¹¹If there is an asymmetric (across players and/or states) BCE outcome distribution that maximizes the designer's expected payoff, then due to the symmetry of the game and the designer's objective there is a "mirror image" of that distribution, which is also a BCE outcome and results in the same designer payoff. A convex combination with equal weights on those two BCE outcome distributions is symmetric, and, since the BCE outcome set is convex, is also a BCE outcome distribution. Since the designer's expected payoff is linear in the probabilities, this distribution will give the same value as the original optimal asymmetric distribution. Additionally, note that Figure 4 represents distributions over action profiles conditional on each state; the unconditional outcome distributions are obtained by multiplying these with the prior.

¹²There are two obedience constraints – one for action E and one for action F – but they reduce to the same inequality due to the symmetry.




Figure 5: BCE set and designer's expected payoff

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 $\frac{2}{3}$.¹³

3.1.2 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

¹³We 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.

(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 $\frac{1}{2} \left(\frac{2}{3} \cdot 2 + \frac{1}{3} \cdot 1 \right) = \frac{5}{6}$, 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 Figure 5 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.¹⁴

In order to satisfy the constraint that firms be willing to Look at their signals¹⁵, 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 and 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 the designer recommends the correct action more frequently on path, makes Firm 1 willing to Look.

Specifically, we calculate that the optimal RCE (γ, β^g, π) corresponds to the direct contingent information structure in Figure 6, where $\alpha \equiv \frac{1}{\sqrt{3}} \approx 0.577$. 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

¹⁴The derivations are in Appendix C.

¹⁵Numerical estimation showed that the optimal RCE does not require mixing between Look and Ignore.

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.

$\pi(\cdot e)$	<i>EF</i>	<i>FE</i>	<i>FF</i>
<i>EF</i>	$\frac{1-\alpha}{2} \approx 0.2115$	0	0
<i>FE</i>	0	$\alpha \approx 0.577$	0
<i>FF</i>	0	0	$\frac{1-\alpha}{2} \approx 0.2115$

$\omega = e$

$\pi(\cdot f)$	<i>EF</i>	<i>FE</i>	<i>FF</i>
<i>EF</i>	$1-\alpha \approx 0.423$	0	0
<i>FE</i>	0	$1-\alpha \approx 0.423$	0
<i>FF</i>	0	0	$2\alpha - 1 \approx 0.154$

$\omega = f$

Figure 6: Optimal direct contingent information structure

We note that the state-wise mirror image of that optimal RCE denoted by $(\gamma, \bar{\beta}^g, \bar{\pi})$, is also an optimal RCE. An equally weighted convex combination of those two, however, is not. Indeed, while the BCE set is always convex, the RCE 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 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 RCEs is not an RCE itself. This non-convexity extends to the outcome sets: the equally weighted convex combination of the two RCE outcome distributions is not an RCE outcome distribution. This is in contract to the BCE outcome set, which is always convex.

3.1.3 Exogenous information

In the standard information design environment, the set of implementable outcomes is decreasing 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 given by (A, \tilde{P}) . When the players start out with no exogenous information, this distribution is not implementable under strategic ignorance. However, if they start out with exactly the exogenous information structure implied by this outcome distribution, then the designer could simply reveal nothing further and there will be a BNE that implements that distribution. Hence, by increasing the exogenous information of the players in this way, an outcome that was previously not implementable becomes implementable in the presence of strategic ignorance. If we further increase the exogenous information of the agents to fully reveal the state, then the designer's optimal outcome distribution becomes again not implementable. Thus, the set of implementable outcome distributions under strategic ignorance is non-monotone in the players' exogenous information.

3.2 Currency attacks

We next consider a model of currency attacks. Here the effect of strategic ignorance is much more dramatic. There are $N \geq 2$ symmetric players deciding whether or not to attack a currency (*attack* or *not*). 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 a :

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

and

$$u_i(n, K; \omega) = 0 \text{ for all } K, \omega.$$

We assume $x \geq 1$, so that the payoff when all players attack is at least as high when the currency is weak as when it is strong.

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

3.2.1 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

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

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 player will choose n .

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

3.2.2 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 $x > 1$, and suppose that Player i deviates to Ignore. When Player i is uninformed, then k is dominant at the action stage: as shown in Figure 7, it gives a strictly positive payoff against any strategy profile mapping states to actions for the other $N - 1$ players, while n gives 0. We can summarize a strategy profile for the other players as (K_W, K_S) denoting the number who play k in each state.

	$K_W = N - 1,$ $K_S > 0$	$K_W < N - 1,$ $K_S = 0$	$K_W = N - 1,$ $K_S = 0$	$K_W < N - 1,$ $K_S > 0$
k	$\frac{x+1}{2}$	$\frac{1}{2}$	$\frac{x-1}{2}$	$\frac{3}{2}$
n	0	0	0	0

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

Figure 7: a is dominant for an uninformed player

In either state, the unique best response for any other player when Player i chooses k is k . The outcome is thus (k, \dots, k) regardless of the designer's recommendations, and

Player i 's resulting payoff is $\frac{1}{2} \cdot x + \frac{1}{2} \cdot 1 > \frac{1}{4}x$. It follows that deviating to Ignore is profitable.

In fact, when $x > 1$ 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.¹⁶

If $x = 1$, then the situation changes. From Figure 7, 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 0; 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 $(W : k, S : n)$. Player i 's unique best response to anything other than the strategy profile of $(W : k, S : n)$ for all opponents is k , and the rest of the argument is the same as in the $x > 1$ case.

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

Under (\mathcal{A}, P^*) , it is an equilibrium for all players to Look at and follow their recommendations, yielding payoff $x/2 = \frac{1}{2}$. If Player 1 deviates to Ignore, then there is a BNE where he plays n and all other players follow their recommendation by playing $(W : k, S : n)$. That BNE gives Player 1 a payoff of $0 < \frac{1}{2}$, so the deviation to Ignore is not profitable.

An interesting feature of the optimal information structure (\mathcal{A}, P^*) is that, as just argued, the constraint that players must be willing to view their signals is slack. In the invest-

¹⁶This 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) ," and so *every* equilibrium outcome at the Look-Ignore stage involves at least one player choosing Ignore. We are still selecting the designer's preferred equilibrium of the dynamic game, but there is only one outcome to select from.

ment game, the designer optimally modified the baseline information structure by raising the players' on-path payoffs and lowering the post-deviation payoffs so that the constraint was just satisfied. In the currency attack game with $x = 1$, the worst post-deviation BNE payoff is constant with respect to the information until a discontinuous downward jump when players become fully informed about the state. Consequently, the constraint is either strictly violated or strictly satisfied. Another qualitative difference is that in the investment game, the designer adjusts by giving the players less precise information about the state, and in the currency attack game she gives them more precise information. A qualitative similarity of the investment and the currency attack games is that the players are better off under strategic ignorance.¹⁷ However, recall from Section 2.3 that this need not be the case in general. In the Online Appendix we provide an example where both players are worse off under the ability to strategically ignore information.

4 Two Counterintuitive Results

In this section we consider two extensions presented in the context of specific examples, which demonstrate the reversal of standard information design results in the presence of strategic ignorance. The first example shows that the designer may do strictly worse by giving recommendations in multiple rounds (once before the Look-Ignore choices have been made and then conditional on the specific Look-Ignore profile that has realised) rather than all at once upfront, before the agents decide whether to Look or Ignore. The second example shows that strategic communication between the players can actually be beneficial for the designer.

4.1 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 this subsection we provide a two-player example that shows that the designer

¹⁷The designer is always (weakly) worse off under strategic ignorance due to the added incentive constraints.

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.

Consider the following state-contingent payoff matrices in Figure 8. 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.

$$\begin{array}{c} \begin{array}{cc} & X & Y \\ E & \boxed{0, 0} & \boxed{1, 1} \\ F & \boxed{2, 2} & \boxed{1, 1} \end{array} \\ \omega = e \end{array}, \begin{array}{c} \begin{array}{cc} & X & Y \\ E & \boxed{1, 0} & \boxed{0, 0} \\ F & \boxed{1, 1} & \boxed{1, 1} \end{array} \\ \omega = f \end{array}$$

Figure 8: State-contingent payoffs

At the prior, the players' expected payoffs are given in Figure 9.

$$\begin{array}{c} \begin{array}{cc} & X & Y \\ E & \boxed{\frac{1}{2}, 0} & \boxed{\frac{1}{2}, \frac{1}{2}} \\ F & \boxed{\frac{3}{2}, \frac{3}{2}} & \boxed{1, 1} \end{array} \\ \Pr(\omega = e) = \frac{1}{2} \end{array}$$

Figure 9: 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

$$\tilde{P}(A, Y|\omega = a) = \tilde{P}(B, Y|\omega = b) = 1.$$

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 above outcome is robust to strategic ignorance, if the designer can provide the direct contingent recommendations all at once. Specifically, the

outcome can be implemented by providing the following direct contingent information structure (\mathcal{A}, P^*) :

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

Now both players' messages reveal the state. 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 $s = (\ell, \ell)$.

Giving punishment recommendations only after a deviation to Ignore. Here, we have in mind the following 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 $(1, \frac{1}{2})$. 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 $(\frac{3}{2}, \frac{3}{2})$. Therefore, conditional on Player 1 choosing to Ignore the initial message, it is a 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. 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 need any information on path.

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 in 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 who has initially chosen to Look only after his opponent has deviated to Ignore means having to satisfy a second Look constraint, and that effect makes the designer worse off.

4.2 Communication between Players

In the standard information design environment without strategic ignorance, the possibility of players communicating their private signals with each other can never be strictly ben-

eficial 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.

The example of this subsection builds upon the complete information game in Figure 10. 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 $(2, 2)$.

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

Figure 10: Complete information game

Next we add two states of nature, which give rise to the payoff matrices in Figure 11. 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.

	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
	$\omega = 2$				

Figure 11: State-contingent payoffs

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 lower payoffs than any other action of Player 2.

At the prior, when both players are uninformed, their expected payoffs are given in Figure 12. The designer gets a payoff of 1 if (X_ω, F_ω) is played in state ω , and 0 otherwise.

	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

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

Figure 12: Expected payoffs at the prior

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. The baseline information structure 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 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 2.¹⁸ The designer achieves her maximal payoff of 1, and the players’ payoffs are $(2, 2)$.

¹⁸We assume communication between the players is in the form of verifiable all-or-nothing disclosure, where the receiver can refuse to be informed.

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 $1 < 2$. Finally, 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.

Interpretation. In this example, given 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 the perfectly informative signal to him only, after which he would want to pass it on to Player 2. As long as Player 2 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 1 and perfectly negatively correlated in state 0. 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 others' actions.

Appendix

A Correlated Equilibrium Equivalence Result

We will start by providing the appropriate definition of correlated equilibrium in our setting. Recall that the designer can provide correlation of strategies only at the action stage and not at the Look-Ignore stage, and only for those agents who choose Look. That limitation implies that agents' Look-Ignore choices must be independent of each other and independent of the state ω , and that the action-stage choices of an agent who chose Ignore must be independent of ω and of the actions of other agents (although the agent may condition on the observed Look-Ignore choices of the other agents s_{-i}).

Therefore, the object of interest is an element

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

where γ denotes the distributions of Look-Ignore recommendations, β^g denotes the distri-

butions of the post-Ignore recommendations, and π denotes the joint distribution of post-Look recommendations and the state.

The timing proceeds as follows:

1. The designer commits to and publicly announces (γ, β^g, π) .
2. Look-Ignore recommendations $s \in S$ are drawn from γ , post-Ignore recommendations from β^g , and (m, ω) from π .
3. The designer privately recommends the realization s_i to each agent i .
4. Each agent i chooses \tilde{s}_i , which is publicly observed.
5. The designer sends the realized draws from step 2 corresponding to the choices in Step 4: $m_i \in \mathcal{A}_i$ to each i such that $\tilde{s}_i = \ell$, and $a_i(\tilde{s}_{-i}) \in A_i$ to each i such that $\tilde{s}_i = g$.
6. Each agent i makes an action choice \tilde{a}_i .

A robust correlated equilibrium is a triple (γ, β^g, π) such that 1) for all $\tilde{s} \in S$, the action recommendations sent in Step 5 are obedient and 2) the Look-Ignore recommendations 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 a robust correlated equilibrium (RCE) of G^* if

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

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

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 $\beta_i^g(a_i|s_{-i}) > 0$

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

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

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

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

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

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

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

We can now prove Theorem 1.

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

Given profile $s \in S$ of Look-Ignore choices, let $\mathcal{A}_{\mathcal{L}(s)} \equiv \times_{i \in \mathcal{L}(s)} \mathcal{A}_i$ and $P_{\mathcal{L}(s)}(m_{\mathcal{L}(s)}|\omega) := \pi(m_{\mathcal{L}(s)}, \omega)/\mu(\omega)$. In $G(\mathcal{A}_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$ consider the following strategy for all 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 $m_i \in \mathcal{A}_i$, and for all 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 $m_i \in \mathcal{A}_i$ and choosing action $a_i \in A_i$ when his opponents play according to $\tilde{\beta}_{-i}^s$ is given by

$$\begin{aligned} & \sum_{a_{-i}, m_{\mathcal{L}(s) \setminus i}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(m_i, m_{\mathcal{L}(s) \setminus i} | \omega) \prod_{j \in \mathcal{L}(s) \setminus i} \tilde{\beta}_j^s(a_j | m_j) \prod_{k \in \mathcal{G}(s)} \tilde{\beta}_k^s(a_k) u_i(a_i, a_{-i}, \omega) \\ &= \sum_{m_{\mathcal{L}(s) \setminus i}, a_{\mathcal{G}(s)}, \omega} \pi(m_i, m_{\mathcal{L}(s) \setminus i}, \omega) \prod_{k \in \mathcal{G}(s)} \beta_k^g(a_k | s_{-k}) u_i(a_i, (m_j(s_{-j}))_{j \in \mathcal{L}(s) \setminus i}, a_{\mathcal{G}(s)}, \omega). \end{aligned} \quad (8)$$

Hence, by (5) we obtain

$$\begin{aligned} & \sum_{a_{-i}, m_{\mathcal{L}(s) \setminus i}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(m_i, m_{\mathcal{L}(s) \setminus i} | \omega) \prod_{j \in \mathcal{L}(s) \setminus i} \tilde{\beta}_j^s(a_j | m_j) \prod_{k \in \mathcal{G}(s)} \tilde{\beta}_k^s(a_k) u_i(m_i(s_i), a_{-i}, \omega) \\ & \geq \sum_{a_{-i}, m_{\mathcal{L}(s) \setminus i}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(m_i, m_{\mathcal{L}(s) \setminus i} | \omega) \prod_{j \in \mathcal{L}(s) \setminus i} \tilde{\beta}_j^s(a_j | m_j) \prod_{k \in \mathcal{G}(s)} \tilde{\beta}_k^s(a_k) u_i(a'_i, a_{-i}, \omega). \end{aligned} \quad (9)$$

for all $i \in \mathcal{L}(s)$, $m_i \in \mathcal{A}_i$, and $a'_i \in A_i$. This establishes the BNE interim incentive compatibility constraint (1) for all $i \in \mathcal{L}(s)$, $m_i \in \mathcal{A}_i$, 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

$$\begin{aligned} & \sum_{a_{-i}, m_{\mathcal{L}(s)}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(m_{\mathcal{L}(s)} | \omega) \prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j | m_j) \prod_{k \in \mathcal{G}(s) \setminus i} \tilde{\beta}_k^s(a_k) u_i(a_i, a_{-i}, \omega) \\ &= \sum_{m_{\mathcal{L}(s)}, a_{\mathcal{G}(s) \setminus i}, \omega} \pi(m_{\mathcal{L}(s)}, \omega) \prod_{k \in \mathcal{G}(s) \setminus i} \beta_k^g(a_k | s_{-k}) u_i(a_i, (m_j(s_{-j}))_{j \in \mathcal{L}(s)}, a_{\mathcal{G}(s) \setminus i}, \omega). \end{aligned} \quad (10)$$

Hence, by (6) we obtain

$$\begin{aligned} & \sum_{a_{-i}, m_{\mathcal{L}(s)}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(m_{\mathcal{L}(s)} | \omega) \prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j | m_j) \prod_{k \in \mathcal{G}(s) \setminus i} \tilde{\beta}_k^s(a_k) u_i(a_i, a_{-i}, \omega) \\ & \geq \sum_{a_{-i}, m_{\mathcal{L}(s)}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(m_{\mathcal{L}(s)} | \omega) \prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j | m_j) \prod_{k \in \mathcal{G}(s) \setminus i} \tilde{\beta}_k^s(a_k) u_i(a'_i, a_{-i}, \omega) \end{aligned} \quad (11)$$

for all $i \in \mathcal{G}(s)$, 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 all $i \in \mathcal{G}(s)$ 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 $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$. Then $v(\tilde{\beta}^s)$ defined as

$$\begin{aligned} v(\tilde{\beta}^s)(a, \omega) &:= \sum_{m_{\mathcal{L}(s)}} \mu(\omega) P_{\mathcal{L}(s)}(m_{\mathcal{L}(s)}|\omega) \left(\prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j|m_j) \prod_{k \in \mathcal{G}(s)} \tilde{\beta}_k^s(a_k) \right) \\ &= \sum_{m_{\mathcal{L}(s)}} \pi(m_{\mathcal{L}(s)}, \omega) \left(\prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j|m_j) \prod_{k \in \mathcal{G}(s)} \tilde{\beta}_k^s(a_k) \right) \end{aligned} \quad (12)$$

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

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

$$\begin{aligned} &\sum_{s_{-i}, a, \omega} \prod_{j \neq i} \gamma_j(s_j) \left(\sum_{m_{\mathcal{L}(s)}} \pi(m_{\mathcal{L}(s)}, \omega) \left(\prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j|m_j) \prod_{k \in \mathcal{G}(s)} \tilde{\beta}_k^s(a_k) \right) \right) u_i(a_i, a_{-i}, \omega) \\ &= \sum_{s_{-i}, a, \omega} \prod_{j \neq i} \gamma_j(s_j) v(\tilde{\beta}^s)(a, \omega) u_i(a_i, a_{-i}, \omega) \\ &\geq \sum_{s_{-i}, a, \omega} \prod_{j \neq i} \gamma_j(s_j) v(\tilde{\beta}^{s'}) (a, \omega) u_i(a_i, a_{-i}, \omega) \\ &= \sum_{s'_{-i}, a, \omega} \prod_{j \neq i} \gamma_j(s_j) \left(\sum_{m_{\mathcal{L}(s')}} \pi(m_{\mathcal{L}(s')}, \omega) \left(\prod_{j \in \mathcal{L}(s')} \tilde{\beta}_j^{s'}(a_j|m_j) \prod_{k \in \mathcal{G}(s')} \tilde{\beta}_k^{s'}(a_k) \right) \right) u_i(a_i, a_{-i}, \omega), \end{aligned} \quad (13)$$

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 $G^*(\mathcal{A}, P)$. 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 $G^*(\mathcal{A}, P)$, that is $\hat{v} \in PBE^*(G^*(\mathcal{A}, P))$.

Notice that for all $a \in A$ and $\omega \in \Omega$

$$\begin{aligned} \widehat{v}(a, \omega) &= \sum_{s \in S} \prod_{i \in \mathcal{I}} \gamma_i(s_i) v(\tilde{\beta}^s)(a, \omega) \\ &= \sum_{s \in S} \prod_{i \in \mathcal{I}} \gamma_i(s_i) \left(\sum_{m_{\mathcal{L}(s)}: (m_j(s-j))_{j \in \mathcal{L}(s)} = a_{\mathcal{L}(s)}} \pi(m_{\mathcal{L}(s)}, \omega) \right) \prod_{k \in \mathcal{G}(s)} \beta_k^g(a_k | s_{-k}) = v(\gamma, \beta^g, \pi)(a, \omega). \end{aligned} \quad (14)$$

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

Next, we prove that $RCE(G^*) \supseteq \cup_{(T,P)} PBE^*(G^*(T, P))$. 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 all $i \in \mathcal{I}$ define $\beta_i^g : S_{-i} \rightarrow \Delta A_i$ 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 $\beta^g = \times_i \beta_i^g$. Define $\pi \in \Delta(\mathcal{A} \times \Omega)$ such that for all $s \in S$

$$\pi(m_{\mathcal{L}(s)}, \omega) = \sum_{t_{\mathcal{L}(s)}} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)} | \omega) \prod_{i \in \mathcal{L}(s)} \tilde{\beta}_i^s(a_i | t_i) \quad (15)$$

for all $a_{\mathcal{L}(s)} \in \times_{i \in \mathcal{L}(s)} A_i$ and $m_{\mathcal{L}(s)}$ such that $(m_j(s-j))_{j \in \mathcal{L}(s)} = a_{\mathcal{L}(s)}$. Notice, this ensures that $\pi(\mathcal{A} \times \{\omega\}) = \mu(\omega)$ for all $\omega \in \Omega$.

Multiplying both sides of (1) by $\tilde{\beta}_i^s(a_i | t_i)$ and summing across t_i we obtain for all $s \in S$,

$i \in \mathcal{L}(s)$, and $a_i, a'_i \in A_i$

$$\begin{aligned}
& \sum_{a_{-i}, \omega} \left(\sum_{t_{\mathcal{L}(s)}} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)} | \omega) \prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j | t_j) \right) \prod_{k \in \mathcal{G}(s)} \tilde{\beta}_k^s(a_k) u_i(a_i, a_{-i}, \omega) \\
&= \sum_{m_{\mathcal{L}(s) \setminus i}, a_{\mathcal{G}(s)}, \omega} \pi(m_{\mathcal{L}(s)}, \omega) \prod_{k \in \mathcal{G}(s)} \beta_k^g(a_k | s_{-k}) u_i(m_i(s_{-i}), (m_j(s_{-j}))_{j \in \mathcal{L}(s) \setminus i}, a_{\mathcal{G}(s)}, \omega) \\
&\geq \sum_{m_{\mathcal{L}(s) \setminus i}, a_{\mathcal{G}(s)}, \omega} \pi(m_{\mathcal{L}(s)}, \omega) \prod_{k \in \mathcal{G}(s)} \beta_k^g(m_k(s_{-k})) u_i(a'_i, (m_j(s_{-j}))_{j \in \mathcal{L}(s) \setminus i}, a_{\mathcal{G}(s)}, \omega) \\
&= \sum_{a_{-i}, \omega} \left(\sum_{t_{\mathcal{L}(s)}} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)} | \omega) \prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j | t_j) \right) \prod_{k \in \mathcal{G}(s)} \tilde{\beta}_k^s(a_k) u_i(a'_i, a_{-i}, \omega) \quad (16)
\end{aligned}$$

which establishes (5).

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

$$\begin{aligned}
& \sum_{a_{-i}, \omega} \left(\sum_{t_{\mathcal{L}(s)}} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)} | \omega) \prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j | t_j) \right) \prod_{k \in \mathcal{G}(s) \setminus i} \tilde{\beta}_k^s(a_k) u_i(a_i, a_{-i}, \omega) \\
&= \sum_{m_{\mathcal{L}(s)}, a_{\mathcal{G}(s) \setminus i}, \omega} \pi(m_{\mathcal{L}(s)}, \omega) \prod_{k \in \mathcal{G}(s) \setminus i} \beta_k^g(a_k | s_{-k}) u_i(a_i, (m_j(s_{-j}))_{j \in \mathcal{L}(s)}, a_{\mathcal{G}(s) \setminus i}, \omega) \\
&\geq \sum_{m_{\mathcal{L}(s)}, a_{\mathcal{G}(s) \setminus i}, \omega} \pi(m_{\mathcal{L}(s)}, \omega) \prod_{k \in \mathcal{G}(s) \setminus i} \beta_k^g(a_k | s_{-k}) u_i(a'_i, (m_j(s_{-j}))_{j \in \mathcal{L}(s)}, a_{\mathcal{G}(s) \setminus i}, \omega) \\
&= \sum_{a_{-i}, \omega} \left(\sum_{t_{\mathcal{L}(s)}} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)} | \omega) \prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j | t_j) \right) \prod_{k \in \mathcal{G}(s) \setminus i} \tilde{\beta}_k^s(a_k) u_i(a'_i, a_{-i}, \omega), \quad (17)
\end{aligned}$$

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

$$\begin{aligned}
& \sum_{s_{-i}, a, \omega} \prod_{j \neq i} \gamma_j(s_j) \left(\sum_{t_{\mathcal{L}(s)}} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)} | \omega) \prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j | t_j) \right) \prod_{k \in \mathcal{G}(s)} \tilde{\beta}_k^s(a_k) u_i(a_i, a_{-i}, \omega) \\
&= \sum_{s_{-i}, m_{\mathcal{L}(s)}, a_{\mathcal{G}(s)}, \omega} \prod_{j \neq i} \gamma_j(s_j) \pi(m_{\mathcal{L}(s)}, \omega) \prod_{k \in \mathcal{G}(s)} \beta_k^g(a_k | s_{-k}) u_i((m_j(s_{-j}))_{j \in \mathcal{L}(s)}, a_{\mathcal{G}(s)}, \omega) \\
&\geq \sum_{s'_{-i}, m_{\mathcal{L}(s')}, a_{\mathcal{G}(s')}, \omega} \prod_{j \neq i} \gamma_j(s'_j) \pi(m_{\mathcal{L}(s')}, \omega) \prod_{k \in \mathcal{G}(s')} \beta_k^g(a_k | s'_{-k}) u_i((m_j(s'_{-j}))_{j \in \mathcal{L}(s')}, a_{\mathcal{G}(s')}, \omega), \\
&= \sum_{s'_{-i}, a, \omega} \prod_{j \neq i} \gamma_j(s_j) \left(\sum_{t_{\mathcal{L}(s')}} \mu(\omega) P_{\mathcal{L}(s')}(t_{\mathcal{L}(s')} | \omega) \prod_{j \in \mathcal{L}(s')} \tilde{\beta}_j^{s'}(a_j | t_j) \right) \prod_{k \in \mathcal{G}(s')} \tilde{\beta}_k^{s'}(a_k) u_i(a_i, a_{-i}, \omega)
\end{aligned} \tag{18}$$

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 a RCE of G^* . Then, $v(\gamma, \beta^g, \pi) \in \Delta(A \times \Omega)$ is a RCE outcome of G^* , that is $v \in RCE(G^*)$. Notice that

$$\begin{aligned}
v(\gamma, \beta^g, v)(a, \omega) &= \sum_{s \in \mathcal{S}} \prod_{i \in \mathcal{I}} \gamma_i(s_i) \left(\sum_{m_{\mathcal{L}(s)}: (m_j(s_{-j}))_{j \in \mathcal{L}(s)} = a_{\mathcal{L}(s)}} \pi(m_{\mathcal{L}(s)}, \omega) \right) \prod_{k \in \mathcal{G}(s)} \beta_k^g(a_k | s_{-k}) \\
&= \sum_{s \in \mathcal{S}} \prod_{i \in \mathcal{I}} \gamma_i(s_i) \sum_{t_{\mathcal{L}(s)}} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)} | \omega) \left(\prod_{j \in \mathcal{L}(s)} \tilde{\beta}_j^s(a_j | t_j) \prod_{k \in \mathcal{G}(s)} \tilde{\beta}_k^s(a_k) \right) = \bar{v}(a, \omega) \tag{19}
\end{aligned}$$

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

B Example: the necessity of ignorance

Consider the following example, 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 Figure 13.

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,

	L	R_e	R_f	P_e	P_f	Q
E	3, 0	1, 1	1, 1	3, 0	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, 0	-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

$\omega = f$

Figure 13: State-contingent payoffs

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 Figure 14.

	L	R_e	R_f	P_e	P_f	Q
E	1.5, 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	1.5, 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

Figure 14: Expected payoffs at the prior

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 example, 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 eqm 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 $3p + (1 - p) = 2p + 1 > 2$.

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 (\mathcal{A}, P) with

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

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 $\frac{1}{2}1 + \frac{1}{2}\frac{1}{2} = \frac{3}{4}$.

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 her recommendation must give her 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 in the previous section cannot be duplicated in a pure Look-Look equilibrium.

C Investment game

Consider the parameterized symmetric direct information structures of Figure 4. 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 is choosing to Look, it is dominant strategy for the player who has chosen to Look to play the action that corresponds to the most likely state given his signal, while the payoff of the player who has chosen Ignore is independent of 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 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 Figure 5 and the area to 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.

References

- Akerlof, George A. 1970. “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism.” *Quarterly Journal of Economics* 84 (3):488–500. [1.2](#)
- Arcuri, Alesandro. 2021. “Hear-No-Evil Equilibrium.” Working paper, Cornell University. [1.2](#)
- Atayev, Atabek. 2022. “Information Asymmetry and Search Intensity.” Working paper, ZEW – Leibniz Centre for European Economic Research in Mannheim. [1.2](#)
- Athey, Susan and Jonathan Levin. 2001. “Information and Competition in U.S. Forest Service Timber Auctions.” *Journal of Political Economy* 109 (2):375–417. [1.2](#), [8](#)
- Ben-Porath, Elchanan and Eddie Dekel. 1992. “Signaling Future Actions and the Potential for Sacrifice.” *Journal of Economic Theory* 57 (1):36–51. [7](#)
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2017. “First-Price Auctions with General Information Structures in Games: Implications for Bidding and Revenue.” *Econometrica* 85 (1):107–143. [4](#)
- Bergemann, Dirk and Stephen Morris. 2016. “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games.” *Theoretical Economics* 11:487–522. [1](#), [1.1](#), [3.1.3](#)
- . 2019. “Information Design: A Unified Perspective.” *Journal of Economic Literature* 57 (1):44–95. [1](#)
- de Oliveira, Henrique and Rohit Lamba. 2019. “Rationalizing Dynamic Choices.” Working paper, Penn State University. [1.2](#)
- Deimen, Inga and Desz  Szalay. 2019. “Delegated Expertise, Authority, and Communication.” *American Economic Review* 109 (4):1349–1374. [1.2](#)
- Doval, Laura and Jeffrey Ely. 2020. “Sequential Information Design.” *Econometrica* 86 (6):2575–2608. [1.2](#)
- Gal-Or. 1987. “First Mover Disadvantages with Private Information.” *Review of Economic Studies* 54 (2):279–292. [1.2](#)

- Golman, Russell, David Hagmann, and George Loewenstein. 2017. “Information Avoidance.” *Journal of Economic Literature* 55 (1):96–135. 6
- Gul, Faruk. 2001. “Unobservable Investment and the Hold-Up Problem.” *Econometrica* 69 (2):343–376. 1.2
- Hirshleifer, Jack. 1971. “The Private and Social Value of Information and the Reward to Inventive Activity.” *American Economic Review* 61 (4):561–574. 1.2
- Kamenica, Emir and Matthew Gentzkow. 2011. “Bayesian Persuasion.” *American Economic Review* 101 (6):2590–2615. 1
- Kessler, Anke S. 1998. “The Value of Ignorance.” *RAND Journal of Economics* 29 (2):339–354. 1.2
- Makris, Miltiadis and Ludovic Renou. 2022. “Information Design in Multi-Stage Games.” *Theoretical Economics* :forthcoming. 1.1, 1.2
- McAdams, David. 2012. “The Value of Ignorance.” *Economic Letter* 114:83–85. 1.2
- Palfrey, Thomas R. 1982. “Risk Advantages and Information Acquisition.” *Bell Journal of Economics* 13 (1):219–224. 1.2
- Roesler, Anne-Katrin and Balázs Szentes. 2017. “Buyer-Optimal Learning and Monopoly Pricing.” *American Economic Review* 107 (7):2072–2080. 1.2
- Rogerson, William P. 1992. “Contractual Solutions to the Hold-Up Problem.” *Review of Economic Studies* 59 (4):777–793. 1.2
- Rothschild, Michael and Joseph E. Stiglitz. 1976. “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information.” *Quarterly Journal of Economics* 90 (4):629–649. 1.2
- Schelling, Thomas C. 1956. “An Essay on Bargaining.” *American Economic Review* 46 (3):281–306. 1.2
- Schelling, Thomas C. 1960. *The Strategy of Conflict*. Cambridge, Massachusetts: Harvard University Press. 7

- Schlee, Edward E. 2001. “The Value of Information in Efficient Risk Sharing Arrangements.” *American Economic Review* 91:509–524. [1.2](#)
- Taneva, Ina. 2019. “Information Design.” *American Economic Journal: Microeconomics* 11 (4):151–185. [1](#), [1](#), [3.1](#), [3.1.1](#)
- Taneva, Ina and Thomas Wiseman. 2023. “Ignoring Harmful Correlations.” Working paper, University of Edinburgh. [2](#)
- Tirole, Jean. 1986. “Procurement and Renegotiation.” *Journal of Political Economy* 94 (2):235–259. [1.2](#)
- van Damme, Eric. 1989. “Stable Equilibria and Forward Induction.” *Journal of Economic Theory* 48 (2):476–496. [7](#)
- Whitmeyer, Mark. 2022. “Opacity Design in Signaling Games.” Working paper, Arizona State University. [7](#)