

Costly monitoring in signaling games*

Tommaso Denti[†]

March 2021

Abstract

Off-path beliefs are a key free variable in equilibrium analysis of dynamic games of incomplete information. Our starting point is the observation that, if we take into account players' incentives to monitor past events in the game, then we can sidestep the question of off-path beliefs. We focus on signaling games where the receiver has to pay a cost to monitor the sender's action. We show that Nash equilibrium is outcome equivalent to any refinement of perfect Bayesian equilibrium that puts restrictions on the receiver's off-path beliefs. We then characterize all Nash equilibria that can arise across all cost functions. As an application, we consider the case of vanishing costs to provide a micro-foundation for restrictions on off-path beliefs in standard signaling games where the receiver observes for free the sender's action.

1 Introduction

Off-path beliefs are a key free variable in equilibrium analysis of dynamic games of incomplete information. What should players believe when confronted with events that were not supposed to happen? The question has motivated a rich and important literature in game theory (e.g., Selten, 1975; Kreps and Wilson, 1982; Kohlberg and Mertens, 1986; Banks and Sobel, 1987; Cho and Kreps, 1987).

Our starting point is the observation that, if we take into account players' incentives to monitor past events in the game, then we can sidestep the question of off-path beliefs. For an intuition, suppose that players have to pay some cost to monitor past events in the game. Then, no matter how small the cost is, players never monitor events that are not supposed to happen: monitoring zero-probability events is not rational, since, ex ante, it has zero value but positive cost. When monitoring is costly, players are never confronted with zero-probability events and can always use Bayes rule to update their beliefs.

*I thank Marco Battaglini, Luciano Pomatto, and seminar participants at Notre Dame. I am especially grateful to Mohammad Ghomi for essential help in proving Lemma 3.

[†]Cornell University. Email: tjd237@cornell.edu

In this paper, we study costly monitoring in signaling games and formalize the intuition that equilibrium analysis does not depend on off-path beliefs. We then characterize all equilibria that can arise across all cost functions. As an application, we consider the case of vanishing costs to provide a micro-foundation for restrictions on off-path beliefs in standard signaling games where the receiver observes for free the sender’s action.

Our model introduces costly monitoring in otherwise standard signaling games. The model features an informed sender and an uninformed receiver. The sender observes the state of nature and chooses an action. In standard signaling games, the receiver observes for free the action of the sender and takes an action in response. We assume instead that monitoring the sender’s action is costly. The receiver can choose from a number of feasible monitoring structures. Each monitoring structure provides different information about the action of the sender and comes with its own cost.

We make two broad assumptions on monitoring costs. First, we assume that there is free disposal of information: if a monitoring structure is feasible, then all monitoring structures that are less informative are also feasible. Second, we assume that monitoring costs are strictly increasing in the Blackwell order: acquiring strictly more information (in the sense of Blackwell, 1951) is strictly more costly.

Our initial finding is that Nash equilibrium is outcome equivalent to any refinement of perfect Bayesian equilibrium that puts restrictions on off-path beliefs. The result formalizes the intuition that, if we account for the receiver’s incentives to monitor the sender’s action, then we can sidestep the question of off-path beliefs.

We then characterize all Nash equilibria that can arise across all cost functions. It is a “robust” equilibrium characterization because it abstracts from specific details of monitoring costs. The robust approach is motivated by the observation that, in many settings, monitoring costs are affected by factors—such as time, effort, and cognitive resources—that are difficult to measure. The analyst therefore may be interested in predictions that are not sensitive to specific details of monitoring costs.

As an application, we consider the case of vanishing costs to provide a micro-foundation for restrictions on off-path beliefs in standard signaling games. The case of vanishing costs can be seen as a perturbation of the standard signaling game where the receiver observes for free the sender’s action. We characterize what perfect Bayesian equilibria of standard signaling games can be approximated by smaller and smaller perturbations. An interesting aspect of the exercise is that off-path beliefs are irrelevant in the perturbations, as shown by the equivalence of Nash equilibrium and refinements of perfect Bayesian equilibrium.

Signaling games with costly monitoring have recently received attention in a number of economic applications—espionage (Solan and Yarovitz, 2004), cheap talk with language barriers (Sobel, 2012; de Clippel and Rozen, 2021), strategic pricing with inattentive consumers

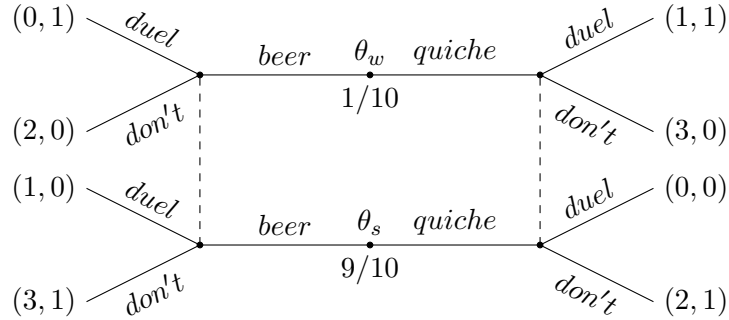


Figure 1: Beer-quiche game (Cho and Kreps, 1987).

(Matejka and McKay, 2012; Ravid, 2020). This paper takes a broader view and do not commit to a specific payoff structure or cost function.

We hope to encourage more applications: monitoring costs seem relevant to many other settings where signaling is a core issue, such as labor markets and advertising. For example, Bartoš, Bauer, Chytilová, and Matějka (2016) find evidence of limited attention to job applicants' CVs in hiring decisions and show that it may lead to discrimination; Acharya and Wee (2020) discuss the implications for wage and unemployment dynamics. Marketing research studies the phenomenon of “ad avoidance:” how media users differentially reduce their exposure to ad content and how, in response, firms design their ads to maximize engagement (see, e.g., Pieters and Wedel, 2004). We postpone a comprehensive review of the literature to the end of the paper.

2 A simple example

The basic ideas in this paper are illustrated by the following simple game, borrowed from Cho and Kreps (1987). The reader should refer to Figure 1 throughout.

The players are Ann, the sender, and Bob, the receiver. Ann can either be *surly* or a *wimp*, denoted by θ_s and θ_w . Ann's type is her private information. Bob ex ante believes that Ann is surly with probability 9/10.

In the morning, Ann chooses her breakfast. The options are *beer* and *quiche*. The surly type prefers beer, the wimpish type prefers quiche. In the afternoon, Ann meets Bob. Bob decides whether to *duel* Ann or not. Bob wants a duel if and only if Ann is a wimp. Bob's problem is that he does not know Ann's type.

In the standard formulation of this game, which is depicted by the extensive form in Figure 1, Bob observes Ann's breakfast before deciding whether to duel her. As shown by Cho and Kreps (1987), there are two classes of equilibria: in one class both types of Ann

choose beer, in the other class both choose quiche. In these pooling equilibria, off-path Bob threatens a duel. Bob is willing to carry out the threat if he believes that Ann is more likely a wimp. The question of off-path beliefs is as follows: after a given zero-probability event, is it reasonable for Bob to revise his beliefs and put more weight on Ann being a wimp? Cho and Kreps's influential work proposes a methodology, the "intuitive criterion," to address these important issues.

In this paper, we study Bob's incentives to monitor Ann's action and relate them to the question of off-path beliefs. Concretely, we ask what happens if Bob has to pay some cost, possibly small, to learn Ann's breakfast. For example, it could be that Bob has to bribe Ann's roommate to tell him whether Ann had beer or quiche. Or maybe Bob has to hire a private investigator to discover Ann's morning habits.

Our starting point is the observation that, if we account for Bob's incentives to monitor Ann's breakfast, then we can sidestep the question of Bob's off-path beliefs. To illustrate, suppose that Bob expects the two types of Ann to have the same breakfast, say quiche. Then Bob must choose not to monitor Ann's breakfast: Bob is already sure that Ann had quiche, so he is not willing to pay the monitoring cost, no matter how small it is. Thus Bob is never confronted with the off-path event of Ann having beer and can always use Bayes rule to revise his beliefs.

In the rest of the paper, we formalize these ideas for general signaling games and monitoring costs. We will characterize equilibrium behavior and use the case of vanishing costs to provide a micro-foundation for restrictions on off-path beliefs in standard signaling games where the receiver observes for free the sender's action. For example, when applied to beer-quiche, our results will show that the surly type always chooses beer with probability one in every equilibrium with costly monitoring. Thus, as the cost of monitoring vanishes, all equilibria converge to the beer-beer pooling equilibrium, the same equilibrium selected by Cho and Kreps's intuitive criterion.

3 Model

Our model introduces costly monitoring in otherwise standard signaling games. The players are Ann, the sender, and Bob, the receiver. At the beginning of the game, nature draws a state θ from a finite set Θ according to a full-support probability distribution $\pi \in \Delta(\Theta)$. Ann observes the realized state and takes an action a from a finite set A . The realized state is also called Ann's type.

Bob can monitor the state and Ann's action at a cost. Bob's monitoring choice is

represented by a function

$$\begin{aligned} P : \Theta \times A &\longrightarrow \Delta(X) \\ (\theta, a) &\longrightarrow P_{(\theta, a)} \end{aligned}$$

where X is a finite set of signals. The quantity $P_{(\theta, a)}(x)$ is the probability of observing signal x when the state is θ and Ann chooses a .

By analogy with the terminology of statistical decision theory, we call P Bob's *experiment*. We emphasize that we use the term experiment in the broad sense of monitoring structure, rather than in the strict sense of scientific procedure. For example, performing an experiment could represent conducting a market research, hiring a private investigator, reading a newspaper, or watching a panel discussion.

We denote by $\text{supp } P_{(\theta, a)}$ the support of $P_{(\theta, a)}$. We denote by $\text{supp } P$ the set of all signals that can be generated with positive probability by some θ and a :

$$\text{supp } P = \{x : P_{(\theta, a)}(x) > 0 \text{ for some } \theta \text{ and } a\}.$$

Let $\mathcal{E} \subseteq \Delta(X)^{\Theta \times A}$ be the set of feasible experiments; a special case of interest is $\mathcal{E} \subseteq \Delta(X)^A$ where Bob can acquire information only about Ann's action.

Bob performs an experiment $P \in \mathcal{E}$ at cost $c(P) \in [0, \infty)$. The monitoring cost could be pecuniary, psychological, or a combination of the two. After observing the outcome of the experiment, Bob takes an action b from a finite set B . We assume that the set of signals is rich enough to contain the set of Bob's actions: $B \subseteq X$.¹

Finally, payoffs are realized. Ann receives payoff $u(\theta, a, b) \in \mathbb{R}$ as a function of the state θ , her own action a , and Bob's action b . Bob receives payoff $v(\theta, a, b, c(P)) \in \mathbb{R}$ as a function of the state θ , Ann's action a , his own action b , and the cost of the experiment P that he performs. For Bob's payoff, we assume the separable form

$$v(\theta, a, b, c(P)) = v(\theta, a, b, 0) - c(P).$$

For short, we write $v(\theta, a, b)$ instead of $v(\theta, a, b, 0)$.

To recap, the timeline of the game is as follows:

1. Nature draws a state $\theta \in \Theta$ according to $\pi \in \Delta(\Theta)$.
2. Ann observes θ and chooses an action $a \in A$.
3. Without observing neither θ nor a , Bob chooses an experiment $P \in \mathcal{E}$.

¹The assumption is de facto equivalent to the weaker requirement that B has smaller cardinality than X .

4. Nature draws a signal $x \in X$ according to $P_{(\theta,a)} \in \Delta(X)$.
5. Bob observes x and chooses an action $b \in B$.
6. Payoffs are realized: $u(\theta, a, b)$ for Ann, $v(\theta, a, b) - c(P)$ for Bob.

Throughout the paper, we maintain two broad assumptions on monitoring costs: (i) there is free disposal of information, (ii) more information is more costly and strictly more information is strictly more costly. To formalize them, we use the classic ranking of experiments due to Blackwell (1951). Given an experiment $P : \Theta \times A \rightarrow \Delta(X)$ and a stochastic kernel $K : X \rightarrow \Delta(X)$, denote by $K \circ P$ the experiment given by

$$(K \circ P)_{(\theta,a)}(x) = \sum_{x'} K_{x'}(x) P_{(\theta,a)}(x').$$

Definition 1. The *Blackwell order* \succeq is a binary relation on $\Delta(X)^{\Theta \times A}$ defined by $P \succeq Q$ if there is a stochastic kernel $K : X \rightarrow \Delta(X)$ such that $Q = K \circ P$.

In words, P dominates Q in the Blackwell order if, via a stochastic kernel K , Q is reproducible from P . This reflects the idea that P conveys more information than Q . We denote by \succ the asymmetric part of \succeq : $P \succ Q$ if $P \succeq Q$ and $Q \not\succeq P$.

Assumption 1. (i) If $P \in \mathcal{E}$ and $P \succeq Q$, then $Q \in \mathcal{E}$. (ii) If $P \succeq Q$, then $c(P) \geq c(Q)$; if $P \succ Q$, then $c(P) > c(Q)$.

Assumptions (i) and (ii) are broad but still restrictive. In particular, (i) implies that there is some flexibility in information acquisition. Since the work of Sims (2003) on rational inattention, models of flexible information acquisition have become increasingly popular. These models usually assume that *all* information structures are feasible; in this paper, this would correspond to $\mathcal{E} = \Delta(X)^{\Theta \times A}$ or $\mathcal{E} = \Delta(X)^A$. Our flexibility assumption is weaker: \mathcal{E} can be smaller than $\Delta(X)^{\Theta \times A}$ or $\Delta(X)^A$.

An important difference from rational inattention is that information costs are independent of beliefs. In the most adopted specification of rational inattention, the cost of information is the expected reduction in the entropy of beliefs. In this paper, it would imply that Bob's monitoring cost depends both on his experiment P and on his beliefs about Ann's behavior. As highlighted by Ravid (2020) in a signaling game with a rationally inattentive receiver, the dependence of monitoring costs on beliefs that are endogenous to the game may generate a number of unappealing equilibria. Here we do not have to deal with these issues because Bob's monitoring cost $c(P)$ depends only on his experiment P .

In some applications, our results will hold under simpler conditions than (i) and (ii). For example, the arguments we outlined in the previous section for the beer-quiche game

require less structure than the one imposed by (i) and (ii). However, for general signaling games, the structure imposed by (i) and (ii) will be needed.

4 Sidestepping off-path beliefs

Our initial finding is that, in signaling games with costly monitoring, Nash equilibrium is outcome equivalent to any refinement of perfect Bayesian equilibrium that puts restrictions on off-path beliefs. To present the result, we first adapt the textbook definitions of Nash and perfect Bayesian equilibrium to the present model.²

A *strategy* for Ann is an action rule

$$\begin{aligned}\sigma : \Theta &\longrightarrow \Delta(A) \\ \theta &\longrightarrow \sigma_\theta.\end{aligned}$$

The quantity $\sigma_\theta(a)$ is the probability that Ann takes action a when the state is θ . We denote by $\text{supp } \sigma_\theta$ the support of σ_θ . The set of Ann's strategies is $\Sigma = \Delta(A)^\Theta$.

For simplicity, we assume that Bob does not randomize over experiments—we postpone to the end of the section a discussion of this assumption. A *strategy* for Bob consists of an experiment $P \in \mathcal{E}$ and an action rule

$$\begin{aligned}\tau : X &\longrightarrow \Delta(B) \\ x &\longrightarrow \tau_x.\end{aligned}$$

The quantity $\tau_x(b)$ is the probability that Bob takes action b after observing signal x . We denote by $\text{supp } \tau_x$ the support of τ_x . The set of Bob's action rules is $T = \Delta(B)^X$. The set of Bob's strategies is $\mathcal{E} \times T$.

Ann's strategy σ is a *best reply* to Bob's strategy (P, τ) if, for all $\theta \in \Theta$, $a \in \text{supp } \sigma_\theta$, and $a' \in A$,

$$\sum_{x,b} u(\theta, a, b) \tau_x(b) P_{(\theta,a)}(x) \geq \sum_{x,b} u(\theta, a', b) \tau_x(b) P_{(\theta,a')}(x).$$

Bob's strategy (P, τ) is a best reply to Ann's strategy σ if (P, τ) is a solution of

$$\max_{P', \tau'} \sum_{\theta, a, x, b} v(\theta, a, b) \tau'_x(b) P'_{(\theta,a)}(x) \sigma_\theta(a) \pi(\theta) - c(P').$$

Definition 2. A strategy profile (σ, P, τ) is a *Nash equilibrium* if σ and (P, τ) are best replies to each other.

²Fudenberg and Tirole (1991b) provide the leading formal definition of perfect Bayesian equilibrium for finite multi-period games with observed actions and independent type.

A *belief system* is a function

$$\begin{aligned} \mu : X &\longrightarrow \Delta(\Theta \times A) \\ x &\longrightarrow \mu_x. \end{aligned}$$

The quantity $\mu_x(\theta, a)$ represents the probability that Bob assigns to nature drawing state θ and Ann choosing action a , given that x is the realized signal. The set of belief systems is $M = \Delta(\Theta \times A)^X$.

An *assessment* (σ, P, τ, μ) consists of a strategy profile (σ, P, τ) and a belief system μ . Let $\pi \times \sigma$ be the induced probability distribution over $\Theta \times A$:

$$(\pi \times \sigma)(\theta, a) = \sigma_\theta(a)\pi(\theta).$$

Let $P_{\pi \times \sigma}$ be the induced probability distribution over X :

$$P_{\pi \times \sigma}(x) = \sum_{\theta, a} P_{(\theta, a)}(x)\sigma_\theta(a)\pi(\theta).$$

We denote by $\text{supp } P_{\pi \times \sigma}$ the support of $P_{\pi \times \sigma}$.

Definition 3. An assessment (σ, P, τ, μ) is a *perfect Bayesian equilibrium* if the following conditions hold:

- The strategy profile (σ, P, τ) is a Nash equilibrium.
- For all $x \in X$, $b \in \text{supp } \tau_x$, and $b' \in B$,

$$\sum_{\theta, a} v(\theta, a, b)\mu_x(\theta, a) \geq \sum_{\theta, a} v(\theta, a, b')\mu_x(\theta, a).$$

- For all $x \in \text{supp } P_{\pi \times \sigma}$, $\theta \in \Theta$, and $a \in A$,

$$\mu_x(\theta, a) = \frac{P_{(\theta, a)}(x)\sigma_\theta(a)\pi(\theta)}{P_{\pi \times \sigma}(x)}.$$

The literature has proposed several refinements of perfect Bayesian equilibrium that put restrictions on off-path beliefs. Here we adopt a general formulation. We represent restrictions on off-path beliefs by a correspondence $F : \Sigma \times \mathcal{E} \times T \rightrightarrows M$. We interpret $F(\sigma, P, \tau)$ as the set of belief systems that are “consistent” with the strategy profile (σ, P, τ) . Different notions of consistency lead to different correspondences F , that is, to different restrictions on off-path beliefs.

We consider correspondences $F : \Sigma \times \mathcal{E} \times T \rightrightarrows M$ that satisfy three conditions:

(i) For all $\sigma \in \Sigma$, $P \in \mathcal{E}$, and $\tau \in T$, the set $F(\sigma, P, \tau)$ is nonempty.

(ii) For all $\mu \in F(\sigma, P, \tau)$, $x \in \text{supp } P_{\pi \times \sigma}$, $\theta \in \Theta$, and $a \in A$,

$$\mu_x(\theta, a) = \frac{P_{(\theta, a)}(x) \sigma_\theta(a) \pi(\theta)}{P_{\pi \times \sigma}(x)}.$$

(iii) For all $\sigma \in \Sigma$ and $P \in \mathcal{E}$, if $\tau_x = \tau'_x$ for all $x \in \text{supp } P$, then $F(\sigma, P, \tau) = F(\sigma, P, \tau')$.

Condition (i) states that F has nonempty values, and (ii) that Bayes rule is satisfied on-path. To interpret (iii), recall that $\text{supp } P$ is the set of all signals that have positive probability for some θ and a :

$$\text{supp } P = \{x : P_{(\theta, a)}(x) > 0 \text{ for some } \theta \text{ and } a\}.$$

Condition (iii) states that restrictions on off-path beliefs are independent of the actions that Bob takes after signals that can never be generated. We denote by \mathcal{F} the set of correspondences $F : \Sigma \times \mathcal{E} \times T \rightrightarrows M$ that satisfy (i)-(iii).

Definition 4. A perfect Bayesian equilibrium (σ, P, τ, μ) is a *F-perfect Bayesian equilibrium* if $\mu \in F(\sigma, P, \tau)$.

We are now ready to state our first result.

Theorem 1. *Let Assumption 1 hold. A probability distribution $\rho \in \Delta(\Theta \times A \times B)$ is the outcome of a Nash equilibrium if and only if, for every $F \in \mathcal{F}$, it is the outcome of a F-perfect Bayesian equilibrium.*

Thus, when monitoring is costly, a probability distribution over states and actions can be induced by a Nash equilibrium if and only if it can be induced by *any* F-perfect Bayesian equilibrium. The theorem shows that the set of equilibrium outcomes does not depend on restrictions on off-path beliefs. This is the sense in which we can sidestep the question of off-path beliefs.

The “if” statement of Theorem 1 is tautological. The non-trivial result is the “only if” statement. The driving force is the following lemma:

Lemma 1. *Under Assumption 1, if (P, τ) is a best reply to σ then*

$$\text{supp } P = \text{supp } P_{\pi \times \sigma}.$$

The lemma implies that, in a Nash equilibrium, even if Ann deviates, Bob is not surprised by the signal he observes. There may be signals that Bob does not expect, that is, it is

possible that $P_{\pi \times \sigma}$ does not have full support. But those signals cannot be induced by a deviation of Ann. Theorem 1 then easily follows.

Lemma 1 is a consequence of optimal monitoring. For an intuition, assume by contradiction that $\text{supp } P \neq \text{supp } P_{\pi \times \sigma}$, that is, assume that there are θ , a , and x such that $P_{(\theta,a)}(x) > 0$ and $P_{\pi \times \sigma}(x) = 0$. If $P_{\pi \times \sigma}(x) = 0$, then $\sigma_\theta(a) = 0$ for all θ and a such that $P_{(\theta,a)}(x) > 0$. Thus x provides information only about an event that is not supposed to happen, an off-path event. From Bob's perspective, x is a costly waste, and therefore it cannot be rational to perform an experiment that generates such signal with positive probability. It follows that (P, τ) is not a best reply to σ .

We conclude the section by discussing the assumption that Bob does not randomize over experiments. First, we remark that, even if we extend our setting and allow Bob to randomize over experiments, Theorem 1 remains true, that is, we can still sidestep off-path beliefs. To illustrate, consider a ‘‘mixed’’ Nash equilibrium where Ann chooses $\sigma \in \Sigma$ and Bob chooses an element of $\Delta(\mathcal{E} \times T)$. If Bob puts positive probability on $(P, \tau) \in \mathcal{E} \times T$, then (P, τ) must be a best reply to σ . Thus Lemma 1 applies and Theorem 1 follows through.

We chose to rule out randomizations over experiments mostly to encourage applications. In applications, the set of feasible experiments can be quite rich: considering mixtures of experiments can be daunting. Of course this raises the question of equilibrium existence, which we address next.

To prove equilibrium existence, we make the following additional assumptions on monitoring costs:

Assumption 2. (i) \mathcal{E} is closed and convex. (ii) c is convex and continuous.

For (i) and (ii), we view experiments as elements of a Euclidean space whose dimension is the cardinality of $\Theta \times A$. The sequence of experiments (P^n) has limit P if, for every $\theta \in \Theta$, $a \in A$, and $x \in X$, the sequence $(P_{(\theta,a)}^n(x))$ has limit $P_{(\theta,a)}(x)$. For $\lambda \in [0, 1]$, the convex combination $\lambda P + (1 - \lambda)Q$ of P and Q is defined by

$$(\lambda P + (1 - \lambda)Q)_{(\theta,a)}(x) = \lambda P_{(\theta,a)}(x) + (1 - \lambda)Q_{(\theta,a)}(x).$$

The next result guarantees the existence of a Nash equilibrium where Bob does not randomize over experiments.

Proposition 1. *Under Assumptions 1 and 2 hold, a Nash equilibrium (σ, P, τ) exists.*

To prove the result, we use a version of the revelation principle. Let \mathcal{E}_B be the set of ‘‘direct’’ experiments whose signals are action recommendation (recall that $B \subseteq X$):

$$\mathcal{E}_B = \{P \in \mathcal{E} : \text{supp } P \subseteq B\}.$$

Lemma 2. *Let Assumption 1 hold. A probability distribution $\rho \in \Delta(\Theta \times A \times B)$ is a Nash equilibrium outcome if and only if it is induced by some $\sigma \in \Sigma$ and $P \in \mathcal{E}_B$ that satisfy the following conditions:*

(i) *For all $\theta \in \Theta$, $a \in \text{supp } \sigma_\theta$, and $a' \in A$,*

$$\sum_{x,b} u(\theta, a, b) P_{(\theta,a)}(b) \geq \sum_{x,b} u(\theta, a', b) P_{(\theta,a')}(b).$$

(ii) *P is a solution of*

$$\max_{P' \in \mathcal{E}_B} \sum_{\theta,a,b} v(\theta, a, b) P'_{(\theta,a)}(b) \sigma_\theta(a) \pi(\theta) - c(P').$$

The lemma follows from standard arguments. Proposition 1 is an immediate consequence of Lemma 2: it is enough to observe that the auxiliary game where Ann chooses $\sigma \in \Sigma$ and Bob chooses $P \in \mathcal{E}_B$ satisfies the usual conditions for the existence of pure-strategy Nash equilibrium in infinite games with continuous payoffs.

Beyond the derivation of Proposition 1, the lemma is of independent interest: it simplifies the search for Nash equilibria. In particular, Bob's optimization problem reduces in (ii) to finite-dimensional convex programming (provided that c is continuous and convex). Denti, Marinacci, and Rustichini (2020) have many examples of cost functions that can give more structure to (ii) (see also Mensch, 2018; Pomatto, Strack, and Tamuz, 2020).

5 Robust equilibrium characterization

We now characterize all Nash equilibria that can arise across all cost functions—the focus on Nash equilibria is justified by Theorem 1. To state the result, we need additional notation. Fix a set \mathcal{E} of experiments that satisfies free disposal of information: if $P \succeq Q$ and $P \in \mathcal{E}$, then $Q \in \mathcal{E}$. Recall that \mathcal{E}_B is the set of experiments that put positive probability only on signals belonging to B (cf. Lemma 2):

$$\mathcal{E}_B = \{P \in \mathcal{E} : \text{supp } P \subseteq B\}.$$

For $P \in \mathcal{E}_B$ and $b \in B$, let $P_b \in \mathbb{R}^{\Theta \times A}$ be the vector given by

$$P_b = (P_{(\theta,a)}(b) : \theta \in \Theta, a \in A).$$

For $p \in \Delta(\Theta \times A)$, denote by $BR(p) \subseteq B$ the set of Bob's rational actions:

$$BR(p) = \arg \max_b \sum_{\theta, a} v(\theta, a, b) p(\theta, a).$$

Let \mathcal{C} be the class of cost functions $c : \mathcal{E} \rightarrow [0, \infty)$ that are strictly increasing in the Blackwell order: if $P \succeq Q$, then $c(P) \geq c(Q)$; if $P \succ Q$, then $c(P) > c(Q)$.

Theorem 2. *A probability distribution $\rho \in \Delta(\Theta \times A \times B)$ is a Nash equilibrium outcome for some $c \in \mathcal{C}$ if and only if it is induced by some $\sigma \in \Sigma$ and $P \in \mathcal{E}_B$ that satisfy the following conditions:*

(i) *For all $\theta \in \Theta$, $a \in \text{supp } \sigma_\theta$, and $a' \in A$,*

$$\sum_b u(\theta, a, b) P_{(\theta, a)}(b) \geq \sum_b u(\theta, a', b) P_{(\theta, a')}(b).$$

(ii) *For all $b \in \text{supp } P_{\pi \times \sigma}$ and $b' \in B$,*

$$\sum_{\theta, a} v(\theta, a, b) \mu_b(\theta, a) \geq \sum_{\theta, a} v(\theta, a, b') \mu_b(\theta, a)$$

where $\mu_b(\theta, a) = P_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta) / P_{\pi \times \sigma}(b)$.

(iii) $\text{supp } P = \text{supp } P_{\pi \times \sigma}$.

(iv) *For all $b, b' \in \text{supp } P_{\pi \times \sigma}$, if $BR(\mu_b) \cap BR(\mu_{b'}) \neq \emptyset$ then P_b and $P_{b'}$ are linearly dependent.*

To put it differently, if $\sigma \in \Sigma$ and $P \in \mathcal{E}_B$ satisfy (i)-(iv), then there is $c \in \mathcal{C}$ such that the signaling game with monitoring cost c has a Nash equilibrium outcome $\rho \in \Delta(\Theta \times A \times B)$ given by $\rho(\theta, a, b) = P_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta)$. Conversely, for every $c \in \mathcal{C}$ and for every Nash equilibrium outcome $\rho \in \Delta(\Theta \times A \times B)$ of the signaling game with monitoring cost c , there are $\sigma \in \Sigma$ and $P \in \mathcal{E}_B$ that generate ρ and satisfy (i)-(iv).

The theorem provides a robust equilibrium characterization across cost functions. We are motivated by the observation that, in many settings, monitoring costs depend on factors—such as time, effort, and cognitive resources—that are difficult to measure. The analyst therefore may be interested in predictions that are not sensitive to the specific details of monitoring costs.

Conditions (i)-(iv) admit a simple interpretation: Think of $P \in \mathcal{E}_B$ as a “direct” experiment that both conveys information and recommends what action to take. Then (i) is Ann's incentive compatibility constraint, and (ii) states that Bob is willing to obey the

recommendation of the direct experiment. By (iii), Bob never receives a recommendation that he does not expect (cf. Lemma 1). To understand (iv), take $b'' \in BR(\mu_b) \cap BR(\mu_{b'})$. Bob could merge the recommendations b and b' into a single recommendation b'' . The merge would reduce monitoring costs, unless b and b' are “replica” of each other—formally, unless P_b and $P_{b'}$ are linearly dependent.

Conditions (i)-(iv) are broad but can have substantial bite. As a proof of concept, the next example revisits Cho and Kreps’s beer-quiche game. Assuming that Bob can acquire information only about Ann’s action, we use Theorem 2 to show that, in every equilibrium with costly monitoring, the surly type of Ann chooses beer with probability one. By contrast, in the standard beer-quiche game where Bob observes Ann’s action, both beer-beer and quiche-quiche are possible equilibrium outcomes.

Example 1. Consider the beer-quiche game of Figure 1. Assume that Bob can acquire information only about Ann’s action: $\mathcal{E} \subseteq \Delta(X)^A$. We claim that, if $\sigma \in \Sigma$ and $P \in \Delta(B)^A$ satisfy (i)-(iv) of Theorem 2, then $\sigma_{\theta_s}(beer) = 1$.

By contradiction, suppose (i)-(iv) and $\sigma_{\theta_s}(beer) < 1$. By (i), $\sigma_{\theta_w}(quiche) = 1$: if the surly chooses quiche with positive probability, then the wimp strictly prefers quiche to beer. Moreover, again by (i), $P_{beer}(duel) > P_{quiche}(duel)$: if the surly chooses quiche with positive probability, then she expects a duel with higher probability after beer. It follows from (iii) that $P_{\pi \times \sigma}(duel) > 0$: Bob expects to receive the recommendation *duel*.

Now we show that Bob disobeys the recommendation *duel*, reaching a contradiction. Because $P_{beer}(duel) > P_{quiche}(duel)$ and $\sigma_{\theta_w}(quiche) = 1$, it follows from Bayes rule that $\mu_{duel}(\theta_w) < 1/10$ (recall that $1/10$ is the prior probability that Ann is a wimp). Thus, in particular, $\mu_{duel}(\theta_w) < 1/2$, which implies $duel \notin BR(\mu_{duel})$. Thus Bob disobeys the recommendation *duel*, contradicting (ii). We conclude that (i)-(iv) imply $\sigma_{\theta_s}(beer) = 1$. \blacktriangle

Overall, the beer-quiche game with costly monitoring features two classes of equilibria. In one class, the types of Ann split up: the surly type chooses beer, the wimp chooses quiche. Intuitively, these separating equilibria arise when the cost of monitoring is large and Bob does not find it optimal to acquire information about Ann’s breakfast, even if he has incentive to do so. For example, (i)-(iv) are satisfied by

$$\sigma_{\theta_s}(beer) = \sigma_{\theta_w}(quiche) = 1 \quad \text{and} \quad P_{beer}(don't) = P_{quiche}(don't) = 1.$$

In words, Bob does not monitor Ann’s breakfast and acts on his prior beliefs that surly is more likely; Ann always has her favorite breakfast, not fearing Bob’s duel.

In the second class of equilibria, the surly type chooses beer and the wimp randomizes between quiche and beer. Intuitively, these semi-separating equilibria arise when the cost of monitoring is not too large: Bob acquires noisy information to make the wimp indifferent

between beer and quiche; the wimp randomizes to give Bob the incentive to monitor Ann's breakfast. For example, (i)-(iv) are satisfied by

$$\sigma_{\theta_s}(beer) = 1, \sigma_{\theta_w}(quiche) = \delta, P_{beer}(duel) = \delta^2, P_{quiche}(duel) = \frac{1}{2} + \delta^2$$

where $\delta > 0$ is sufficiently small. Note that, as δ goes to zero, these equilibria converge to the beer-beer pooling equilibrium of the standard signaling game where Bob observes Ann's action. We will come back to this limit result in the next section when we study the case of vanishing monitoring costs.

We conclude this section by discussing the proof of Theorem 2. Recall that a version of the revelation principle holds (cf. Lemma 2): for a fixed $c \in \mathcal{C}$, finding the Nash equilibria is equivalent to finding all $\sigma \in \Sigma$ and $P \in \mathcal{E}_B$ such that

- (1) for all $\theta \in \Theta$, $a \in \text{supp } \sigma_\theta$, and $a' \in A$,

$$\sum_{x,b} u(\theta, a, b)P_{(\theta,a)}(b) \geq \sum_{x,b} u(\theta, a', b)P_{(\theta,a')}(b),$$

- (2) P is a solution of

$$\max_{P' \in \mathcal{E}_B} \sum_{\theta,a,b} v(\theta, a, b)P'_{(\theta,a)}(b)\sigma_\theta(a)\pi(\theta) - c(P').$$

Condition (1) corresponds to (i) of Theorem 2. In the proof of Theorem 2, we show that (2) implies (ii)-(iv). Moreover, we show that if (ii)-(iv) hold, there exists $c \in \mathcal{C}$ such that (2) holds.

6 Vanishing costs

Small monitoring costs can be seen as perturbations of the standard signaling game where Bob observes for free Ann's action. What equilibria of standard signaling games survive these perturbations? We now address this refinement question.

The structure

$$\Gamma = (\Theta, \pi, A, B, u, v)$$

determines a standard signaling game where Bob, before moving, observes Ann's action. An *assessment* (σ, τ, μ) in Γ consists of an action rule $\sigma : \Theta \rightarrow A$ for Ann, an action rule $\tau : A \rightarrow \Delta(B)$ for Bob, and a belief system $\mu : A \rightarrow \Delta(\Theta)$. Let $\sigma_\pi \in \Delta(A)$ be the induced probability distribution over Ann's actions:

$$\sigma_\pi(a) = \sum_{\theta} \sigma_\theta(a)\pi(\theta).$$

We denote by $\text{supp } \sigma_\pi$ the support of σ_π .

Definition 5. An assessment (σ, τ, μ) in Γ is a *perfect Bayesian equilibrium* (PBE) if the following conditions are satisfied:

- For all $\theta \in \Theta$, $a \in \text{supp } \sigma_\theta$, and $a' \in A$,

$$\sum_b u(\theta, a, b)\tau_a(b) \geq \sum_b u(\theta, a', b)\tau_{a'}(b).$$

- For all $a \in A$, $b \in \text{supp } \tau_a$, and $b' \in B$,

$$\sum_{\theta} v(\theta, a, b)\mu_a(\theta) \geq \sum_{\theta} v(\theta, a, b')\mu_a(\theta).$$

- For all $\theta \in \Theta$ and $a \in \text{supp } \sigma_\pi$,

$$\mu_a(\theta) = \frac{\sigma_\theta(a)\pi(\theta)}{\sigma_\pi(a)}.$$

Throughout the section, we assume that $\mathcal{E} = \Delta(X)^A$. For every cost function $c : \Delta(X)^A \rightarrow [0, \infty)$, the structure

$$\Gamma^c = (\Gamma, \Delta(X)^A, c) = (\Theta, \pi, A, B, u, v, \Delta(X)^A, c)$$

determines a signaling game with costly monitoring where Bob can flexibly acquire information about Ann's action. Denote by \mathcal{C} the class of cost functions $c : \Delta(X)^A \rightarrow [0, \infty)$ that are strictly increasing in the Blackwell order: if $P \succeq Q$, then $c(P) \geq c(Q)$; if $P \succ Q$, then $c(P) > c(Q)$. Next is the main definition of this section.

Definition 6. A perfect Bayesian equilibrium $(\sigma^*, \tau^*, \mu^*)$ of Γ is *consistent with costly monitoring* if, for every $\epsilon > 0$, there are a cost function $c \in \mathcal{C}$ and a Nash equilibrium (σ, P, τ) of Γ^c such that the following conditions hold:

- For all $P \in \Delta(X)^A$, $c(P) \leq \epsilon$.
- For all $\theta \in \Theta$, $a \in A$, and $b \in B$,

$$\left| \sum_x \tau_x(b)P_a(x)\sigma_\theta(a) - \tau_a^*(b)\sigma_\theta^*(a) \right| \leq \epsilon.$$

The definition is in the spirit of the refinement literature: a PBE that is inconsistent with costly monitoring may be deemed less likely to occur and therefore discarded to obtain sharper predictions. Most existing refinements of PBE focus on the property of the belief system μ^* when Bayes rule does not apply. An interesting aspect of the exercise we propose is that off-path beliefs play no role in Γ^c (cf. Theorem 1).

The next theorem, the main result of the section, operationalizes Definition 6.

Theorem 3. *A perfect Bayesian equilibrium $(\sigma^*, \tau^*, \mu^*)$ of Γ is consistent with costly monitoring if and only if, for every $\epsilon > 0$, there are $\sigma \in \Sigma$ and $P \in \Delta(B)^A$ that satisfy the following conditions:*

(i) *For all $\theta \in \Theta$, $a \in \text{supp } \sigma_\theta$, and $a' \in A$,*

$$\sum_b u(\theta, a, b) P_a(b) \geq \sum_b u(\theta, a', b) P_{a'}(b).$$

(ii) *For all $b \in \text{supp } P_{\sigma_\pi}$ and $b' \in B$,*

$$\sum_{\theta, a} v(\theta, a, b) \mu_b(\theta, a) \geq \sum_{\theta, a} v(\theta, a, b') \mu_b(\theta, a)$$

where $\mu_b(\theta, a) = P_a(b) \sigma_\theta(a) \pi(\theta) / P_{\sigma_\pi}(b)$.

(iii) $\text{supp } P = \text{supp } P_{\sigma_\pi}$.

(iv) *For all $b, b' \in \text{supp } P_{\sigma_\pi}$, if $BR(\mu_b) \cap BR(\mu_{b'}) \neq \emptyset$ then P_b and $P_{b'}$ are linearly dependent.*

(v) *For all $\theta \in \Theta$, $a \in A$, and $b \in B$, $|P_a(b) \sigma_\theta(a) - \tau_a^*(b) \sigma_\theta^*(a)| \leq \epsilon$.*

The theorem has bite because it bypasses the cost of monitoring. In principle, to find what PBE of Γ are consistent with costly monitoring, we would need to search over all cost functions $c \in \mathcal{C}$ and all Nash equilibria of Γ^c . Thanks to Theorem 3, we only need to search over all $\sigma \in \Sigma$ and $P \in \mathcal{E}_B$ that satisfy (i)-(iv). Note that (i) and (ii) per se put no restrictions: the two conditions are trivially satisfied by $\sigma = \sigma^*$ and $P = \tau^*$. If σ_π^* has full support, that is, if Ann takes all actions with positive probability, then also (iii) is trivially satisfied by $\sigma = \sigma^*$ and $P = \tau^*$.

The proof of Theorem 3 builds on Theorem 2. The “only if” statement of Theorem 3 logically follows from the “only if” statement of Theorem 2. For the “if” statement, an extra step is needed: The “if” statement of Theorem 2 guarantees the existence of a cost function c . The “if” statement of Theorem 3 requires, in addition, that the cost function c is arbitrarily small.

We now illustrate the content of Theorem 3 through a series of examples. The first example goes back to Cho and Kreps’s beer-*quiche* game and shows that the only PBE of Γ that is consistent with costly monitoring is the beer-beer pooling equilibrium.

Example 1 (continued). Consider the beer-*quiche* game of Figure 1. Recall that Γ has two classes of PBE: Ann always has *quiche* or always has beer.

Using Theorem 2 we have already shown that, in every equilibrium with costly monitoring, the surly type chooses beer with probability one. Therefore the *quiche-quiche* pooling PBE of Γ is not consistent with costly monitoring, since it cannot be approximated by any equilibrium of Γ^c for any choice of c .

We claim that the beer-beer PBE is consistent with costly monitoring. To prove it, we use Theorem 3. For every $\delta \in (0, 1)$, define $\sigma \in \Sigma$ and $P \in \Delta(B)^A$ by

$$\sigma_{\theta_s}(beer) = 1, \sigma_{\theta_w}(quiche) = \delta, P_{beer}(duel) = \delta^2, P_{quiche}(duel) = \frac{1}{2} + \delta^2.$$

It is easy to verify that, for every $\epsilon > 0$, we can take δ sufficiently small for σ and P to satisfy (i)-(v). Thus the beer-beer PBE is consistent with costly monitoring. \blacktriangle

The example shows that consistency with costly monitoring is a nontrivial equilibrium refinement. Applying Theorem 3 to the beer-*quiche* game, we concluded that the *quiche-quiche* PBE is not consistent with costly monitoring.

There is a simple intuition for this selection result: For the surly type of Ann to have *quiche* with high probability, Bob must threaten a duel when Ann has beer. To implement the threat, Bob must monitor Ann’s action. For Bob to have an incentive to monitor Ann’s action, Ann must have beer with positive probability. But only the surly type of Ann can have beer with positive probability; when Ann is a wimp, having *quiche* and not dueling is the best she can hope for. Thus, whether Ann has beer or *quiche*, the best action for Bob is not to duel. Hence Bob does not have an incentive to monitor Ann’s action, and cannot implement the threat.

Cho and Kreps (1987) also propose to select the beer-beer pooling equilibrium. They deem the *quiche-quiche* pooling equilibrium “counterintuitive.” As they argue, it is unreasonable for Bob to believe that Ann is a wimp if she deviates to beer, given that the wimp achieves in the *quiche-quiche* equilibrium the best she can hope for. Cho and Kreps formalize these ideas for general signaling games and propose an influential equilibrium refinement, the “intuitive criterion.”

To provide a formal comparison with consistency with costly monitoring, we now define

the intuitive criterion. For $\Theta' \subseteq \Theta$ and $a \in A$, define $BR(\Theta', a) \subseteq B$ by

$$BR(\Theta', a) = \left\{ b : b \in BR(p) \text{ for some } p \in \Delta(\Theta \times A) \text{ such that } \sum_{\theta' \in \Theta'} p(\theta', a) = 1 \right\}.$$

In words, $BR(\Theta', a)$ is the set of actions that are optimal for Bob when he is convinced that the state belongs to Θ' and Ann takes action a .

Definition 7 (Cho and Kreps, 1987). Let (σ, τ, μ) be a PBE of Γ . For every $a \notin \text{supp } \sigma_\pi$, define $\Theta_a \subseteq \Theta$ by

$$\Theta_a = \left\{ \theta : \sum_{a', b} u(\theta, a', b) P_{a'}(b) \sigma_\theta(a') > \max_{b \in BR(\Theta, a)} u(\theta, a, b) \right\}.$$

The PBE fails the *intuitive criterion* if there are $\theta \notin \Theta_a$ and $a \notin \text{supp } \sigma_\pi$ such that

$$\sum_{a', b} u(\theta, a', b) P_{a'}(b) \sigma_\theta(a') < \min_{b \in BR(\Theta \setminus \Theta_a, a)} u(\theta, a, b).$$

The intuitive criterion puts restrictions on Bob's off-path beliefs. To illustrate, suppose that Ann does not play action a in equilibrium, that is, $a \notin \text{supp } \sigma_\pi$. The set Θ_a contains the states where Ann has no reason to deviate with a : her equilibrium payoff is better than anything she can get with a . The intuitive criterion prescribes that, given a , Bob should put probability zero on the state belonging to Θ_a .

As Cho and Kreps show, in the beer-quiche game the intuitive criterion rules out the PBE where Ann always has quiche. Thus, in the beer-quiche game, the predictions of the intuitive criterion coincide with the predictions of consistency with costly monitoring. For general signaling games, there is not a simple relation between the intuitive criterion and consistency with costly monitoring, as we show in the next two examples. In the first example, the intuitive criterion is a stronger refinement; in the second example, the intuitive criterion is a weaker refinement.

Example 2. Consider the signaling game in Figure 2. The game Γ has two classes of PBE, a pooling class and a separating class. In the pooling PBE, Ann always goes left; given left, Bob goes down; given right, Bob goes down with probability at least $1/2$. In the separating PBE, Ann goes right when the state is high, and goes left when the state is low; given left, Bob goes down; given right, Bob goes up.

The separating PBE trivially satisfies the intuitive criterion, since Ann takes all actions with positive probability. Using Theorem 3, it is also easy to see that the separating PBE is consistent with costly monitoring.




Figure 2: Signaling game for Example 2.




Figure 3: Settlement game (Banks and Sobel, 1987).

The situation is different for the pooling PBE: the pooling PBE fails the intuitive criterion but is consistent with costly monitoring. From the perspective of the intuitive criterion, it is unreasonable for Bob to put positive probability on the state being low when Ann goes right—when the state is low, left is a strictly dominant action for Ann. To prove that the pooling PBE is consistent with costly monitoring, define $\sigma \in \Sigma$ and $P \in \Delta(B)^A$ by

$$\sigma_\theta(a) = \begin{cases} 1 & \text{if } a = \textit{left} \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad P_a(b) = \begin{cases} 1 & \text{if } b = \textit{down} \\ 0 & \text{otherwise.} \end{cases}$$

For every $\epsilon > 0$, σ and P satisfy (i)-(v) of Theorem 3. Thus the pooling PBE is consistent with costly monitoring.

Intuitively, if Ann always goes left, Bob has no incentive to monitor her action: he can just go down. But down is precisely the action of Bob that makes going right suboptimal for Ann. Thus Bob does not need to monitor Ann's action to implement the threat of going down when she goes right. ▲

Example 3. Consider the signaling game in Figure 3, which is the settlement game from Banks and Sobel (1987). Ann, the defendant, is either guilty ($\theta = \theta_g$) or innocent ($\theta = \theta_i$). To avoid trial, Ann makes a settlement offer to Bob, the plaintiff. The offer can be either high or low. Bob decides whether to accept or reject the offer.

The game Γ has two classes of PBE. In the first class of PBE, Ann always makes the low offer, and Bob accepts both offers. In the second class of PBE, Ann always makes the high offer; Bob accepts the high offer and reject the low offer. It is easy to see that both classes of PBE satisfy the intuitive criterion.

We claim that the PBE where Ann always makes the high offer is not consistent with costly monitoring. We prove it by contradiction using Theorem 3. By contradiction, suppose that, for every $\epsilon > 0$, there are $\sigma \in \Sigma$ and $P \in \Delta(B)^A$ that satisfy (i)-(v) of Theorem 3. For ϵ sufficiently small, $\sigma_\theta(\text{high})$ must be close to one for both $\theta = \theta_g$ and $\theta = \theta_i$. Thus, by (i), $P_{\text{low}}(\text{reject}) > 0$. Together with (iii), this implies that $P_{\sigma_\pi}(\text{reject}) > 0$. By (ii), $\text{reject} \in BR(\mu_{\text{reject}})$, which in turn implies $\sigma_{\theta_g}(\text{low}) > 0$. By (i), $\sigma_{\theta_g}(\text{low}) > 0$ implies that

$$-3P_{\text{low}}(\text{accept}) - 11P_{\text{low}}(\text{reject}) \geq -5P_{\text{high}}(\text{accept}) - 11P_{\text{high}}(\text{reject}).$$

It follows from simple algebra that

$$-3P_{\text{low}}(\text{accept}) - 6P_{\text{low}}(\text{reject}) > -5P_{\text{high}}(\text{accept}) - 6P_{\text{high}}(\text{reject}).$$

Together with (i), this implies that $\sigma_{\theta_i}(\text{low}) = 1$, and this contradicts the fact that $\sigma_{\theta_i}(\text{low})$ is close to one. Overall, we conclude that the PBE where Ann always makes the high offer is not consistent with costly monitoring.

Using Theorem 3, it is easy to see that the PBE where Ann always makes the low offer is consistent with costly monitoring. Define $\sigma \in \Sigma$ and $P \in \Delta(B)^A$ by

$$\sigma_\theta(a) = \begin{cases} 1 & \text{if } a = \text{low} \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad P_a(b) = \begin{cases} 1 & \text{if } b = \text{accept} \\ 0 & \text{otherwise.} \end{cases}$$

For every $\epsilon > 0$, σ and P satisfy (i)-(v) of Theorem 3. Thus the PBE where Ann always makes the low offer is consistent with costly monitoring. \blacktriangle

The literature has proposed several alternative refinements to the intuitive criterion. For example, the settlement game of Figure 3 motivated Banks and Sobel (1987) to introduce “divinity.” We discuss divinity in the online appendix. In particular, we show that also divinity is neither stronger nor weaker than consistency with costly monitoring. A (generically) stronger refinement than the intuitive criterion and divinity is “stability” (Kohlberg and Mertens, 1986). An open question is whether stability is (generically) stronger than




Figure 4: Non-generic signaling game for Example 4.

consistency with costly monitoring.

So far we have discussed the restrictions that consistency with costly monitoring puts on the perfect Bayesian equilibria of Γ where some action of Ann is off the equilibrium path. The next example shows that consistency with costly monitoring puts restrictions also on PBE where Ann plays all actions with positive probability.

Example 4. Consider the signaling game in Figure 4. In all PBE of Γ , Ann goes left when the state is high, and goes right when the state is low. Bob is always indifferent, thus any specification of Bob's play is part of a PBE.

We claim that a perfect Bayesian equilibrium $(\sigma^*, \tau^*, \mu^*)$ of Γ is consistent with costly monitoring if and only if $\tau_{left}^* = \tau_{right}^*$. The intuition for the result is straightforward: whatever Ann does, Bob has no incentive to monitor her action, thus his play must be independent of her play.

To formally prove the result, we can use Theorem 3. Suppose first $(\sigma^*, \tau^*, \mu^*)$ is a PBE of Γ that is consistent with costly monitoring. Then, for every $\epsilon > 0$, there are $\sigma \in \Sigma$ and $P \in \Delta(B)^A$ that satisfy (i)-(v) of Theorem 3. By (i), $\sigma_{\theta_h}(left) = 1$ and $\sigma_{\theta_l}(right) = 1$. Thus, by (iv), we must have $P_{left} = P_{right}$. Taking the limit as $\epsilon \rightarrow 0$, we conclude that $\tau_{left}^* = \tau_{right}^*$.

Conversely, suppose that $(\sigma^*, \tau^*, \mu^*)$ is a PBE of Γ such that $\tau_{left}^* = \tau_{right}^*$. It is easy to check that, for every $\epsilon > 0$, $\sigma = \sigma^*$ and $P = \tau^*$ satisfy (i)-(v) of Theorem 3. We conclude that $(\sigma^*, \tau^*, \mu^*)$ is consistent with costly monitoring. \blacktriangle

The example shows that it is possible for a PBE of Γ where Ann plays all actions with positive probability to be inconsistent with costly monitoring. An important difference from the examples discussed before is that Example 4 is not generic: Bob is indifferent among all actions, a non-generic property of Γ . An open question is whether there exists a non-generic example of a PBE where all Ann's actions have positive probability and the PBE




Figure 5: Signaling game for Example 5.

is inconsistent with costly monitoring.

Finally, we observe that the notion of consistency with costly monitoring can be extended to the Nash equilibria (σ^*, τ^*) of Γ in an obvious way:

Definition 8. A Nash equilibrium (σ^*, τ^*) of Γ is *consistent with costly monitoring* if, for every $\epsilon > 0$, there are a cost function $c \in \mathcal{C}$ and a Nash equilibrium (σ, P, τ) of Γ^c such that the following conditions hold:

- For all $P \in \Delta(B)^A$, $c(P) \leq \epsilon$.
- For all $\theta \in \Theta$, $a \in A$, and $b \in B$,

$$\left| \sum_x \tau_x(b) P_x(x) \sigma_\theta(a) - \tau_a^*(b) \sigma_\theta^*(a) \right| \leq \epsilon.$$

Inspecting the proof in the appendix, it is easy to see that also Theorem 3 extends in an obvious way to the Nash equilibria (σ^*, τ^*) of Γ .³ The extension is interesting because, as the next example shows, Γ may have Nash equilibria that are consistent with costly monitoring and do not correspond to any perfect Bayesian equilibrium.

Example 5. Consider the signaling game in Figure 5. The game Γ has a Nash equilibrium (σ^*, τ^*) where Ann always goes left and Bob always goes down. This Nash equilibrium does not correspond to any PBE of Γ : it is not sequentially rational for Bob to go down after right. The game Γ has a unique PBE: Ann goes right when the state is high, and goes left when the state is low; Bob goes down on the left, and goes up on the right.

The Nash equilibrium (σ^*, τ^*) does not correspond to any PBE, but it is consistent with costly monitoring. The intuition for the result is straightforward: if Ann always goes

³The proof of Theorem 3 in the appendix uses only the fact that (σ^*, τ^*) is Nash equilibrium and does not make any reference to the strategy profile (σ^*, τ^*) being part of a PBE $(\sigma^*, \tau^*, \mu^*)$.

left, Bob has no incentive to monitor her action; if Bob does not monitor Ann’s action, the sequential rationality constraint that would force him to go up after the deviation to right is not binding. More formally, it is easy to check that, for every $\epsilon > 0$, $\sigma = \sigma^*$ and $P = \tau^*$ satisfy (i)-(v) of Theorem 3. Thus (σ^*, τ^*) is consistent with costly monitoring. \blacktriangle

Using the equilibrium existence result for signaling games with costly monitoring (cf. Proposition 1), it is easy to prove that any signaling game Γ has a Nash equilibrium that is consistent with costly monitoring.

Proposition 2. *Every signaling game Γ has a Nash equilibrium that is consistent with costly monitoring.*

An open question is whether every signaling game Γ has a perfect Bayesian equilibrium that is consistent with costly monitoring.

7 Related literature and concluding remarks

Signaling games with costly monitoring have recently received attention. Solan and Yariv (2004) study “games with espionage,” a special case of our framework where the sender has no private information. They prove many interesting properties of games with espionage. For example, they characterize all distributions over actions that can arise in a game with espionage across monitoring costs. The result is in the same spirit of our Theorem 2. Ultimately, however, the two equilibrium characterizations are quite different, since the two papers make different assumptions on monitoring costs. In particular, Solan and Yariv do not assume that the cost of monitoring is strictly increasing in the Blackwell order. For this reason, the outcome equivalence between Nash equilibrium and perfect Bayesian equilibrium does not hold in their framework (their solution concept is PBE).

Sobel (2012) and de Clippel and Rozen (2021) study cheap talk games where the receiver has to pay a cost to decode the message of the sender. Their setups can be seen as special cases of ours where the action of the sender is payoff irrelevant. An important difference with our paper is that, in Sobel (2012) and de Clippel and Rozen (2021), the sender can put effort to reduce or increase the monitoring cost of the receiver. This reflects the natural idea that the sender, by being more or less eloquent, can affect how easy is for the receiver to understand what she says. It could be interesting to extend our framework to allow the sender to affect the monitoring cost of the receiver. For example, the sender’s action could be bi-dimensional $a = (a_1, a_2)$ where a_1 is not observed, a_2 is observed, and a_2 affects the cost of monitoring a_1 .

Bilancini and Boncinelli (2018) study signaling games with costly monitoring where the monitoring decision is binary: the receiver either observes the action of the sender or he

does not. Because of the rigidity in the monitoring choice, the outcome equivalence between Nash equilibrium and perfect Bayesian equilibrium does not hold in their framework (their solution concept is PBE). The relevance of off-path beliefs motivate Bilancini and Boncinelli to develop a number of equilibrium refinements for signaling games with costly monitoring.

Ravid (2020) studies a signaling game between an informed seller and an inattentive buyer. The seller, who is informed of the quality of the good, makes an offer to the buyer. The buyer has to pay costly attention to learn both the quality good and the offer of the seller. Ravid shows that attention limitation may lead to inefficiencies. In particular, in the case of vanishing costs, the buyer overpays for low-quality goods, underpays for high-quality goods, and earns strictly positive profits. Matejka and McKay (2012) also study strategic pricing with inattentive buyers. Their framework features multiple sellers competing for the attention of a unit mass of buyers.

A main difference with our paper is the specification of the cost of information. Matejka and McKay (2012) and Ravid (2020) follow the literature on rational inattention and assume that the cost of information is proportional to the expected reduction in the entropy of beliefs. In this specification, the cost of information depends both on the experiment P and on the receiver's beliefs about the strategy σ chosen by the sender. The dependence on the receiver's belief about σ , which is an endogenous object, may give rise to a number of unappealing equilibria. Ravid (2020) finds a clever solution to the issue via a refinement in the spirit of Selten's perfect equilibrium (the issue is not binding in Matejka and McKay, 2012). The issue that Ravid has to deal with does not arise in our framework, since the cost of information depends only on the experiment P chosen by the receiver. See Denti, Marinacci, and Rustichini (2020) for a broader discussion of the relation between beliefs and cost of information.

We study the case of vanishing costs to provide a micro-foundation for restrictions on off-path beliefs in standard signaling games where the receiver has knowledge of the sender's action. The same motivation underlies Fudenberg and He (2018). Fudenberg and He embed an otherwise standard signaling game in a learning model where a large population of agents interacts repeatedly. Their framework also features costly experimentation, but it is the sender—to be more precise, the agents who play in the role of the sender—who experiments. In Fudenberg and He (2018), the sender has an incentive to experiment with different actions to learn the strategy of the receiver. The cost of experimentation is the implicit cost of not taking the myopically optimal action. In our paper, the cost of experimentation is explicit; the receiver has an incentive to experiment to learn the action taken by the sender.

A number of recent papers have used costly information acquisition to relax common knowledge assumptions in games. Yang (2015), Denti (2020), and Morris and Yang (2019) study coordination games with costly information acquisition. They use the case of van-

ishing costs to revisit the selection results from global games. Ravid, Roesler, and Szentes (2020) study a bilateral trading game where the buyer can acquire information about the quality of the good before receiving an offer from the seller (the offer from the seller is observed for free). They show that, when the cost of information vanishes, equilibria converge to the worst free-learning equilibrium.

Appendix

A Proofs

A.1 Proofs of the results in Section 4

Proof of Lemma 1. We prove the contrapositive: if $\text{supp } P \neq \text{supp } P_{\pi \times \sigma}$, the (P, τ) is not a best reply to σ .

Fix some $x^* \in \text{supp } P_{\pi \times \sigma}$. Define $K : X \rightarrow \Delta(X)$ by

$$K_{x'}(x) = \begin{cases} 1 & \text{if } x = x' \text{ and } x' \in \text{supp } P_{\pi \times \sigma} \\ 0 & \text{if } x \neq x' \text{ and } x' \in \text{supp } P_{\pi \times \sigma} \\ 1 & \text{if } x = x^* \text{ and } x' \notin \text{supp } P_{\pi \times \sigma} \\ 0 & \text{if } x \neq x^* \text{ and } x' \notin \text{supp } P_{\pi \times \sigma}. \end{cases}$$

Take $Q = K \circ P$. We now show that (Q, τ) is a better reply to σ than (P, τ) .

Claim 1. If $Q_{(\theta, a)}(x) > 0$ for some $\theta \in \Theta$ and $a \in A$, then $Q_{\pi \times \sigma}(x) > 0$.

Proof of the claim. Suppose that $Q_{(\theta, a)}(x) > 0$ for some θ and a . Then $x \in \text{supp } P_{\pi \times \sigma}$, which implies that $Q_{(\theta', a')}(x) \geq P_{(\theta', a')}(x)$ for all θ' and a' . It follows from $P_{\pi \times \sigma}(x) > 0$ that $Q_{\pi \times \sigma}(x) > 0$. \square

Claim 2. $P \succ Q$.

Proof of the claim. By definition $P \succeq Q$. To verify that $Q \not\succeq P$, by the equivalence theorem of Blackwell (1951), it is enough to find conjectures $\hat{\pi} \in \Delta(\Theta)$ and $\hat{\sigma} \in \Sigma$, a finite set of actions \hat{B} , and a payoff function $\hat{v} : \Theta \times A \times \hat{B} \rightarrow \mathbb{R}$ such that

$$\sum_x \max_{\hat{b}} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}) P_{(\theta, a)}(x) \hat{\sigma}_{\theta}(a) \hat{\pi}(\theta) > \sum_x \max_{\hat{b}} \sum_{\theta, a} \hat{v}(a, \theta, \hat{b}) Q_{(\theta, a)}(x) \hat{\sigma}_{\theta}(a) \hat{\pi}(\theta).$$

So, let $\hat{\pi} \times \hat{\sigma}$ have full support, let $\hat{B} = \{\hat{b}_1, \hat{b}_2\}$, and define \hat{v} by

$$\hat{v}(\theta, a, \hat{b}_1) = \begin{cases} 1 & \text{if } \sigma_\theta(a) = 0 \\ 0 & \text{otherwise,} \end{cases}$$

$$\hat{v}(\theta, a, \hat{b}_2) = \max_{x \in \text{supp } Q_{\hat{\pi}}} \sum_{(\theta, a): \sigma(a|\theta)=0} \frac{Q_{(\theta, a)}(x) \hat{\sigma}_\theta(a) \hat{\pi}(\theta)}{Q_{\hat{\pi} \times \hat{\sigma}}(x)}.$$

In words, \hat{b}_1 is a bet on the event $\{(\theta, a) : \sigma_\theta(a) = 0\}$ —which is nonempty because $\text{supp } P \neq \text{supp } P_{\pi \times \sigma}$ —and \hat{b}_2 is a safe action whose payoff is independent of θ and a .

For short, write $\hat{v}(\hat{b}_2)$ instead of $\hat{v}(\theta, a, \hat{b}_2)$. We observe that $\hat{v}(\hat{b}_2) \in (0, 1)$. Indeed, $\hat{v}(\hat{b}_2) > 0$ because the set $\{(\theta, a) : \sigma_\theta(a) = 0\}$ is nonempty. Moreover, $\text{supp } Q = \text{supp } Q_{\pi \times \sigma}$ by Claim 1. Thus, for every x such that $Q_{\hat{\pi}}(x) > 0$, there are a and θ such that $Q_{(\theta, a)}(x) > 0$ and $\sigma_\theta(a) > 0$. This shows that $\hat{v}(\hat{b}_2) < 1$.

For every x , we have

$$\begin{aligned} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}_1) Q_{(\theta, a)}(x) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) &= \sum_{(\theta, a): \sigma(a|\theta)=0} Q_{(\theta, a)}(x) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) \\ &\leq \hat{v}(\hat{b}_2) Q_{\hat{\pi} \times \hat{\sigma}}(x) = \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}_2) Q_{(\theta, a)}(x) \hat{\sigma}_\theta(a) \hat{\pi}(\theta). \end{aligned}$$

It follows that

$$\sum_x \max_{\hat{b}} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}) Q_{(\theta, a)}(x) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) = \hat{v}(\hat{b}_2).$$

Obviously, we have that

$$\sum_x \max_{\hat{b}} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}) P_{(\theta, a)}(x) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) \geq \hat{v}(\hat{b}_2).$$

Since $\text{supp } P \neq \text{supp } P_{\pi \times \sigma}$, we can find θ , a , and x such that $P_{(\theta, a)}(x) > 0$ and $P_{\pi \times \sigma}(x) = 0$. Since $P_{(\theta, a)}(x) > 0$, we have $P_{\hat{\pi} \times \hat{\sigma}}(x) > 0$, given that $\hat{\pi} \times \hat{\sigma}$ has full support. Since $P_{\pi \times \sigma}(x) = 0$, we have

$$\sum_{(\theta, a): \sigma(a|\theta)=0} P_{(\theta, a)}(x) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) = \sum_{\theta, a} P_{(\theta, a)}(x) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) = P_{\hat{\pi} \times \hat{\sigma}}(x).$$

We obtain that

$$\begin{aligned} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}_1) P_{(\theta, a)}(x) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) &= P_{\hat{\pi} \times \hat{\sigma}}(x) \\ &> \hat{v}(\hat{b}_2) P_{\hat{\pi} \times \hat{\sigma}}(x) = \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}_2) P_{(\theta, a)}(x) \hat{\sigma}_\theta(a) \hat{\pi}(\theta), \end{aligned}$$

where we use the fact that $\hat{v}(\hat{b}_2) \in (0, 1)$. Overall, we conclude that

$$\sum_x \max_{\hat{b}} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}) P_{(\theta, a)}(x) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) > \hat{v}(\hat{b}_2) = \sum_x \max_{\hat{b}} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}) Q_{(\theta, a)}(x) \hat{\sigma}_\theta(a) \hat{\pi}(\theta).$$

By Blackwell's equivalence theorem, this shows that $Q \not\preceq P$, as desired. \square

Claim 3. For all $\theta \in \Theta$, $a \in A$, and $x \in X$,

$$P_{(\theta, a)}(x) \sigma_\theta(a) = Q_{(\theta, a)}(x) \sigma_\theta(a).$$

Proof of the claim. If $\sigma_\theta(a) = 0$, then the result is trivial. Suppose therefore that $\sigma_\theta(a) > 0$. By definition of Q , we have that

$$\begin{aligned} Q_{(\theta, a)}(x) \sigma_\theta(a) &= \sum_{x'} K_{x'}(x) P_{(\theta, a)}(x') \sigma_\theta(a) \\ &= \sum_{x' \in \text{supp } P_{\pi \times \sigma}} 1_{\{x\}}(x') P_{(\theta, a)}(x') \sigma_\theta(a) + \sum_{x' \notin \text{supp } P_{\pi \times \sigma}} 1_{\{x\}}(x^*) P_{(\theta, a)}(x') \sigma_\theta(a). \end{aligned}$$

If $\sigma_\theta(a) > 0$ and $P_{\pi \times \sigma}(x') = 0$, then it must be that $P_{(\theta, a)}(x') = 0$. Thus

$$\sum_{x' \notin \text{supp } P_{\pi \times \sigma}} 1_{\{x\}}(x^*) P_{(\theta, a)}(x') \sigma_\theta(a) = 0.$$

We obtain that

$$Q_{(\theta, a)}(x) \sigma_\theta(a) = \begin{cases} P_{(\theta, a)}(x) \sigma_\theta(a) & \text{if } P_{\pi \times \sigma}(x) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

If $P_{\pi \times \sigma}(x) = 0$, then $P_{(\theta, a)}(x) \sigma_\theta(a) = 0$. The desired result follows. \square

Since $P \succ Q$ (see Claim 2), it follows from Assumption 1 that $Q \in \mathcal{E}$ and $c(P) > c(Q)$. Moreover, by Claim 3,

$$\sum_{\theta, a, x, b} v(\theta, a, b) \tau_x(b) P_{(\theta, a)}(x) \sigma_\theta(a) \pi(\theta) = \sum_{\theta, a, x, b} v(\theta, a, b) \tau_x(b) Q_{(\theta, a)}(x) \sigma_\theta(a) \pi(\theta).$$

Thus, when Ann plays σ , Bob does strictly better with (Q, τ) rather than with (P, τ) . We conclude that (P, τ) is not a best reply to σ . ■

Proof of Theorem 1. The “if” direction is tautological. To show the “only if” direction, let (σ, P, τ) be a Nash equilibrium. Fix some $\mu \in F(\sigma, P, \tau)$. For every $x \notin \text{supp } P$, fix some $b_x^* \in B$ such that

$$b_x^* \in \arg \max_{b \in B} \sum_{\theta, a} v(\theta, a, b) \mu_x(a, \theta).$$

Define $\tau^* : X \rightarrow \Delta(B)$ by

$$\tau_x^*(b) = \begin{cases} \tau_x(b) & \text{if } x \in \text{supp } P \\ 1 & \text{if } x \notin \text{supp } P \text{ and } b = b_x^* \\ 0 & \text{if } x \notin \text{supp } P \text{ and } b \neq b_x^*. \end{cases}$$

The strategy profiles (σ, P, τ) and (σ, P, τ^*) induce the same state-action distribution.

We claim that (σ, P, τ^*, μ) is a F -perfect Bayesian equilibrium. Since $\tau_x = \tau_x^*$ for all $x \in \text{supp } P$, it is clear that σ and (P, τ^*) are best replies to each other, given that σ and (P, τ) are best replies to each other. Thus (σ, P, τ^*) is a Nash equilibrium.

Now take $x \in \text{supp } P$ and $b \in \text{supp } \tau_x$. By Lemma 1, $x \in \text{supp } P_{\pi \times \sigma}$. Since (P, τ) is a best reply to σ , b is a solution of

$$\max_{b'} \sum_{\theta, a} v(\theta, a, b') P_{(\theta, a)}(x) \sigma_\theta(a) \pi(\theta).$$

Moreover, since μ obeys Bayes rule on-path, for all θ and a ,

$$\mu_x(\theta, a) = \frac{P_{(\theta, a)}(x) \sigma_\theta(a) \pi(\theta)}{P_{\pi \times \sigma}(x)}.$$

We conclude that b is a solution of

$$\max_{b'} \sum_{\theta, a} v(\theta, a, b') \mu_x(\theta, a). \tag{1}$$

If $x \notin \text{supp } P$ and $b \in \text{supp } \tau_x$, then $b = b_x^*$ and (1) holds by construction.

Finally, since $\tau_x = \tau_x^*$ for all $x \in \text{supp } P$, $F(\sigma, P, \tau) = F(\sigma, P, \tau^*)$. Thus $\mu \in F(\sigma, P, \tau^*)$. We conclude that (σ, P, τ^*, μ) is a F -perfect Bayesian equilibrium. ■

Proof of Lemma 2. The result follows by combining the two claims below.

Claim 4. If (P, τ) is a best reply to σ , then $\tau \circ P$ is a solution of

$$\max_{Q \in \mathcal{E}_B} \sum_{\theta, a, b} v(\theta, a, b) Q_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta) - c(Q). \quad (2)$$

Proof of the claim. Let (P, τ) be a best reply to σ . Define $P^* = \tau \circ P$. Observe that $c(P) \geq c(P^*)$ (being c is increasing in the Blackwell order and $P \succeq P^*$). Moreover,

$$\sum_{\theta, a, b} v(\theta, a, b) P_{(\theta, a)}^*(b) \sigma_\theta(a) \pi(\theta) = \sum_{\theta, a, x, b} v(\theta, a, b) \tau_x(b) P_{(\theta, a)}(x) \sigma_\theta(a) \pi(\theta).$$

We obtain that

$$\begin{aligned} & \sum_{\theta, a, b} v(\theta, a, b) P_{(\theta, a)}^*(b) \sigma_\theta(a) \pi(\theta) - c(P^*) \\ & \geq \sum_{\theta, a, x, b} v(\theta, a, b) \tau_x(b) P_{(\theta, a)}(x) \sigma_\theta(a) \pi(\theta) - c(P). \end{aligned}$$

Take any $\tau^* : X \rightarrow \Delta(B)$ such that $\tau_b^*(b) = 1$ for all $b \in B$. Since (P, τ) is a best reply to σ , we have that, for all $Q \in \mathcal{E}_B$,

$$\begin{aligned} & \sum_{\theta, a, x, b} v(\theta, a, b) \tau_x(b) P_{(\theta, a)}(x) \sigma_\theta(a) \pi(\theta) - c(P) \\ & \geq \sum_{\theta, a, x, b} v(\theta, a, b) \tau_x^*(b) Q_{(\theta, a)}(x) \sigma_\theta(a) \pi(\theta) - c(Q) \\ & = \sum_{\theta, a, b} v(\theta, a, b) Q_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta) - c(Q). \end{aligned}$$

We conclude that P^* is a solution of (2). \square

Claim 5. For $P \in \mathcal{E}_B$ and $\tau \in T$, if P is a solution of (2) and $\tau_b(b) = 1$ for all $b \in B$, then (P, τ) is a best reply to σ .

Proof of the claim. Let P be a solution of (2), and choose τ such that $\tau_b(b) = 1$ for all $b \in B$. To verify that (P, τ) is a best reply to σ , consider an alternative strategy (P', τ') . Define $Q' = \tau' \circ P'$. Since $B \subseteq X$, the experiment Q' is well defined. In addition, since

$P' \succeq Q'$, we have $Q' \in \mathcal{E}_B$ and $c(P') \geq c(Q')$. Thus

$$\begin{aligned}
& \sum_{\theta,a,b} v(\theta, a, b) P_{(\theta,a)}(b) \sigma_\theta(a) \pi(\theta) - c(P) \\
& \geq \sum_{\theta,a,b} v(\theta, a, b) Q'_{(\theta,a)}(b) \sigma_\theta(a) \pi(\theta) - c(Q') \\
& = \sum_{\theta,a,x,b} v(\theta, a, b) \tau'_x(b) P'_{(\theta,a)}(x) \sigma_\theta(a) \pi(\theta) - c(Q') \\
& \geq \sum_{\theta,a,x,b} v(\theta, a, b) \tau'_x(b) P'_{(\theta,a)}(x) \sigma_\theta(a) \pi(\theta) - c(P').
\end{aligned}$$

where the first inequality follows from P being a solution of (2). We conclude that (P', τ') is not a profitable deviation: (P, τ) is a best reply to σ . \square

■

Proof of Proposition 1. Consider the auxiliary game where Ann chooses $\sigma \in \Sigma$ to maximize

$$\sum_{\theta,a,b} u(\theta, a, b) P_{(\theta,a)}(b) \sigma_\theta(a) \pi(\theta), \tag{3}$$

and Bob choose $P \in \mathcal{E}_B$ to maximize

$$\sum_{\theta,a,b} v(\theta, a, b) P_{(\theta,a)}(b) \sigma_\theta(a) \pi(\theta) - c(P). \tag{4}$$

The set Σ is compact and convex, and the quantity (3) is continuous in (σ, P) and affine in σ . By Assumption 2, the set \mathcal{E}_B is compact and convex, and the quantity (4) is continuous in (σ, P) and concave in P . By standard argument (see, e.g., Fudenberg and Tirole, 1991a, Theorem 1.2), the auxiliary game has a pure Nash equilibrium (σ^*, P^*) . The pair (σ^*, P^*) satisfies (i) and (ii) of Lemma 2. It follows from Lemma 2 that the signaling game with costly monitoring has a Nash equilibrium. \blacksquare

A.2 Proof of Theorem 2

Proof of the “only if” statement

Fix $c \in \mathcal{C}$. Let (σ^*, P^*, τ^*) be a Nash equilibrium of the signaling game with monitoring cost c . Define $\sigma = \sigma^*$ and $P = \tau^* \circ P^*$. We now show that $\sigma \in \Sigma$ and $P \in \mathcal{E}_B$ satisfy (i)-(iv).

Condition (i) is satisfied because σ^* is a best reply to (P^*, τ^*) . To verify (ii), take b

such that $P_{\pi \times \sigma}(b) > 0$. For all $x \in \text{supp } P_{\pi \times \sigma}^*$, define $\mu_x^* \in \Delta(\Theta \times A)$ by

$$\mu_x^*(\theta, a) = \frac{P_{(\theta, a)}^*(x) \sigma_\theta^*(a) \pi(\theta)}{P_{\pi \times \sigma}^*(x)}.$$

Simple algebra shows that

$$\mu_b = \sum_{x \in \text{supp } P_{\pi \times \sigma}^*} \frac{\tau_x^*(b) P_{\pi \times \sigma}^*(x)}{P_{\pi \times \sigma}(b)} \mu_x^*.$$

Since (P^*, τ^*) is a best reply to σ^* , if $P_{\pi \times \sigma}^* > 0$ and $\tau_x^*(b) > 0$, then $b \in BR(\mu_x^*)$. Thus $b \in BR(\mu_b)$. This shows that (ii) is satisfied.

To verify (iii), take θ, a , and b such that $P_{(\theta, a)}(b) > 0$. Since $P_{(\theta, a)}(b) = \sum_x \tau_x^*(b) P_{(\theta, a)}^*(x)$, there is x such that $P_{(\theta, a)}^*(x) > 0$ and $\tau_x^*(b) > 0$. By Lemma 1, $P_{\pi \times \sigma}^*(x) > 0$. Thus

$$P_{\pi \times \sigma}(b) = \sum_x \tau_x^*(b) P_{\pi \times \sigma}^*(x) > 0.$$

We conclude that (iii) holds.

Now we verify (iv). It follows from standard arguments that P is a solution of

$$\max_{Q \in \mathcal{E}_B} \sum_{\theta, a, b} v(\theta, a, b) Q_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta) - c(Q).$$

Since c is strictly increasing in the Blackwell order, it must be that, for all $Q \in \mathcal{E}_B$ such that $P \succ Q$, we have that

$$\sum_{\theta, a, b} v(\theta, a, b) P_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta) > \sum_{\theta, a, b} v(\theta, a, b) Q_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta).$$

Then (iv) follows from the next result, which concludes the proof of the ‘‘only if’’ statement of Theorem 2.

Claim 6. Let $\sigma \in \Sigma$ and $P \in \mathcal{E}_B$. Suppose that, for all $Q \in \mathcal{E}_B$ such that $P \succ Q$,

$$\sum_{\theta, a, b} v(\theta, a, b) P_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta) > \sum_{\theta, a, b} v(\theta, a, b) Q_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta).$$

Then, for all $b, b' \in \text{supp } P_{\sigma_\pi}$, if $BR(\mu_b) \cap BR(\mu_{b'}) \neq \emptyset$ then the vectors P_b and $P_{b'}$ are linearly dependent.

Proof. By contraposition, suppose there are $b^1, b^2 \in \text{supp } P_{\pi \times \sigma}$ such that $BR(\mu_{b^1}) \cap BR(\mu_{b^2}) \neq \emptyset$ and the vectors P_{b^1} and P_{b^2} are linearly independent. Take $b^* \in BR(\mu_{b^1}) \cap BR(\mu_{b^2})$ and

define the stochastic kernel $K : B \rightarrow \Delta(B)$ by

$$K_{b'}(b) = \begin{cases} 1 & b = b^* \text{ and } b' \in \{b^1, b^2\} \\ 1 & b = b' \text{ and } b' \notin \{b^1, b^2\} \\ 0 & \text{otherwise.} \end{cases}$$

Take $Q = K \circ P$. By definition, $P \succeq Q$. In addition, because $b^* \in BR(\mu_{b^1}) \cap BR(\mu_{b^2})$,

$$\sum_{\theta, a, b} v(\theta, a, b) P_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta) = \sum_{\theta, a, b} v(\theta, a, b) Q_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta).$$

It remains to show that $Q \not\succeq P$. By Blackwell's equivalence theorem, it is enough to find conjectures $\hat{\pi} \in \Delta(\Theta)$ and $\hat{\sigma} \in \Sigma$, a finite set of actions \hat{B} , and a payoff function $\hat{v} : \Theta \times A \times \hat{B} \rightarrow \mathbb{R}$ such that

$$\sum_b \max_{\hat{b}} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}) P_{(\theta, a)}(b) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) > \sum_b \max_{\hat{b}} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}) Q_{(\theta, a)}(b) \hat{\sigma}_\theta(a) \hat{\pi}(\theta). \quad (5)$$

So, choose $\hat{\pi}$ and $\hat{\sigma}$ such that $\hat{\pi} \times \hat{\sigma}$ has full support on $\Theta \times A$. For $b \in \text{supp } P_{\hat{\pi} \times \hat{\sigma}}$, define $\hat{\mu}_b \in \Delta(\Theta \times A)$ by

$$\hat{\mu}_b(\theta, a) = \frac{P_{(\theta, a)}(b) \hat{\sigma}_\theta(a) \hat{\pi}(\theta)}{P_{\hat{\pi} \times \hat{\sigma}}(b)}.$$

Since P_{b^1} and P_{b^2} are linearly independent, $\hat{\mu}_{b^1} \neq \hat{\mu}_{b^2}$. Thus, by the hyperplane separation theorem, we can find $\hat{B} = \{\hat{b}^1, \hat{b}^2\}$ and $\hat{v} : \Theta \times A \times \hat{B} \rightarrow \mathbb{R}$ such that

$$\begin{aligned} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}^1) \hat{\mu}_{b^1}(\theta, a) &> \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}^1) \hat{\mu}_{b^2}(\theta, a) \\ \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}^2) \hat{\mu}_{b^1}(\theta, a) &< \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}^2) \hat{\mu}_{b^2}(\theta, a). \end{aligned}$$

By construction, for all $b \notin \{b^1, b^2, b^*\}$, we have

$$\max_{\hat{b}} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}) P_{(\theta, a)}(b) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) = \max_{\hat{b}} \sum_a \hat{v}(a, \hat{b}) Q_a(b) \hat{\sigma}_\theta(a) \hat{\pi}(\theta).$$

Moreover, it follows from $BR(\hat{\mu}_{b^1}) \cap BR(\hat{\mu}_{b^2}) = \emptyset$ that

$$\begin{aligned}
& \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}^1) P_{(\theta, a)}(b^1) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) + \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}^2) P_{(\theta, a)}(b^2) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) \\
& + 1_{B \setminus \{b_1, b_2\}}(b^*) \max_{\hat{b}} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}) P_{(\theta, a)}(b^*) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) \\
> & \max_{\hat{b}} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}) (P_{(\theta, a)}(b^1) + P_{(\theta, a)}(b^2)) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) \\
& + 1_{B \setminus \{b_1, b_2\}}(b^*) \max_{\hat{b}} \sum_{\theta, a} \hat{v}(\theta, a, b^*) P_{(\theta, a)}(b^*) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) \\
\geq & \max_{\hat{b}} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}) (P_{(\theta, a)}(b^1) + P_{(\theta, a)}(b^2) + 1_{B \setminus \{b_1, b_2\}}(b^*) P_{(\theta, a)}(b^*)) \hat{\sigma}_\theta(a) \hat{\pi}(\theta) \\
= & \max_{\hat{b}} \sum_{\theta, a} \hat{v}(\theta, a, \hat{b}) Q_{(\theta, a)}(b^*) \hat{\sigma}_\theta(a) \hat{\pi}(\theta).
\end{aligned}$$

We conclude that (5) holds, as desired. \square

Proof of the “if” statement

We prove the “if” statement of Theorem 2 through a number of intermediate lemmas.

Lemma 3. *Let $C \subseteq \mathbb{R}^n$ be a polyhedral convex set, and let $S \subseteq C$ be a finite subset. Suppose there is no hyperplane that supports C at two distinct points of S . Then there exists a compact, strictly convex set $K \subseteq \mathbb{R}^n$ such that $S \subseteq K \subseteq C$.*

Proof. The following argument was suggested to me by Mohammad Ghomi, who has studied related, albeit much harder, problems (see, e.g., Ghomi, 2004).

Let $H_1, \dots, H_m \subseteq \mathbb{R}^n$ be closed half-spaces whose intersection is C . Without loss of generality, suppose that $\partial H_i \cap C \neq \emptyset$ for all $i = 1, \dots, m$. By hypothesis, if $\partial H_i \cap S \neq \emptyset$ then there exists a unique $\xi_i \in S$ such that $\xi_i \in \partial H_i \cap S$. If $\partial H_i \cap S = \emptyset$, pick an arbitrary element $\xi_i \in \partial H_i \cap C$.

Denote by $B(\xi, r) \subseteq \mathbb{R}^n$ the ball centered in $\xi \in \mathbb{R}^n$ with radius $r > 0$:

$$B(\xi, r) = \{\xi' \in \mathbb{R}^n : \|\xi - \xi'\| \leq r\}$$

where $\|\cdot\|$ is the Euclidean norm. For every $i = 1, \dots, m$ and $r > 0$, pick $\xi_{(i,r)} \in \mathbb{R}^n$ such that the ball centered in $\xi_{(i,r)}$ with radius r is contained by H_i and is tangent to ∂H_i at ξ_i :

$$\xi_i \in B(\xi_{(i,r)}, r) \subseteq H_i.$$

Furthermore, pick $r_i > 0$ sufficiently large so that $B(\xi_{(i,r_i)}, r_i) \supseteq S$.

Define $K \subseteq \mathbb{R}^n$ by

$$K = \bigcap_{i=1}^m B(\xi_{(i,r_i)}, r_i).$$

The set K is compact, contains S , and is contained by $\bigcap_i H_i = C$. In addition, K is strictly convex, being the intersection of finitely many strictly convex sets. The desired result follows. \blacksquare

Lemma 4. *Let $p_1, \dots, p_n \in \Delta(\Theta \times A)$ be such that*

(i) *for all $i \neq j$, $BR(p_i) \cap BR(p_j) = \emptyset$;*

(ii) *the convex hull of $\{p_1, \dots, p_n\}$ is $\Delta(\Theta \times A)$.*

Then there exists a strictly convex function $\phi : \Delta(\Theta \times A) \rightarrow \mathbb{R}$ such that

$$\begin{aligned} \phi(p) &\geq \max_b \sum_{\theta, a} v(\theta, a, b) p(\theta, a), & p \in \Delta(\Theta \times A), \\ \phi(p_i) &= \max_b \sum_{\theta, a} v(\theta, a, b) p_i(\theta, a), & i = 1, \dots, n, \\ \phi(p) &\leq \max_{\theta, a, b} v(\theta, a, b), & p \in \Delta(\Theta \times A). \end{aligned}$$

Proof. Let n be the cardinality of $\Theta \times A$. For this proof, we identify $\Delta(\Theta \times A)$ with a subset of \mathbb{R}^{n-1} under the usual embedding. Define $\psi : \Delta(\Theta \times A) \rightarrow \mathbb{R}$ by

$$\psi(p) = \max_b \sum_{\theta, a} v(\theta, a, b) p(\theta, a).$$

Let $\text{epi}\psi$ be the epigraph of ψ , which is polyhedral convex set:

$$\text{epi}\psi = \bigcap_b \left\{ (p, t) \in \Delta(\Theta \times A) \times \mathbb{R} : \sum_{\theta, a} v(\theta, a, b) p(\theta, a) \leq t \right\}.$$

By (i), there is no hyperplane (of dimension $n-1$) that supports $\text{epi}\psi$ at two distinct points $(p_i, \phi(p_i))$ and $(p_j, \phi(p_j))$. Thus, by Lemma 3, there exists a compact, strictly convex set $K \subseteq \text{epi}\psi$ such that $(p_i, \phi(p_i)) \in K$ for all $i = 1, \dots, n$.

Define $\phi : \Delta(\Theta \times A) \rightarrow \mathbb{R}$ by

$$\phi(p) = \min\{t : (p, t) \in K\}.$$

By (ii), the quantity $\phi(p)$ is well defined for all $p \in \Delta(\Theta \times A)$. Because $K \subseteq \text{epi}\psi$, we have $\phi(p) \geq \psi(p)$ for all $p \in \Delta(\Theta \times A)$. Because $(p_i, \phi(p_i)) \in K$, we have $\phi(p_i) \leq \psi(p_i)$ for all

$i = 1, \dots, n$. Because K is strictly convex, ϕ is strictly convex. Finally, since ϕ is convex and the convex hull $\{p_1, \dots, p_n\}$ is $\Delta(\Theta \times A)$,

$$\max_p \phi(p) = \max_i \phi(p_i) = \max_{i,b} \sum_{\theta,a} v(\theta, a, b) p_i(\theta, a) \leq \max_{\theta,a,b} v(\theta, a, b).$$

The desired result follows. ■

Lemma 5. *Let $p_1, \dots, p_n \in \Delta(\Theta \times A)$ be such that, for all $i \neq j$, $BR(p_i) \cap BR(p_j) = \emptyset$. Then there exists a strictly convex function $\phi : \Delta(\Theta \times A) \rightarrow \mathbb{R}$ such that*

$$\begin{aligned} \phi(p) &\geq \max_b \sum_{\theta,a} v(\theta, a, b) p(\theta, a), & p \in \Delta(\Theta \times A), \\ \phi(p_i) &= \max_b \sum_{\theta,a} v(\theta, a, b) p_i(\theta, a), & i = 1, \dots, n, \\ \phi(p) &\leq 1 + \max_{\theta,a,b} v(\theta, a, b), & p \in \Delta(\Theta \times A). \end{aligned}$$

Proof. Let $p_{(\theta,a)} \in \Delta(\Theta \times A)$ be the Dirac measure concentrated on $(\theta, a) \in \Theta \times A$. Let p_{n+1}, \dots, p_{n+m} be an enumeration of the set

$$\{p_{(\theta,a)} : p_{(\theta,a)} \notin \text{co}\{p_1, \dots, p_n\}\}.$$

By construction, the convex hull of $\{p_1, \dots, p_{n+m}\}$ is $\Delta(\Theta \times A)$.

For every index $i = n+1, \dots, n+m$ and $t > 0$, define the vector $v_{(i,t)} \in \mathbb{R}^{\Theta \times A}$ by

$$v_{(i,t)}(\theta, a) = \begin{cases} 1 + \max_{\theta,a,b} v(\theta, a, b) & \text{if } p_i = p_{(\theta,a)} \\ -t & \text{otherwise.} \end{cases}$$

Observe that

$$\sum_{a,\theta} v_{(i,t)}(a, \theta) p_i(a, \theta) > \max_b \sum_{a,\theta} v(a, \theta, b) p_i(a, \theta).$$

Moreover, since p_i does not belong to the convex hull of $\{p_1, \dots, p_n\}$, we can find $t_i > 0$ sufficiently large so that, for all $j \neq i$,

$$\sum_{a,\theta} v_{(i,t_i)}(a, \theta) p_j(a, \theta) > \max_b \sum_{a,\theta} v(a, \theta, b) p_j(a, \theta).$$

Without loss of generality, assume that $B \cap \{n+1, \dots, n+m\} = \emptyset$. Define $B^* =$

$B \cup \{n+1, \dots, n+m\}$ and $v^* : \Theta \times A \times B^* \rightarrow \mathbb{R}$ by

$$v^*(\theta, a, b^*) = \begin{cases} v(\theta, a, b) & \text{if } b^* = b \in B \\ v_{(i, t_i)}(\theta, a) & \text{if } b^* = i \in \{n+1, \dots, n+m\}. \end{cases}$$

Observe that $BR^*(p_i) = BR(p_i)$ for all $i = 1, \dots, n$. Moreover, $BR^*(p_i) = \{i\}$ for all $i = n+1, \dots, n+m$. By Lemma 4, there exists a strictly convex function $\phi : \Delta(\Theta \times A) \rightarrow \mathbb{R}$ such that

$$\begin{aligned} \phi(p) &\geq \max_{b^*} \sum_{\theta, a} v^*(\theta, a, b^*) p(\theta, a), & p \in \Delta(\Theta \times A), \\ \phi(p_i) &= \max_{b^*} \sum_{\theta, a} v^*(\theta, a, b^*) p_i(\theta, a), & i = 1, \dots, n+m, \\ \phi(p) &\leq \max_{\theta, a, b^*} v^*(\theta, a, b^*), & p \in \Delta(\Theta \times A). \end{aligned}$$

Observe that

$$\begin{aligned} \max_{b^*} \sum_{\theta, a} v^*(\theta, a, b^*) p(\theta, a) &\geq \max_b \sum_{\theta, a} v(\theta, a, b) p(\theta, a), & p \in \Delta(\Theta \times A), \\ \max_{b^*} \sum_{\theta, a} v^*(\theta, a, b^*) p_i(\theta, a) &= \max_b \sum_{\theta, a} v(\theta, a, b) p_i(\theta, a), & i = 1, \dots, n, \\ \max_{\theta, a, b^*} v^*(\theta, a, b^*) &= 1 + \max_{\theta, a, b} v(\theta, a, b). \end{aligned}$$

The desired result follows. ■

Lemma 6. *Suppose that $\sigma \in \Sigma$ and $P \in \mathcal{E}_B$ satisfy the following conditions:*

(i) *for all $b \in \text{supp } P_{\pi \times \sigma}$ and $b' \in B$,*

$$\sum_{\theta, a} v(\theta, a, b) \mu_b(\theta, a) \geq \sum_{\theta, a} v(\theta, a, b') \mu_b(\theta, a)$$

where, for all $\theta \in \Theta$ and $a \in A$, $\mu_b(\theta, a) = P_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta) / P_{\pi \times \sigma}(b)$;

(ii) *$\text{supp } P = \text{supp } P_{\pi \times \sigma}$;*

(iii) *for all $b, b' \in \text{supp } P_{\pi \times \sigma}$, if $BR(\mu_b) \cap BR(\mu_{b'}) \neq \emptyset$ then P_b and $P_{b'}$ are linearly dependent.*

Then there exists a cost function $c \in \mathcal{C}$ such that P is a solution of

$$\max_{Q \in \mathcal{E}_B} \sum_{\theta, a} v(\theta, a, b) Q_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta) - c(Q).$$

In addition, we can choose c such that

$$c(Q) \leq 3 + \max_{\theta, a, b} |v(\theta, a, b)|, \quad Q \in \mathcal{E}.$$

Proof. Let $\sigma^* \in \Sigma$ be a fully mixed strategy of Ann: $\sigma^*(a|\theta) > 0$ for all θ and a . Define $v^* : \Theta \times A \times B \rightarrow \mathbb{R}$ by

$$v^*(\theta, a, b) = \frac{v(\theta, a, b)\sigma_\theta(a)}{\sigma_\theta^*(a)}.$$

For every b such that $P_{\pi \times \sigma^*}(b) > 0$, define $\mu_b^* \in \Delta(\Theta \times A)$ by

$$\mu_b^*(\theta, a) = \frac{P_{(\theta, a)}(b)\sigma_\theta^*(a)\pi(\theta)}{P_{\pi \times \sigma^*}(b)}.$$

Let $BR^*(\mu_b^*)$ be the set of optimal actions given utility function v^* and belief μ_b^* :

$$BR^*(\mu_b^*) = \arg \max_{b'} \sum_{\theta, a} v^*(\theta, a, b')\mu_b^*(\theta, a).$$

Claim 7. Let $b, b' \in \text{supp } P_{\pi \times \sigma}$ such that $BR(\mu_b) \cap BR(\mu_{b'}) = \emptyset$. If σ^* is sufficiently close to σ , then $BR^*(\mu_b^*) \cap BR^*(\mu_{b'}^*) = \emptyset$.

Proof of the claim. By contradiction, suppose there exist a sequence $(\sigma^n)_{n=1}^\infty$ of fully mixed strategies of Ann with limit σ such that, for every n , $BR^n(\mu_b^n) \cap BR^n(\mu_{b'}^n) \neq \emptyset$. Let b^n be an element of $BR^n(\mu_b^n) \cap BR^n(\mu_{b'}^n)$. Since B is finite, we can assume without loss of generality that $b^n = b^1$ for all n . Since $\mu_b^n \rightarrow \mu_b$, $\mu_{b'}^n \rightarrow \mu_{b'}$, and $u^n \rightarrow u$, it must be that $b^1 \in BR(\mu_b) \cap BR(\mu_{b'})$. This contradicts the hypothesis that $BR(\mu_b) \cap BR(\mu_{b'}) = \emptyset$. \square

From now on, we assume that σ^* is sufficiently close to σ so that, for all $b, b' \in \text{supp } P_{\pi \times \sigma}$ such that $BR(\mu_b) \cap BR(\mu_{b'}) = \emptyset$,

$$BR^*(\mu_b^*) \cap BR^*(\mu_{b'}^*) = \emptyset. \tag{6}$$

The assumption is justified by the claim above.

Claim 8. If $\mu_b^* \neq \mu_{b'}^*$, then $BR^*(\mu_b^*) \cap BR^*(\mu_{b'}^*) = \emptyset$.

Proof of the claim. Take $b, b' \in \text{supp } P_{\pi \times \sigma^*}$. By (iii), we have $b, b' \in \text{supp } P_{\pi \times \sigma}$. If $BR(\mu_b) \cap BR(\mu_{b'}) \neq \emptyset$, then $\mu_b^* = \mu_{b'}^*$ because of (ii). If $BR(\mu_b) \cap BR(\mu_{b'}) = \emptyset$, then $BR^*(\mu_b^*) \cap BR^*(\mu_{b'}^*) = \emptyset$ by (6). The desired result follows. \square

By Lemma 5, we can find a strictly convex function $\phi : \Delta(\Theta \times A) \rightarrow \mathbb{R}$ such that

$$\phi(p) \geq \max_b \sum_{\theta, a} v^*(\theta, a, b) p(\theta, a), \quad p \in \Delta(\Theta \times A), \quad (7)$$

$$\phi(p_b^*) = \max_{b'} \sum_{\theta, a} v^*(\theta, a, b') p_b^*(\theta, a), \quad b \in \text{supp } P_{\pi \times \sigma^*}, \quad (8)$$

$$\phi(p) \leq 1 + \max_{\theta, a, b} v^*(\theta, a, b), \quad p \in \Delta(\Theta \times A). \quad (9)$$

Define $c : \mathcal{E} \rightarrow [0, \infty)$ by

$$c(Q) = \sum_x \phi(\nu_x^*) Q_{\pi \times \sigma^*}(x) - \phi(\pi \times \sigma^*)$$

where $\nu_x^*(\theta, a) = Q_{(\theta, a)}(x) \sigma_\theta^*(a) \pi(\theta) / Q_{\pi \times \sigma^*}(x)$. Since ϕ is strictly convex and $\pi \times \sigma^*$ has full support, c is strictly increasing in the Blackwell order: $c \in \mathcal{C}$. In addition, possibly by making σ^* even closer to σ , we obtain that

$$c(Q) \leq 2 + \max_{\theta, a, b} |v^*(\theta, a, b)| \leq 3 + \max_{\theta, a, b} |v(\theta, a, b)|, \quad Q \in \mathcal{E}.$$

Claim 9. P is a solution of

$$\max_{Q \in \mathcal{E}_B} \sum_{\theta, a} v(\theta, a, b) Q_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta) - c(Q).$$

Proof of the claim. By (ii), we have

$$\sum_{\theta, a} v(\theta, a, b) P_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta) = \sum_b \max_{b'} v(\theta, a, b') P_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta).$$

By (8), we have

$$\begin{aligned} c(P) + \phi(\pi \times \sigma^*) &= \sum_b \max_{b'} \sum_{\theta, a} v^*(\theta, a, b') \mu_b^*(\theta, a) P_{\pi \times \sigma^*}(b) \\ &= \sum_b \max_{b'} v(\theta, a, b') P_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta). \end{aligned}$$

By (7), we have

$$\begin{aligned} c(Q) + \phi(\pi \times \sigma^*) &\geq \sum_b \max_{b'} \sum_{\theta, a} v^*(\theta, a, b') \nu_b^*(\theta, a) Q_{\pi \times \sigma^*}(b) \\ &= \sum_b \max_{b'} v(\theta, a, b') Q_{(\theta, a)}(b) \sigma_\theta(a) \pi(\theta). \end{aligned}$$

Overall, we obtain that

$$\begin{aligned}
\sum_{\theta,a} v(\theta, a, b)P_{(\theta,a)}(b)\sigma_\theta(a)\pi(\theta) - c(P) &= \sum_b \max_{b'} v(\theta, a, b')P_{(\theta,a)}(b)\sigma_\theta(a)\pi(\theta) - c(P) \\
&= \phi(\pi \times \sigma^*) \\
&\leq \sum_b \max_{b'} v(\theta, a, b')Q_{(\theta,a)}(b)\sigma_\theta(a)\pi(\theta) - c(Q) \\
&\leq \sum_{\theta,a} v(\theta, a, b)Q_{(\theta,a)}(b)\sigma_\theta(a)\pi(\theta) - c(Q).
\end{aligned}$$

The desired result follows. \square

The last claim concludes the proof of the lemma. \blacksquare

The “if” statement of Theorem 2 follows from Lemmas 2 and 6.

A.3 Proofs of the results in Section 6

Proof of Theorem 3. The “only if” direction follows immediately from Theorem 2. To prove the “if” direction, for every $n = 1, 2, \dots$, take $\sigma^n \in \Sigma$ and $P^n \in \Delta(B)^A$ that satisfy (i)-(v) for $\epsilon = 1/n$. By Lemma 6, there exists a strictly convex function $f^n : \Delta(X)^A \rightarrow [0, \infty)$ such that P^n is a solution of

$$\max_{Q \in \Delta(B)^A} \sum_{\theta,a} v(\theta, a, b)Q_a(b)\sigma_\theta^n(a)\pi(\theta) - f^n(Q). \quad (10)$$

In addition, for all $Q \in \mathcal{E}$,

$$f^n(Q) \leq 3 + \max_{\theta,a,b} |v(\theta, a, b)|. \quad (11)$$

For every $Q \in \mathcal{E}$, define $v^n(Q) \in \mathbb{R}$ by

$$v^n(Q) = \sum_b \max_{b'} \sum_{\theta,a} v(\theta, a, b')Q_a(b)\sigma_\theta^n(a)\pi(\theta) - \max_{b'} \sum_{\theta,a} v(\theta, a, b')\sigma_\theta^n(a)\pi(\theta).$$

Define also $g^n(Q) \in \mathbb{R}$ by

$$g^n(Q) = \max\{v^n(Q) - v^n(P), 0\}.$$

The function $g^n : \mathcal{E} \rightarrow \mathbb{R}$ is increasing in the Blackwell order, but not necessarily strictly increasing.

Claim 10. P^n is a solution of

$$\max_{Q \in \mathcal{E}_B} \sum_{\theta, a} v(\theta, a, b) Q_a(b) \sigma_\theta^n(a) \pi(\theta) - g^n(Q).$$

Proof of the claim. To ease the exposition, assume without loss of generality that

$$\max_{b'} \sum_{\theta, a} v(\theta, a, b') \sigma_\theta^n(a) \pi(\theta) = 0.$$

If $v^n(P^n) \geq v^n(Q)$, then

$$\begin{aligned} \sum_{\theta, a, b} v(\theta, a, b) Q_a(b) \sigma_\theta^n(a) \pi(\theta) - g^n(Q) &\leq v^n(Q) \\ &\leq v^n(P^n) = \sum_{\theta, a, b} v(\theta, a, b) P_a^n(b) \sigma_\theta^n(a) \pi(\theta) - g^n(P^n) \end{aligned}$$

where the last equality holds by (ii). If instead $v^n(P^n) < v^n(Q)$, then

$$\begin{aligned} \sum_{\theta, a, b} v(\theta, a, b) Q_a(b) \sigma_\theta^n(a) \pi(\theta) - g^n(Q) &\leq v^n(Q) - (v^n(Q) - v^n(P^n)) \\ &= v^n(P^n) = \sum_{\theta, a, b} v(\theta, a, b) P_a^n(b) \sigma_\theta^n(a) \pi(\theta) - g^n(P^n) \end{aligned}$$

where the last equality holds by (ii). The desired result follows. \square

Claim 11. For every $Q \in \mathcal{E}$, $g^n(Q) \rightarrow 0$.

Proof of the claim. Define

$$v^*(Q) = \sum_x \max_b \sum_{\theta, a} v(\theta, a, b) Q_a(x) \sigma_\theta^*(a) \pi(\theta) - \max_b \sum_{\theta, a} v(\theta, a, b) \sigma_\theta^*(a) \pi(\theta).$$

By (v), $\sigma^n \rightarrow \sigma^*$, which implies $v^n(Q) \rightarrow v^*(Q)$. In addition, $v^n(P^n) \rightarrow v^*(\tau)$. Thus

$$\lim_{n \rightarrow \infty} w^n(Q) = \max\{v^*(Q) - v^*(\tau), 0\} = 0$$

where the last equality follows from τ^* being a best reply to σ^* . \square

Define the cost function $c^n : \mathcal{E} \rightarrow [0, \infty)$ by

$$c^n(Q) = \frac{1}{n} f^n(Q) + \frac{n-1}{n} g^n(Q).$$

Since f^n and g^n are increasing in the Blackwell order, with f_n strictly increasing, c^n is strictly increasing: $c^n \in \mathcal{C}$. By (10) and Claim 10, P^n is a solution of

$$\max_{Q \in \mathcal{E}_B} \sum_{\theta, a} v(\theta, a, b) Q_a(b) \sigma_\theta^n(a) \pi(\theta) - c^n(Q).$$

By (11) and Claim 11, $c^n(Q) \rightarrow 0$. The desired result follows from Lemma 2. \blacksquare

Proof of Proposition 2. Take $c \in \mathcal{C}$ such that c is convex and continuous. For every $n = 1, 2, \dots$, define $c^n : \mathcal{E} \rightarrow [0, \infty)$ by $c^n(P) = c(P)/n$. The cost function c^n is convex, continuous, and belongs to \mathcal{C} . By Proposition 1, the game Γ^{c^n} has a Nash equilibrium (σ^n, P^n, τ^n) . Without loss of generality, suppose that the sequences $(\sigma^n)_{n=1}^\infty$, $(P^n)_{n=1}^\infty$, and $(\tau^n)_{n=1}^\infty$ converges to $\sigma \in \Sigma$, $P \in \mathcal{E}$, and $\tau \in T$. Define $\sigma^* = \sigma$ and $P^* = \tau \circ P$.

Claim 12. (σ^*, τ^*) is a Nash equilibrium of Γ .

Proof of the claim. Take $\theta \in \Theta$ and $a \in A$ such that $\sigma^*(a|\theta) > 0$. For n sufficiently large, $\sigma^n(a|\theta) > 0$. Since σ^n is a best reply to (P^n, τ^n) , we have

$$\sum_{x, b} u(\theta, a, b) \tau_x^n(b) P_a^n(x) \geq \sum_{x, b} u(\theta, a', b) \tau_x^n(b) P_a^n(x), \quad a' \in A.$$

Taking the limit as $n \rightarrow \infty$, we obtain

$$\sum_b u(\theta, a, b) \tau_a^*(b) \geq \sum_b u(\theta, a', b) \tau_{a'}^*(b), \quad a' \in A.$$

This shows that σ^* is a best reply to τ^* .

Let Q be an experiment that fully reveals Ann's action. Since (P^n, τ^n) is a best reply to σ^n , we have

$$\begin{aligned} & \sum_{\theta, a, x, b} v(\theta, a, b) \tau_x^n(b) P_a^n(x) \sigma_\theta^n(a) \pi(\theta) - \frac{1}{n} c(P^n) \\ & \geq \max_{\tau'} \sum_{\theta, a, x, b} v(\theta, a, b) \tau_x'(b) Q_a(x) \sigma_\theta^n(a) \pi(\theta) - \frac{1}{n} c(Q) \\ & = \sum_a \max_b \sum_\theta v(\theta, a, b) \sigma_\theta^n(a) \pi(\theta) - \frac{1}{n} c(Q). \end{aligned}$$

Taking the limit as $n \rightarrow \infty$, we obtain

$$\sum_{\theta, a, b} v(\theta, a, b) \tau_a^*(b) \sigma_\theta^*(a) \pi(\theta) \geq \sum_a \max_b \sum_\theta v(\theta, a, b) \sigma_\theta^*(a) \pi(\theta).$$

This shows that τ^* is a best reply to σ^* . □

Thus (σ^*, τ^*) is a Nash equilibrium of Γ consistent with costly monitoring. ■

B Divinity

We now compare divinity (Banks and Sobel, 1987) and consistency with costly monitoring. To define divinity, we need additional notation. Let (σ, τ) be a strategy profile in Γ . For $\theta \in \Theta$, $a \in A$, and $\beta \in \Delta(B)$, define $L_a(\theta, \beta) \subseteq [0, 1]$ by

$$L_a(\theta, \beta) = \begin{cases} \{1\} & \text{if } \sum_b u(\theta, a, b)\beta(b) > \sum_{a,b} u(\theta, a, b)\tau_a(b)\sigma_\theta(a) \\ [0, 1] & \text{if } \sum_b u(\theta, a, b)\beta(b) = \sum_{a,b} u(\theta, a, b)\tau_a(b)\sigma_\theta(a) \\ \{0\} & \text{otherwise.} \end{cases}$$

For $a \in A$ and $\beta \in \Delta(B)$, define $\mathcal{D}_a(\beta) \subseteq \Delta(\Theta)$ by

$$\mathcal{D}_a(\beta) = \bigcup_{k>0} \bigcap_{\theta} \{p \in \Delta(\Theta) : \exists l_a(\theta, \beta) \in L_a(\theta, \beta) \text{ s.t. } p(\theta) = kl_a(\theta, \beta)\pi(\theta)\}.$$

For $a \in A$ and $\mathcal{B} \subseteq \Delta(B)$, define $\mathcal{D}_a(\mathcal{B}) \subseteq \Delta(\Theta)$ by

$$\mathcal{D}_a(\mathcal{B}) = \text{co} \left(\bigcup_{\beta \in \mathcal{B}} \mathcal{D}_a(\beta) \right)$$

where co is convex hull. For $p \in \Delta(\Theta)$ and $a \in A$, define $\overline{BR}(p, a) \subseteq \Delta(B)$ by

$$\overline{BR}(p, a) = \arg \max_{\beta \in \Delta(B)} \sum_{\theta, a, b} v(\theta, a, b)\beta(b).$$

For $D \subseteq \Delta(\Theta)$ and $a \in A$, define $\overline{BR}(D, a) \subseteq \Delta(B)$ by

$$\overline{BR}(D, a) = \bigcup_{p \in D} \overline{BR}(p, a).$$

Finally, for $a \in A$ and $n = 0, 1, \dots$, define

$$D_a^0 = \Delta(\Theta) \quad \text{and} \quad \mathcal{B}_a^0 = \Delta(B)$$

$$D_a^{n+1} = \begin{cases} \mathcal{D}_a(\mathcal{B}^n) & \text{if } \mathcal{D}_a(\mathcal{B}^n) \neq \emptyset \\ D_a^n & \text{if } \mathcal{D}_a(\mathcal{B}^n) = \emptyset \end{cases}$$




Figure 6: Signaling game for Example 6.

$$\mathcal{B}_a^{n+1} = \overline{BR}(D_a^{n+1}, a).$$

Definition 9 (Banks and Sobel, 1987). A perfect Bayesian equilibrium (σ, τ, μ) is *divine* if, for all $a \notin \text{supp } \sigma_\pi$,

$$\mu_a \in \bigcap_n D_a^n.$$

Like the intuitive criterion, divinity puts restrictions on Bob's off-path beliefs. We refer the reader to the original paper of Banks and Sobel for more discussion and interpretation. Here we focus on the relation between divinity and consistency with costly monitoring.

Like the intuitive criterion, divinity, as an equilibrium refinement, is neither weaker nor stronger than consistency with costly monitoring. It is easy to check that, in Example 2, divinity is a stronger refinement. Next we provide an example where divinity is a weaker refinement.

Example 6. Consider the signaling game in Figure 6, which is a variation on games in Banks and Sobel (1987, Figure 3) and in Cho and Kreps (1987, Figure IV). The game Γ has a class of PBE where Ann always goes left; given left, Bob chooses middle; given right, Bob randomizes between middle and down, with probability of down at least $1/2$.

It is easy to see that all PBE where Ann always goes left are divine. Indeed, notice that

$$\mathcal{D}_{\text{right}}(\text{up}) = \{\theta_l\} \quad \text{and} \quad \mathcal{D}_{\text{right}}\left(\frac{1}{2}\text{mid} + \frac{1}{2}\text{down}\right) = \{\theta_h\}.$$

This implies that $D_{\text{right}}^1 = \Delta(\Theta)$:

$$\Delta(\Theta) \supseteq D_{\text{right}}^1 \supseteq \text{co} \left\{ \mathcal{D}_{\text{right}}(\text{up}) \cup \mathcal{D}_{\text{right}}\left(\frac{1}{2}\text{mid} + \frac{1}{2}\text{down}\right) \right\} = \Delta(\Theta).$$

Furthermore, since

$$B_{right}^1 = BR(\Delta(\Theta), right) = \{\beta : \text{supp } \beta \subseteq \{up, mid\} \text{ or } \text{supp } \beta \subseteq \{mid, down\}\},$$

we conclude that $D_{right}^n = \Delta(\Theta)$ for all n . Thus divinity puts no restrictions on Bob's off-path beliefs: all left-left PBE are divine.

We claim that no PBE where Ann always goes left is consistent with costly monitoring. To prove it, we use Theorem 3. By contradiction, suppose that, for every $\epsilon > 0$, there are σ and P that satisfy (i)-(v) of Theorem 3. For ϵ sufficiently small, both $\sigma_{\theta_h}(left)$ and $\sigma_{\theta_l}(left)$ are close to one. We consider four cases depending on whether $\sigma_{\theta_h}(left)$ or $\sigma_{\theta_l}(left)$ are exactly one:

- Case 1. Suppose that $\sigma_{\theta_h}(left) = 1$ and $\sigma_{\theta_l}(left) = 1$. By (ii), the support of P_{σ_π} must be $\{mid\}$. By (iii), the support of P_{right} must be $\{mid\}$. It follows from (i) that $\sigma_{\theta_h}(left) = 0$ and $\sigma_{\theta_l}(left) = 0$: contradiction.
- Case 2. Suppose that $\sigma_{\theta_h}(left) = 1$ and $\sigma_{\theta_l}(left) \in (0, 1)$. By (i), we must have

$$P_{right}(middle) \leq \frac{1}{2} \quad \text{and} \quad P_{right}(down) = \frac{1}{3}.$$

Thus, $P_{right}(up) \geq 1/6$. By (iii), $P_{\sigma_\pi}(up) > 0$. It follows from $\mu_{up}(\theta_h, right) = 0$ that $up \notin BR(\mu_{up})$. This contradicts (ii).

- Case 3. Suppose that $\sigma_{\theta_h}(left) \in (0, 1)$ and $\sigma_{\theta_l}(left) = 1$. By (i), we must have

$$P_{right}(middle) = \frac{1}{2} \quad \text{and} \quad P_{right}(down) \geq \frac{1}{3}.$$

By (iii), $P_{\sigma_\pi}(down) > 0$. It follows from $\mu_{down}(\theta_l, right) = 0$ that $down \notin BR(\mu_{down})$. This contradicts (ii).

- Case 4. Suppose that $\sigma_{\theta_h}(left) \in (0, 1)$ and $\sigma_{\theta_l}(left) \in (0, 1)$. By (i), we must have

$$P_{right}(middle) = \frac{1}{2} \quad \text{and} \quad P_{right}(down) = \frac{1}{3}.$$

By (iii), $P_{\sigma_\pi}(up) > 0$ and $P_{\sigma_\pi}(down) > 0$. By (ii), we must have

$$\begin{aligned} 4\mu_{up}(\theta_h, right) + \mu_{up}(\theta_l, right) &\geq 3 \\ 4\mu_{down}(\theta_l, right) + \mu_{down}(\theta_h, right) &\geq 3. \end{aligned}$$

This implies that

$$\frac{\mu_{up}(\theta_h, right)}{\mu_{up}(\theta_l, right)} > 1 > \frac{\mu_{down}(\theta_h, right)}{\mu_{down}(\theta_l, right)}.$$

However, this contradicts the fact that

$$\frac{\mu_{up}(\theta_h, right)}{\mu_{up}(\theta_l, right)} = \frac{\sigma_{\theta_h}(right)}{\sigma_{\theta_l}(right)} = \frac{\mu_{down}(\theta_h, right)}{\mu_{down}(\theta_l, right)}.$$

Overall, we conclude that no PBE where Ann always goes left is consistent with costly monitoring.

Note that Γ has also a PBE where Ann always goes right and Bob always chooses middles. Using Theorem 3, it is easy to check that this PBE is consistent with costly monitoring. \blacktriangle

References

- Sushant Acharya and Shu Lin Wee. Rational inattention in hiring decisions. *American Economic Journal: Macroeconomics*, 12:1–40, 2020.
- Jeffrey S Banks and Joel Sobel. Equilibrium selection in signaling games. *Econometrica*, 55:647–661, 1987.
- Vojtěch Bartoš, Michal Bauer, Julie Chytilová, and Filip Matějka. Attention discrimination: Theory and field experiments with monitoring information acquisition. *American Economic Review*, 106:1437–75, 2016.
- Ennio Bilancini and Leonardo Boncinelli. Signaling with costly acquisition of signals. *Journal of Economic Behavior & Organization*, 145:141–150, 2018.
- David Blackwell. Comparison of experiments. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*, pages 93–102. University of California Press, 1951.
- In-Koo Cho and David M Kreps. Signaling games and stable equilibria. *Quarterly Journal of Economics*, 102:179–221, 1987.
- Geoffroy de Clippel and Kareen Rozen. Communication, perception, and strategic obfuscation. 2021.
- Tommaso Denti. Unrestricted information acquisition. 2020.
- Tommaso Denti, Massimo Marinacci, and Aldo Rustichini. Experimental cost of information. 2020.

- Drew Fudenberg and Kevin He. Learning and type compatibility in signaling games. *Econometrica*, 86:1215–1255, 2018.
- Drew Fudenberg and Jean Tirole. *Game theory*. MIT press, 1991a.
- Drew Fudenberg and Jean Tirole. Perfect bayesian equilibrium and sequential equilibrium. *journal of Economic Theory*, 53(2):236–260, 1991b.
- Mohammad Ghomi. Optimal smoothing for convex polytopes. *Bulletin of the London Mathematical Society*, 36:483–492, 2004.
- Elon Kohlberg and Jean-Francois Mertens. On the strategic stability of equilibria. *Econometrica*, 54:1003–1037, 1986.
- David M Kreps and Robert Wilson. Sequential equilibria. *Econometrica*, 50:863–894, 1982.
- Filip Matejka and Alisdair McKay. Simple market equilibria with rationally inattentive consumers. *American Economic Review P&P*, 102:24–29, 2012.
- Jeffrey Mensch. Cardinal representations of information. 2018.
- Stephen Morris and Ming Yang. Coordination and continuous stochastic choice. 2019.
- Rik Pieters and Michel Wedel. Attention capture and transfer in advertising: Brand, pictorial, and text-size effects. *Journal of Marketing*, 68:36–50, 2004.
- Luciano Pomatto, Philipp Strack, and Omer Tamuz. The cost of information. 2020.
- Doron Ravid. Ultimatum bargaining with rational inattention. *American Economic Review*, 110:2948–63, 2020.
- Doron Ravid, Anne-Katrin Roesler, and Balazs Szentes. Learning before trading: on the inefficiency of ignoring free information. 2020.
- Reinhard Selten. Re-examination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4:25–55, 1975.
- Christopher A Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50:665–690, 2003.
- Joel Sobel. Complexity versus conflict in communication. In *2012 46th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6, 2012.
- Eilon Solan and Leeat Yariv. Games with espionage. *Games and Economic Behavior*, 47:172–199, 2004.

Ming Yang. Coordination with flexible information acquisition. *Journal of Economic Theory*, 158:721–738, 2015.