Strategic Ignorance and Information Design

Ina Taneva
University of Edinburgh
ina.taneva@ed.ac.uk

Thomas Wiseman*
University of Texas at Austin
wiseman@austin.utexas.edu

December 2021

Abstract

We study information design in strategic settings when agents can publicly refuse to view their private signals. The requirement that agents must be willing to view their signals represents additional constraints for the designer, comparable to participation constraints in mechanism design. Ignoring those constraints 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 characterize implementable distributions over states and actions. Requiring robustness to strategic ignorance undoes two standard information design results: providing information conditional on players' choices rather than all at once may hurt the designer, and communication between players may help her.

1 Introduction

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. The information structure comprises distributions of joint signal realizations conditional on each possible realization of the state. Through the release of information, the designer incentivizes the agents to take actions that will benefit her. An implicit assumption in the information design literature is that players will agree to get informed according to the information structure chosen by the designer. Crucially, that setting does not permit players to refuse to observe the signals and to credibly signal this choice to the other players. In many strategic environments, however, an agent may benefit

*We thank Ben Brooks, Erik Madsen, Elliot Lipnowski, and Leeat Yariv for helpful comments.

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 — one where some or all agents act only given their prior information.

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, $\omega \in \{A, B\}$, corresponding to whether the agency needs solution A or solution B. The row player is the small company, which has three possible actions. Action A represents a choice to invest, ahead of the auction, in technology that will let it provide solution A at a low cost and hence a low bid; action B is the equivalent choice for solution B; 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.

	H	L
A	$3, 0$	$1, 1$
M	$2, 2$	$0, 0$
B	$0, 0$	$-2, 1$

$\omega = a$

	H	L
A	$0, 0$	$-2, 1$
M	$2, 2$	$0, 0$
B	$3, 0$	$1, 1$

$\omega = b$

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 (A, L) is played in state a or (B, L) is played in state b, 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 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 blindly investing in either solution, as shown in Figure 2. The large company's best response is high bid H, so by ignoring the report the small company increases its payoff from 1 to 2. The agency, though, gets payoff 0.

	H	L
A	1.5	-0.5
M	2	0
B	1.5	-0.5

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

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

We argue that in modeling information design, it is important to incorporate incentives to accept information as well as the designer's incentive to provide it. 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, and so 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. There are many economic settings where committing to ignorance is valuable, as we discuss below.

The requirement that agents must be incentivized to view their signals 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 some settings, including 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.

A given distribution is implementable if it is the outcome of a sequential equilibrium of the two-stage game. This requires three things: 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; 2) players who have chosen Ignore in the first stage follow second-stage strategies that are independent of the state and the other player's information; and 3) the Look-Ignore choices in the first stage constitute an equilibrium given the continuation equilibria chosen in 1).

In Theorem 1, we characterize the implementable distributions over actions and states under strategic ignorance in general finite environments. We show that it is without loss of generality to restrict the designer to *direct 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: direct information structures are sufficient for the designer only if we allow randomization at the Look-Ignore stage. Surprisingly, in some cases the designer's optimal outcome can be achieved via action recommendations only in an equilibrium where some players Ignore their signals with positive probability.[1] We further elaborate on this in Section 2.2.

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 information structure in order to satisfy the Look constraints. Method 1 is to raise the on-path payoff of the player(s) whose Look constraint is violated. Method 2 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.

We formulate the designer's problem using the approach of Bergemann and Morris (2016) and Taneva (2019), which exploits the equivalence between the set of Bayes correlated equilibria (BCEs) and the set of all BNE outcomes across all information structures. As in those papers, we assume that the designer costlessly commits to an information structure without observing the state and that the agents cannot communicate with each

[1]We thank Elliot Lipnowski for pushing us to investigate this question.

other, and we restrict attention to the best equilibrium for the designer.[2] Our other key assumption is that each agent's choice of whether to Look 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.

In order to further explore the impact of strategic ignorance, we consider a couple of extensions. 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 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 (2021)). 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,[3] 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

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

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

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. As in the standard information design environment, allowing communication may also harm the designer because she cannot prevent the players from sharing information that she would prefer to remain private.

The setting where players can communicate also provides an exception to the result that providing direct information in the form of recommended actions is sufficient for the designer. In the example in the previous paragraph, the designer may not want Player 1 to know the state. She can achieve that goal (again, subject to the appropriate incentive constraints) by giving each player a coded signal that is uninformative on its own but that reveals the state when combined with the other signal. For example, each players 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. The two signals combined correspond to an action recommendation, but neither does on its own.

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 in the unique sequential equilibrium of the dynamic game all players remain uninformed.

1.1 Relation to Literature

Ignoring free and payoff-relevant information is never beneficial in the standard single-receiver environment of Kamenica and Gentzkow (2011). However, when beliefs directly enter the agent's utility function, avoidance of decision-relevant costless information can occur as shown by Lipnowski and Mathevet (2018). Other mechanisms such as time inconsistency (Meng and Wang (2021)), temptation and self-control (Carrillo and Mariotti (2000)), present bias (Benabou and Tirole (2002)), pro-social preferences (Grossman and van der Weele (2017)) and enjoyment from suspense (Ely, Frankel, and Kamenica (2015)) can also lead to information avoidance in single-receiver settings. Additionally, departures from Bayesian updating can explain an ex-ante preference for less informative experiments, as shown in Jakobsen (2021).

In contrast, in this paper we focus on the *strategic* rationale behind ignoring free and payoff-relevant information, and the implications for the third party provider of this information, the information designer. This rationale exists even though the agents in our model are standard Bayesian-updating expected utility maximizers and it is purely by strategic motives that are completely absent in the single-agent environment. When an agent publicly commits not to observe the information provided by the designer, then other players cannot rely on his knowledge when deciding on their own actions. This commitment to ignoring information may thus lead to a better outcome for the agent by changing the optimal behavior of the other players.

Strategic ignorance can take different forms. In some games, committing to remain uninformed of the previous action choice of an opponent can be beneficial for reversing the opponent's first-mover advantage (Schelling (1960), van Damme (1989), Ben-Porath and Dekel (1992)). 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 information avoidance about a payoff-relevant state can be used to prevent a situation of asymmetric information and the resulting adverse selection problems of Akerlof (1970) or to preserve incentives for efficient risk-sharing as in Hirshleifer (1971) and Rothschild and Stiglitz (1976). 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). Similarly, a public commitment to information avoidance can be used to convincingly strengthen one's bargaining position (Schelling (1956)).[4] More recent papers have pointed out that agents may strategically choose to remain less than perfectly informed about a payoff relevant state in the context of procurement costs (Kessler (1998)), private-values in second-price auctions (McAdams (2012)), and buyer valuations in bilateral trade (Roesler and Szentes (2017)).

[4]Goldman, 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.

In this paper, we focus on general finite environments with a common payoff-relevant state and an information structure according to which players can learn about the state's realization. The information structure, which is provided with commitment by the designer, contains information about the state and about the information received by other players, as well as a component of pure correlation. Players first publicly decide whether to get informed according to the given information structure or remain ignorant, and then actions are taken simultaneously in the resulting game of incomplete information. Thus, in our setting, information avoidance is not about previously taken actions by the other players, but about the given information structure. Choosing to strategically ignore information results in both remaining uninformed about the state and unable to coordinate one's actions with those of other players.

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 incorporating the designer's strategic provision of information 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, for example, 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.[5]

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 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.

Due to the multi-stage nature of the interaction, 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 (2021)). The main conceptual difference with these papers is that we do not allow the designer to provide information

[5]A related implication is that a collusive agreement (corresponding to a designer who seeks to maximize players' payoffs) on what types of information to obtain and observe may not be sustainable. Bergemann, 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.

more than once. Specifically, we do not allow the designer to send additional messages to players contingent on their action choices in the first stage. Instead, the players' choices in the first "Look-Ignore stage" determine how the fixed information structure provided ends up being transformed into the informational environment that governs the second, action stage. Importantly, as outlined in the introduction, conditioning the information sent to players on their "Look-Ignore" 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; $A \equiv A_1 \times \ldots \times A_N$ 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 $u : A \times \Omega \to \mathbb{R}^N$, where agent i's payoff function $u_i : A \times \Omega \to \mathbb{R}$ depends on the action profile and the (ex ante unknown) state. The designer has a utility function $u^D : A \times \Omega \to \mathbb{R}$, so that her payoff also depends on the agents' actions and the state. The agents and the designer share a common full-support prior μ over Ω. Let $G = ((A, u), \mu)$ be the *basic game*.

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

At the *Look-Ignore stage*, each agent makes a choice $s_i \in S_i \equiv \{\ell, g\}$: whether to *Look* ($\ell$) at his signal and learn the realization of t_i, or to *Ignore* (g) it and remain uninformed. The Look-Ignore choices are public and simultaneous. Given a profile $s \in S \equiv \{\ell, g\}^N$ of realized choices from the Look-Ignore stage, let $\mathcal{L}(s) \coloneqq \{i : s_i = \ell\}$ and $\mathcal{G}(s) \coloneqq \mathcal{I} \setminus \mathcal{L}(s)$. 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 $i \in \mathcal{G} \coloneqq \mathcal{I} \setminus \mathcal{L}$ 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, $\left(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)}\right)$ and the prior μ. An agent who chooses to Ignore his signal uses $\left(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)}\right)$ and μ to form beliefs about t_{-i} and does not update his beliefs about the state.

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

Our solution concept for a dynamic game G^* is sequential equilibrium. In particular, continuation play in the action stage $G\left(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)}\right)$ 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 $\left(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)}\right)$ is common knowledge. Agent i who has chosen Look and observed s and t_i updates his beliefs about ω and t by using Bayes' rule. An agent who has chosen Ignore observes only s so he updates his beliefs about t only.

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

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

for all $a_i' \in A_i$;

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

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

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

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

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

Definition 2. *A strategy profile $\left(\gamma, \left(\tilde{\beta}^s\right)_s\right)$ is a sequential equilibrium of $G^*(T,P)$ if for each $s \in S$, $\tilde{\beta}^s$ is a BNE of $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$, and for each $i \in \mathcal{I}$ and $s_i \in \{\ell, g\}$ with $\gamma_i(s_i) > 0$,*

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

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 sequential equilibrium outcome of $G^*(T,P)$.

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

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

2.1 Characterization

The designer maximizes over a very large space, the set of all information structures (T, P). We show in Theorem 1, however, that without loss of generality we can restrict the designer to selecting a *direct* information structure. In a direct 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 information structure $T_i = \mathscr{A}_i$ for each agent i, where $\mathscr{A}_i \equiv A_i^{|S_{-i}|}$ denotes the set of agent i's (pure) mappings from S_{-i} to A_i. We will denote a generic element of $\mathscr{A}_i$ 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 $\mathscr{A} \equiv \mathscr{A}_1 \times \ldots \times \mathscr{A}_N$. Say that an outcome $v \in \Delta(A \times \Omega)$ is *implementable with direct recommendations* if there exists a conditional message distribution $P : \Omega \to \Delta(\mathscr{A})$ such that v is a sequential equilibrium outcome of $G^*(\mathscr{A}, P)$.

Theorem 1. *An outcome v is a sequential equilibrium outcome if and only if it is implementable with direct recommendations.*

Theorem 1 follows immediately from Theorem 3, which states that the outcome of any sequential equilibrium can be achieved in a BCE with appropriate restrictions on the correlation structure. 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 sequential equilibria 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. Therefore, we get a more limited result: we can view the designer as optimizing over the set of correlated equilibria that satisfy the corresponding constraints on the feasible correlation structure. In Appendix A we provide the appropriate definition of correlated equilibrium for this purpose in our setting and present the equivalence result.

In the rest of the paper, we exploit Theorem 1 in order to characterize the solution to designer's problem. We maximize the designer's payoff over direct information structures $(\mathscr{A}, P)$, Look-Ignore strategies γ, and post-Ignore contingent (mixed) strategies $\beta_i^g : S_{-i} \to \Delta A_i$, subject to the specified behavior being a sequential equilibrium. Conceptually, think of the designer as optimizing over direct information structures $(\mathscr{A}, P)$ by choosing P, and then nature optimizes on the designer's behalf over the set of sequential equilibria of $G^*(\mathscr{A}, P)$.

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

Theorem 2. *A sequential equilibrium 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 Player 2 also is completely uninformed. On path, however, the designer 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.

Importantly, the reason why equilibria in which there is strict randomization between Look and Ignore are necessary is because we are restricting attention to direct information structures. If we were to enlarge the signal space of the information structures to include information about whether the messages of the other players are informative or uninformative, we could restrict attention to pure strategy Look equilibria by the agents and the designer could do the randomization between informative and uninformative messages. Relatedly, we can without loss of generality disregard equilibria in which any agent plays Ignore with certainty, as this is simply equivalent to the designer choosing a completely uninformative message for that agent and the agent choosing Look with certainty.

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 Appendix C.

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 an unknown state of the world. 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, where $w > u > 0$ and each state is equally likely.

The designer wants the project to fail. In particular, she gets a payoff of 1 if (B, B) is

	A	B
A	w, w	$u, 0$
B	$0, u$	$0, 0$

$\omega = a$

	A	B
A	$0, 0$	$0, u$
B	$u, 0$	w, w

$\omega = b$

Figure 3: Investment game

played in state a or (A, A) is played in state b, and 0 otherwise.

3.1.1 Baseline

In the baseline information design environment, where agents automatically observe their private signals from the designer, the designer's optimal direct information structure is $(A, \tilde{P})$ with

$$\tilde{P}\left(A, A|\omega = a\right) = \tilde{P}\left(B, B|\omega = b\right) = \frac{u}{w+u},$$

$$\tilde{P}\left(B, B|\omega = a\right) = \tilde{P}\left(A, A|\omega = b\right) = \frac{w}{w+u}.$$

The designer's payoff is $w/(w+u) > 0.5$, and each firm's payoff is $wu/(w+u)$.

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 $u/(w+u) < 0.5$. 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 w with probability $u/(w+u)$, and switching yields u with probability $w/(w+u)$.[6]

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 A and B. Firm 2's best response is to choose the opposite of the

[6] We note the role of advantageous equilibrium selection here. Under $m^{baseline}$, 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.

project that the designer recommended: now that Firm 1 cannot coordinate by following the designer's recommendation, Firm 2 just wants to pick the project that is more likely to succeed. Under $(A, \tilde{P})$, the project that the designer recommends is more likely to be the wrong one, so Firm 2 will pick the other project.

In that BNE, Firm 1 gets an expected payoff of

$$\frac{1}{2}\left(\frac{w}{w+u}w + \frac{u}{w+u}u\right),$$

which is strictly greater than the payoff $wu/(w+u)$ 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.

In order to satisfy the constraint that firms be willing to Look at their signals, the designer's optimal adjustment is to reduce the probability that Firm 2 can choose the correct project if Firm 1 deviates to Ignore. That is, to reduce the probability of recommending the wrong project, π, from $\pi = w/(w+u)$. Under the new optimal information structure, the value of π is $(2w-u)/(3w-u)$, which is lower than $w/(w+u)$ but still greater than $\frac{1}{2}$. Formally, the best that the designer can do when we require robustness to strategic ignorance comes from the direct information structure $(\mathscr{A}, P^*)$ with

$$P^*\left(AB, AB|\omega = a\right) = P^*\left(BA, BA|\omega = b\right) = \frac{w}{3w-u},$$

$$P^*\left(BA, BA|\omega = a\right) = P^*\left(AB, AB|\omega = b\right) = \frac{2w-u}{3w-u}.$$

The first term in the message is the recommendation after the other firm chooses Look, and the second term is the recommendation after the other firm chooses Ignore.

The probability of recommending the wrong project, π, is the designer's payoff, so $(\mathscr{A}, P^*)$ is the information structure that maximizes π subject to the constraint that a firm's payoff when both follow their recommendations exceeds the worst BNE payoff after unilaterally deviating to Ignore. For $\pi > \frac{1}{2}$, that worst BNE is still Firm 1 (the deviator) randomizing uniformly between A and B, and Firm 2 choosing the opposite of the project that was recommended on-path:

$$\frac{2w-u}{3w-u} = \max\left[\pi : (1-\pi)u \geq \frac{1}{2}\left(\pi w + (1-\pi)u\right)\right].$$

Under $(\mathscr{A}, P^*)$, it is an equilibrium for both firms to Look at and follow their signals: a firm gets the same payoff from choosing Look as from deviating to Ignore. Relative to the baseline case, the ability to Ignore a signal improves the firms' payoff from $wu/(w+u)$ to $w^2/(3w-u)$. The designer's payoff falls from $w/(w+u)$ to $(2w-u)/(3w-u)$. In summary, the designer's response to the threat of strategic ignorance is a modest adjustment of the information structure. At values $u = 1$ and $w = 1.5$, for example, π falls from 0.6 to 0.57.

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 ($a(ttack)$ or $n(ot)$). The currency may be either weak or strong with equal probability. If the currency is weak, then one player is enough for a successful attack, and so attacking is strictly dominant. If the currency is strong, then the attack succeeds if and only if at least two players attack. We capture that setting with the following payoff function, where player i's payoff depends on the state $\omega \in \{W(eak), S(trong)\}$, his own action, and the number K of other players who play a:

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

where $x \geq 1$, and

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

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

3.2.1 Baseline

At the prior, a 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 $(n, \ldots, n)$ to be an equilibrium. Formally, the designer's optimal information structure with *single* action recommentations, $(A, \tilde{P})$, is

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

The obedience constraint binds for a player who gets recommendation n: the updated probability of state S is $0.5/(0.5 + 0.25) = 2/3$, so both a 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 suppose that $x > 1$, and suppose that Player i deviates to no-look. When Player i is uninformed, then a is dominant: as shown in Figure 4, 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 a in each state.

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

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

Figure 4: a is dominant for an uninformed player

In either state, the unique best response for any other player when Player i chooses a is a. The outcome is thus $(a, \ldots, a)$ 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 $(a, \ldots, a)$ 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.[7]

If $x = 1$, then the situation changes. From Figure 4, we see that now an uninformed Player i's expected payoff from playing a against a strategy of (a in state W, n in state S)

[7] This example illustrates the distinction between equilibrium selection and strategic ignorance. "Always play $(a, \ldots, a)$" 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 $(a, \ldots, a)$," 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.

by each other player (that is, $K_W = N-1, K_S = 0$) is 0; both a 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 $(W : a, S : n)$. Player i's unique best response to anything other than the strategy profile of $(W : a, S : n)$ for all opponents is a, 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 $(W : a, S : n)$ becomes a feasible strategy. The new optimal direct information structure, $(\mathscr{A}, P^*)$, is to recommend action a 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 $(\mathscr{A}, P^*)$, 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 a and each other players follow their recommendation by playing $(W : a, S : n)$. 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 $(\mathscr{A}, P^*)$ is that, as just argued, the constraint that players must be willing to view their signals is slack. In the investment game, the designer optimally modified the baseline information structure by slightly raising the players' on-path payoffs and lowering the post-deviation payoffs until 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.[8] However, recall from Section 2.3 that this need not be the case in general. In Appendix C we provide an example where both players are worse off under the ability to strategically ignore information.

4 Discussion and Conclusion

We have shown that the ability of agents to publicly refuse information has important effects for information design in strategic settings. Requiring robustness to strategic ignorance significantly alters optimal information structures and the ensuing outcomes in

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

leading economic applications, and it undoes standard qualitative 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. 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 BCE 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

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

where γ denotes the Look-Ignore strategies, β^g denotes the post-Ignore strategies, and v denotes the post-Look strategies. For each $s \in S$, let $v(m_{\mathcal{L}(s)}, \omega) := \sum_{m_{\mathcal{G}(s)}} v(m_{\mathcal{L}(s)}, m_{\mathcal{G}(s)}, \omega)$, and let $v((m_i(s_{-i}))_{i \in \mathcal{L}(s)}, \omega)$ denote the corresponding projection of $v(m_{\mathcal{L}(s)}, \omega)$. Let $a_{\mathcal{G}} := (a_i)_{i \in \mathcal{G}}$.

Definition 4. (γ, β^g, v) *is a BCE of* G^* *if*

1. *(***Consistency with the prior***)* $v(\mathscr{A} \times \{\omega\}) = \mu(\omega)$ *for all* $\omega \in \Omega$*;*

2. *(***Obedience for agent*** i ***who chose Look****) for every* $s \in S$*,* $i \in \mathcal{L}(s)$*,* $m_i \in \mathscr{A}_i$*, and* $a_i' \in A_i$

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

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

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

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

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

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

Definition 5. *Given a BCE distribution* $(\gamma, \beta^g; v)$*, let* $\tilde{v}(\gamma, \beta^g; v) \in \Delta(A \times \Omega)$ *defined as*

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

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

Theorem 3. $\cup_{(T,P)} SE\left(G^*\left(T,P\right)\right) = BCE\left(G^*\right)$.

Proof. First we prove that $BCE\left(G^*\right) \subseteq \cup_{(T,P)} SE\left(G^*\left(T,P\right)\right)$. Take any $\tilde{v}(\gamma, \beta^g; v) \in BCE\left(G^*\right)$. Consider the information structure $(\mathscr{A}, P)$ with $P(m|\omega) := v(m,\omega)/\mu(\omega)$ for all $m \in \mathscr{A}, \omega \in \Omega$.

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

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

for all $m_i \in \mathscr{A}_i$, 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 $m_i \in \mathscr{A}_i$ and choosing action $a_i \in A_i$ when his opponents play according to $\tilde{\beta}_{-i}^s$ is given by

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

Hence, by (5) we obtain

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

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

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

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

Hence, by (6) we obtain

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

for all $i \in \mathcal{G}(s)$, a_i such that $\beta^g_i(a_i|s_{-i}) > 0$, 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 $\tilde{\beta}^s_i(a_i) > 0$.

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

$$v(\tilde{\beta}^s)(a,\omega) := \sum_{m_{\mathcal{L}(s)}} \mu(\omega) P_{\mathcal{L}(s)}(m_{\mathcal{L}(s)}|\omega) \left(\prod_{j\in\mathcal{L}(s)} \tilde{\beta}^s_j(a_j|m_j) \prod_{k\in\mathcal{G}(s)} \tilde{\beta}^s_k(a_k) \right)$$
$$= \sum_{m_{\mathcal{L}(s)}} v(m_{\mathcal{L}(s)}, \omega) \left(\prod_{j\in\mathcal{L}(s)} \tilde{\beta}^s_j(a_j|m_j) \prod_{k\in\mathcal{G}(s)} \tilde{\beta}^s_k(a_k) \right) \quad (12)$$

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

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

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

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

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

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

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

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

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

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

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

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

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

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 $v(\mathscr{A} \times \{\omega\}) = \mu(\omega)$ 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 $a_i, a_i' \in A_i$

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

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

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

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

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

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

Hence, (γ, β^g, v) is a BCE of G^*. Then, $\tilde{v}(\gamma, \beta^g, v) \in \Delta(A \times \Omega)$ is a BCE outcome of G^*, that is $\tilde{v} \in BCE(G^*)$. Notice that

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

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

□

B Example: the necessity of ignorance

	L	R_a	R_b	P_a	P_b	P
A	$3, 0$	$1, 1$	$1, 1$	$3, 0$	$3, -1$	$1, -1$
M	$2, 2$	$0, 0$	$0, 0$	$2, 0$	$2, 0$	$0, 0$
B	$0, 0$	$-2, 1$	$-2, 1$	$-2, 0$	$-2, -1$	$-2, -1$
M'	$0, 2$	$-1, 0$	$-1, 0$	$-1, 3$	$-1, 1$	$0, 2$

$\omega = a$

	L	R_a	R_b	P_a	P_b	P
A	$0,0$	$-2,1$	$-2,1$	$-2,-1$	$-2,-1$	$-2,-1$
M	$2,2$	$0,0$	$0,0$	$2,0$	$2,0$	$0,0$
B	$3,0$	$1,1$	$1,1$	$3,-1$	$3,-1$	$1,-1$
M'	$0,2$	$-1,0$	$-1,0$	$-1,1$	$-1,3$	$0,2$

$\omega = b$

Consider the following example with $\Omega = \{a, b\}$. Each state is equally likely, so that the prior is $\mu(a) = \mu(b) = \frac{1}{2}$. There are two players with action sets $A_1 = \{A, M, B, M'\}$ and $A_2 = \{L, R_a, R_b, P_a, P_b, P\}$. The players' state contingent payoffs are given by the above payoff matrices, where the first entry in each cell corresponds to the payoff of Player 1, whose possible action choices are represented by the rows, and the second entry corresponds to Player 2's payoff, whose possible action choices are represented by the columns.

The designer gets a payoff of 1 if (A, R_a) is played in state a, or if (B, R_b) is played in state b, and payoff 0 otherwise. The actions R_a and R_b 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 a, A is dominant for Player 1, and R_a is the unique best response for Player 2. In state b, B is dominant for Player 1, and R_b is the unique 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

	L	R_a	R_b	P_a	P_b	P
A	$1.5,0$	$-\frac{1}{2},1$	$-\frac{1}{2},1$	$\frac{1}{2},-1$	$\frac{1}{2},-1$	$-\frac{1}{2},-1$
M	$2,2$	$0,0$	$0,0$	$2,0$	$2,0$	$0,0$
B	$1.5,0$	$-\frac{1}{2},1$	$-\frac{1}{2},1$	$\frac{1}{2},-1$	$\frac{1}{2},-1$	$-\frac{1}{2},-1$
M'	$0,2$	$-1,0$	$-1,0$	$-1,2$	$-1,2$	$0,2$

Suppose it is common knowledge that Player 1 knows the state and that Player 2's beliefs equal the prior. In state a, A is dominant for Player 1, and in state b, B is dominant. In both cases, either R_a or R_b 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_a in state $\omega = a$ and R_b in state $\omega = b$.

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 A and B, M weakly dominates M' for Player 1: M' is a weak best response for Player 1 if and only if Player 2 plays P with probability 1. P is a best response to M'. So (M', P) is an eqm with payoff $(0, 2)$. Next, suppose Player 2 assigns belief $p > \frac{1}{2}$ to state ω. Then P is not a best response to M': P 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', P) is not an eqm. Instead, M is dominant against state-contingent strategies of Player 2 and the unique equilibrium is (M, L), giving payoff $(2, 2)$.

A mixed Look-Ignore outcome

Suppose the designer's information structure is given by $(\mathscr{A}, P)$ with

$$P(AA, R_aL|\omega = a) = P(BB, R_bL|\omega = b) = 1$$

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 (A or B), 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 (A or B). Player 2's post-Ignore strategy is $\beta_2^g(R_a|\ell) = \beta_2^g(R_a|\ell) = \frac{1}{2}$, where he randomizes between R_a or R_b, 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 $\beta_2^g(P|g) = 1$. At the prior, M' is a best response to P , and P 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 $\frac{1}{2}2 + \frac{1}{2}0 = 1$. 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 Example: the harm of ignorance

Consider the following symmetric game, where each state $\omega \in \{0, 1\}$ is equally likely, so that $\mu(1) = \mu(2) = \frac{1}{2}$. The players' state contingent payoffs are given by the following payoff matrices, where the row player is Player 1 and the column player is Player 2:

	X	Y	A_1	B_1	A_2	B_2
X	$0, 0$	$0.1, 0.1$	$1.1, 0.12$	$1.12, 0.14$	$-1.1, -0.2$	$-1.12, -0.2$
Y	$0.1, 0.1$	$0.15, 0.15$	$1, 0.18$	$1.1, 0.16$	$1, -0.2$	$1.1, -0.2$
A_1	$0.12, 1.1$	$0.18, 1$	$1.11, 1.11$	$1.111, 1.1$	$1.1, 0$	$1.1, 0$
B_1	$0.14, 1.12$	$0.16, 1.1$	$1.1, 1.111$	$1.11, 1, 11$	$1.11, 0$	$1.11, 0$
A_2	$-0.2, -1.1$	$-0.2, 1$	$0, 1.1$	$0, 1.11$	$0, 0$	$0, 0$
B_2	$-0.2, -1.12$	$-0.2, 1.1$	$0, 1.1$	$0, 1.11$	$0, 0$	$0, 0$

,

Payoffs in $\omega = 1$

	X	Y	A_1	B_1	A_2	B_2
X	$0, 0$	$0.1, 0.1$	$-1.1, -0.2$	$-1.12, -0.2$	$1.1, 0.12$	$1.12, 0.14$
Y	$0.1, 0.1$	$0.15, 0.15$	$1, -0.2$	$1.1, -0.2$	$1, 0.18$	$1.1, 0.16$
A_1	$-0.2, -1.1$	$-0.2, 1$	$0, 0$	$0, 0$	$0, 1.1$	$0, 1.11$
B_1	$-0.2, -1.12$	$-0.2, 1.1$	$0, 0$	$0, 0$	$0, 1.1$	$0, 1.11$
A_2	$0.12, 1.1$	$0.18, 1$	$1.1, 0$	$1.1, 0$	$1.11, 1.11$	$1.111, 1.1$
B_2	$0.14, 1.12$	$0.16, 1.1$	$1.11, 0$	$1.11, 0$	$1.1, 1.111$	$1.11, 1.11$

Payoffs in $\omega = 2$

At the prior, expected payoffs are

	X	Y	A_1	B_1	A_2	B_2
X	$0, 0$	$0.1, 0.1$	$0, -0.04$	$0, -0.03$	$0, -0.04$	$0, -0.03$
Y	$0.1, 0.1$	$0.15, 0.15$	$1, -0.01$	$1.1, -0.02$	$1, -0.01$	$1.1, -0.02$
A_1	$-0.04, 0$	$-0.01, 1$	$0.555, 0.555$	$0.5555, 0.55$	$0.55, 0.55$	$0.55, 0.555$
B_1	$-0.03, 0$	$-0.02, 1.1$	$0.55, 0.5555$	$0.555, 0.555$	$0.555, 0.55$	$0.555, 0.555$
A_2	$-0.04, 0$	$-0.01, 1$	$0.55, 0.55$	$0.55, 0.555$	$0.555, 0.555$	$0.5555, 0.55$
B_2	$-0.03, 0$	$-0.02, 1.1$	$0.555, 0.55$	$0.555, 0.555$	$0.55, 0.5555$	$0.555, 0.555$

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

so that Y is strictly dominant for each player.

The designer gets a payoff equal to the sum of the payoffs of the two players. In the baseline information design environment, where agents automatically observe their private signals from the designer, the designer can achieve her maximum feasible payoff of 2.22 with a perfectly informative information structure that recommends action $A\omega$ in state ω to each player:

- After (ℓ, ℓ): In this case the state is common knowledge. In state ω, action $A\omega$ strictly dominates every action except $B\omega$. The unique best response to any mixing between $A\omega$ and $B\omega$ is $A\omega$. Thus, the unique BNE is $(A\omega, A\omega)$, and the payoffs are $u(\ell, \ell) = (1.11, 1.11)$.

- After (g, ℓ): In this case it is common knowledge that Player 2 knows the state and that Player 1's beliefs are given by the prior. As above, in state ω, action $A\omega$ strictly dominates every action except $B\omega$ for Player 2. Thus, Player 2 has four undominated strategies: $A1A2$,$A1B2$, $B1A2$, and $B1B2$, where the first element denotes the action in state 1 and the second element denotes the action in state 2. Player 1's expected payoffs against those strategies are as follows:

	$A1A2$	$A1B2$	$B1A2$	$B1B2$
X	1.1	1.11	1.11	1.12
Y	1	1.05	1.05	1.1
A_1	0.555	0.555	0.5555	0.5555
B_1	0.55	0.55	0.555	0.555
A_2	0.555	0.5555	0.555	0.5555
B_2	0.55	0.555	0.55	0.555

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

Player 1's unique best response against any of those four strategies is X. Player 2's best response to X is $B1B2$. Thus, the unique BNE is $(X, B1B2)$, and the payoffs are $u(g, \ell) = (1.12, 0.14)$.

- After (ℓ, g): This case is symmetric to the preceding one.
- After (g, g): In this case it is common knowledge that both players' beliefs are given by the prior distribution, and, hence, Y is strictly dominant. Thus, the unique BNE is (Y, Y), and the payoffs are $u(g, g) = (0.15, 0.15)$.

Equilibrium at the Look-Ignore Stage

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

	ℓ	g
ℓ	$1.11, 1.11$	$0.14, 1.12$
g	$1.12, 0.14$	$0.15, 0.15$

Ignore is strictly dominant, so the outcome is that both players choose Ignore and wind up with payoff 0.15. Without the possibility of strategic ignorance, they would get payoff 1.11.

Interpretation

The example has the flavor of a prisoners' dilemma, where Look corresponds to Cooperate, and Ignore corresponds to Defect. Intuitively, an informed player wants to match the state with either $A\omega$ or $B\omega$, but which is better depends on the action of the other player. An uninformed player effectively commits to playing X against either $A\omega$ or $B\omega$, and the informed player's best reponse is $B\omega$, which is slightly better for the uninformed player than $A\omega$ would be.

An informed player matched with an informed player chooses $A\omega$. Thus, given that Player 2 is informed, Player 1 prefers to be uninformed: the gain from getting Player 2 to switch from $A\omega$ to $B\omega$ outweighs the loss from not being able to exactly best respond to $B\omega$. Player 2 loses more than Player 1 gains, because playing X instead of $A\omega$ gives the informed opponent a low payoff.

Against an uninformed player, choosing Ignore effectively commits a player to playing Y, and so the opponent's response will also be Y. Thus, given that Player 2 is uninformed, Player 1 prefers to be uninformed: the benefit from getting Player 2 to switch from X to Y outweighs the loss from not being able to exactly best respond to X. Player 2 loses more than Player 1 gains, because playing Y instead of $B\omega$ gives the opponent a lower payoff.

An information structure robust to strategic ignorance

Suppose that the designer provides the following direct information structure:

$$\Pr(A1A1, A1A1, \omega = 1) = \Pr(A2A2, A2A2, \omega = 2) = p = 21/22,$$
$$\Pr(A2A2, A2A2, \omega = 1) = \Pr(A1A1, A1A1, \omega = 2) = 1 - p = 1/22.$$

where the first term in the message is the recommendation of what action to play after the other player chooses Look, and the second term is the recommendation for after the other player chooses Ignore. That is, in state ω, the designer recommends action $A\omega$ to both players with probability $p = 21/22$, and otherwise recommends action $A\omega'$. Additionally, a player's recommendation is the same irrespective of the Look-Ignore choice of the other player.

- After (ℓ, ℓ): Both players follow the message recommendation. The expected payoff is $1.11p \approx 1.060$, i.e., $u(\ell, \ell) = (1.060, 1.060)$.
- After (g, ℓ): The uninformed Player 1 plays Y and the informed Player 2 follows the received action recommendation for the case when his opponents has chosen Ignore. The expected payoff for the uninformed player is 1, while for the informed player it is $0.18p - 0.2(1 - p) \approx 0.163$. Hence, $u(g, \ell) = (1, 0.163)$.
- After (ℓ, g): This case is symmetric to the preceding one.
- After (g, g): Y is strictly dominant for both players and the payoffs are $u(g, g) = (0.15, 0.15)$.

Under these continuation payoffs, the expected payoffs at the Look-Ignore stage are:

	ℓ	g
ℓ	$1.060, 1.060$	$0.163, 1$
g	$1, 0.163$	$0.15, 0.15$

.

Now Look is dominant at the Look-Ignore stage.

Discussion

For an uninformed Player 1, playing Y is safe while X is risky if Player 2 sometimes gets it wrong (that is, plays $A\omega'$ or $B\omega'$ instead of $A\omega$ or $B\omega$). Under perfect information, an informed Player 2 never gets it wrong, so Player 1 chooses X, leading Player 2 to choose $B\omega$ instead of $A\omega$, and Player 1 benefits.

Adding a little noise to the signals ($p < 1$) makes X too risky. Now Player 1 prefers Y, so Player 2 plays $A\omega$ regardless of whether or not Player 1 chooses Look. When Player 1's Look-Ignore choice does not change Player 2's action, then Player 1 cannot possibly gain from ignoring his signal. He can only lose from not being able to match his action to the state.

References

Akerlof, George A. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 84 (3):488–500. 1.1

Arcuri, Alesandro. 2021. "Hear-No-Evil Equilibrium." Working paper, Cornell University. 1.1

Ben-Porath, Elchanan and Eddie Dekel. 1992. "Signaling Future Actions and the Potential for Sacrifice." *Journal of Economic Theory* 57 (1):36–51. 1.1

Benabou, Roland and Jean Tirole. 2002. "Self-Confidence and Personal Motivation." *Quarterly Journal of Economics* 117 (3):871–915. 1.1

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. 5

Bergemann, Dirk and Stephen Morris. 2016. "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games." *Theoretical Economics* 11:487–522. 1

———. 2019. "Information Design: A Unified Perspective." *Journal of Economic Literature* 57 (1):44–95. 1

Carrillo, Juan D. and Thomas Mariotti. 2000. "Strategic Ignorance as a Self-Disciplining Device." *Review of Economic Studies* 67 (3):529–544. 1.1

de Oliveira, Henrique and Rohit Lamba. 2019. "Rationalizing Dynamic Choices." Working paper, Penn State University. 1.1

Doval, Laura and Jeffrey Ely. 2020. "Sequential Information Design." *Econometrica* 86 (6):2575–2608. 1.1

Ely, Jeffrey, Alexander Frankel, and Emir Kamenica. 2015. "Suspense and Surprise." *Journal of Political Economy* 123 (1):215–260. 1.1

Goldman, Russell, David Hagmann, and George Loewenstein. 2017. "Information Avoidance." *Journal of Economic Literature* 55 (1):96–135. 4

Grossman, Zachary and Joël J. van der Weele. 2017. "Self-Image and Willful Ignorance in Social Decisions." *Journal of the European Economic Association* 15 (1):173–217. 1.1

Gul, Faruk. 2001. "Unobservable Investment and the Hold-Up Problem." *Econometrica* 69 (2):343–376. 1.1

Hirshleifer, Jack. 1971. "The Private and Social Value of Information and the Reward to Inventive Activity." *American Economic Review* 61 (4):561–574. 1.1

Jakobsen, Alexander. 2021. "Coarse Bayesian Updating." Working paper, University of Calgary. 1.1

Kamenica, Emir and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* 101 (6):2590–2615. 1, 1.1

Kessler, Anke S. 1998. "The Value of Ignorance." *RAND Journal of Economics* 29 (2):339–354. 1.1

Lipnowski, Elliot and Laurent Mathevet. 2018. "Disclosure to a Psychological audience." *AEJ: Microeconomics.* 10(4):67–93. 1.1

Makris, Miltiadis and Ludovic Renou. 2021. "Information Design in Multi-Stage Games." Working paper, University of Essex. 1, 1.1

McAdams, David. 2012. "The Value of Ignorance." *Economic Letter* 114:83–85. 1.1

Meng, Delong and Siyu Wang. 2021. "Persuading a Time-Inconsistent Receiver." Working paper, Wichita State University. 1.1

Palfrey, Thomas R. 1982. "Risk Advantages and Information Acquisition." *Bell Journal of Economics* 13 (1):219–224. 1.1

Roesler, Anne-Katrin and Balázs Szentes. 2017. "Buyer-Optimal Learning and Monopoly Pricing." *American Economic Review* 107 (7):2072–2080. 1.1

Rogerson, William P. 1992. "Contractual Solutions to the Hold-Up Problem." *Review of Economic Studies* 59 (4):777–793. 1.1

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. 1.1

Schelling, Thomas C. 1956. "An Essay on Bargaining." *American Economic Review* 46 (3):281–306. 1.1

Schelling, Thomas C. 1960. *The Strategy of Conflict*. Cambridge, Massachussets: Harvard University Press. 1.1

Taneva, Ina. 2019. "Information Design." *American Economic Journal: Microeconomics* 11 (4):151–185. 1, 1, 3.1

Tirole, Jean. 1986. "Procurement and Renegotiation." *Journal of Political Economy* 94 (2):235–259. 1.1

van Damme, Eric. 1989. "Stable Equilibria and Forward Induction." *Journal of Economic Theory* 48 (2):476–496. 1.1