

Information Validates the Prior: A Theorem on Bayesian Updating and Applications[†]

By NAVIN KARTIK, FRANCES XU LEE, AND WING SUEN*

We develop a result on expected posteriors for Bayesians with heterogeneous priors, dubbed information validates the prior (IVP). Under familiar ordering requirements, Anne expects a (Blackwell) more informative experiment to bring Bob's posterior mean closer to Anne's prior mean. We apply the result in two contexts of games of asymmetric information: voluntary testing or certification, and costly signaling or falsification. IVP can be used to determine how an agent's behavior responds to additional exogenous or endogenous information. We discuss economic implications. (JEL C11, D82, D84)

Bayesian agents revise their beliefs upon receiving new information. From an ex ante point of view, however, one cannot expect information to systematically alter one's beliefs in any particular direction. More precisely, a fundamental property of Bayesian updating is that beliefs are a martingale: an agent's expectation of his posterior belief is equal to his prior belief. But what about an agent's expectation of *another* agent's posterior belief when their current beliefs differ? Relatedly, should agents expect new information to systematically affect their disagreement, and if so, how? These questions are not only of intrinsic interest, but tackling them is useful for economics with asymmetric information.

Information Validates the Prior.—In Section I, we develop the following general result. Let $\Omega \subset \mathbb{R}$ be possible states. Bayesians Anne (A) and Bob (B) have mutually known but different priors over Ω , with means m_A and m_B . A signal s will be drawn from a known information structure or experiment \mathcal{E} . Let m_B^s denote Bob's posterior mean after observing the signal. Let $E_A^{\mathcal{E}}[m_B^s]$ be Anne's ex ante expectation of Bob's posterior mean under experiment \mathcal{E} .

*Kartik: Department of Economics, Columbia University (email: nkartik@columbia.edu); Lee: Quinlan School of Business, Loyola University Chicago (email: francesxu312@gmail.com); Suen: Faculty of Business and Economics, The University of Hong Kong (email: wsuen@econ.hku.hk). Larry Samuelson was coeditor for this article. A previous version of this paper was titled "A Theorem on Bayesian Updating and Applications to Communication Games." We thank Nageeb Ali, Ilan Guttman, Rick Harbaugh, Keri Hu, Alessandro Lizzeri, Steve Matthews, Paula Onuchic, Ariel Pakes, Andrea Prat, Mike Riordan, Satoru Takahashi, Jianrong Tian, Enrico Zanardo, Jidong Zhou, and various seminar and conference audiences for comments. We have benefited from thoughtful suggestions made by the editor and three referees. Sarba Raj Bartaula, Bruno Furtado, Dilip Ravindran, and Teck Yong Tan provided excellent research assistance. Kartik gratefully acknowledges financial support from the NSF (Grant SES-1459877).

[†]Go to <https://doi.org/10.1257/aeri.20200284> to visit the article page for additional materials and author disclosure statement(s).

Consider two experiments \mathcal{E} and $\tilde{\mathcal{E}}$, with \mathcal{E} more informative than $\tilde{\mathcal{E}}$ in the sense of Blackwell (1951, 1953). Theorem 1 establishes that under familiar ordering requirements,¹

$$(1) \quad m_A \leq (\geq) m_B \Rightarrow E_A^{\mathcal{E}}[m_B^s] \leq (\geq) E_A^{\tilde{\mathcal{E}}}[m_B^s].$$

In words, if Anne has a lower (respectively, higher) prior mean than Bob, then Anne predicts that a more informative experiment will, on average, reduce (respectively, raise) Bob's posterior mean by a larger amount than a less informative experiment. Put differently, Anne expects more information to further validate her prior in the sense of bringing Bob's posterior mean closer to her prior mean. Of course, Bob expects just the reverse. We refer to the result as “information validates the prior,” IVP hereafter.

IVP has an implication about expected disagreement. A particular signal can lead to larger posterior disagreement than prior disagreement. Nevertheless, IVP implies (Corollary 1) that when beliefs' disagreement is quantified by the difference in means, both Anne and Bob predict that a more informative experiment will, on average, reduce their posterior disagreement by more. At the extreme, both predict zero posterior disagreement under a fully informative experiment—even though each one's expectation of the other's posterior mean may be very different.

Applications.—IVP is a statistical result, which we believe is of intrinsic interest. Our paper shows why it is also instrumentally useful in games of asymmetric information. After all, even in common-prior environments, private information can endow Anne, an informed agent, with a different belief about a fundamental than Bob, an uninformed agent. In equilibrium, Anne anticipates how her actions will affect Bob's belief. Anne's strategic incentives may depend on how she expects new information to affect Bob, for which IVP is a useful tool. The new information can be exogenous or endogenous, for example, owing to the behavior of still other agents.

We develop these points in two contexts.

Voluntary Testing.—Section II studies voluntary testing or certification. An agent has some private information about his true ability or product quality. He can choose to undertake a costly test that provides an independent public signal of quality. Gross of the testing cost, the agent's payoff is the market's posterior expectation of his quality.

Proposition 1 establishes that under familiar informational assumptions, the more (Blackwell) informative the test, the less the agent will choose to get tested. The logic we elucidate, using IVP, is that because of pooling, the marginal type who takes the test expects more informative tests to reduce the benefit of getting tested.

Proposition 1 offers economic insights. When market information is of concern, a trade-off must be resolved between more informative tests and participation (cf. Harbaugh and Rasmusen 2018). But in settings where information only affects

¹The priors must be likelihood-ratio ordered, and experiments must satisfy the monotone likelihood-ratio property (MLRP). These assumptions are unrestrictive if the state is binary.

surplus division, ex ante efficiency improves with better tests because they reduce the deadweight loss from testing. We further discuss implications for a monopolist certifier's test choice, generalizing Lizzeri's (1999) observation about the emergence of uninformative tests.

Costly Signaling.—IVP is also useful in games of asymmetric information even when the information asymmetry is eliminated in equilibrium by an agent's own behavior. Section III deploys IVP in a sender-receiver game with lying costs (Kartik 2009) but explains how the logic also applies to canonical applications like education signaling (Spence 1973). Our concern is how exogenous information the receiver obtains affects the sender's signaling.

Proposition 2 establishes that, under reasonable conditions, better exogenous information reduces the sender's benefit from falsification to appear more favorable. Intuitively, the sender expects any favorable receiver belief he induces to get neutralized more by better exogenous information. Costly falsification becomes less attractive. Consequently, better exogenous information reduces wasteful signaling—every sender type is better off, even under full separation. We further discuss how this implies a strategic complementarity when there are multiple, possibly opposed, senders with bounded signal spaces.

Other Applications.—We believe IVP will be useful in other contexts too. Indeed, the logic of IVP unifies aspects of mechanisms in some existing papers that study models with heterogeneous priors under specific information structures. See, for example, Yildiz (2004), Che and Kartik (2009), Van den Steen (2010, proposition 5), Hirsch (2016, proposition 8), and Sethi and Yildiz (2012, proposition 5). We ourselves have used IVP to study information acquisition prior to disclosure (Kartik, Lee, and Suen 2017).

I. Information Validates the Prior

Fix any finite (multi-)set of states $\Omega \equiv \{\omega_1, \dots, \omega_L\} \subset \mathbb{R}$, with $\omega_1 \leq \dots \leq \omega_L$. We denote a generic element of Ω by either ω or ω_l . Fix a measurable space of signals, (S, \mathcal{S}) , endowed with a σ -finite reference measure. An *experiment* is $\mathcal{E} \equiv \{P_\omega\}_{\omega \in \Omega}$, where P_ω is a probability measure over signals in state ω . We only consider experiments for which each P_ω is absolutely continuous with respect to the reference measure, so that a Radon-Nikodym derivative exists, denoted $p(s|\omega)$; this is the probability density or mass function.²

Experiment \mathcal{E} is (Blackwell) *more informative* than experiment $\tilde{\mathcal{E}} \equiv \{\tilde{P}_\omega\}_{\omega \in \Omega}$, whose signal may be denoted $\tilde{s} \in S$ for clarity, if there is a Markov kernel $Q(\cdot|s)$ such that for each $\omega \in \Omega$ and every $\Sigma \in \mathcal{S}$,

$$\tilde{P}_\omega(\Sigma) = \int_S Q(\Sigma|s) dP_\omega(s).$$

²This permits the familiar definition of monotone likelihood ratio on the following page. Lehrer and Wang's (2020) "strong stochastic dominance" should be applicable more generally.

This definition captures the statistical notion that $\tilde{\mathcal{E}}$ does not provide any information beyond \mathcal{E} : state by state, the distribution of signals in $\tilde{\mathcal{E}}$ can be generated by taking signals from \mathcal{E} and transforming them through the state-independent kernel $Q(\cdot)$. As established by Blackwell (1953), this notion is equivalent to that of \mathcal{E} being more valuable than $\tilde{\mathcal{E}}$ for decision-making in canonical senses.

An experiment is an *MLRP experiment* if there is a total order on S , denoted \succeq (with asymmetric relation \succ), such that the MLRP holds:

$$s' \succ s \text{ and } \omega' > \omega \Rightarrow p(s'|\omega')p(s|\omega) \geq p(s'|\omega)p(s|\omega').$$

Let $\beta \in \Delta\Omega$ denote a belief, with $\beta(\omega)$ the probability ascribed to state ω . We say that a belief β' *likelihood-ratio (LR) dominates* belief β , written $\beta \leq_{LR} \beta'$, if for all $\omega' > \omega$,

$$\beta'(\omega')\beta(\omega) \geq \beta(\omega')\beta'(\omega).$$

A pair of beliefs are *likelihood-ratio ordered* if one likelihood-ratio dominates the other.

Anne (*A*) and Bob (*B*) are Bayesians with full-support priors on Ω , denoted by β_A and β_B respectively, with expectations or means m_A and m_B .³ Given an experiment \mathcal{E} , we denote individual i 's posterior mean by $m_i^{\mathcal{E}}$, computed by Bayes' rule. Let $E_i^{\mathcal{E}}[m_j^{\mathcal{E}}]$ denote i 's ex ante expectation of j 's posterior mean.

If the individuals have the same prior, then $E_A^{\mathcal{E}}[m_B^{\mathcal{E}}] = E_A^{\mathcal{E}}[m_A^{\mathcal{E}}] = m_A$ by iterated expectation. Under different priors, the same conclusion holds when \mathcal{E} is fully informative (i.e., every signal reveals the state) because prior differences become irrelevant. On the other hand, if \mathcal{E} is uninformative (i.e., no signal provides any information), then $E_A^{\mathcal{E}}[m_B^{\mathcal{E}}] = m_B$. For the intermediate cases between fully informative and uninformative experiments, there is certain monotonicity under some conditions.

THEOREM 1: *Let $\beta_A \leq_{LR} \beta_B$ and \mathcal{E} and $\tilde{\mathcal{E}}$ be MLRP experiments. If \mathcal{E} is more informative than $\tilde{\mathcal{E}}$, then*

$$m_A \leq E_A^{\mathcal{E}}[m_B^{\mathcal{E}}] \leq E_A^{\tilde{\mathcal{E}}}[m_B^{\tilde{\mathcal{E}}}] \leq m_B; \text{ and}$$

$$m_A \leq E_B^{\tilde{\mathcal{E}}}[m_A^{\tilde{\mathcal{E}}}] \leq E_B^{\mathcal{E}}[m_A^{\mathcal{E}}] \leq m_B.$$

Theorem 1 says that under its ordering requirements,⁴ each individual i predicts that a more informative experiment will, on average, bring the other's posterior mean closer to i 's prior. In this sense more information is expected to further validate one's prior, or more succinctly, *information validates the prior*. It bears emphasis that by relabeling states (cf. footnote 3), information validates the prior in the sense of not

³That is, $m_i \equiv \sum_{\omega} \omega \beta_i(\omega)$. We could just as well take m_i to be $\sum_{\omega} h(\omega) \beta_i(\omega)$ for any increasing $h: \Omega \rightarrow \mathbb{R}$. (By "increasing," we mean "weakly increasing," similarly for related terminology throughout this paper unless made explicit otherwise.) This amounts to relabeling each state ω as $h(\omega)$.

⁴The priors' likelihood-ratio ordering may be viewed as without loss—the states can be relabeled—so long as the experiments' MLRP and beliefs' means are understood with respect to the states' relabeling.

only the mean state but the expectation of any increasing function of the state, such as the probability of any $\{\omega_l, \dots, \omega_L\}$.

The proof of Theorem 1 is in Appendix A. For an illustration, suppose Bayesian updating takes the canonical linear form: the posterior mean is a convex combination of the prior mean and the signal, as is the case for any exponential family of signals with conjugate prior (e.g., normal-normal). It then holds that for any signal $s \in \mathbb{R}$ under experiment \mathcal{E} , $m_j^s = (1 - \alpha^\mathcal{E})m_j + \alpha^\mathcal{E}s$ for some $\alpha^\mathcal{E} \in [0, 1]$. Hence,

$$E_i^\mathcal{E}[m_j^s] = (1 - \alpha^\mathcal{E})m_j + \alpha^\mathcal{E}m_i.$$

The weight $\alpha^\mathcal{E}$ is larger when experiment \mathcal{E} is more informative, which implies Theorem 1's conclusions.

To shed light on the ordering requirements in Theorem 1, it is useful to decompose its conclusions. Each individual i expects the other's posterior mean to

- (i) (*Direction*) move toward i 's prior mean under any experiment; for example, $E_A^\mathcal{E}[m_B^s] \leq m_B$;
- (ii) (*Undershooting*) not move past i 's prior mean under any experiment; for example, $m_A \leq E_A^\mathcal{E}[m_B^s]$;
- (iii) (*Intensity*) move more when the experiment is more informative; for example, $E_A^\mathcal{E}[m_B^s] \leq E_A^{\mathcal{E}'}[m_B^s]$.

Note that the third point implies the previous two given the observations before the theorem about extreme experiments. The directional result (point i) does not require priors to be likelihood-ratio ordered; given the experiment's MLRP, the priors' ordering by first-order stochastic dominance (FOSD) is sufficient and is in fact essentially necessary.⁵ Online Appendix B.1 shows that even with likelihood-ratio ordered priors, the directional result can fail for a non-MLRP experiment. The undershooting result (point ii) does not require any assumption on the experiment; the priors' likelihood-ratio ordering is sufficient and is in fact essentially necessary.⁶ The intensity result (point iii) does not require any assumption on the less informative experiment; the priors' likelihood-ratio ordering and the more informative experiment's MLRP are sufficient. This is verified by studying the theorem's proof, which is fairly straightforward because the experiments are Blackwell ranked. We leave to future research the question of whether less restrictive comparisons of experiments would suffice.

Francetich and Kreps (2014) prove that conditional on any event being true, a Bayesian's expected posterior on that event (ignorant of the truth) is larger than her prior; see also Good (1965). It can be shown that for a binary state, that result

⁵Suppose $m_A \leq m_B$ with $\beta_A \not\leq_{FOSD} \beta_B$. Then for some index k , $\sum_{l \geq k} \beta_B(\omega_l) < \sum_{l \geq k} \beta_A(\omega_l)$. For an MLRP-experiment \mathcal{E} that only reveals whether $\omega < \omega_k$ or $\omega \geq \omega_k$, it can be verified that $E_A^\mathcal{E}[m_B^s] > m_B$ so long as $\omega_1 < \omega_L$.

⁶If $\beta_A \not\leq_{LR} \beta_B$, then so long as $\omega_1 < \dots < \omega_L$ there is an MLRP-experiment \mathcal{E} such that $E_A^\mathcal{E}[m_B^s] < m_A$. Specifically, let l be any index such that $\beta_B(\omega_{l+1})/\beta_B(\omega_l) < \beta_A(\omega_{l+1})/\beta_A(\omega_l)$, and consider \mathcal{E} that fully reveals every state except $\{\omega_l, \omega_{l+1}\}$, which are pooled together. Onuchic and Ray (2019, proposition 4) note a related point.

and Theorem 1 are equivalent. Note that any priors and experiments satisfy the theorem's ordering assumptions with a binary state. More generally though, neither result implies the other.

IVP has a noteworthy implication about expected disagreement. Consider measuring disagreement between beliefs by the distance in their means.⁷ Typically, experiments can generate signals for which posterior disagreement is larger than prior disagreement.⁸ But can individuals expect, *ex ante*, to disagree more after observing more information?

COROLLARY 1: *Let $\beta_A \leq_{LR} \beta_B$ and \mathcal{E} and $\tilde{\mathcal{E}}$ be MLRP experiments. If \mathcal{E} is more informative than $\tilde{\mathcal{E}}$, then for either i and j ,*

$$E_i^{\mathcal{E}}[|m_i^s - m_j^s|] \leq E_i^{\tilde{\mathcal{E}}}[|m_i^{\tilde{s}} - m_j^{\tilde{s}}|].$$

Corollary 1 says that—subject to the ordering hypotheses—more information reduces expected disagreement when disagreement is measured by the (absolute) difference in means. We omit a proof, as the corollary is equivalent to Theorem 1 because any signal from an MLRP experiment preserves the prior likelihood-ratio ordering of beliefs and thus also the ordering by their means; moreover, for any experiment \mathcal{E} and individual i , $E_i^{\mathcal{E}}[m_i^s] = m_i$.

Corollary 1 and IVP relate to the literature on merging of opinions initiated by Blackwell and Dubins (1962). That literature establishes asymptotic merging under general conditions; we are concerned instead with updating after just one observation.

In the following sections, we apply IVP to study games with a common prior but asymmetric information.

II. Voluntary Testing

Model.—An agent has unknown ability or product quality $q \in Q \subset \mathbb{R}$, where Q is a finite set. There is a full-support prior $\pi(q)$, whose expectation is denoted π^e . The agent privately receives information, referred to as his type, $t \in T \equiv [t, \bar{t}] \subset \mathbb{R}$ from the density $f(t|q)$ satisfying the strict MLRP. Without loss, we assume $E[q|t] = t$. The agent then chooses whether to take a test or not. If he takes the test, a result or signal $s \in S$ is drawn from density $g(s|q)$ with MLRP; s is conditionally independent of t . After observing whether the agent took the test—and if so, the test result—a decision-maker or market forms belief $\delta \in \Delta Q$ and hires the agent at a wage (or buys the product at a price) $w = \delta^e$, the market's posterior expectation of quality. Taking the test costs the agent $c \geq 0$, and so the agent's von Neumann-Morgenstern payoff is $w - c$ if he took the test and w if he did not.

Equilibrium.—Since the test is conditionally independent of the agent's information, we can think of the market as updating in two steps: it first forms an “interim

⁷The distance between individuals' expectations is an interesting measure of disagreement but obviously coarse and not without limitations. Zanardo (2017) studies disagreement axiomatically; he provides a result in the spirit of our Corollary 1 for a family of disagreement measures that include the Kullback-Leibler divergence.

⁸Yet no signal can polarize beliefs in the sense of FOSD (Baliga, Hanany, and Klibanoff 2013, theorem 1).

belief” about quality by updating its prior with what the agent’s action reveals about his private information t ; thereafter, if the agent took the test, it updates the interim belief using the test result s . Accordingly, denote the market’s interim belief when the agent takes the test (before observing the test result) as δ_+ and when the agent does not as δ_- . The agent’s private belief is denoted $\beta(t)$. Plainly, the agent of type t will take the test if and only if (modulo indifference)

$$E_{s|t}[\delta^e(s; \delta_+)] - c > \delta_-^e,$$

where $\delta^e(s; \delta_+)$ is the market’s posterior expectation upon observing test result s given interim belief δ_+ , and δ_-^e is the no-test expected quality. The left-hand side of the above inequality is strictly increasing in t : for any $t_H > t_L$,

$$E_{s|t_H}[\delta^e(s; \delta_+)] - E_{s|t_L}[\delta^e(s; \delta_+)] = \sum_q [\beta(q|t_H) - \beta(q|t_L)] E_{s|q}[\delta^e(s; \delta_+)]$$

is strictly positive because $E_{s|q}[\delta^e(s; \delta_+)]$ is increasing in q by the test’s MLRP, and $\beta(\cdot|t_H)$ strictly first-order stochastically dominates $\beta(\cdot|t_L)$ by strict MLRP of the agent’s private information.⁹

Therefore, all (weak perfect Bayesian) equilibria are described by a cutoff $t^* \in T$ such that the agent takes the test if $t > t^*$ and does not if $t < t^*$. Let $\delta_-(t^*)$ and $\delta_+(t^*)$ denote the interim beliefs computed from Bayes’ rule given any interior cutoff t^* . Consistent with the expected gain from testing strictly increasing in type, we restrict attention to equilibria in which if no type gets tested ($t^* = \bar{t}$), then the off-path interim belief upon testing is $\delta_+(\bar{t}) \equiv \beta(\bar{t})$; similarly, if all types get tested ($t^* = \underline{t}$), then the no-test off-path belief is $\delta_-(\underline{t}) \equiv \beta(\underline{t})$. Since intervals of signals preserve MLRP structure (Milgrom 1981, theorem 4), it holds that for any $t^* \in T$, $\delta_-(t^*)$ and $\delta_+(t^*)$ are likelihood-ratio ordered with all $\beta(t)$.

LEMMA 1: A cutoff t^* is an equilibrium cutoff if and only if

(i) t^* is interior and

$$(2) \quad E_{s|t^*}[\delta^e(s; \delta_+(t^*))] = \underbrace{E[t|t < t^*]}_{=\delta_-^e(t^*)} + c;$$

(ii) $t^* = \underline{t}$ and $E_{s|\underline{t}}[\delta^e(s; \pi)] \geq \underline{t} + c$; or

(iii) $t^* = \bar{t}$ and $\bar{t} \leq \pi^e + c$.

We omit the routine proof. The characterization implies that an equilibrium exists. There may be multiple equilibria. A key observation is that at any interior equilibrium cutoff t^* ,

$$\delta_-^e(t^*) = E[t|t < t^*] < t^* \leq E_{s|t^*}[\delta^e(s; \delta_+(t^*))] \leq \delta_+^e(t^*) = E[t|t > t^*],$$

⁹More precisely, our conclusion holds if $E_{s|q}[\delta^e(s; \delta_+)]$ is strictly increasing in q . Accounting for the possibility that this function may only be weakly increasing would not materially affect what follows. The possibility is ruled out when $g(s|q)$ has the strict MLRP and δ_+ is nondegenerate.

with at least one of the weak inequalities being strict. The weak inequalities follow from IVP, as Theorem 1’s ordering hypotheses are satisfied. Intuitively, as the cutoff type is pooled with all higher types, it expects the test to cause a downward revision from the market’s interim quality expectation. The cutoff type is nevertheless willing to take the test because foregoing it would lead to an even worse expectation, $\delta_-^e(t^*)$. It follows that $t^* - \delta_-^e(t^*) \leq c$; equality holds only with a fully informative test. With such a test, the setting is effectively one of costly voluntary disclosure (Jovanovic 1982, Verrecchia 1983). If the test is uninformative, then $\delta_+^e(t^*) - \delta_-^e(t^*) = c$; it is as if the agent can simply take a pure money-burning action.

Comparative Statics.—What happens when the test becomes (Blackwell) more informative? IVP assures that the cutoff type expects the “disagreement” between its private belief and the market’s interim belief δ_+ to shrink by more. The cutoff type’s expected benefit of taking the test is thus lower. On the other hand, the payoff from not taking the test, δ_-^e , is unaffected. Consequently, fewer types should take the test. Formally, we have the following result.

PROPOSITION 1: *If the test becomes more informative, then the smallest and largest equilibrium cutoffs increase.*

PROOF:

Consider two test distributions, \underline{g} and \bar{g} , with the latter more informative. Define $L(t; \underline{g}) \equiv E_{s|t}^{\underline{g}}[\delta^e(s; \delta_+(t))]$ and $R(t) \equiv \delta_-^e(t) + c$; these correspond to the two sides of equation (2). IVP (Theorem 1) implies that $L(t; \underline{g}) \geq L(t; \bar{g})$ for all t . We argue below that the largest equilibrium cutoff increases; an analogous argument applies to the smallest.

The conclusion is trivial if $\bar{t} \leq \pi^e + c$, as Lemma 1 implies the largest equilibrium cutoff is \bar{t} regardless of the test. So suppose $\bar{t} > \pi^e + c$, or equivalently, $L(\bar{t}; \underline{g}) > R(\bar{t})$ for any \underline{g} . If $L(t; \underline{g}) > R(t)$ for all $t > \underline{t}$, the only equilibrium cutoff under \underline{g} is \underline{t} , and the conclusion follows. So suppose $L(t; \underline{g}) = R(t)$ for some interior t . Continuity and $L(\cdot; \underline{g}) \geq L(\cdot; \bar{g})$ imply that the largest intersection of $L(\cdot; \bar{g})$ and $R(\cdot)$ is at least as large as that of $L(\cdot; \underline{g})$ and $R(\cdot)$. ■

Equilibria with extremal cutoffs are stable in the sense of best-response dynamics. Proposition 1’s comparative static extends to other stable equilibria, but it can reverse for unstable equilibria, as is a common theme in games with multiple equilibria.

Ex ante, the agent prefers a larger equilibrium cutoff because his ex ante expected wage is π^e in any equilibrium, and a larger cutoff reduces the likelihood of having to take the test. Focusing on the largest equilibrium is thus justified on ex ante utilitarian grounds if the agent’s wage is purely redistributive and information has no efficiency benefit. Indeed, Proposition 1 implies that from this perspective more informative tests and their largest equilibrium cutoffs are socially preferable—only indirectly, because they lower the deadweight loss from testing.

Proposition 1’s lesson that more informative tests reduce testing participation echoes and broadens a point made by Harbaugh and Rasmusen (2018) that fully

informative tests do not maximize participation. Unlike them, we do not require the agent to be perfectly informed about quality, and we compare any pair of Blackwell-ranked tests. Their focus was on solving for the optimal test to minimize the market's mean squared error.

Certification by an Intermediary.—View the agent's cost c as the price charged by a monopolist testing firm or certifying intermediary who must offer a single test. Momentarily take c as given. The monopolist then seeks to maximize the fraction of agents who take the test, that is, to minimize the equilibrium cutoff. It follows from Proposition 1 that (focusing on an extremal equilibrium) the monopolist will choose the least informative test available if such a test exists—even if all available tests are costless to perform, and a fortiori, if more informative tests are more costly. In particular, if the uninformative test is available (and least costly), the monopolist will choose that. Consequently, when the monopolist can choose any test and any price c , and we select the profit-maximizing equilibrium cutoff, it is optimal to choose an uninformative test and price $c = \pi^e - \underline{t}$. All types then get tested, and the monopolist makes the maximum possible profit subject to the agent's minimum payoff of \underline{t} .¹⁰

This discussion generalizes themes from Lizzeri (1999), who assumed a perfectly informed agent, endogenous pricing, and availability of all tests. Our analysis clarifies that the economic force favoring less informative tests does not turn on any of these conditions—rather, even with a partially informed agent, demand in a profit-maximizing equilibrium at any price is higher when the monopolist chooses a less informative test.

III. Costly Signaling

Our second application shows how IVP is useful in games with asymmetric information even when there may be no “disagreement” in equilibrium because of (full) separation; rather, what is crucial is how information affects disagreement *off the equilibrium path*.

A. Model

We consider communication with lying costs or costly falsification, following Kartik, Ottaviani, and Squintani (2007) and Kartik (2009) but adding exogenous information. A sender and receiver share a common nondegenerate prior about a state $\omega \in \{0, 1\}$. The sender privately learns his type $t \in [0, 1]$, normalized to equal his private belief that $\omega = 1$, drawn from a density $f(t|\omega)$. Our normalization implies that $f(\cdot)$ has the MLRP. The sender sends a report $r \geq 0$ at a cost $c(r, t)$, elaborated below. The receiver forms a belief based on both r and an additional signal $s \in \mathbb{R}$ that, conditional on the state ω , is drawn independently of t or r from a density $g(s|\omega) > 0$. Without loss, we assume $g(\cdot)$ satisfies the MLRP. The signal s can either be the receiver's private information or publicly observed after the sender

¹⁰If the type distribution under the prior has a strictly decreasing density, then $\delta_+^e(t) - \delta_-^e(t)$ is strictly increasing and hence \underline{t} is the unique cutoff equilibrium when the test is uninformative and $c = \pi^e - \underline{t}$.

has acted. Denote the receiver's posterior expectation of the state by $E[\omega|r, s]$. The sender's von Neumann-Morgenstern payoff is linear in this expectation; specifically, his payoff is

$$E[\omega|r, s] - c(r, t).$$

The receiver's belief updating can be analyzed in two steps: based on the sender's report r , she forms an interim belief $\pi(r) \in [0, 1]$ about the state (all beliefs refer to the probability of $\omega = 1$) and then uses this interim belief to further update from the signal s to form a posterior belief about the state $\beta(s; \pi(r))$. By Bayes' rule,

$$(3) \quad \beta(s; \pi) = \frac{\pi g(s|1)}{\pi g(s|1) + (1 - \pi)g(s|0)}.$$

The expected payoff of a type- t sender from report r is thus

$$E_{s|t}[\beta(s; \pi(r))] - c(r, t),$$

where $E_{s|t}$ denotes t 's expectation over the signal s . The first term in the above display is strictly increasing in $\pi(r)$ because a higher interim belief raises the receiver's posterior after any signal.

Assume the cost function $c(r, t)$ is smooth with $\partial c(r, t)/\partial r = 0$ for all t , that is, the marginal cost of lying when telling the truth is zero. Furthermore, $\partial^2 c(r, t)/\partial r^2 > 0 > \partial^2 c(r, t)/\partial r \partial t$ for all t, r , that is, the marginal cost of sending a higher report is increasing in the report and decreasing in the type.

A Single-Crossing Condition.—Due to the receiver's exogenous information, the above cost assumptions do not guarantee a suitable single-crossing property. Consider the sender's indifference curves in the space of his report r and the receiver's interim belief π . For any type t , these indifference curves slope upward for $r > t$: costly lying requires a higher interim-belief compensation. The requisite single-crossing property is that—no matter the exogenous experiment—these indifference curves are flatter for higher types, meaning that higher types are more willing to inflate their report to induce a higher interim belief. Lemma A.1 in Appendix A establishes that this property is assured by the following assumption that we will maintain:

$$(4) \quad \text{for } t < 1 \text{ and } t < r: \quad \frac{\partial^2 c(r, t)/\partial r \partial t}{\partial c(r, t)/\partial r} \leq -\frac{1}{1 - t}.$$

Condition (4) can be interpreted as saying that higher types have a sufficiently large marginal cost advantage relative to marginal cost. All our cost assumptions are satisfied by, for example, $c(r, t) = (r - t)^2$.

Another Interpretation.—Although our model is posed as a communication game, it can also be viewed as adding exogenous information to a variation of the

Spence (1973) signaling model. In this interpretation, a worker possesses a private trait, his type t (e.g., intelligence), which is indicative of binary job productivity. An employer also observes a signal about that productivity (e.g., through an interview). The worker's wage depends on the employer's posterior on productivity given the worker's schooling level r and the employer's signal s . The marginal cost of schooling decreases in the characteristic t . While we assume that higher-type workers intrinsically prefer acquiring more education, our analysis also applies if all workers prefer less education; see Section IIID.

B. Equilibrium

We focus on the least-cost (fully) separating equilibrium, or LCSE, as is standard. In such an equilibrium the sender's pure strategy $\rho : [0, 1] \rightarrow \mathbb{R}_+$ is strictly increasing, with the "Riley condition" $\rho(0) = 0$. For $r \in [0, \rho(1)]$, the receiver's interim belief upon observing r is $\pi(r) = \rho^{-1}(r)$. Without loss, we stipulate that for $r > \rho(1)$, $\pi(r) = 1$. Standard arguments imply that the LCSE strategy $\rho(\cdot)$ must satisfy the differential equation

$$(5) \quad \frac{\partial c(\rho(t), t)}{\partial r} \rho'(t) = \frac{\partial E_{s|t}[\beta(s; t)]}{\partial \pi}.$$

Equation (5) obtains from the binding local upward incentive compatibility constraints. The left-hand side is type t 's marginal cost of mimicking a slightly higher type; the right-hand side is the marginal benefit, which comes from inducing a higher receiver interim belief. This benefit is affected by the sender's belief about the exogenous signal s . Since $\rho'(t) > 0$ and equation (5)'s right-hand side is strictly positive, any solution has $\rho(t) > t$ for $t > 0$.

Condition (4), which yields the requisite single-crossing property, guarantees that equation (5) and the boundary condition $\rho(0) = 0$ are not only necessary in an LCSE but also sufficient (i.e., global incentive compatibility is assured). As standard arguments imply that this boundary-value problem has a unique solution, we state the following result without proof.

LEMMA 2: *There is a unique LCSE.*

C. Comparative Statics

How does (Blackwell) more informative exogenous information affect the sender's LCSE signaling strategy? The strategy is determined by the local upward incentive constraints. Intuitively when type t mimics a slightly higher type $t + \varepsilon$, it creates disagreement: type t views the receiver's interim belief $t + \varepsilon$ as higher than the truth. IVP implies that t expects a more informative exogenous signal to correct that interim belief by more, and so t 's gain from inducing that interim belief is lower under a more informative experiment. Formally, we have the following result.

LEMMA 3: *If \tilde{s} is drawn from a more informative experiment than s , then for any $t < 1$,*

$$\frac{\partial E_{\tilde{s}|t}[\beta(\tilde{s};t)]}{\partial \pi} \leq \frac{\partial E_{s|t}[\beta(s;t)]}{\partial \pi}.$$

PROOF:

IVP (Theorem 1) implies that for any $t < 1$ and small $\varepsilon > 0$,

$$E_{\tilde{s}|t}[\beta(\tilde{s};t + \varepsilon)] - E_{\tilde{s}|t}[\beta(\tilde{s};t)] \leq E_{s|t}[\beta(s;t + \varepsilon)] - E_{s|t}[\beta(s;t)],$$

because $E_{\tilde{s}|t}[\beta(\tilde{s};t)] = E_{s|t}[\beta(s;t)] = t$. (Theorem 1's ordering hypotheses hold because both the sender's and receiver's information have the MLRP.) The result follows from dividing both sides of the inequality above by ε and taking $\varepsilon \rightarrow 0$. ■

Lemma 3 implies that more informative exogenous information reduces the right-hand side of equation (5). Since each type expects a smaller marginal benefit from inducing a higher belief in the receiver, the solution ρ to the differential equation (5) with boundary condition $\rho(0) = 0$ is pointwise lower. It is therefore intuitive, and proved in Appendix A, that the following result holds.

PROPOSITION 2: *In the LCSE, every sender type bears a lower signaling cost when the receiver's exogenous information is more informative.*

Consequently, every sender type is better off when the receiver's information improves, because in the LCSE every type t expects the receiver's posterior to be t regardless of the exogenous information distribution. Plainly, the receiver is also better off (given any von Neumann-Morgenstern payoff function) with better exogenous information.

D. Discussion and Implications

A very similar analysis, with the same conclusion as Proposition 2, would apply if we had instead assumed à la Spence (1973) that all types would choose $r = 0$ under complete information. Specifically, we could have instead assumed the signaling cost function $c(r,t)$ satisfies $\partial c / \partial r > 0$, $\partial^2 c / \partial r^2 > 0$, and $\partial^2 c / \partial r \partial t < 0$ and required the inequality in (4) to hold for all $t < 1$ and $r > 0$.

One reason we study costly lying is that the analysis also extends to a bounded report space, which is natural there. Suppose the sender's report must be in $[0, 1]$, his type space. Then, given our original assumptions on $c(\cdot)$, there is no fully separating equilibrium; in particular, recall that for a strictly increasing strategy ρ , a solution to equation (5) entails $\rho(t) > t$ for all $t > 0$. In a previous version of this paper (Kartik, Lee, and Suen 2019), we analyzed a salient equilibrium that extends the LCSE to this case: there is separation up to some cutoff type $t^* < 1$ and pooling on $r = 1$ thereafter. The sender's strategy in the separating region is unchanged: $\rho(0) = 0$, and equation (5) holds. The cutoff t^* is determined by its indifference. More informative exogenous information now not only lowers every type's signaling

cost but also raises the cutoff t^* (again due to IVP, similar to the logic that reduces the benefit of pooling in Section II). The larger separating region means the sender's signaling is more informative.

One can adapt the analysis to study multi-sender signaling when each sender gets a conditionally independent signal. Suppose some senders are upward biased, others are downward biased, and each chooses a report in $[0, 1]$.¹¹ This model applies to competing persuaders with falsification costs in various domains: lobbyists, media, legal parties, etc. From each sender's perspective, the other senders' reports are endogenous experiments. The logic underlying Proposition 2 implies that senders' strategies are *strategic complements*: in equilibria where each sender uses a cutoff strategy as described above, each sender reveals more and bears a lower signaling cost when other senders reveal more. The receiver thus learns more with more senders; importantly, beyond the obvious direct benefit of adding a sender, there is an indirect benefit of existing senders revealing more. Each sender's ex ante welfare is also higher, as his equilibrium signaling cost decreases while the receiver's average posterior is unaffected. Furthermore, if any sender's cost increases in a suitable sense (for example, if sender i 's cost is $k_i c(r, t)$ and $k_i > 0$ increases), then not only does that sender reveal more but so do all other senders.

Although it may seem unsurprising that better receiver information reduces a sender's incentive to incur falsification costs, we emphasize that Proposition 2 relies on IVP. Online Appendix B.2 shows that if the sender's payoff is not linear in the receiver's posterior, or if the receiver's information violates MLRP (in a multi-state extension of the model), then better exogenous information can *increase* the sender's marginal benefit from mimicking a higher type, leading to higher equilibrium signaling costs.

Frank (1985, section III) also suggests that better exogenous information can reduce dissipative signaling. Weiss (1983) studies when exogenous information allows for separating equilibria even absent any heterogeneity in the direct costs of signaling; his focus is not on comparative statics. Daley and Green (2014) emphasize the stability of nonseparating equilibria when there is "double crossing" of appropriate indifference curves, contrary to the single crossing assured by our condition (4). Truyts (2015) shows that better exogenous information can exacerbate dissipative signaling when signaling is noisy.

IV. Conclusion

IVP is instructive more broadly than for just the two applications with asymmetric information developed in this paper. In an earlier version (Kartik, Lee, and Suen 2019), we used IVP to study voluntary disclosure with either concealment or disclosure costs.

We close by briefly commenting on a domain with symmetric information. Consider Bayesian persuasion (Kamenica and Gentzkow 2011). Alonso and Câmara (2016) develop a general analysis under heterogeneous priors; see also Onuchic & Ray (2019). When the sender's preferences are state independent and

¹¹ A downward-biased sender's payoff is $-E[\omega|r, s] - c(r, t)$. The relevant strategy for this sender is given by $\rho(1) = 1$, the analog of equation (5) with the right-hand side's sign flipped, and all types below some cutoff pooling on report $r = 0$. For any $t \in (0, 1)$, $\rho(t) < t$, that is, a downward-biased sender deflates his report.

concave in the receiver’s posterior expectation, there is no scope for beneficial persuasion under common priors. Alonso and Câmara (2016, section 4.3) show that this observation does not hold generically with heterogeneous priors and at least three states. Our Theorem 1 delivers additional insights. For example, if the priors are likelihood-ratio ordered, then among MLRP experiments the sender prefers less informative experiments when facing a favorable receiver, that is, one whose prior dominates (respectively, is dominated by) the sender’s if the sender’s utility is increasing (respectively, decreasing) in the receiver’s posterior expectation. Consequently, in that scenario, an uninformative experiment is optimal if all and only MLRP experiments are available.¹²

APPENDIX A. PROOFS

PROOF OF THEOREM 1:

We prove the statement about A ’s expectations; B ’s are analogous. It is sufficient to establish $E_A^{\mathcal{E}}[m_B^s] \leq E_A^{\tilde{\mathcal{E}}}[m_B^{\tilde{s}}]$, as the other two inequalities are implied by the observations in the paragraph preceding Theorem 1.

Since $\beta_A \leq_{LR} \beta_B$, it follows that for any \tilde{s} , $\beta_A^{\tilde{s}} \leq_{LR} \beta_B^{\tilde{s}}$ and hence $\beta_A^{\tilde{s}} \leq_{FOSD} \beta_B^{\tilde{s}}$. (Here, $\beta_i^{\tilde{s}}$ denotes i ’s posterior after observing signal \tilde{s} .) Experiment \mathcal{E} ’s MLRP implies that its distribution of signals is increasing in the state in the sense of FOSD. Consequently, conditional on any \tilde{s} , the probability distribution ascribed by B to signals s dominates in FOSD ascribed by A . To prove that, let $p(s|\omega)$ denote the signal density of experiment \mathcal{E} and $\tilde{p}(\tilde{s}|\omega)$ that of the garbling $\tilde{\mathcal{E}}$. For notational simplicity, we proceed assuming there is a nonnegative kernel $q(\tilde{s}|s)$ with $\int_{\tilde{s}} q(\tilde{s}|s) d\tilde{s} = 1$ for all s , such that for any state ω ,

$$\tilde{p}_i(\tilde{s}) = \int_s q(\tilde{s}|s)p(s|\omega)ds.$$

For $i = A, B$, let $p_i(s) \equiv \sum_{\omega} p(s|\omega)\beta_i(\omega)$ and $\tilde{p}_i(\tilde{s}) \equiv \sum_{\omega} \tilde{p}(\tilde{s}|\omega)\beta_i(\omega)$. The conditional density of any signal s from \mathcal{E} given any signal \tilde{s} from $\tilde{\mathcal{E}}$, using prior β_i , is

$$\hat{q}_i(s|\tilde{s}) = \frac{q(\tilde{s}|s)p_i(s)}{\tilde{p}_i(\tilde{s})}.$$

Therefore,

$$\frac{\hat{q}_B(s|\tilde{s})}{\hat{q}_A(s|\tilde{s})} = \frac{\tilde{p}_A(\tilde{s}) p_B(s)}{\tilde{p}_B(\tilde{s}) p_A(s)}.$$

¹²Following our discussion after Theorem 1, priors ordered by FOSD are sufficient for optimality of an uninformative experiment (Onuchic © Ray 2019, also note this point restricting attention to a subset of MLRP experiments) but not for the broader preference ranking of MLRP experiments by their informativeness.

As \mathcal{E} satisfies MLRP and $\beta_A \leq_{LR} \beta_B$, the ratio $p_B(s)/p_A(s)$ is increasing.¹³ Therefore, for any \tilde{s} , $\hat{q}_A(\cdot|\tilde{s}) \leq_{LR} \hat{q}_B(\cdot|\tilde{s})$, which implies $\hat{q}_A(\cdot|\tilde{s}) \leq_{FOSD} \hat{q}_B(\cdot|\tilde{s})$.

As the mean m_B^s is increasing in s (by \mathcal{E} 's MLRP), it follows that for any \tilde{s} ,

$$(A1) \quad E_A^{\mathcal{E}}[m_B^s|\tilde{s}] \leq E_B^{\mathcal{E}}[m_B^s|\tilde{s}].$$

Experiment \mathcal{E} being more informative than $\tilde{\mathcal{E}}$ implies that for any s and \tilde{s} , $m_B^s = m_B^{s,\tilde{s}}$, where $m_B^{s,\tilde{s}}$ denotes B 's posterior mean having observed (s, \tilde{s}) . As $E_B^{\mathcal{E}}[m_B^{s,\tilde{s}}|\tilde{s}] = m_B^{\tilde{s}}$ by iterated expectation, it follows from (A1) that for any \tilde{s} ,

$$(A2) \quad E_A^{\mathcal{E}}[m_B^s|\tilde{s}] \leq m_B^{\tilde{s}}.$$

Taking the expectation in (A2) over \tilde{s} using the prior β_A yields

$$E_A^{\mathcal{E}}[m_B^s] = E_A^{\tilde{\mathcal{E}}}[E_A^{\mathcal{E}}[m_B^s|\tilde{s}]] \leq E_A^{\tilde{\mathcal{E}}}[m_B^{\tilde{s}}],$$

where the equality is by iterated expectation. ■

LEMMA A.1: For $t < 1$ and $t < r$, $\frac{\partial c(r,t)/\partial r}{\partial E_{s|t}[\beta(s;\pi)]/\partial \pi}$ strictly decreases in t .

PROOF:

Differentiating (3) and manipulating, the sender's marginal rate of substitution between r and π is given by

$$\frac{\partial c(r,t)/\partial r}{\partial E_{s|t}[\beta(s;\pi)]/\partial \pi} = \pi(1-\pi) \frac{\partial c(r,t)/\partial r}{E_{s|t}[\beta(s;\pi)(1-\beta(s;\pi))]}.$$

Differentiating with respect to t , noting that $E_{s|t}[\cdot] = t[\cdot] + (1-t)[\cdot]$, and rearranging, we see that for $t < m$, the marginal rate of substitution strictly decreases in t if

$$\frac{\partial^2 c(r,t)/\partial r \partial t}{\partial c(r,t)/\partial r} < \frac{E_{s|1}[\beta(s;\pi)(1-\beta(s;\pi))] - E_{s|0}[\beta(s;\pi)(1-\beta(s;\pi))]}{tE_{s|1}[\beta(s;\pi)(1-\beta(s;\pi))] + (1-t)E_{s|0}[\beta(s;\pi)(1-\beta(s;\pi))]}.$$

Since both $E_{s|1}[\beta(s;\pi)(1-\beta(s;\pi))]$ and $E_{s|0}[\beta(s;\pi)(1-\beta(s;\pi))]$ are strictly positive, for $t < 1$ the right-hand side of the above inequality is strictly greater than

$$\frac{0 - E_{s|0}[\beta(s;\pi)(1-\beta(s;\pi))]}{0 + (1-t)E_{s|0}[\beta(s;\pi)(1-\beta(s;\pi))]} = -\frac{1}{1-t}.$$

¹³To confirm this, observe that

$$\frac{p_B(s)}{p_A(s)} = \frac{\sum_{\omega} \beta_B(\omega)p(s|\omega)}{\sum_{\omega} \beta_A(\omega)p(s|\omega)} = \sum_{\omega} \frac{\beta_A(\omega)p(s|\omega)}{\sum_{\omega'} \beta_A(\omega')p(s|\omega')} \frac{\beta_B(\omega)}{\beta_A(\omega)} = \sum_{\omega} \beta_A(\omega|s) \frac{\beta_B(\omega)}{\beta_A(\omega)}.$$

Fix any $s' > s$. The experiment's MLRP implies $\beta_A(\cdot|s) \leq_{FOSD} \beta_A(\cdot|s')$. The LR ordering of priors implies $\beta_B(\omega)/\beta_A(\omega)$ is increasing. Therefore, $p_B(s')/p_A(s') \geq p_B(s)/p_A(s)$.

So if (4) holds, the marginal rate of substitution strictly decreases in t for $t < 1$ and $t < r$. ■

PROOF OF PROPOSITION 2:

If we show that $\rho(t)$ decreases pointwise when the right-hand side of equation (5) decreases for all t , then the result follows from Lemma 3. Accordingly, let $\tilde{\rho}(t)$ and $\rho(t)$ be two solutions to equation (5), with $\tilde{\rho}(0) = \rho(0) = 0$, where $\tilde{\rho}$ solves equation (5) with a pointwise lower right-hand side. For any $t > 0$, if $\tilde{\rho}(t) = \rho(t)$ then $\rho'(t) \geq \tilde{\rho}'(t) > 0$. This implies that at any touching point, ρ must touch $\tilde{\rho}$ from below. Consequently, by continuity,

$$\rho(t') \geq \tilde{\rho}(t') \text{ for } t' > 0 \Rightarrow \rho(t) \geq \tilde{\rho}(t) \text{ for all } t \geq t'.$$

Now suppose, to contradiction, that $\tilde{\rho}(\hat{t}) > \rho(\hat{t})$ for some for $\hat{t} > 0$. It must hold that $\tilde{\rho}(t) > \rho(t)$ for all $t \in (0, \hat{t})$. Since $\partial^2 c / \partial r^2 > 0$ and $\tilde{\rho}$ corresponds to a lower right-hand side of equation (5), it follows from equation (5) that $\tilde{\rho}'(t) < \rho'(t)$ for all $t \in (0, \hat{t})$. But then

$$\tilde{\rho}(\hat{t}) - \rho(\hat{t}) = \int_0^{\hat{t}} [\tilde{\rho}'(t) - \rho'(t)] dt < 0,$$

a contradiction. ■

REFERENCES

- Alonso, Ricardo, and Odilon Câmara. 2016. "Bayesian Persuasion with Heterogeneous Priors." *Journal of Economic Theory* 165: 672–706.
- Baliga, Sandeep, Eran Hanany, and Peter Klibanoff. 2013. "Polarization and Ambiguity." *American Economic Review* 103 (7): 3071–83.
- Blackwell, David. 1951. "Comparison of Experiments." In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 1, edited by Jerzy Neyman, 93–102. Berkeley: University of California Press.
- Blackwell, David. 1953. "Equivalent Comparisons of Experiments." *Annals of Mathematical Statistics* 24 (2): 265–72.
- Blackwell, David, and Lester Dubins. 1962. "Merging of Opinions with Increasing Information." *Annals of Mathematical Statistics* 33 (3): 882–86.
- Che, Yeon-Koo, and Navin Kartik. 2009. "Opinions as Incentives." *Journal of Political Economy* 117 (5): 815–60.
- Daley, Brendan, and Brett Green. 2014. "Market Signaling with Grades." *Journal of Economic Theory* 151 (1): 114–45.
- Francetich, Alejandro, and David Kreps. 2014. "Bayesian Inference Does Not Lead You Astray...On Average." *Economics Letters* 125 (3): 444–46.
- Frank, Robert H. 1985. "The Demand for Unobservable and Other Nonpositional Goods." *American Economic Review* 75 (1): 101–16.
- Good, Irving J. 1965. "A List of Properties of Bayes-Turing Factors." *NSA Technical Journal* 10 (2): 1–6.
- Harbaugh, Rick, and Eric Rasmusen. 2018. "Coarse Grades: Informing the Public by Withholding Information." *American Economic Journal: Microeconomics* 10 (1): 210–35.
- Hirsch, Alexander V. 2016. "Experimentation and Persuasion in Political Organizations." *American Political Science Review* 110 (1): 68–84.
- Jovanovic, Boyan. 1982. "Truthful Disclosure of Information." *Bell Journal of Economics* 13 (1): 36–44.
- Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* 101 (6): 2590–2615.

- Kartik, Navin.** 2009. "Strategic Communication with Lying Costs." *Review of Economic Studies* 76 (4): 1359–95.
- Kartik, Navin, Frances Xu Lee, and Wing Suen.** 2017. "Investment in Concealable Information by Biased Experts." *RAND Journal of Economics* 48 (1): 24–43.
- Kartik, Navin, Frances Xu Lee, and Wing Suen.** 2019. "A Theorem on Bayesian Updating and Applications to Communication Games." <http://www.columbia.edu/~nk2339/Papers/KLS-disclosure.pdf>.
- Kartik, Navin, Marco Ottaviani, and Francesco Squintani.** 2007. "Credulity, Lies, and Costly Talk." *Journal of Economic Theory* 134 (1): 93–116.
- Lehrer, Ehud, and Tao Wang.** 2020. "Strong Stochastic Dominance." <http://www.math.tau.ac.il/~lehrer/Papers/StrongStochasticDominance.pdf>.
- Lizzeri, Alessandro.** 1999. "Information Revelation and Certification Intermediaries." *RAND Journal of Economics* 30 (2): 214–31.
- Milgrom, Paul R.** 1981. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics* 12 (2): 380–91.
- Onuchic, Paula** © **Debraj Ray.** 2019. "Conveying Value via Categories." <https://debrajray.com/wp-content/uploads/2019/10/OnuchicRay.pdf>.
- Sethi, Rajiv, and Muhamet Yildiz.** 2012. "Public Disagreement." *American Economic Journal: Microeconomics* 4 (3): 57–95.
- Spence, Michael.** 1973. "Job Market Signaling." *Quarterly Journal of Economics* 87 (3): 355–74.
- Truyts, Tom.** 2015. "Noisy Signaling with Exogenous Information." Unpublished.
- Van den Steen, Eric.** 2010. "Culture Clash: The Costs and Benefits of Homogeneity." *Management Science* 56 (10): 1718–38.
- Verrecchia, Robert E.** 1983. "Discretionary Disclosure." *Journal of Accounting and Economics* 5 (1): 179–94.
- Weiss, Andrew.** 1983. "A Sorting-cum-Learning Model of Education." *Journal of Political Economy* 91 (3): 420–42.
- Yildiz, Muhamet.** 2004. "Waiting to Persuade." *Quarterly Journal of Economics* 119 (1): 223–48.
- Zanardo, Enrico.** 2017. "How to Measure Disagreement?" Unpublished.

