Signaling via Research

Christian Baker [*]

January, 2024
DRAFT - IN PROGRESS

Abstract

Firms hire researchers of unknown type in a competitive job market. Researchers are able to signal their type via the choice of a costly experiment. Unlike standard signaling scenarios, researchers face a problem of information design rather than optimal effort. While the realized information of the experiment is observable, the experimental process is not. We show that this can limit the informativeness of the equilibrium research for pooling and fully-separating equilibria: except in a special case of isodivergent experiments, there always exists an equilibrium with observable experiments that is more informative.

1 Introduction

A number of specialized job markets determine job offers based upon the quality of research produced by a prospective worker. For instance, a university hiring faculty will examine an applicant's portfolio of research in order to infer the potential contribution of that applicant's future research. Researchers entering such a job market will attempt to assemble a portfolio that will signal their value. However, the firm making the job offer may not be able to discriminate between experiments ex post when they differ only in their counterfactual outcomes, forcing it to make an offer based on the

[*]cbaker8128@gmail.com

experiment's result, rather than its design. To what extent does signalling only via observable evidence alter the informativeness of the research design?

This signaling problem is the focus of this paper. It has an important difference from the classic signaling problem (Spence, 1973): Instead of a worker choosing an effort level, she must choose an experiment. She must decide not only how much information to gather, but also how to gather it. For instance, an investigative journalist conducting a survey must not only choose the number of hours to devote to an investigation, but also which witnesses to question and what documents to seek. The journalist may conclude her investigation immediately if she finds just a few witnesses that reveal a scandal. However, this result would not reveal how she would have continued the investigation if the witnesses had failed to provide relevant information. Hence, a newspaper considering hiring the journalist cannot depend on the counterfactual portion of her research.

Our analysis utilizes a model in which firms in a competitive labor market wish to hire a researcher. The value to the firm of a researcher depends on her type, which is drawn from a finite set and is private information. Prior to receiving a wage offer, the researcher can conduct any experiment about an unknown state to signal her type. These experiments are costly, and the marginal cost of a more informative experiment (in the sense of Blackwell (1953)) depends on her type. A key assumption in our model is that neither the researcher nor the firms are influenced directly by the outcome of the experiment. When the experiment itself is observable, this reduces to the standard job market signalling problem of Spence (1973), where experiment cost is analogous to effort level.

However, signalling instead only via the outcome of the researcher's experiment, i.e. based on the posterior belief induced by the experiment, eliminates degrees of freedom from the design of the wage contract. To see why, suppose that there are two possible states of the world and consider an equilibrium in which a lower-cost researcher's experiment may either decrease the belief that the state of the world is 1 (rather than 0) to μ_1 or increase it to μ_2, and a higher-cost researcher's experiment has possible results μ_1' and μ_2'. A wage function which specifies rewards for these four results pins down the expected payoff for the experiments being performed by the two types of researchers. However, it also specifies the expected payoff for performing an experiment which can shift beliefs to either μ_1 or μ_2', as well as for one which has possible results μ_1' or μ_2.

Moreover, if an equilibrium is fully-separating so that each type is offered their expected value to the firm, then when there are only two states it is impossible for a researcher to choose an experiment with three or more outcomes. This is because the researcher would receive the same payoff regardless of which outcome is realized, and would prefer to choose an experiment that only includes the two least expensive outcomes.

Theorem 1 show that these considerations restrict the set of experiments that may be conducted in fully-separating or pooling equilibrium with observable outcomes, compared to when experiments are observable. In particular, it is the most informative experiments that cannot be supported by an equilibrium with observable outcomes. This has implications for the externalities of the signaling game. Signaling via outcomes rather than experiments may lead to less useful research being contributed to society.

An important application is to the preregistration of research. Preregistration involves specifying the hypotheses, methods, and analysis protocol in advance of conducting a study. The primary aim of preregistration is to improve the validity of statistical inference in confirmatory research by reducing problematic research practices such as p-hacking and post hoc analysis. However, preregistration also commits the researcher to fully revealing the experiment. Our results imply that this expands the set of possible equilibria to include more informative experiments.

The paper is organized as follows. Section 2 presents the signaling model, and Section 3 demonstrates several examples of equilibria. We present our main results in Section 4. Section 5 concludes.

Related Literature This paper builds on the classic signaling literature (Spence, 1973). The key difference is that instead of costly effort, costly experiments are used to signal type. Another strand of the signaling literature looks noisy signaling. de Haan et al. (2011) looks at a Spence model with two types where the receiver observes effort plus some normally distributed noise. This eliminates all pooling equilibria except for the zero-effort one, since the receiver has to have reasonable beliefs about all possible signal realizations. It also eliminates all but one separating equilibrium, where the bad type chooses zero effort and the receiver chooses a cutoff strategy. Heinsalu (2018) uses a similar model but with dynamics.

Also related is the literature on incentives in information design problems. Yoder (2022) considers the problem of designing incentives for heterogeneous researchers. In his model, principal contracts with a researcher with unknown costs to conduct an experiment This model can be seen as a screening counterpart to our signaling model, as the principal cares directly about the experiment.

Our model follows the literature on flexible costly information acquisition (Caplin and Dean, 2013; Matějka and McKay, 2015) to model the cost of information. The researcher can select any experiment at a cost that is proportional to the expected reduction of entropy it induces.

2 Model

Firms in a competitive labor market wish to hire researchers. Heterogeneous researchers choose costly experiments to signal their type. After completing their experiments, researchers enter a job market. Conditional on the belief induced by the outcome of an experiment, firms simultaneously make wage offers to the researcher. Finally, the researcher decides whether to work for a firm, and if so which one. Neither firms nor researchers care directly about the experiments.

Researchers There is finite state space Ω with $|\Omega| \geq 2$, and a common prior $\mu_0 \in \Delta(\Omega)$. The type of a researcher is drawn from the type space $\Theta = \{\theta_1, \dots, \theta_N\} \subset \mathbb{R}_{++}$ according to a full-support probability distribution f. We assume that $\theta_i < \theta_j$ for $i < j$. The type of a researcher is not observable by the firms, but researchers know their own type.

In order to signal their type, researchers can choose a finite experiment, i.e. a set of signal labels and a conditional distribution over signals in each state. Due to having a common prior, each signal can be identified with the posterior belief that it induces. Hence, we will model experiments as unconditional distributions over posterior beliefs that satisfy Bayes plausibility: the expectation of the posteriors μ is equal to the prior. Thus the researcher may choose any experiment in the set

$$\mathcal{I}(\mu_0) = \{p \in \Delta(\Delta(\Omega)) : \mathbb{E}_p[\mu] = \mu_0\}$$

The cost to a type θ researcher of running experiment p follows a posterior-separable form:

$$C(p, \theta) = \theta \left(H(\mu_0) - \mathbb{E}_p[H(\mu)] \right),$$

where $H : \Delta(\Omega) \to \mathbb{R}$ is continuously differentiable and strictly concave. The function H is a measure of uncertainty about the state. For instance, H could be taken to be Shannon entropy,

$$H(\mu) = -\sum_{\omega \in \Omega} \mu(\omega) \ln(\mu(\omega)).$$

The cost of an experiment is then proportional to the expected reduction in uncertainty.

Researchers are offered wages $w : \Delta(\Omega) \to \mathbb{R}_+$ as a function of realized posteriors. The ex ante expected payoff to a type θ_i researcher who chooses experiment $p \in \mathcal{I}(\mu_0)$ is then

$$u(p, \theta_i) = \mathbb{E}_p[w(\mu)] - C(p, \theta).$$

Firms Firms do not directly care about the experiments performed by the researcher. The value to a firm of hiring a type θ_i researcher is $V_i = v - \theta_i k$, where $v, k > 0$ are interpreted as being derived from an experiment that the firm wishes to run in some other context, for which the firm will pay the cost. Note that V_i is decreasing in i. We will assume that $V_N > 0$, and that for all i, $V_i < V_N + \theta_i H(\mu_0)$.

After observing the outcome of an experiment and forming posterior belief μ, the firm forms a belief $m(\theta_i|\mu)$ that the researcher is type θ_i. Firms compete in hiring, and offer wages equal to the expected value of the researcher, so that

$$w(\mu) = \sum_i m(\theta_i|\mu)V_i.$$

Equilibrium The timing is then: (i) Researcher chooses experiment p, (ii) signals are realized yielding posterior μ, and (iii) firms make offer $w(\mu)$. We will look at pure-strategy perfect Bayesian equilibria.

An *equilibrium with observable outcomes* $(\{p_i^*\}, m^*, w^*)$ consists of experiment strategies $p_i^* \in \mathcal{I}(\mu_0)$, beliefs $m \in \Delta(\Theta \times \Delta\Omega)$, and wages $w : \Delta\Omega \to \mathbb{R}_+$ such that:

1. For all $p \in \mathcal{I}(\mu_0)$ and $i \in \Theta$,

$$u(p_i^*, \theta_i) \geq u(p, \theta_i). \tag{IC}$$

2. For all $\mu \in \Delta\Omega$ satisfying $p_i^*(\mu) > 0$ for some i,

$$m(\theta_i|\mu) = \frac{p_i(\mu)f(\theta_i)}{\sum_j p_j(\mu)f(\theta_j)}.$$

3. For all $\mu \in \Delta\Omega$,

$$w(\mu) = \sum_i m(\theta_i|\mu)V_i.$$

The first condition is incentive compatibility, requiring all types of researchers choose an experiment that maximizes their expected payoff, given the wage function. The second condition stipulates that firms form beliefs about the type of a researcher using Bayes' rule, whenever possible. The final condition is that the wage offers are equal to the expected value of a researcher. This is a result of the firms engaging in Bertrand competition. Define $\mathcal{E}^o$ to be the set of all equilibria with observable outcomes.

It is important to note that the second condition only pins down the beliefs of the firms for outcomes that can result from some experiment chosen in the equilibrium. If some posterior μ is not in the support of any experiment then $m(\cdot|\mu)$ can not be derived via Bayes' rule. However, the firms' beliefs at these off-path posteriors are still important because they determine the wages that the researchers face when considering alternative strategies. The lowest wage that can be offered at any posterior is V_N, corresponding to firms assuming that a researcher is high-cost after observing an off-equilibrium outcome.

The two simplest classes of equilibria are *pooling equilibria* and *fully-separating equilibria*. In a pooling equilibrium every type of researcher chooses the same experiment, so that $p_i^* = p_j^*$, for all $i, j \in \Theta$. This in turn implies that every type of researcher would be offered the same wage, equal to the expected value of a randomly selected worker, $\sum_i V_i f(\theta_i)$. In a fully-separating equilibrium no two researchers choose experiments with overlapping supports, so that $\cap_{i\in\Theta}\text{supp}(p_i^*) = \emptyset$. Note that this condition is stronger than merely requiring that no two researchers choose exactly the same experiment: there can be no overlap in the distributions of posteriors. Hence in a fully-separating equilibrium every realizable outcome is fully revealing, and so the wage offered to a type θ_i researcher is V_i. We will use $\mathcal{E}_P^o$ and $\mathcal{E}_{FS}^o$ to denote the sets of all pooing and fully-separating equilibria with observable outcomes.

Observable Experiments To understand how the market for researchers differs from more traditional markets, it is useful to consider the case where firms can directly observe the experiments that are chosen by the researchers, rather than just the outcomes of those experiments. As before, researchers choose Bayes plausible experiments $p \in \mathcal{I}(\mu_0)$. But now firms form their beliefs about a researcher's type as a function of the experiment p itself.

An *equilibrium with observable experiments* consists of strategies $\hat{p}_i \in \mathcal{I}(\mu_0)$, beliefs $\hat{m} \in \Delta(\Theta \times \mathcal{I}(\mu_0))$, and wages $\hat{w} : \mathcal{I}(\mu_0) \to \mathbb{R}_+$ such that

1. For all $p \in \mathcal{I}(\mu_0)$ and $i \in \Theta$,

$$\hat{w}(\hat{p}_i) - C(\hat{p}_i, \theta_i) \geq \hat{w}(p) - C(p, \theta_i).$$

2. For all $i \in \Theta$,

$$\hat{m}(\theta_i|\hat{p}_i) = f(\theta_i) / \sum_{j|\hat{p}_j = \hat{p}_i} f(\theta_j).$$

3. For all $p \in \mathcal{I}(\mu_0)$,

$$\hat{w}(p) = \sum_i \hat{m}(\theta_i|p) V_i.$$

The interpretation of these equilibrium conditions is the same as before. The sole difference is that beliefs $\hat{m}$ about researcher types are a function of experiments, not posteriors. Since equilibria with observable experiments lack this stochastic element, they are analogous to the standard signaling problem. Define $\mathcal{E}^e$ to be the set of all equilibria with observable experiments.

Note that any set of researcher strategies in a pooling or fully-separating equilibrium with only observable outcomes can also be supported in an equilibrium with observable experiments via a straightforward transformation of the belief and wage functions. In this sense, observability of the full experiments expands the set of pooling and fully-separating equilibria.

3 Examples

To illustrate the structure of equilibria, we will take as an example the case where there are two types of researchers and two states. Take as parameter values $v = 3$, $k = 0.25$, $\theta_1 = 1$, $\theta_2 = 3$, $f(\theta_1) = 0.85$, and $\mu_0 = 0.5$. Hence the value of a type θ_1 researcher is $V_1 = 2.75$, and the value of a type θ_2 researcher is $V_2 = 2.25$. We will take H to be Shannon entropy.

We can frame a researcher's incentive compatibility constraint as

$$p_i^* \in \arg\max_{p \in \mathcal{I}(\mu_0)} \mathbb{E}_p[w(\mu) + \theta_i H(\mu)]. \tag{1}$$

It is a widely established result (Aumann et al., 1995; Kamenica and Gentzkow, 2011) that the value of the objective function at a solution to this problem is the value of the concavification of $w(\mu) + \theta_i H(\mu)$ at the prior μ_0[1]. Hence to graphically illustrate equilibria, we can plot the function $w(\mu) + \theta_i H(\mu)$ for each i and consider their concavifications.

3.1 Pooling

In any pooling equilibrium we must have that $p_1^* = p_2^*$, and that $w^*(\mu) = f(\theta_1)V_1 + f(\theta_2)V_2$ for all $\mu \in \text{supp}(p_i^*)$. Figure 1 illustrates an equilibrium in which the experiment has two outcomes in its support, $\mu_1 = 0.35$ and $\mu_2 = 0.85$, so that $p_i^*(\mu_1) = 0.3$. Regardless of which outcome is observed, the researcher is paid $w(\mu_1) = w(\mu_2) = 2.675$. Off-path beliefs are given by $m^*(\theta_2|\mu) = 1$ for all $\mu \notin \text{supp}(p_i^*)$, so that $w(\mu) = V_2$ for all off-path beliefs.

The blue and red dots in Figure 1 show the values of the objective function (1). Note that even though the value of this function is greater for the high-cost researcher, the net payoff $u(p_1^*, \theta_1)$ will be greater for the low-cost researcher.

3.2 Fully-separating

In any fully-separating equilibrium we have that $m^*(\theta_1|\mu) = 1$ for all $\mu \in \text{supp}(p_1^*)$, and $m^*(\theta_2|\mu) = 0$ for all $\mu \in \text{supp}(p_2^*)$. Type θ_1 researchers are paid V_1, and type θ_2 researchers are paid V_2, regardless of the outcome of the experiment. Figure 2 illustrates a fully-separating equilibrium in which type θ_1 chooses a binary experiment p_1^* with posteriors $\mu_1 = 0.15$ and $\mu_2 = 0.8$. Regardless of which posterior is realized, the type θ_1 researcher gets a wage V_1. The type θ_2 researcher chooses an uninformative

[1] The concavification of a function $f : X \to \mathbb{R}$ is the smallest concave function that lies weakly above f.

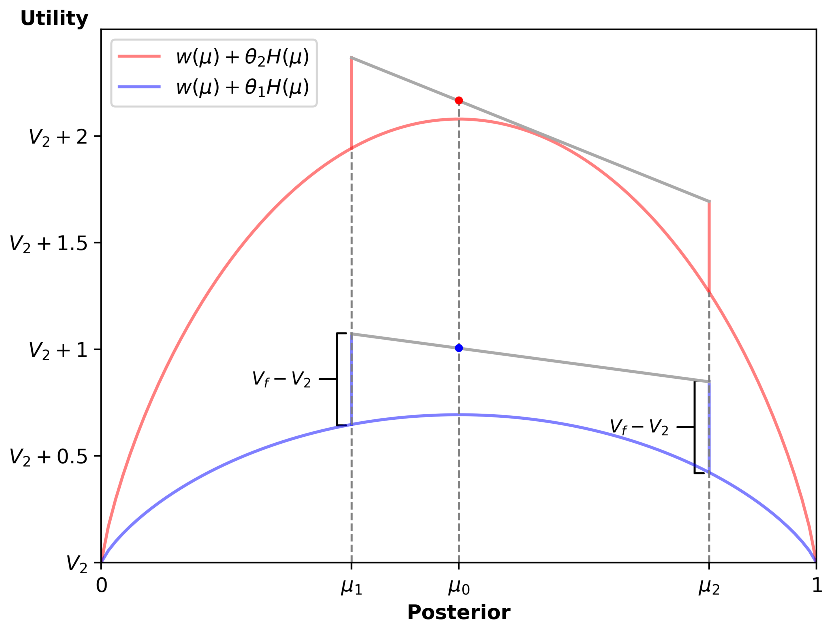

Figure 1: Pooling Equilibrium, $V_f = f(\theta_1)V_1 + f(\theta_2)V_2$

experiment with $p_2^*(\mu_0) = 1$, getting paid a wage V_2. Off-path beliefs are given by $m^*(\theta_2|\mu) = 1$ for all $\mu \notin \text{supp}(p_1^*)$, so that $w(\mu) = V_2$ for all off-path beliefs.

This demonstrates a general feature of fully-separating equilibria: the highest-cost researcher will choose an uninformative experiment, whose support contains only the prior μ_0. This is because in a fully-separating equilibrium the type θ_2 researcher must be paid V_2, while the firms can never offer a wage below V_2 to incentivize the researcher. In order to prevent the type θ_2 researcher from deviating to p_1^*, the posteriors μ_1 and μ_2 must be far enough away from the prior so that the gray line connecting the vertical red lines in Figure 2 passes below the red curve at the prior.

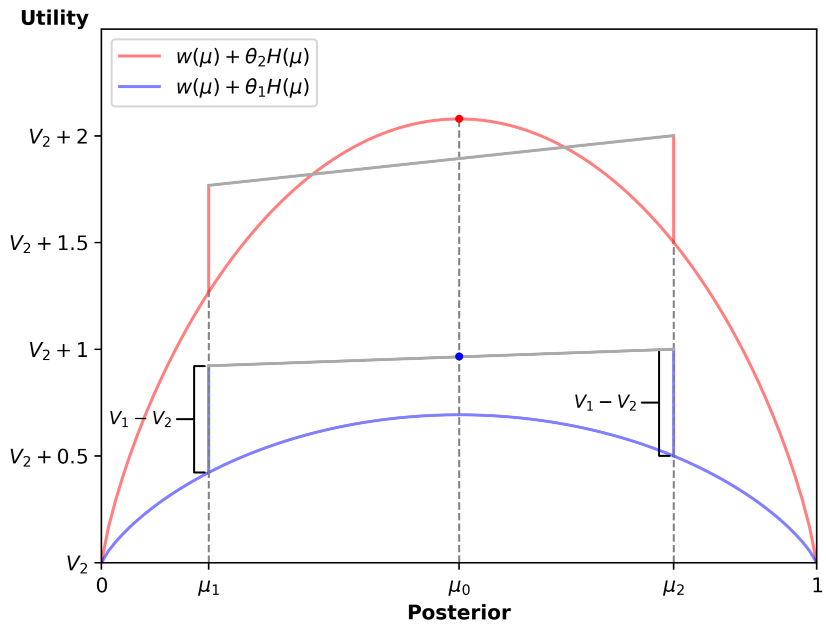

Figure 2: Fully Separating Equilibrium

4 Analysis

4.1 Researcher's Problem

We will first examine the researcher's problem to determine the set of compatible wage schedules. Given a wage schedule w, define the function

$$M_{i,w}(\mu) = w(\mu) + \theta_i H(\mu).$$

In equilibrium researcher i chooses an experiment such that

$$p_i \in \underset{p \in \mathcal{I}(\mu_0)}{\arg\max}\, \mathbb{E}_p[M_{i,w}(\mu)]. \tag{2}$$

The value of the maximization problem, when it exists, is the value of the concavification of $M_{i,w}(\mu)$ at the prior μ_0.

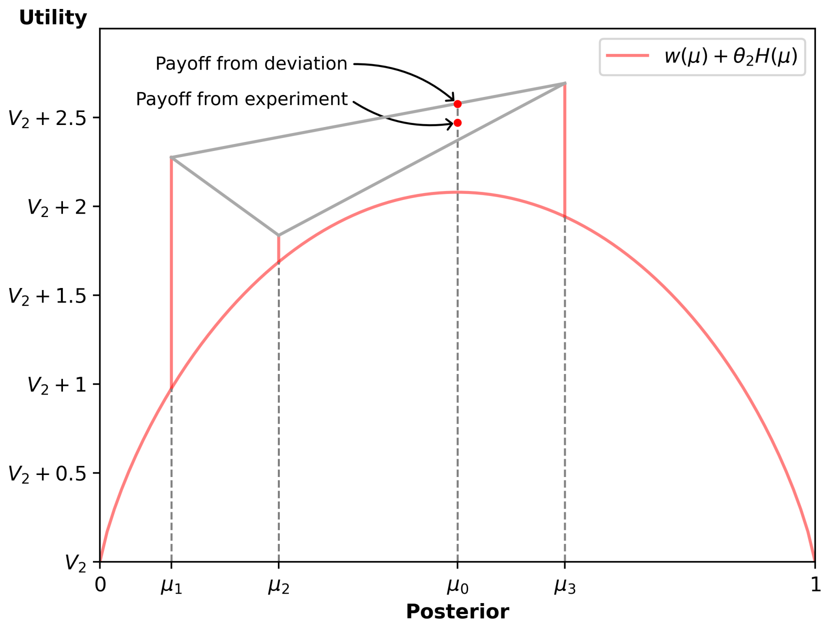

(a) $L_{i,w,p}$ does not lie in hyperplane

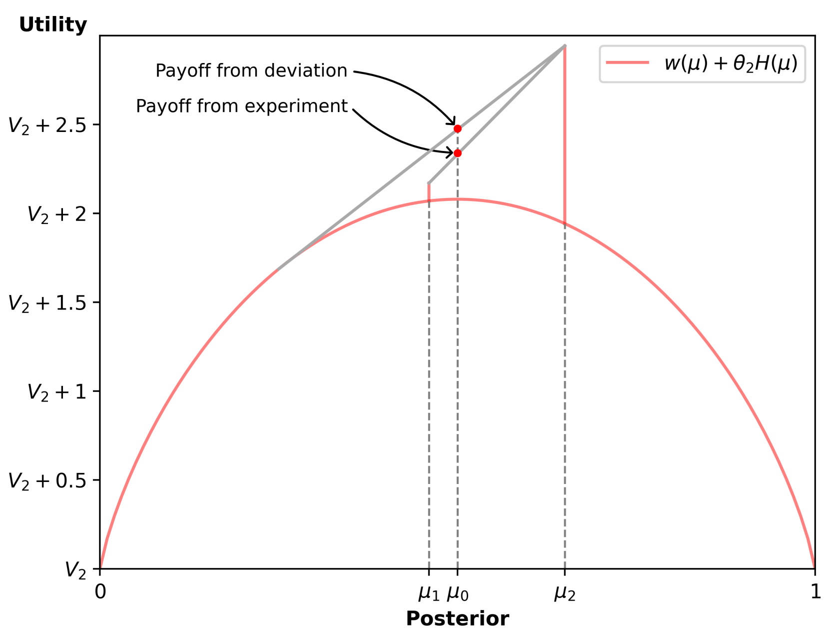

(b) Hyperplane spanned by $L_{i,w,p}$ does not lie above graph of $M_{i,w}$

Figure 3: Profitable deviations when the conditions in Proposition 1 are not met

For a given experiment p and wage schedule w, define

$$L_{i,w,p} = \{(\mu, M_{i,w}(\mu)|\mu \in \operatorname{supp}(p)\}.$$

Now consider dividing the set of feasible experiments $\mathcal{I}(\mu_0)$ based on whether or not their support is contained in that of p. The points of $L_{i,w,p}$ must lie in a hyperplane in order to prevent deviations to experiments with supports contained in that of p. In such deviations, the researcher could benefit by choosing a strict subset of the posteriors in the support of p that span the hyperplane with the highest value at the prior. This is illustrated in Figure 3a. Moreover, this hyperplane must lie weakly above $\operatorname{gr}(M_{i,w})$ in order to prevent deviations to experiments with supports not contained in that of p. In such deviations, the researcher will substitute a posterior μ outside the support of p such that $(\mu, M_{i,w}(\mu))$ is above the hyperplane. This is illustrated in Figure 3b. More generally, we can derive the following result:

Proposition 1. *A set of experiments $\{p_i\}$ is IC for a wage schedule w if and only if for all i, L_{i,w,p_i} lies in a hyperplane weakly above $gr(M_{i,w})$.*

Binary Case In the case of binary states and only two researchers, we can use Proposition 1 to bound the size of the supports of the experiments. In an equilibrium, the support of an experiment contains at most two non-overlapping posteriors and two overlapping posteriors.

Proposition 2. *Suppose that $|\Theta| = 2$ and $N = 2$. In any equilibrium*

$$|supp(p_1^*) \cap supp(p_2^*)| \leq 2$$

and for $i \in \{1, 2\}$,

$$|supp(p_i^*) \setminus supp(p_{-i}^*)| \leq 2.$$

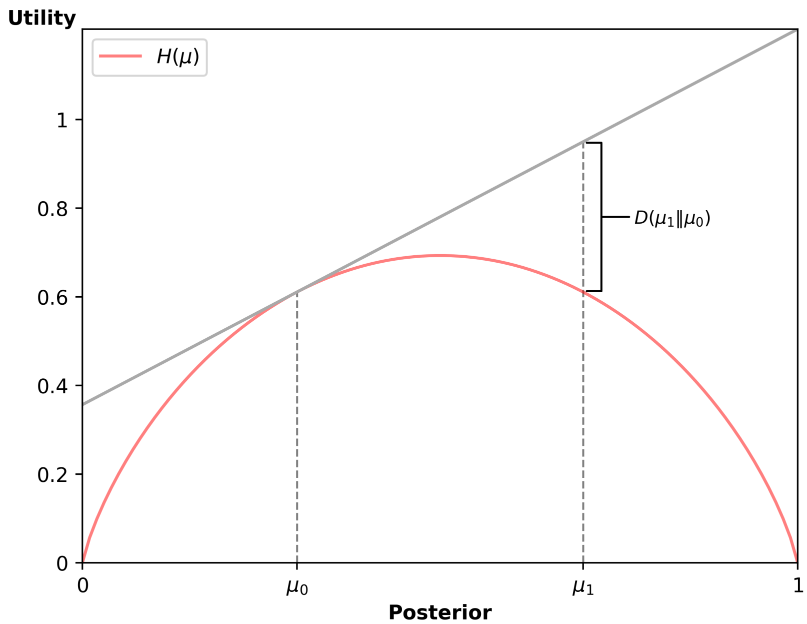

Figure 4: Divergence measure for entropy.

4.2 Divergence

Given a prior belief μ_0, define the divergence measure associated with H to be

$$D(\mu\|\mu_0) \equiv H(\mu_0) - H(\mu) + \nabla H(\mu_0) \cdot (\mu - \mu_0).$$

Geometrically, $D(\mu\|\mu_0)$ is the distance at μ between the graph of H, $\mathrm{gr}(H)$, and the hyperplane tangent to $\mathrm{gr}(H)$ at μ_0, as illustrated in Figure 4. We can rewrite the cost of an experiment as the expected cost of obtaining the posteriors in the support of the experiment,

$$C(p, \theta) = \theta\, \mathbb{E}_p[D(\mu\|\mu_0)].$$

Thus $D(\mu\|\mu_0)$ can be thought of as the cost of generating the posterior μ.

When each posterior in the support of the experiment has the same cost of being generated, we will call the experiment *isodivergent*.

Definition 1. An experiment $p \in \mathcal{I}(\mu_0)$ is *isodivergent* if there exists a constant

$a \in \mathbb{R}$ such that for all $\mu \in \operatorname{supp}(p)$, $D(\mu \| \mu_0) = a$.

We will show that it is non-isodivergence that limits the informativeness of equilibria with observable outcomes. A typical non-isodivergent experiment is one that obtains near certainty with low probability and is relatively uninformative otherwise. An example of this kind of experiment is a scientific researcher looking for a breakthrough.

4.3 Informativeness

We can judge the informativeness of an equilibrium $(\{p_i^*\}, m^*, w^*)$ based on the aggregate experiment defined by

$$p^*(\mu) = \sum_i f(\theta_i) p_i^*(\mu).$$

If the aggregate experiment of an equilibrium Blackwell dominates the aggregate experiment of another equilibrium, we will say that the first equilibrium is more informative. Define $\bar{\mathcal{E}}_{FS}^o \subset \mathcal{E}_{FS}^o$ to be the set of fully-separating equilibria that are maximally informative, i.e. their aggregate experiments are not strictly dominated by the aggregate experiment of any other equilibrium in $\mathcal{E}_{FS}^o$. Similarly define $\bar{\mathcal{E}}_P^o$ to be the set of maximally informative pooling equilibria.

The equilibrium shown in Figure 2 is not maximally informative. We could slightly lower μ_1 and slightly raise μ_2 in order to construct a more informative equilibrium. This process could be continued until the line connecting $(\mu_1, M_{1,w}(\mu_1))$ and $(\mu_2, M_{1,w}(\mu_2))$ is tangent to the curve $V_2 + \theta_1 H(\mu)$. To generalize this idea, define

$$M_i(\mu) = V_N + \theta_i H(\mu).$$

Proposition 3. *Suppose* $(\{p_i^*\}, m^*, w^*) \in \bar{\mathcal{E}}_{FS}^o$*. Then for all* i*,* L_{i,w^*,p_i^*} *lies in a hyperplane tangent to gr*(M_i) *at some* $\mu_i \in \Delta(\Omega)$*. The equilibrium payoff to researcher* i *is* $V_N + \theta_i D(\mu_0 \| \mu_i)$*.*

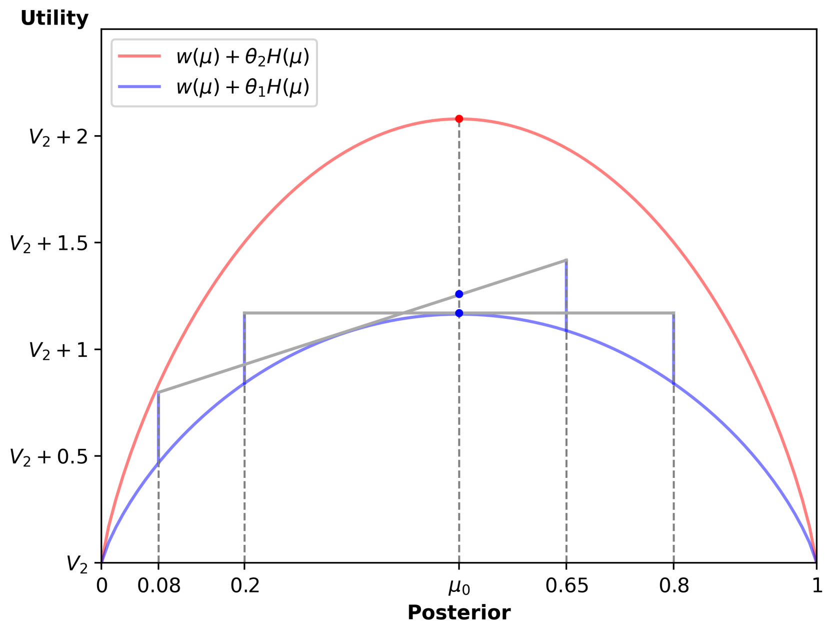

Figure 5: Two Equilibria in $\bar{\mathcal{E}}^o_{FS}$, one isodivergent and one not

Proposition 4. *Suppose $(\{p'_i\}, m', w') \in \bar{\mathcal{E}}^o_P$. Then L_{N,w',p'_N} lies in a hyperplane tangent to $gr(M_N)$ at some $\mu' \in \Delta\Omega$. The equilibrium payoff to researcher i is $V_N + \theta_i D(\mu_0 \| \mu')$.*

Our main result relates the informativeness of equilibria with observable outcomes to those with observable experiments. The result implies that lack of observability of the experiment can limit the informativeness of equilibrium research. Except in the non-generic case of isodivergent experiments, there always exists an equilibrium with observable experiments that is more informative.

Theorem 1. *Take any $(\{p^*_i\}, m^*, w^*) \in \bar{\mathcal{E}}^o_{FS} (\in \bar{\mathcal{E}}^o_P)$. There exists an equilibrium with observable experiments $(\{\hat{p}_i\}, \hat{m}, \hat{w}) \in \mathcal{E}^e_{FS} (\in \mathcal{E}^e_P)$ that is strictly more informative than $(\{p^*_i\}, m^*, w^*)$ if and only if at least one experiment in $\{p^*_i\}$ is not isodivergent.*

The idea for the proof is as follows. In a maximally informative equilibrium, the hyperplane containing the points L_{i,w,p_i} will be tangent to the curve $V_N + \theta_i H(\mu)$. If an experiment p^*_i is isodivergent, then this point of tangency will be at μ_0 and researcher i will have a net payoff of V_N. This is illustrated in Figure 5 for researcher 1 by the experiment with posteriors at 0.2 and 0.8. Even with observable experiments researcher i could never be incentivized to choose a more informative, and thus more costly, experiment, because she always has the choice of choosing an uninformative experiment and receiving at least V_N.

If an experiment p_i^* is not isodivergent, then the cost of p_i will be less than $V_i - V_N$, yielding a net payoff greater than V_N. This is illustrated in Figure 5 for researcher 1 by the experiment with posteriors at 0.08 and 0.65. We can then construct an otherwise identical equilibrium with observable experiments in which researcher i chooses some more informative experiment $\hat{p}_i$ with cost $V_i - V_N$.

5 Conclusion

This paper analyzes the signaling problem that arises in job markets for researchers. We compare the case where the chosen experiment is observable to the case where only the realized outcome of the experiment is observable. We show that this can limit the informativeness of the equilibrium research conducted in pooling and fully-separating equilibria: except in a special case of isodivergent experiments, there always exists an equilibrium with observable experiments that is more informative.

Our analysis has focused on the case of pooling and fully-separating equilibria. Future work should consider partially-separating equilibria, in which the wage offered to researchers would differ depending on the outcome of the experiment.

A Proofs

A.1 Proof of Proposition 1

Proof.

($\Longrightarrow$) Suppose that a given set of experiments $\{p_i\}$ solves (2) for a wage schedule w, and fix any i. Then p_i solves

$$\max_{p \in \mathcal{I}(\mu_0)} \mathbb{E}_p[w(\mu) + \theta_i H(\mu)].$$

If $w(\mu) + \theta_i H(\mu)$ is upper semi-continuous, then the result follows immediately from e.g. Kamenica (2017). However, can make no continuity assumptions on w. So define $w' : \Delta\Omega \to \mathbb{R}$ to be the upper envelope of w, i.e.

$$w'(\mu) = \limsup_{\mu_n \to \mu} w(\mu_n).$$

By construction, the function w' is upper semi-continuous.

First, we will show that p_i is also IC for w'. If not, then there exists some p_i' such that for some $\varepsilon > 0$,

$$\mathbb{E}_{p_i'}[w'(\mu) + \theta_i H(\mu)] - \varepsilon > \mathbb{E}_{p_i}[w'(\mu) + \theta_i H(\mu)].$$

Since the objective function is linear, it is without loss to assume that p_i' is an extreme experiment, and thus has finite support. So for any $\delta > 0$ we can take a set of sequences $\mu_n^j \to \mu^j$ for every $\mu^j \in \text{supp}(p_i')$ such that $\lim_{\mu_n^j \to \mu^j} w(\mu_n^j) > w'(\mu^j) - \delta$. Since μ_0 is in the interior of the convex hull of $\text{supp}(p_i')$, for large enough n we have that μ_0 will also be in the interior of the convex hull of $\{\mu_n^j\}_j$. Hence the set of experiments with a support of $\{\mu_n^j\}_j$ will be non-empty. Moreover, since p_i' is an extreme point, for sufficiently large n there will be a unique experiment p_n on $\{\mu_n^j\}_j$. The sequence $\{p_n\}_n$ converges in distribution to p_i'. So for sufficiently small δ we have that

$$\begin{aligned}\lim_{p_n \to p'_i} \mathbb{E}_{p_n}[w(\mu) + \theta_i H(\mu)] &> \mathbb{E}_{p'_i}[w'(\mu) + \theta_i H(\mu)] - \varepsilon \\ &> \mathbb{E}_{p_i}[w'(\mu) + \theta_i H(\mu)] \\ &> \mathbb{E}_{p_i}[w(\mu) + \theta_i H(\mu)].\end{aligned}$$

This contradicts the fact that p_i is optimal under w. Therefore, p_i is also optimal for w'.

Next, we will show that $w(\mu) = w'(\mu)$ at every $\mu \in \text{supp}(p_i)$. Suppose by way of contradiction that for some $\mu \in \text{supp}(p_i)$ we had that $w'(\mu) > w(\mu)$. Now take an extreme experiment p'_i such that $\mu \in \text{supp}(p'_i) \subset \text{supp}(p_i)$. Due to the linearity of the objective function, p'_i will be optimal for w'. Then there exists some $\varepsilon > 0$ such that

$$\mathbb{E}_{p'_i}[w'(\mu) + \theta_i H(\mu)] - \varepsilon > \mathbb{E}_{p'_i}[w(\mu) + \theta_i H(\mu)].$$

As before, for any $\delta > 0$ we can take a set of sequences $\mu^j_n \to \mu^j$ for every $\mu^j \in \text{supp}(p'_i)$ such that $\lim_{\mu^j_n \to \mu^j} w(\mu^j_n) > w'(\mu^j) - \delta$, and take $\{p_n\}_n$ with $\text{supp}(p_n) = \{\mu^j_n\}_j$. The sequence $\{p_n\}_n$ converges in distribution to p'_i. So for sufficiently small δ we have that

$$\lim_{p_n \to p'_i} \mathbb{E}_{p_n}[w(\mu) + \theta_i H(\mu)] > \mathbb{E}_{p'_i}[w'(\mu) + \theta_i H(\mu)] - \varepsilon.$$

This contradicts the fact that p_i is optimal under w. Therefore, $w(\mu) = w'(\mu)$ at every $\mu \in \text{supp}(p_i)$

In Kamenica and Gentzkow (2011) it is shown that $\max_{p \in \mathcal{I}(\mu_0)} \mathbb{E}_p[w'(\mu) + \theta_i H(\mu)]$ exists, and that the support of any solution spans an affine subspace supported by the surface of $\text{gr}(M_{i,w'})$. Hence in order for p_i to be optimal under w', the set L_{i,w',p_i} must span an affine subspace that is supported by $\text{gr}(M_{i,w'})$. And since $w'(\mu) = w(\mu)$ for all $\mu \in \text{supp}(p_i)$, we have that L_{i,w,p_i} must span an affine subspace supported by $\text{gr}(M_{i,w'})$. Define ℓ to be the associated linear function. It follows that for all $\mu \in \Delta\Omega$ we have that $\ell(\mu) \geq M_{i,w'}(\mu) \geq M_{i,w}(\mu)$. Therefore, L_{i,w,p_i} must lie in a hyperplane weakly above $\text{gr}(M_{i,w})$.

($\Longleftarrow$) Now suppose that for a given set of experiments $\{p_i\}$ and wage schedule w, L_{i,w,p_i} lies in a hyperplane weakly above $\text{gr}(M_{i,w})$ for all i. For a given i, the payoff to any experiment is bounded above by the concavification of $M_{i,w}(\mu)$ at the prior μ_0. Since L_{i,w,p_i} lies in a hyperplane weakly above $\text{gr}(M_{i,w})$, its value at μ_0 is above the concavification of $M_{i,w}(\mu)$ at μ_0, and hence above the value of any other experiment. Therefore $\{p_i\}$ is IC for w.

□

A.2 Proof of Proposition 2

Proof. Let $(\{p_i^*\}, m^*, w^*)$ be an equilibrium. Define $A = \text{supp}(p_1^*) \cap \text{supp}(p_2^*)$ and $B_i = \text{supp}(p_i^*) \setminus \text{supp}(p_{-i}^*)$.

First we will show that $|A| \leq 2$. Suppose that A were to have at least three distinct posteriors, $\mu_1 < \mu_2 < \mu_3$. Then for some $t \in (0,1)$ we can write $\mu_2 = t\mu_1 + (1-t)\mu_3$. Define $h_{ij} = (\mu_j, w^*(\mu_j) + \theta_i H(\mu_j))$. By Proposition 1, for $i \in \{1,2\}$ the points h_{i1}, h_{i2}, and h_{i3} must lie in a hyperplane. Thus h_{i2} can be written as a linear combination of h_{i1} and h_{i3}, so that $h_2 = th_1 + (1-t)h_3$. It follows that for $i \in \{1,2\}$,

$$w^*(\mu_2) + \theta_i H(\mu_2) = t(w^*(\mu_1) + \theta_i H(\mu_1)) + (1-t)(w^*(\mu_3) + \theta_i H(\mu_3)),$$

and so

$$tw(\mu_1) + (1-t)w(\mu_3) - w(\mu_2) = \theta_i H(\mu_2) - t\theta_i H(\mu_1) - (1-t)\theta_i H(\mu_3).$$

Taking the right hand side of this equation for $i \in \{1,2\}$, equating the expressions, and rearranging yields

$$(\theta_1 - \theta_2)H(\mu_2) = (\theta_1 - \theta_2)tH(\mu_1) + (\theta_1 - \theta_2)(1-t)H(\mu_3).$$

Since $\mu_2 = t\mu_1 + (1-t)\mu_3$, we now have that $H(t\mu_1 + (1-t)\mu_3) = tH(\mu_1) + (1-t)H(\mu_3)$. This violates the convexity of H. Therefore, $|A| \leq 2$.

Next we will show that for any given i, $|B_i| \leq 2$. Note that for all $\mu \in B_i$ we have that $w(\mu) = V_i$, since the realization of a posterior in B_i perfectly reveals the researcher's type. Suppose that B_i were to have at least three distinct posteriors, $\mu_1 < \mu_2 < \mu_3$. Then for some $t \in (0, 1)$ we can write $\mu_2 = t\mu_1 + (1-t)\mu_3$. Write $h_j = (\mu_j, V_i + \theta_i H(\mu_j))$. By Proposition 1, the points h_1, h_2, and h_3 must lie in a hyperplane. Thus h_2 can be written as a linear combination of h_1 and h_3, i.e. $h_2 = th_1 + (1-t)h_3$. This implies that

$$V_i + \theta_i H(\mu_2) = t(V_i + \theta_i H(\mu_1)) + (1-t)(V_i + \theta_i H(\mu_3)).$$

It follows that $H(t\mu_1 + (1-t)\mu_3) = tH(\mu_1) + (1-t)H(\mu_3)$, contradicting the strict convexity of H. Therefore, $|B_i| \leq 2$. □

A.3 Proof of Theorem 1

Proof. In a maximally informative equilibrium, the hyperplane containing the points L_{i,w,p_i} will be tangent to the curve $V_N + \theta_i H(\mu)$. If an experiment p_i^* is isodivergent, then this point of tangency will be at μ_0 and researcher i will have a net payoff of V_N. Even with observable experiments researcher i could never be incentivized to choose a more informative, and thus more costly, experiment, because she always has the choice of choosing an uninformative experiment and receiving at least V_N.

If an experiment p_i^* is not isodivergent, then the cost of p_i will be less than $V_i - V_N$, yielding a net payoff greater than V_N. We can then construct an otherwise identical equilibrium with observable experiments in which researcher i chooses some more informative experiment $\hat{p}_i$ with cost $V_i - V_N$.

□

References

Aumann, R. J., M. Maschler, and R. E. Stearns (1995). *Repeated games with incomplete information.* MIT press.

Blackwell, D. (1953). Equivalent comparisons of experiments. *The annals of mathematical statistics*, 265–272.

Caplin, A. and M. Dean (2013). Behavioral implications of rational inattention with shannon entropy. Technical report, National Bureau of Economic Research.

de Haan, T., T. Offerman, and R. Sloof (2011). Noisy signaling: theory and experiment. *Games and Economic Behavior 73*(2), 402–428.

Heinsalu, S. (2018). Dynamic noisy signaling. *American Economic Journal: Microeconomics 10*(2), 225–49.

Kamenica, E. (2017). Information economics. *Journal of Political Economy 125*(6), 1885–1890.

Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *American Economic Review 101*(6), 2590–2615.

Matějka, F. and A. McKay (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review 105*(1), 272–98.

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics 87*(3), 355–374.

Yoder, N. (2022). Designing incentives for heterogeneous researchers. *Journal of Political Economy 130*(8), 2018–2054.