

Escudé, Matteo; Sinander, Ludvig

Article

Slow persuasion

Theoretical Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Escudé, Matteo; Sinander, Ludvig (2023) : Slow persuasion, Theoretical Economics, ISSN 1555-7561, The Econometric Society, New Haven, CT, Vol. 18, Iss. 1, pp. 129-162, <https://doi.org/10.3982/TE5175>

This Version is available at:

<https://hdl.handle.net/10419/296407>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/4.0/>

Slow persuasion

MATTEO ESCUDÉ

Department of Economics and Finance, LUISS

LUDVIG SINANDER

Department of Economics and Nuffield College, University of Oxford

What are the value and form of optimal persuasion when information can be generated only slowly? We study this question in a dynamic model in which a “sender” provides public information over time subject to a graduality constraint, and a decision maker takes an action in each period. Using a novel “viscosity” dynamic programming principle, we characterize the sender’s equilibrium value function and information provision. We show that the graduality constraint inhibits information provision relative to unconstrained persuasion. The gap can be substantial, but closes as the constraint slackens. Contrary to unconstrained persuasion, less-than-full information may be provided even if players have aligned preferences but different prior beliefs.

KEYWORDS. Persuasion, information, gradual, dynamic, viscosity, constrained.

JEL CLASSIFICATION. C61, C65, C73, D82, D83.

1. INTRODUCTION

Information production takes time. Experts seeking to influence behavior can provide information no faster than it can be produced. In this paper, we study the impact of such graduality constraints on strategic communication.

Consider, for example, a firm undertaking a risky project. Management produces information about the soundness of the investment with a view to persuading investors to back the enterprise. This typically cannot be achieved instantaneously, as evidence of profitability accumulates only gradually over the months as the project is taken from conception to market.

Alternatively, consider a funding body that wishes to influence policy-making (or individual behavior) by commissioning research. Consensus on scientific questions is

Matteo Escudé: mescude@luiss.it

Ludvig Sinander: ludvig.sinander@economics.ox.ac.uk

We are grateful to Giacomo Calzolari, Eddie Dekel, Piero Gottardi, Alessandro Pavan and Bruno Strulovici for guidance and comments, and to three anonymous referees, Gregorio Curello, Théo Durandard, Piotr Dworzak, Jeff Ely, Alex Frug, George Georgiadis, Duarte Gonçalves, Yingni Guo, Navin Kartik, Frédéric Koessler, Elliot Lipnowski, Ludmila Matysková, Konstantin Milbradt, Marco Ottaviani, Eduardo Perez-Richet, Francisco Poggi, Sven Rady, Marciano Siniscalchi, Alex Smolin, Pierpaolo Soravia, Quitzé Valenzuela-Stokey, and audiences at EUI, Northwestern, the 2019 Barcelona GSE Summer Forum, and the 2019 Summer School of the Econometric Society for comments and suggestions. A previous version of this paper was titled “Strategic research funding.”

built over dozens of studies, each of which takes time to conduct and be peer-reviewed. This constrains how rapidly information can be produced.

How much can an expert benefit from strategic information provision in such an environment? How does the graduality constraint affect the amount of information provided? In this paper, we answer these questions by characterizing the value of graduality-constrained information provision and by delineating how graduality shapes the extent of information production.

The value that an expert can derive from strategic information provision is a key concern of the literature on Bayesian persuasion. In the well-known model of [Kamenica and Gentzkow \(2011\)](#), an uninformed “sender” can costlessly produce public information about the unknown state of the world by designing a public signal, and a (symmetrically uninformed) decision maker observes the signal’s realization and takes an action.

Our *slow persuasion* model augments this setting with a graduality constraint. Since graduality is inherently dynamic, so is our model. In each period, the sender can produce at most a *small* amount of public information about the unknown (and fixed) state of the world. The players’ shared belief about the state evolves over time as evidence accumulates. Informed by this belief, the decision maker takes an action in each period, which together with the state determines (flow) payoffs. This defines a two-player stochastic game in continuous time, which we study by describing the set of Markov perfect equilibria with the common belief as state variable. We characterize equilibrium payoffs and behavior (in particular, information provision).

The persuasion literature is distinguished from the earlier cheap-talk literature by its *commitment assumption*: once the signal has been chosen, its realization cannot be concealed from the decision maker. Thus the two players remain symmetrically informed throughout their interaction, which eliminates communication frictions and so permits a sharper focus on the value of (constrained) information provision. We maintain the commitment assumption within each period: whatever information is produced becomes public. (We do not assume commitment *over time*. But given our focus on Markov perfect equilibria, nothing would change if we did.)

Our first result characterizes the value of slow persuasion. The sender’s *value function* describes, at each public belief, her expected discounted continuation payoff in equilibrium. Proposition 1 asserts that the value function is a version of the concave envelope (the value in unconstrained persuasion) that accounts for graduality and impatience. In particular, the value is strictly convex and below the concave envelope whenever the latter is affine (and exceeds the sender’s flow payoff).

To see why, consider “splitting” the current public belief across two posteriors by producing information. In unconstrained persuasion, the split can be effected immediately, so that its value equals the average of payoffs at the two posteriors. When persuasion is slow, it takes time for the belief to reach (one of) the target posteriors, which has two effects. First, the sender impatiently discounts the payoffs she anticipates once one of the posteriors has been reached. Second, until the target posteriors are reached, the actions chosen by the decision maker will be those she considers optimal at the beliefs prevailing in the interim.

The value of slow persuasion increases when the graduality constraint is slackened (allowing faster information production) and is well approximated by the concave envelope for a sufficiently slack constraint (Corollary 1). Outside of this special case, the value and concave envelope may differ substantially.

We next study behavior. Using our characterization of the value, we obtain a description of the sender's equilibrium information-provision strategy. To assess the impact of the graduality constraint, we derive comparative statics (Proposition 2): as the constraint tightens, the sender optimally provides Blackwell-less information over the course of the relationship. (She *could* provide the same aggregate information, albeit more slowly; we prove that she prefers not to.) This holds no matter what the players' preferences.

In slow persuasion, Blackwell-less information is provided overall than in the unconstrained benchmark (Proposition 3). In general, the gap can be large. However, in the special case of fast information arrival, long-run information provision is well approximated by the prediction of the unconstrained-persuasion model (also Proposition 3).¹ The seemingly similar "slow" limit lacks this continuity property: a sufficiently *tight* constraint does *not* generally lead to negligible aggregate information provision.

We conclude by studying the case in which the conflict of interest between the players arises not from different preferences over actions conditional on the state, but rather from their having different prior beliefs. This case is salient in organizations, where preferences may already have been aligned using contracts, and for policy questions such as climate-change mitigation, where disagreement is primarily over the extent of anthropogenic climate change, rather than over the correct policy response to a given amount of warming.

In unconstrained persuasion, a purely belief-based conflict of this sort amounts to no conflict at all: full information is provided, and *ex post* optimal actions are taken (see, e.g., [Alonso and Câmara \(2016\)](#)). We show by contrast that when persuasion is constrained to be slow, only partial information is produced over the long run if the constraint is tight enough (Proposition 5). It follows in particular that instituting contracts to align preferences conditional on the state is in general insufficient to encourage information production.

Clearly, prior disagreement harms the sender when persuasion is slow, since then only partial evidence is available in the medium term, leading players' priors to color their posterior beliefs and (thus) their views on what actions are best. We show that this welfare effect is monotone: the greater the prior difference, the lower is the sender's value (Proposition 6).

The key to our results is Proposition 1 (the value characterization), from which the remaining propositions are derived. To prove it, we follow the natural route of studying the value as a solution of the Hamilton–Jacobi–Bellman (HJB) equation of the sender's best-reply problem. This poses a technical challenge: the value function may have kinks, in which case it fails to solve the HJB equation (a second-order differential equation)

¹Formally, long-run beliefs (the set of posteriors to which the belief martingale converges) are close to those chosen in the unconstrained problem.

in the classical sense. Such kinks are an unavoidable byproduct of two features of the sender's problem: her ability to freeze the public belief by halting information provision, and the discontinuities in her flow payoff that occur whenever the decision maker switches from one action to another.

To overcome this hurdle, we extend existing results from the mathematics literature to prove a novel dynamic programming principle (Theorem 1): the value function is a *viscosity solution* of the HJB equation. This permits us to use the powerful theory of viscosity solutions of differential equations to characterize the value function in Proposition 1. We view this as a methodological contribution, and believe that our viscosity approach will prove useful for the study of other stochastic models in continuous time.

1.1 Related literature

This paper belongs to the Bayesian persuasion literature initiated by Kamenica and Gentzkow (2011), Rayo and Segal (2010) and Aumann and Maschler (1995) (see Kamenica (2019) and Bergemann and Morris (2019) for surveys). We contribute in particular to the growing strand of that literature which examines the impact of constraints on (or costs of) information production (e.g., Gentzkow and Kamenica (2014), le Treust and Tomala (2019) and Doval and Skreta (2021)).² Whereas these papers consider static settings, we study a dynamic constraint: graduality.

We also contribute to the larger strand on dynamic persuasion. In some important papers on this topic (Ely (2017), Renault, Solan, and Vieille (2017), Ball (2022)), the state of the world evolves over time. We instead consider a fixed state, but impose a graduality constraint.

Several recent papers study dynamic persuasion models with a fixed state in which the decision maker chooses when to take a game-ending action (Au (2015), Ely and Szydlowski (2020), Orlov, Skrzypacz, and Zryumov (2020), Bizzotto, Rüdiger, and Vigier (2021), Smolin (2021), Che, Kim, and Mierendorff (2021)). Our decision maker instead selects freely among her actions in each period. This is an important difference: whereas commitment over time has no value for the sender in our setting, it is strictly valuable in these models. The last two papers feature graduality constraints; the remainder do not.

More closely related are Henry and Ottaviani (2019) and Siegel and Strulovici (2020, §6), who study models of graduality-constrained information provision about a fixed state in which the *sender* chooses when to stop irreversibly, whereupon the decision maker takes an action.³ Our decision maker instead acts in every period, earning the sender a flow payoff. Until she stops, these papers' sender incurs a flow cost $c > 0$, absent in our model. Both papers focus on welfare-improving institutional design, whereas we

²Some of this work is surveyed by Kamenica, Kim, and Zapechelnyuk (2021). Calzolari and Pavan (2006a,b), Pavan and Calzolari (2009), Rosar (2017), Georgiadis and Szentos (2020), Dworzak (2020), Doval and Skreta (2021) and Boleslavsky and Kim (2021) study persuasion problems subject to *incentive* constraints.

³Brocas and Carrillo (2007) study a similar setting in discrete time, with no discounting. See also Fershtman and Pavan (2022, §4 and §5).

study equilibrium behavior in a fixed game. Siegel and Strulovici do not have results analogous to any of ours, but their environment exhibits similar properties to the special case of our model with only two or three actions.

Henry and Ottaviani further assume that there are only two actions and that the sender's payoff is independent of the state. In this special case, the dependence of information provision on the speed at which evidence accumulates that we emphasize (Proposition 2) is entirely absent if (as in our model) $c = 0$. It becomes important as soon as the sender's payoff is allowed to be state-dependent or there are more than two actions (see our two- and three-action examples on p. 138 and in the Online Supplemental Material, Appendix H.2 (available in a supplementary file on the journal website, <http://econtheory.org/supp/5175/supplement.pdf>). The authors characterize the sender's value function, and show that information provision is close to that in unconstrained persuasion when the cost c is small, results analogous to a specialization of our Propositions 1 and 3.⁴ Their two-action environment is tractable enough to avoid the need for viscosity methods.

Our material on purely belief-based conflicts of interest contributes to the literature on strategic communication with heterogeneous beliefs (e.g., Che and Kartik (2009) and van den Steen (2009)). In unconstrained persuasion, full information is provided when the conflict of interest is rooted in differing beliefs alone (e.g., Alonso and Câmara (2016)). We show that slow persuasion overturns this conclusion.

Viscosity solutions of differential equations were introduced by Crandall and Lions (1983). We give a brief exposition and some references in supplemental Appendix K. Viscosity methods have begun to be used in economic theory (e.g., Nikandrova and Pansc (2018), Ke and Villas-Boas (2019), Keller and Rady (2020), Zhong (2022) and Barilla and Gonçalves (2022)). Our sender's best-reply problem is nonstandard due the discontinuities in her flow payoff and her ability to freeze the state variable (the public belief), precluding off-the-shelf use of standard results.

Our technical contribution is related to Kuvalekar and Lipnowski (2020), who also adapt arguments from the viscosity literature to deal with discontinuities in flow payoffs. Their goal is to establish smooth pasting for their model, a property which fails in our setting (as the value may have kinks). We instead derive a dynamic programming principle.

1.2 Roadmap

We formulate the model in the next section, and pin down the decision-maker's equilibrium behavior in Section 3. We then study the sender's best-reply problem, characterizing her value function in Section 4 and her equilibrium information provision in Section 5. We conclude in Section 6 by considering the case in which the conflict of interest arises from differing beliefs alone.

⁴Henry and Ottaviani also show that when c is small *and* the sender is patient, her value is close to the concave envelope. This is analogous to our Corollary 1.

2. MODEL

Our model is a stochastic game in continuous time. The players are a *decision maker* and a *sender*, who take respective actions a_t and λ_t in each period $t \in \mathbf{R}_+$. Flow payoffs depend on the decision-maker's action a_t and on the state variable p_t . The sender's action λ_t affects the stochastic evolution of p_t . In particular, λ_t is the rate at which public information arrives, and p_t is the common belief, which evolves according to Bayes's rule.

2.1 State and payoffs

There is a binary state $\theta \in \{0, 1\}$. The sender and the decision maker have a common prior belief p_0 that the state is $\theta = 1$. (We will drop the common-prior assumption in Section 6.) Time $t \in \mathbf{R}_+$ is continuous. The decision maker takes an action $a \in \mathcal{A}$ at each moment, where \mathcal{A} is a finite set.

When the decision maker takes action a and the common belief is p , the sender's and decision-maker's respective flow payoffs are $f_S(a, p)$ and $f_D(a, p)$, both continuous in p . Expected utility ($f_S(a, \cdot)$ and $f_D(a, \cdot)$ affine) is a natural special case. The sender and decision maker discount flow payoffs at rates $r > 0$ and $r_D > 0$, respectively.

This abstract setting nests the examples in the [Introduction](#) as follows. In the first example, an investor (decision maker) decides in each period how much financing $a \in \mathcal{A} \subseteq \mathbf{R}_+$ to contribute to a firm. The net return on the firm's project is $\alpha\theta - 1$, where $\alpha > 1$, and thus the investor's flow payoff is $f_D(a, p) := a(\alpha p - 1)$.⁵ The firm's management (the sender) cares only about financing, so $f_S(a, p) := a$. In the [Introduction's](#) second example, a policymaker sets policy a in each period, and her and the sender's policy preferences depend on a persistent, unknown state θ .

2.2 The sender's information provision

At each instant, the sender can costlessly permit a small amount of public information to arrive. In particular, she chooses $\lambda_t \in [0, \bar{\lambda}]$, and everyone observes the process

$$dX_t = \theta \lambda_t dt + \sigma \sqrt{\lambda_t} d\tilde{B}_t,$$

where \tilde{B} is a standard Brownian motion and $\sigma > 0$. The parameter $\bar{\lambda} > 0$ quantifies the slackness of the graduality constraint $\lambda_t \leq \bar{\lambda}$. Our assumption that the noise is Brownian rules out information arriving in discrete lumps.

The signal process may be microfounded as follows. Write $\Lambda_t := \int_0^t \lambda_s ds$ for cumulative information-production effort. Total effort Λ produces evidence, summarized by a "score": write Y_Λ for the cumulative score, and assume that today's score \dot{Y}_Λ has mean θ , but is subject to i.i.d. noise. Since white noise is the rate of change of a random walk,

⁵Her expected discounted payoff is $\mathbf{E}(\int_0^\infty e^{-r_D t} a_t (\alpha\theta - 1) dt) = \int_0^\infty e^{-r_D t} f_D(a_t, p_t) dt$.

we may write $dY_\Lambda = \theta d\Lambda + \sigma d\bar{B}_\Lambda$, where \bar{B} is a standard Brownian motion. Then the evolution over time of the cumulative score $X_t := Y_{\Lambda_t}$ follows

$$dX_t = \theta d\Lambda_t + \sigma d\bar{B}_{\Lambda_t} = \theta \frac{d\Lambda_t}{dt} dt + \sigma \sqrt{\frac{d\Lambda_t}{dt}} d\tilde{B}_t = \theta \lambda_t dt + \sigma \sqrt{\lambda_t} d\tilde{B}_t,$$

where \tilde{B} is a(nother) standard Brownian motion.⁶

As the players observe $(X_t)_{t \in \mathbf{R}_+}$, their common belief p_t is updated according to Bayes’s rule. By a well-known result from filtering theory, the belief evolves as

$$dp_t = \sqrt{\lambda_t} \frac{p_t(1 - p_t)}{\sigma} dB_t,$$

where B is a standard Brownian motion according to the common belief.⁷ See, for example, Bolton and Harris (1999, Lemma 1) for a heuristic derivation. The belief process $(p_t)_{t \in \mathbf{R}_+}$ is a martingale with a.s. continuous sample paths.

2.3 Strategies and equilibrium

We focus on Markov perfect equilibria, in which players’ behavior depends on the past only through the current state p_t . These equilibria are natural and standard, and avoid technical issues that can arise in continuous time. It is worth noting, however, that there may be other equilibria (suitably defined), and that these may differ qualitatively from the ones that we study.

DEFINITION 1. A function $[0, 1] \rightarrow \mathbf{R}^n$ is *piecewise continuous* if and only if its discontinuities form a discrete subset of $(0, 1)$.⁸

Note that a piecewise continuous function is continuous at 0 and 1. Recall that a discrete subset of $(0, 1)$ is at most countable.

DEFINITION 2. A *strategy* of the sender (decision maker) is a $[0, \bar{\lambda}]$ -valued (\mathcal{A} -valued) stochastic process $(\lambda_t)_{t \in \mathbf{R}_+}$ ($(a_t)_{t \in \mathbf{R}_+}$) adapted to the filtration generated by $(p_t)_{t \in \mathbf{R}_+}$ and actions. A (*pure*) *Markov strategy of the sender* is a measurable map $\Lambda : [0, 1] \rightarrow [0, \bar{\lambda}]$. A *Markov strategy of the decision maker* is a piecewise continuous map $A : [0, 1] \rightarrow \Delta(\mathcal{A})$, where $\Delta(\mathcal{A})$ denotes the set of all probability distributions over the (finite) set \mathcal{A} . We identify a Markov strategy Λ (A) with the strategy (stochastic process) that it induces via $\lambda_t := \Lambda(p_t)$ (a_t distributed according to $A(p_t)$).

⁶As is well known, the “time-changed” process $t \mapsto d\bar{B}_{\Lambda_t}$ has the same law as the *scaled* process $t \mapsto \sqrt{d\Lambda_t/dt} d\tilde{B}_t$, where \tilde{B} is a standard Brownian motion.

⁷See, for example, Papanicolaou (2016, §4.2.2). The result features assumptions about the process $(\lambda_t)_{t \in \mathbf{R}_+}$, which are satisfied by equilibrium processes in our model. The process B is given by $dB_t = dX_t - p_t dt$ and $B_0 = 0$. It is *not* a Brownian motion according to the “objective” law of X under either $\theta = 0$ or $\theta = 1$, but it is a Brownian motion from the point of view of an observer with belief p_t , as can be seen from the Girsanov theorem (e.g., Karatzas and Shreve (1991, §3.5)).

⁸A subset of $(0, 1)$ is *discrete* if and only if each of its members is an isolated point: it lives in a neighborhood that contains no other members.

The restriction to piecewise continuous strategies is a mild assumption on the decision-maker’s tie-breaking that has no payoff consequences for her, provided her preferences are nondegenerate in a weak sense. (See supplemental Appendix J. This is obvious for expected-utility preferences.) The role of piecewise continuity is to ensure that the sender’s best-reply problem satisfies a dynamic programming principle (Theorem 1 in Section 4.2 below).

DEFINITION 3. A strategy $(\lambda_t)_{t \in \mathbf{R}_+}$ of the sender is a *best reply* at $\bar{p} \in [0, 1]$ to a Markov strategy $A : [0, 1] \rightarrow \Delta(\mathcal{A})$ of the decision maker if and only if it maximizes

$$\mathbf{E} \left(\int_0^\infty e^{-rt} \left[\int_{\mathcal{A}} f_S(a, p_t) A(da|p_t) \right] dt \right),$$

where $dp_t = \sqrt{\lambda_t} \frac{p_t(1-p_t)}{\sigma} dB_t$ and $p_0 = \bar{p}$,

over all $[0, \bar{\lambda}]$ -valued processes $(\lambda_t)_{t \in \mathbf{R}_+}$ adapted to the filtration generated by $(p_t)_{t \in \mathbf{R}_+}$.⁹

A strategy $(a_t)_{t \in \mathbf{R}_+}$ of the decision maker is *undominated* if and only if there is no other strategy $(a'_t)_{t \in \mathbf{R}_+}$ that yields the same expected payoff and has $f_D(a'_t, p) > f_D(a_t, p)$ a.s. for some $p \in [0, 1]$. A strategy $(a_t)_{t \in \mathbf{R}_+}$ of the decision maker is a *best reply* at $\bar{p} \in [0, 1]$ to a strategy $(\lambda_t)_{t \in \mathbf{R}_+}$ of the sender if and only if it is undominated and maximizes

$$\mathbf{E} \left(\int_0^\infty e^{-rDt} f_D(a_t, p_t) dt \right), \quad \text{where } dp_t = \sqrt{\lambda_t} \frac{p_t(1-p_t)}{\sigma} dB_t \text{ and } p_0 = \bar{p},$$

over all \mathcal{A} -valued processes adapted to the filtration generated by $(p_t)_{t \in \mathbf{R}_+}$.

We rule out dominated strategies of the decision maker as uninteresting. Accommodating them merely complicates the statements of some results.

DEFINITION 4. A *Markov perfect equilibrium (MPE)* is a pair of Markov strategies that are best replies to each other at each $\bar{p} \in [0, 1]$.

2.4 Unconstrained (static) benchmark

In unconstrained persuasion, the sender flexibly provides information once and for all, with no graduality constraint. If the sender induces belief p , then the decision maker takes an action $A(p)$ that maximizes $f_D(\cdot, p)$, giving the sender a payoff of $u(p) := f_S(A(p), p)$. Assume that the decision maker breaks ties such that u is upper semicontinuous.¹⁰

Kamenica and Gentzkow (2011) studied this model, and showed the following. The sender is able to induce all and only distributions of beliefs whose mean is p_0 (“splits”

⁹In principle, the sender can use a process that is adapted to the filtration generated by $(p_t)_{t \in \mathbf{R}_+}$ and actions. But when the decision maker uses a Markov strategy, actions contain no additional information.

¹⁰Without this assumption, an optimal policy may fail to exist.

of the prior). The sender's value at prior p_0 is $(\text{cav } u)(p_0)$, where $\text{cav } u$ is the concave envelope of u (the smallest concave function that majorizes u). The sender has an optimal policy that induces either two beliefs (if $(\text{cav } u)(p_0) > u(p_0)$) or one belief (if $(\text{cav } u)(p_0) = u(p_0)$).

As we shall see, unconstrained persuasion typically provides a poor approximation to our sender's value and information provision, but the approximation is good if the graduality constraint is slack ($\bar{\lambda}$ large).

3. MYOPIC BEHAVIOR BY THE DECISION MAKER

To avoid uninteresting technicalities, we will focus on MPEs in which the decision-maker's tie-breaking is well behaved. Say that a Markov strategy $A : [0, 1] \rightarrow \Delta(\mathcal{A})$ of the decision maker is *regular* if and only if it breaks ties such that the sender's induced payoff $u(p) := \int_{\mathcal{A}} f_S(a, p) A(da|p)$ is upper semi-continuous. Regular Markov strategies exist—indeed, *any* Markov strategy of the decision maker need be modified only on a discrete (hence at most countable) set of beliefs p to be made regular, and this modification leaves the decision-maker's flow payoff unchanged.¹¹

Call a Markov strategy $A : [0, 1] \rightarrow \Delta(\mathcal{A})$ of the decision maker *myopic* if and only if at each belief $p \in [0, 1]$, every action in the support of $A(\cdot|p)$ maximizes $f_D(\cdot, p)$.

OBSERVATION 1. A regular strategy of the decision maker is part of a MPE if and only if it is myopic.

That is, all and only myopic behavior can be supported in a MPE, modulo tie-breaking. It follows that our analysis below of the sender's behavior in MPEs carries over to a simpler model with a myopic decision maker, or alternatively, a sequence of short-lived decision makers.

PROOF. If the sender uses a Markov strategy, then since the decision maker cannot affect the evolution of the state, a regular strategy of hers is a best reply if and only if it is myopic. The “only if” part follows.

For the “if” part, fix a regular and myopic Markov strategy $A : [0, 1] \rightarrow \Delta(\mathcal{A})$ of the decision maker. We show in Section 5.1 that (given that A is regular,) there is a Markov strategy Λ of the sender, which is a best reply to A at every $\bar{p} \in [0, 1]$. Since A is myopic, it is a best reply to Λ at every $\bar{p} \in [0, 1]$. So (Λ, A) is a MPE. \square

Observation 1 implies that Markov perfect equilibria exist, and further that a sender-preferred MPE exists. It also follows that the decision-maker's behavior can differ across MPEs only at beliefs at which she is exactly indifferent, which in turn implies that generically, under mild conditions, the MPE is partially unique.¹²

¹¹Recall that a Markov strategy is piecewise continuous by definition, and that the decision-maker's flow payoff $f_D(a, \cdot)$ is continuous.

¹²In particular, for Lebesgue-a.e. expected-utility f_D (viewed as vectors in $\mathbf{R}^{2|\mathcal{A}|}$), the decision-maker's strategy differs across MPEs only on a countable set of beliefs. We will see in Section 5.1 that the sender's best reply is generically partially unique.

In light of Observation 1, it remains only to characterize the sender's best reply to a given regular and myopic strategy of the decision maker. We will in fact characterize her best reply to an arbitrary regular Markov strategy. We proceed in two steps, studying the sender's value function Section 4, and then her best reply in Section 5.

4. THE SENDER'S VALUE FUNCTION

Fix a regular Markov strategy $A : [0, 1] \rightarrow \Delta(\mathcal{A})$ of the decision maker. The sender's induced preferences over beliefs are given by

$$u(p) := \int_{\mathcal{A}} f_S(a, p) A(da|p).$$

Note that u is piecewise continuous and upper semicontinuous since $f_S(a, \cdot)$ is continuous and A is piecewise continuous and regular. We will study the sender's problem given an arbitrary piecewise continuous and upper semi-continuous flow payoff $u : [0, 1] \rightarrow \mathbf{R}$.

The sender's best-reply problem, with (discounted) value function v , is

$$\begin{aligned} v(p_0) = \sup_{(\lambda_t)_{t \in \mathbf{R}_+}} \mathbf{E} \left(r \int_0^\infty e^{-rt} u(p_t) dt \right) \\ \text{subject to } dp_t = \sqrt{\lambda_t} \frac{p_t(1-p_t)}{\sigma} dB_t, \end{aligned} \tag{BRP}$$

where $(\lambda_t)_{t \in \mathbf{R}_+}$ is chosen among all $[0, \bar{\lambda}]$ -valued processes adapted to the filtration generated by $(p_t)_{t \in \mathbf{R}_+}$, and p_0 is given.

To understand the sender's incentives in (BRP) and to motivate our solution technique, we begin with an illustrative example.

4.1 A two-action example

There is a risky action $a = 1$ that yields a benefit of 3 to both players (only) in state $\theta = 1$. Taking the risky action costs the sender and decision maker 1 and 2, respectively, so that expected utilities at belief $p \in [0, 1]$ are

$$f_S(1, p) = 3p - 1 \quad \text{and} \quad f_D(1, p) = 3p - 2.$$

There is also a safe action $a = 0$ giving both players a certain payoff of zero.

The decision-maker's (unique regular) myopic strategy is $A = \mathbf{1}_{[2/3, 1]}$. The sender's induced flow payoff u is depicted in Figure 1.

Most of the sender's best-reply problem is easily solved. When $p_t \in (0, 2/3)$, she finds it strictly optimal to provide information (at full tilt, i.e., $\lambda_t = \bar{\lambda}$), since her flow payoff can only improve. When $p_t \in (2/3, 1)$, it is weakly optimal to provide information, since u is affine on this region.¹³

¹³She could provide information only while $p_t > 2/3$. Then at each instant that she does, the belief changes by a mean-zero random increment dp_t , which since u is affine on $[2/3, 1]$ yields expected payoff $\mathbf{E}(u(p_t + dp_t)) = u(p_t)$, the same as from not providing information.




FIGURE 1. The sender's value function in the two-action example.

The sender faces a nontrivial trade-off at $2/3$, however. By stopping information provision, she can lock in a moderate payoff of 1 forever. If she instead continues, then she may increase her payoff toward 2 (if p_t rises), but may equally suffer a flow payoff of zero in the near future (if p_t declines). The optimal resolution of this trade-off depends on how rapidly evidence can accumulate: if quickly ($\bar{\lambda}$ high), then she provides information, and if slowly ($\bar{\lambda}$ low), then she does not.

The value function is easily computed in either case: it has the shape depicted in Figure 1a if evidence accumulates quickly, and the shape in Figure 1b if slowly. The kink in the latter case arises from the sender choosing to freeze the state variable p_t once it hits the discontinuity point $2/3$ of the flow payoff u (induced by the decision maker switching actions).

For comparison, the value function in the unconstrained benchmark is the concave envelope $cav u$, which is affine and strictly higher.

4.2 The HJB equation and viscosity solutions

The Hamilton–Jacobi–Bellman (HJB) equation for the sender's best-reply problem is the following differential equation in an unknown function $w : [0, 1] \rightarrow \mathbf{R}$:

$$w(p) = \sup_{\lambda \in [0, \bar{\lambda}]} \left\{ u(p) + \frac{1}{r} \left(\sqrt{\lambda} \frac{p(1-p)}{\sigma} \right)^2 \frac{w''(p)}{2} \right\},$$

or equivalently

$$w(p) = u(p) + \bar{\lambda} \frac{p^2(1-p)^2}{2r\sigma^2} \max\{0, w''(p)\}. \tag{HJB}$$

In well-behaved problems, a *dynamic programming principle* holds: the value function v is a *classical solution* of (HJB), meaning that v is twice continuously differentiable and that v and v'' satisfy (HJB) at every $p \in (0, 1)$. The familiar interpretation is that the value v is today's flow payoff u plus the expected rate of change of the value, discounted

by r . In the latter term, $\sqrt{\lambda}p(1-p)/\sigma$ is the rate of information arrival, and $v''(p)/2$ is the (local) value of information.

Our sender's problem is not so well behaved. Since the sender can freeze the state variable p (by setting $\lambda = 0$) and her flow payoff u may have discontinuities (arising from action switches by the decision maker), the value function may have kinks, as seen in the two-action example of the previous section. Since v'' does not exist at kinks, v cannot be a classical solution of (HJB): the right-hand side is ill-defined.¹⁴

To be able to use (HJB) to study the value function when the latter may have kinks, we require a broader notion of "solution" of a differential equation. Let u^* (u_*) denote the upper (lower) semicontinuous envelope of u , that is, the pointwise smallest (largest) upper (lower) semicontinuous function that majorizes (minorizes) u . The envelopes u^* and u_* differ only on a discrete set since u is piecewise continuous, and we have $u_* \leq u = u^*$ since u is upper semicontinuous.

DEFINITION 5. $w : [0, 1] \rightarrow \mathbf{R}$ is a *viscosity subsolution (supersolution)* of (HJB) if and only if it is upper (lower) semicontinuous, and for any twice continuously differentiable $\phi : (0, 1) \rightarrow \mathbf{R}$ and local minimum $p \in (0, 1)$ of $\phi - w$ (of $w - \phi$),

$$w(p) \leq u^*(p) + \bar{\lambda} \frac{p^2(1-p)^2}{2r\sigma^2} \max\{0, \phi''(p)\}$$

$$\left(w(p) \geq u_*(p) + \bar{\lambda} \frac{p^2(1-p)^2}{2r\sigma^2} \max\{0, \phi''(p)\} \right).$$

w is a *viscosity solution* of (HJB) if and only if it is both a sub- and a supersolution.

REMARK 1. It is without loss of generality to restrict attention at each $p \in (0, 1)$ to functions ϕ that satisfy $\phi(p) = w(p)$ and for which $\phi - w$ ($w - \phi$) has a *strict global* minimum at p .

A brief exposition of the theory of viscosity solutions is given in supplemental Appendix K. Observe that if w is a viscosity solution of (HJB) and is twice continuously differentiable on a neighborhood of $p \in (0, 1)$, then it satisfies (HJB) in the classical sense at p .¹⁵

Although the value function need not satisfy (HJB) in the classical sense, it *does* satisfy (HJB) in the viscosity sense.

THEOREM 1 (Dynamic programming principle). *Assume that u is piecewise continuous. Then v is a viscosity solution of (HJB), with boundary condition $v = u$ on $\{0, 1\}$.*

¹⁴Even in the absence of a kink (as in Figure 1a), v cannot be a classical solution of (HJB) unless u is continuous. For whenever u jumps, v'' must also jump to balance (HJB), in which case v fails to be twice continuously differentiable.

¹⁵Since we may choose a twice continuously differentiable ϕ that coincides with w on a neighborhood of p , so $\phi - w$ and $w - \phi$ are locally minimized at p and $\phi''(p) = w''(p)$.

The proof is in Appendix A.

We view Theorem 1 as a technical contribution. It extends a well-known theorem from the optimal control literature in which the flow payoff u is assumed to be continuous. That is an unacceptable hypothesis in economic applications such as ours, where u depends on the endogenous strategic behavior of other players. Theorem 1 may prove useful for studying other models of strategic interaction in continuous-time stochastic environments.

4.3 Characterization of the sender's value function

We shall characterize the sender's value v in terms of its local convexity, defined as follows.

DEFINITION 6. $w : [0, 1] \rightarrow \mathbf{R}$ is *locally (strictly) convex* at $p \in (0, 1)$ if and only if

$$w(p) \leq (<) \gamma w(p') + (1 - \gamma)w(p'')$$

for all $p' < p < p''$ sufficiently close to p , where γ is such that $\gamma p' + (1 - \gamma)p'' = p$. It is *locally (strictly) concave* at p if and only if the reverse (strict) inequality holds.

By way of illustration, the function v depicted in Figure 1b is locally strictly concave at $2/3$ (but is concave on no open neighborhood of $2/3$).

Let $C \subseteq (0, 1)$ be the beliefs at which v is locally strictly convex, and let $D \subseteq (0, 1)$ be the (discrete) set of beliefs at which u is discontinuous.

PROPOSITION 1 (Value function). v is continuous and satisfies $u \leq v \leq \text{cav } u$. On C , we have $v < \text{cav } u$, and v is once continuously differentiable. On $C \setminus D$, we have further that v is twice continuously differentiable and satisfies

$$v(p) = u(p) + \bar{\lambda} \frac{p^2(1-p)^2}{2r\sigma^2} v''(p) \quad \text{at each } p \in C \setminus D. \tag{\partial}$$

On $(0, 1) \setminus C$, we have $v = u$. On $\{0, 1\}$, we have $u = v = \text{cav } u$.

Proposition 1 is summarized in Table 1, where C^k means “continuous and k times continuously differentiable.” The “smooth pasting” property in the entry for the region

TABLE 1. Summary of Proposition 1. C^k means “continuous and k times continuously differentiable.”

Region	Properties of v			
$C \setminus D$	locally strictly convex	$u \leq v < \text{cav } u$	C^2	equation (∂)
$C \cap D$	locally strictly convex	$u \leq v < \text{cav } u$	C^1	smooth pasting
$(0, 1) \setminus C$		$u = v \leq \text{cav } u$	C^0	
$\{0, 1\}$		$u = v = \text{cav } u$	C^0	

$C \cap D$ is the following consequence of continuous differentiability on C : for any sequence $(p_n)_{n \in \mathbf{N}}$ of beliefs in $C \setminus D$ converging to $p \in C \cap D$,¹⁶ we have $\lim_{n \rightarrow \infty} v'(p_n) = v'(p)$. Remark that since v is only C^0 on $(0, 1) \setminus C$, it may have (locally concave) kinks in this region—we saw an example of this in Figure 1b.

The characterization of v in Proposition 1 is a generalization of the concave envelope $\text{cav } u$. Both are upper envelopes of u that exceed u when convex and coincide with u when concave. But whereas $\text{cav } u$ is affine whenever it exceeds u , v is strictly convex when it exceeds u , due to impatience. The differential equation (∂) pins down the exact form of this strict convexity.

Proposition 1 permits us to solve for the value function. Given a candidate C' for C , (∂) may be solved in closed form on each maximal interval of C' up to constants. There is at most one collection of constants that ensures the properties demanded by Proposition 1 (at least if C' comprises *finitely* many maximal intervals), and if there is one then $C' = C$. We give some details in supplemental Appendix H.

TWO-ACTION EXAMPLE (SECTION 4.1), CONTINUED. Proposition 1 implies that the value function must have either the strictly convex shape in Figure 1a or the convex-affine shape in Figure 1b. It further rules out one of the two; for example, for $\bar{\lambda}$ small, the convex candidate violates $u \leq v$ at $2/3$. (See supplemental Appendix H.1 for details.)

The proof relies heavily on Theorem 1. In particular, all three lemmata below are derived using the fact that v is a viscosity solution of (HJB).

PROOF OF PROPOSITION 1. By Theorem 1, v is a viscosity solution of (HJB) satisfying $v = u$ on $\{0, 1\}$. It follows that v is continuous. The fact that $u = \text{cav } u$ on $\{0, 1\}$ follows from upper semicontinuity of u . We have $u \leq v$ because for any $p \in [0, 1]$, the value $u(p)$ is attainable (by setting $\lambda = 0$ forever), so must be lower than the optimal value $v(p)$.

To show that $v \leq \text{cav } u$, take any $p \in [0, 1]$, and consider the auxiliary problem in which the sender may choose any $[0, 1]$ -valued process $(p_t)_{t \in \mathbf{R}_+}$ satisfying $\mathbf{E}(p_t) = p$ for every $t \in \mathbf{R}_+$. The value $V(p)$ of this problem must exceed $v(p)$ since any belief process the sender can induce in her best-reply problem is available in the auxiliary problem. And we have $V(p) = (\text{cav } u)(p)$ since the auxiliary problem consists of a sequence of independent unconstrained persuasion problems (one for each instant t), in each of which the optimal value is $(\text{cav } u)(p)$ since u is upper semicontinuous.

For $(0, 1) \setminus C$, we show in Appendix B that

LEMMA 1. *On $(0, 1) \setminus C$, we have $v = u$.*

Now for C . In Appendix C, we prove that

LEMMA 2. *v is continuously differentiable on C .*

¹⁶Every point in $C \cap D$ can be reached by such a sequence, since every element of D is isolated by piecewise continuity of u .

To show that $v < \text{cav } u$ on C , take $p \in C$ and $p' < p < p''$ sufficiently close to p , and let $\gamma \in (0, 1)$ satisfy $\gamma p' + (1 - \gamma)p'' = p$. We have

$$\begin{aligned} v(p) &< \gamma v(p') + (1 - \gamma)v(p'') \quad \text{since } p \in C \\ &\leq \gamma(\text{cav } u)(p') + (1 - \gamma)(\text{cav } u)(p'') \quad \text{since } v \leq \text{cav } u \\ &\leq (\text{cav } u)(p) \quad \text{since } \text{cav } u \text{ is concave.} \end{aligned}$$

Finally, consider $C \setminus D$. We have

LEMMA 3. *On $C \setminus D$, v is twice continuously differentiable and satisfies (∂).*

This lemma is proved in Appendix D. □

We prove in Appendix E that letting $\bar{\lambda} \rightarrow \infty$ in Proposition 1 yields the following.

COROLLARY 1. *As $\bar{\lambda}$ increases, v increases pointwise. As $\bar{\lambda} \rightarrow \infty$, v converges uniformly to $\text{cav } u$.*

Thus when the sender is able to provide information rapidly, her equilibrium value is well approximated by the unconstrained-persuasion model. Beyond this case, the approximation is typically poor, as evidenced by the two-action example (Figure 1b, p. 139).

5. EQUILIBRIUM INFORMATION PROVISION

Having characterized the sender's value function (Proposition 1), we are ready to study her equilibrium behavior. We first show that she provides more information the less stringent the graduality constraint (Proposition 2, Section 5.2). We then establish (Proposition 3, Section 5.3) that less information is provided than in the unconstrained benchmark, but that the latter model provides a good approximation if the graduality constraint is sufficiently slack. Finally, we highlight the role of discontinuities in the sender's flow payoff u : any discontinuity, if paired with a sufficiently tight constraint, leads less-than-full information to be provided in equilibrium (Proposition 4, Section 5.4).

5.1 Induced beliefs in the long run

As in Section 4, fix a regular Markov strategy $A : [0, 1] \rightarrow \Delta(\mathcal{A})$.

COROLLARY 2. *The following Markov strategy is a best reply of the sender:*

$$\Lambda^*(p) = \begin{cases} \bar{\lambda} & \text{if } v(p) > u(p) \\ 0 & \text{otherwise.} \end{cases}$$

PROOF. When $v = u$, setting $\lambda = 0$ clearly attains the value v in the sender's problem (BRP). When $v > u$, $\lambda > 0$ must be optimal. By inspection of the sender's problem, $\lambda = \bar{\lambda}$ is optimal whenever $\lambda > 0$ is. \square

The strategy Λ^* provides information at full tilt when it is strictly valuable, and provides none otherwise. It is partially unique.¹⁷

Under this strategy, the belief p_t evolves according to $dp_t = \sqrt{\bar{\lambda}/\sigma^2} p_t(1 - p_t) dB_t$ until it hits the (closed) set $\{v = u\}$, then remains constant. The belief process $(p_t)_{t \in \mathbf{R}_+}$ converges a.s. by the martingale convergence theorem (e.g., Theorem 3.15 in Karatzas and Shreve (1991), ch. 1). Write F for the distribution of the limiting random variable. The support of F is the set of beliefs that the sender induces (with positive probability) in the long run. Note that F has mean p_0 since each p_t does, by the bounded convergence theorem (e.g., Theorem 16.5 in Billingsley (1995)).

COROLLARY 3. Fix a prior p_0 . A best reply of the sender induces the beliefs $\{p^-, p^+\}$ in the long run, where

$$p^- := \sup\{p \in [0, p_0] : v(p) = u(p)\} \quad \text{and} \quad p^+ := \inf\{p \in [p_0, 1] : v(p) = u(p)\}.$$

PROOF. The best reply Λ^* in Corollary 2 obviously induces $\{p^-, p^+\}$. \square

Provided $v(p_0) > u(p_0)$ (the interesting case), the long-run beliefs $\{p^-, p^+\}$ that we study are generically the unique ones consistent with a best reply—see supplemental Appendix I for a discussion. In general, they are the least extreme beliefs induced by some best reply (by footnote 17).

5.2 Comparative statics

A slacker graduality constraint leads to more extreme beliefs.

PROPOSITION 2 (Comparative statics). Fix any prior p_0 . As $\bar{\lambda}$ increases, p^- decreases and p^+ increases.

PROOF. Write $\{v > u\}$ for the set of beliefs at which the strategy Λ^* from Corollary 2 provides information. As $\bar{\lambda}$ increases, v increases pointwise by Corollary 1, so $\{v > u\}$ increases in the sense of set inclusion, and thus

$$p^- = \sup([0, p_0] \setminus \{v > u\})$$

decreases and p^+ similarly increases. \square

TWO-ACTION EXAMPLE (SECTION 4.1), CONTINUED. Let the prior be $p_0 \in (0, 2/3)$. The sender induces the long-run beliefs $\{0, 1\}$ if information can be produced quickly (Figure 1a, p. 139), and the less extreme $\{0, 2/3\}$ if not (Figure 1b).

¹⁷Precisely: any best reply must have $\Lambda > 0$ on $\{v > u\}$, $\Lambda = \bar{\lambda}$ a.e. on $\{v > u\}$, and $\Lambda = 0$ a.e. on $\{v = u\} \setminus K$, where $K \subseteq (0, 1)$ is the set on which u is locally (weakly) convex. Anything is optimal on $\{v = u\} \cap K$. (In Figure 1b on p. 139, we have $\{v = u\} \cap K = (2/3, 1)$.)

Proposition 2 states that more information is provided in the long run when evidence accumulates faster. More extreme long-run beliefs are achieved by providing information (at full tilt) for longer a.s. Consequently, slackening the graduality constraint leads Blackwell-more information to be generated. Thus the decision-maker's payoff improves, provided she values information in the sense that $p \mapsto f_D(a, p)$ is convex for each $a \in \mathcal{A}$ (a property satisfied by expected utility).

5.3 Slow versus unconstrained persuasion

We next characterize how equilibrium information provision compares with that in the (static) unconstrained-persuasion benchmark. It is well known (see [Kamenica and Gentzkow \(2011\)](#)) that given p_0 , an optimal policy in the unconstrained persuasion problem induces the beliefs $\{P^-, P^+\}$, where

$$P^- := \sup\{p \in [0, p_0] : (\text{cav } u)(p) = u(p)\}$$

$$P^+ := \inf\{p \in [p_0, 1] : (\text{cav } u)(p) = u(p)\}.$$

It follows that either one or two beliefs are induced, depending on p_0 .

Given Proposition 2, it is intuitive that less information will be provided in the dynamic model than in the benchmark, since the latter corresponds (informally) to the case $\bar{\lambda} = \infty$ of instantaneous information arrival. The following result shows that information provision is close to that in the unconstrained benchmark when evidence accumulates fast (high $\bar{\lambda}$).

PROPOSITION 3 (Slow vs. unconstrained). *Fix a prior p_0 . For any $\bar{\lambda} > 0$, we have $P^- \leq p^- \leq p^+ \leq P^+$. As $\bar{\lambda} \rightarrow \infty$, $p^- \rightarrow P^-$ and $p^+ \rightarrow P^+$.*

TWO-ACTION EXAMPLE (SECTION 4.1), CONTINUED. Let the prior be $p_0 \in (0, 2/3)$. If the graduality constraint is tight (Figure 1b, p. 139), then less information is provided than with no constraint: $p^- = 0 = P^-$ and $p^+ = 2/3 < 1 = P^+$. For a sufficiently slack constraint (Figure 1a), there is no gap.¹⁸

Although it is intuitive that long-run induced beliefs should converge to the unconstrained ones as $\bar{\lambda} \rightarrow \infty$, the result is not obvious. To see why, observe that the analogous result for the “slow limit” $\bar{\lambda} \rightarrow 0$ is false. The natural static benchmark here is the trivial model in which no information is available, so that the belief stays put at the prior p_0 . It is *not* true that p^- and p^+ converge to p_0 as $\bar{\lambda} \rightarrow 0$: indeed, in the two-action example, the (uniquely optimal) long-run induced beliefs are $\{p^-, p^+\} = \{0, 2/3\}$ for every low value of $\bar{\lambda} > 0$.¹⁹

¹⁸Examples also exist in which there is a gap for all $\bar{\lambda} > 0$, which closes only in the limit.

¹⁹The key formal difference between the two limits is that the convergence of v to $\text{cav } u$ as $\bar{\lambda} \rightarrow \infty$ is uniform by Corollary 1 (p. 143), whereas the convergence of v to u as $\bar{\lambda} \rightarrow 0$ is merely pointwise (unless u is continuous).

PROOF OF PROPOSITION 3. To emphasize dependence on parameters, write $v_{\bar{\lambda}}$, $p_{\bar{\lambda}}^-$ and $p_{\bar{\lambda}}^+$ for the value and long-run beliefs. Let $\{v_{\bar{\lambda}} > u\}$ be the set of beliefs at which the strategy Λ^* in Corollary 2 provides information.

For the first part, fix a $\bar{\lambda} > 0$. Clearly, $P^- \leq P^+$. By Proposition 1, $u(p) < v_{\bar{\lambda}}(p)$ implies $u(p) < (\text{cav } u)(p)$, so that $\{v_{\bar{\lambda}} > u\} \subseteq \{\text{cav } u > u\}$. Therefore,

$$P^- = \sup([0, p_0] \setminus \{\text{cav } u > u\}) \leq \sup([0, p_0] \setminus \{v_{\bar{\lambda}} > u\}) = p_{\bar{\lambda}}^-,$$

and similarly $p_{\bar{\lambda}}^+ \leq P^+$.

Now for the second part. Since $p_{\bar{\lambda}}^-$ decreases monotonically as $\bar{\lambda}$ increases by Proposition 2, and lives in the compact set $[P^-, p_0]$, it converges to some limit $p_{\infty}^- \in [P^-, p_0]$. We wish to show that $p_{\infty}^- = P^-$, so suppose toward a contradiction that $p_{\infty}^- > P^-$. On the one hand,

$$(\text{cav } u)(p_{\bar{\lambda}}^-) \rightarrow (\text{cav } u)(p_{\infty}^-) > u(p_{\infty}^-) \tag{1}$$

by continuity of $\text{cav } u$ and $p_{\infty}^- > P^-$ (recall the definition of P^-). On the other hand, (recalling the definition of $p_{\bar{\lambda}}^-$)

$$|(\text{cav } u)(p_{\bar{\lambda}}^-) - u(p_{\bar{\lambda}}^-)| = |(\text{cav } u)(p_{\bar{\lambda}}^-) - v_{\bar{\lambda}}(p_{\bar{\lambda}}^-)| \rightarrow 0$$

since $v_{\bar{\lambda}}$ converges *uniformly* to $\text{cav } u$ by Corollary 1 (p. 143), by a standard property of uniform convergence (e.g., Theorem 7.11 in Rudin (1976)). It follows by upper semi-continuity of u that

$$(\text{cav } u)(p_{\bar{\lambda}}^-) \rightarrow \lim_{\bar{\lambda} \rightarrow \infty} u(p_{\bar{\lambda}}^-) \leq u(p_{\infty}^-)$$

a contradiction with (1). A similar argument shows that $p_{\bar{\lambda}}^+ \rightarrow P^+$. □

5.4 The role of discontinuities

The **two-action example** suggests that discontinuities in the sender's flow payoff u generate a particularly stark dampening effect of the graduality constraint on information provision. The following proposition expresses this idea: any discontinuity, when paired with a sufficiently tight constraint, will preclude full information from being provided in equilibrium.

PROPOSITION 4. *Fix a prior p_0 . If the sender's flow payoff u has an interior discontinuity, then $\{p^-, p^+\} \neq \{0, 1\}$ provided $\bar{\lambda} > 0$ is small enough.*

PROOF. Let u be discontinuous at $p \in (0, 1)$. Since u is piecewise continuous, it is continuous on a left-neighborhood (p', p) of p and on a right-neighborhood (p, p'') .

Fix a value of $\bar{\lambda} > 0$ such that $\{p^-, p^+\} = \{0, 1\}$. (Recall that p^-, p^+ and v all depend on $\bar{\lambda}$, albeit our notation leaves this implicit.) Then $v > u$ on $(0, 1)$ by definition of p^- and p^+ . It follows by Proposition 1 that on the intervals (p', p) and (p, p'') , the value v (is everywhere locally strictly convex, and thus) satisfies equation (∂) on page 141. By inspection, this implies that v is strictly convex on (p', p) and (p, p'') . And since v pastes

smoothly at p , it follows v is convex on all of (p', p'') . To summarize: if $\{p^-, p^+\} = \{0, 1\}$, then v must be a convex majorant of u on (p', p'') .

It thus suffices to show that v fails to be a convex majorant of u on (p', p'') if $\bar{\lambda} > 0$ is small enough. And that follows from the pointwise (monotone) convergence of v to u as $\bar{\lambda} \downarrow 0$ and the fact that u is discontinuous at p . \square

6. BELIEF-BASED CONFLICT OF INTEREST

In this section, we drop the common-prior assumption. The resulting differences in beliefs may engender a conflict of interest even if the players' preferences are aligned conditional on the state of the world.

Such *belief-based* conflicts are pervasive. They arise, for instance, where preexisting contracts have largely aligned all parties' interests conditional on the state, as in many organizational settings. Disagreements about the best course of action remain ubiquitous in such environments, but originate in agents' differing assessments of the evidence.

Some policy problems also feature well-aligned preferences. For example, disagreements over actions to mitigate climate change usually play out as debates about the likely extent of (anthropogenic) global warming, rather than about what policies are desirable in any given physical scenario.

In this section, we show that a purely belief-based conflict may preclude full information being provided in slow-persuasion equilibrium, contrary to the unconstrained-persuasion model. We further argue that belief disagreement harms the sender: the greater the prior gap, the lower her value.

We extend our preceding results to heterogeneous priors (maintaining arbitrary preferences) in the next section, then specialize in Section 6.2 to the case of ex post aligned preferences.

6.1 Equilibrium characterization

The model is as in Section 2, except that the priors $p_0, p_{D,0} \in (0, 1)$ of the sender and decision maker may differ. The priors are commonly known (i.e., the players agree to disagree). Write p_t and $p_{D,t}$ for the sender's and decision-maker's beliefs at time t .

The model remains tractable because we need not keep track of the decision-maker's belief, as it may be backed out from the sender's belief and the priors via Bayes's rule.

OBSERVATION 2. The decision-maker's time- t belief is $p_{D,t} = \phi(p_t, p_0, p_{D,0})$, where

$$\phi(p, p_0, p_{D,0}) := \frac{p}{p + (1-p)\frac{p_0}{1-p_0} \bigg/ \frac{p_{D,0}}{1-p_{D,0}}}.$$

In light of Observation 2, MPEs have all of the same properties as in the common-prior case. For the same reason as in Observation 1 (Section 3), a regular Markov strategy

$A : [0, 1] \rightarrow \Delta(\mathcal{A})$ is part of a MPE if and only if it is myopic. Given a regular Markov strategy A , the sender's induced flow payoff is now

$$u(p) := \int_{\mathcal{A}} f_S(a, p) A(da | \phi(p, p_0, p_{D,0}))$$

since the decision-maker's belief is $\phi(p, p_0, p_{D,0})$ when the sender's is p . (Note that u depends on the priors.) It remains true that u is piecewise continuous and upper semi-continuous. Given u , the sender's best-reply problem is unchanged, noting again that p_t is the *sender's* belief.

All of our preceding results therefore remain valid: the sender's value function is a generalized concave envelope (Proposition 1), she provides more information the faster evidence accumulates and the more patient she is (Proposition 2), and her information provision is well approximated by unconstrained persuasion when information can be generated quickly (Proposition 3).

6.2 Belief-based conflict of interest

We now specialize the model of the previous section by assuming that interests are aligned ex post: $f_S = f_D = f$. (In the special case of expected utility, this is equivalent to players having the same preferences conditional on the state θ .) Whatever conflict remains arises from the prior disagreement alone.

In the unconstrained-persuasion benchmark, the sender provides full information when the conflict is purely belief-based (e.g., [Alonso and Câmara \(2016\)](#)). This follows from two observations. First, the sender would be better-off if she were in charge of choosing the action, and were this the case, then her payoff would be highest under full information.²⁰ Second, the sender attains this upper bound on her payoff by providing full information: the decision maker then chooses as the sender would, since the players' posteriors always agree (they both assign probability 1 to the true state).

When persuasion is constrained to be slow, this argument remains approximately valid if the graduality constraint is loose. In particular, near-full information is provided in equilibrium by Proposition 3.

Otherwise, the argument breaks down. The long-run benefit of providing full information must then be weighed against its potential cost over the short run. In fact, outside of trivial cases, the cost is sure to dominate whenever the constraint is tight enough. To formalize this, call an action $a \in \mathcal{A}$ *redundant* if and only if it is never strictly optimal, that is,

$$f(a, p) \leq \max_{a' \in \mathcal{A} \setminus \{a\}} f(a', p) \quad \text{for every } p \in [0, 1].$$

²⁰This is true for expected-utility preferences f , and more generally for preferences f that value information in the sense that $p \mapsto f(a, p)$ is convex for each $a \in \mathcal{A}$.

PROPOSITION 5. *If the actions \mathcal{A} can be totally ordered so that $f = f_S = f_D$ is strictly single-crossing,²¹ and at least two actions are nonredundant, then for any fixed priors $p_0 \neq p_{D,0}$, we have $\{p^-, p^+\} \neq \{0, 1\}$ provided $\bar{\lambda} > 0$ is small enough.*

Any expected-utility preference f satisfies strict single-crossing.²² We prove Proposition 5 in Appendix F.

TWO-ACTION EXAMPLE (SECTION 4.1), CONTINUED. Modify the example so that the decision maker shares the sender's preference f_S , but has a prior $p_{D,0} = 1/5$ different from the sender's $p_0 = 1/2$. The decision-maker's (unique regular) Markov strategy remains $A = \mathbf{1}_{[2/3, 1]}$,²³ so our preceding analysis remains valid. Thus when evidence accumulates slowly (Figure 1b, p. 139), the sender induces the imperfectly-informative long-run beliefs $\{p^-, p^+\} = \{0, 2/3\}$.

In unconstrained persuasion, the sender's value is invariant to the decision-maker's prior. This is because she provides full information, leading both players' posterior beliefs always to agree (assigning probability 1 to the true state), so that the decision maker chooses actions just like the sender would. By contrast, prior disagreement harms the sender in slow persuasion, because it induces a conflict of interest. Under natural conditions, the sender is better off the smaller the disagreement.

PROPOSITION 6. *If the actions \mathcal{A} can be totally ordered so that $f = f_S = f_D$ is strictly single-crossing, then for any fixed prior p_0 , $v(p_0)$ increases pointwise as $p_{D,0} < p_0$ increases toward p_0 (as $p_{D,0} > p_0$ decreases toward p_0).*

The proof is in Appendix G.

APPENDIX A: PROOF OF THEOREM 1 (P. 140)

The boundary condition $v = u$ on $\{0, 1\}$ holds by inspection of the sender's best-reply problem (BRP), since these are absorbing states. For the remainder, let v_\star and v^\star be the lower and upper semicontinuous envelopes of v . In Theorem 1(a), we show that piecewise continuity of u suffices for v_\star to be a viscosity supersolution. In Theorem 1(b), we prove that v^\star is a viscosity subsolution, without requiring piecewise continuity. In both cases, we adapt a standard argument. Finally, we establish in Theorem 1(c) that v

²¹That is, there is a total order \geq on \mathcal{A} such that for $a' > a$ and $p' > p$, $f(a', p) \geq f(a, p)$ implies $f(a', p') > f(a, p')$. Totality can be weakened: it is enough that \geq be a partial order such that (\mathcal{A}, \geq) is a lattice and $f(\cdot, p)$ is quasisupermodular for each $p \in [0, 1]$.

²²An expected-utility f has $f(a, p) = (1 - p)u_0(a) + pu_1(a)$ for some $u_0, u_1 : \mathcal{A} \rightarrow \mathbf{R}$. Define \geq by $a' \geq a$ if and only if $u_1(a') - u_0(a') \geq u_1(a) - u_0(a)$. Then $f(a', p) \geq f(a, p)$ implies $f(a', p') > f(a, p')$ for any $a' > a$ and $p' > p$. This \geq is a total order (in particular, antisymmetric) once \mathcal{A} is pruned of strictly dominated actions and duplicates.

²³In terms of the sender's belief p , the decision-maker's payoff from action $a = 1$ is $3\phi(p, 1/2, 1/5) - 1 = 3\frac{p}{4-3p} - 1$, which exceeds zero (the payoff of $a = 0$) if and only if $p \geq 2/3$.

is continuous, so that $v_\star = v^\star = v$.²⁴ We will make occasional use of the fact that $v \geq u$, which holds since the value u is attainable (by never providing any information).

First, the supersolution property of v_\star , which relies on piecewise continuity of u :

THEOREM 1(a). *If u is piecewise continuous, then v_\star is a viscosity supersolution of (HJB).*

PROOF. We follow the standard argument (e.g., Pham (2009), Proposition 4.3.1), which assumes that u is continuous. We sketch the steps that are unchanged, and emphasize the juncture at which a new argument is needed to accommodate merely piecewise continuity of u .

Take any $p \in (0, 1)$ and any twice continuously differentiable $\phi : (0, 1) \rightarrow \mathbf{R}$ such that $v_\star - \phi$ has a local minimum at p . In light of Remark 1, we may assume without loss of generality that $v_\star(p) - \phi(p) = 0$ and that $v_\star - \phi \geq 0$ (i.e., p is a *global* minimum of $v_\star - \phi$.) We wish to show that

$$v_\star(p) \geq u_\star(p) + \bar{\lambda} \frac{p^2(1-p)^2}{2r\sigma^2} \max\{0, \phi''(p)\}.$$

We have $v_\star \geq u_\star$ since $v \geq u$, so what must be shown is that

$$v_\star(p) \geq u_\star(p) + \bar{\lambda} \frac{p^2(1-p)^2}{2r\sigma^2} \phi''(p). \tag{2}$$

By definition of v_\star and since ϕ is continuous with $v_\star(p) - \phi(p) = 0$, we may find a sequence $(p_n)_{n \in \mathbf{N}}$ in $(0, 1)$ converging to p along which

$$\gamma_n := v(p_n) - \phi(p_n)$$

vanishes. Choose any strictly positive sequence $(h_n)_{n \in \mathbf{N}}$ in \mathbf{R} such that $h_n \rightarrow 0$ and $\gamma_n/h_n \rightarrow 0$ as $n \rightarrow \infty$.

Consider the “full tilt forever” strategy, which sets $\lambda_t = \bar{\lambda}$ a.s. no matter what happens. Write P_s^n for the induced (belief) process when the initial condition is $P_0 = p_n$. Since $v(p_0)$ is the optimal value, it must exceed the expected discounted payoff obtained by using the “full tilt forever” control process until time h_n , then reverting to optimal behavior:

$$v(p_n) \geq \mathbf{E} \left(r \int_0^{h_n} e^{-rt} u(P_t^n) dt + e^{-rh_n} v(P_{h_n}^n) \right) \quad \forall n \in \mathbf{N}.$$

(The integrand on the right-hand side is in fact measurable, so that the expectation is well-defined; see, e.g., Pham (2009, Theorem 3.3.1).) Using $v - \phi \geq v_\star - \phi \geq 0$ and the definition of γ_n yields

$$\phi(p_n) + \gamma_n \geq \mathbf{E} \left(r \int_0^{h_n} e^{-rt} u(P_t^n) dt + e^{-rh_n} \phi(P_{h_n}^n) \right) \quad \forall n \in \mathbf{N}. \tag{3}$$

²⁴As we explain in supplemental Appendix K.2, it is typical to replace the last step with an appeal to a comparison principle. Since we are not aware of a comparison principle that requires only piecewise continuity of u , we prove continuity directly instead.

Since ϕ is twice continuously differentiable and P_t^n evolves as

$$dP_t^n = \sqrt{\lambda} \frac{P_t^n(1 - P_t^n)}{\sigma} dB_t$$

where $(B_t)_{t \in \mathbf{R}_+}$ is a standard Brownian motion, we may apply Itô's lemma to the process $(e^{-rt} \phi(P_t^n))_{t \in \mathbf{R}_+}$ to obtain, for each $n \in \mathbf{N}$,

$$e^{-rh_n} \phi(P_{h_n}^n) = \phi(p_n) - r \int_0^{h_n} e^{-rt} \phi(P_t^n) dt + \frac{1}{2} \int_0^{h_n} \left(\sqrt{\lambda} \frac{P_t^n(1 - P_t^n)}{\sigma} \right)^2 e^{-rt} \phi''(P_t^n) dt.$$

Substituting in (3) and rearranging slightly yields

$$\begin{aligned} \frac{\gamma_n}{h_n} &\geq \mathbf{E} \left(r \frac{1}{h_n} \int_0^{h_n} e^{-rt} u(P_t^n) dt - r \frac{1}{h_n} \int_0^{h_n} e^{-rt} \phi(P_t^n) dt \right. \\ &\quad \left. + r \frac{1}{h_n} \int_0^{h_n} e^{-rt} \lambda \frac{(P_t^n)^2(1 - P_t^n)^2}{2r\sigma^2} \phi''(P_t^n) dt \right) \quad \forall n \in \mathbf{N}. \end{aligned} \tag{4}$$

We will obtain (2) as the limit of this inequality as $n \rightarrow \infty$.

Since ϕ is twice continuously differentiable and the sample paths of $(P_t^n)_{t \in \mathbf{R}_+}$ are continuous a.s., the mean-value theorem may be applied path-by-path to the second and third terms inside the expectation in (4) to conclude that they converge a.s. to, respectively, $-\phi(p)$ and

$$\lambda \frac{p^2(1 - p)^2}{2r\sigma^2} \phi''(p).$$

It remains to show that the first term converges a.s. to a limit that exceeds $u_*(p)$. If p is a continuity point of u , then u is continuous on a neighborhood of p by piecewise continuity, so the same mean-value theorem argument implies that the first term converges a.s. to $u(p) \geq u_*(p)$, as desired.

Suppose instead that p is a discontinuity point of u ; this requires an additional argument relative to the standard proof. By piecewise continuity, u is continuous on a left- and a right-neighborhood of p . Thus for any sufficiently small $\varepsilon > 0$, we may apply the mean-value theorem on either side of p to obtain the existence of a $p_-^\varepsilon \in (p - \varepsilon, p)$ and a $p_+^\varepsilon \in (p, p + \varepsilon)$ such that

$$\frac{1}{\varepsilon} \int_{p-\varepsilon}^p u = u(p_-^\varepsilon) \quad \text{and} \quad \frac{1}{\varepsilon} \int_p^{p+\varepsilon} u = u(p_+^\varepsilon),$$

so that

$$\frac{1}{2\varepsilon} \int_{p-\varepsilon}^{p+\varepsilon} u \geq \min\{u(p_-^\varepsilon), u(p_+^\varepsilon)\}.$$

The left-hand side converges as $\varepsilon \downarrow 0$, and the right-hand side converges to $\min\{u(p-), u(p+)\}$. Thus

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{2\varepsilon} \int_{p-\varepsilon}^{p+\varepsilon} u \geq \min\{u(p-), u(p+)\} \geq u_*(p).$$

As with the second and third terms in (4), we may apply this argument to a.e. path of the first term since a.e. sample path of $(P_t^n)_{t \in \mathbf{R}_+}$ is continuous. Thus the first term in (4) converges a.s. to a limit that exceeds $u_*(p)$.

Next, observe that all three terms inside the expectation in (4) are bounded off D by a constant independent of n because ϕ , ϕ'' and u are continuous off D . Furthermore, the set D is null under the occupancy measure of $(P_t^n)_{t \in \mathbf{R}_+}$, for every $n \in \mathbf{N}$.²⁵ It follows by the bounded convergence theorem that the right-hand side of (4) converges to a limit exceeding

$$u_*(p) - \phi(p) + \bar{\lambda} \frac{p^2(1-p)^2}{2r\sigma^2} \phi''(p).$$

The left-hand side of (4) vanishes by construction of $(h_n)_{n \in \mathbf{N}}$. Thus

$$0 \geq u_*(p) - \phi(p) + \bar{\lambda} \frac{p^2(1-p)^2}{2r\sigma^2} \phi''(p).$$

Using $\phi(p) = v_*(p)$ and rearranging yields the desired inequality (2). □

Unlike the supersolution property of v_* , the subsolution property of v^* holds for general u .

THEOREM 1 (b). *v^* is a viscosity subsolution of (HJB).*

PROOF. Again, we follow the standard line of reasoning (e.g., Pham (2009, Proposition 4.3.2), noting the errata (Pham (2012))) for the case in which u is continuous. Where continuity of u is usually invoked, we shall make do with the (definitional) upper semi-continuity of u^* .

Take any $p \in (0, 1)$ and any twice continuously differentiable $\phi : (0, 1) \rightarrow \mathbf{R}$ such that $\phi - v^*$ has a local minimum at p . By Remark 1, we may assume without loss that $\phi(p) - v^*(p) = 0$. Suppose that the viscosity subsolution property fails at p :

$$\phi(p) = v^*(p) > u^*(p) + \bar{\lambda} \frac{p^2(1-p)^2}{2r\sigma^2} \max\{0, \phi''(p)\}.$$

We shall derive a contradiction.

By Remark 1 again, we may assume that $\phi - v^*$ has a strict global minimum at p . For $\eta > 0$, write

$$B_\eta := \{q \in (0, 1) : |q - p| < \eta\}$$

for the open ball of radius η around p , and ∂B_η for its boundary. Define

$$k_\eta := \min_{q \in \partial B_\eta} |\phi(q) - v^*(q)|,$$

²⁵This is because the occupancy measure is absolutely continuous with respect to Lebesgue measure, and D is Lebesgue-null since it is discrete.

noting that it is strictly positive for $\eta > 0$ because the minimum of $\phi - v^*$ at p is strict. Since ϕ and ϕ'' are continuous and u^* is upper semicontinuous, we may find an $\eta > 0$ and an $\varepsilon \in (0, k_\eta]$ small enough that

$$\phi(q) \geq u^*(q) + \bar{\lambda} \frac{q^2(1-q)^2}{2r\sigma^2} \max\{0, \phi''(q)\} + \varepsilon \quad \text{for all } q \in B_\eta. \tag{5}$$

By definition of v^* and since ϕ is continuous with $\phi(p) - v^*(p) = 0$, we may find a sequence $(p_n)_{n \in \mathbf{N}}$ in B_η converging to p along which

$$\gamma_n := \phi(p_n) - v(p_n)$$

vanishes. Let $(\lambda_t^n)_{t \in \mathbf{R}_+}$ be an $\varepsilon/2$ -best reply in the sender's best-reply problem with prior $p_0 = p_n$, and write $(P_t^n)_{t \in \mathbf{R}_+}$ for the belief process induced by this strategy. Let τ_n be the first exit time of $(P_t^n)_{t \in \mathbf{R}_+}$ from B_η . Using $(\lambda_t^n)_{t \in \mathbf{R}_+}$ only until time τ_n and then reverting to optimal behavior is even better, so certainly attains value at least $v(p_n) - \varepsilon/2$:

$$v(p_n) - \frac{\varepsilon}{2} \leq \mathbf{E} \left(r \int_0^{\tau_n} e^{-rt} u(P_t^n) dt + e^{-r\tau_n} v(P_{\tau_n}^n) \right).$$

(It is nontrivial but true that the integrand is measurable, so that the expectation is well-defined; see, e.g., [Pham \(2009, Theorem 3.3.1\)](#).) Subtracting $\phi(p_n)$ from both sides and using the fact that

$$\phi - v \geq \phi - v^* \geq k_\eta \geq \varepsilon \quad \text{on } \partial B_\eta$$

yields

$$-\gamma_n - \frac{\varepsilon}{2} \leq \mathbf{E} \left(-e^{-r\tau_n} \varepsilon + r \int_0^{\tau_n} e^{-rt} u(P_t^n) dt + e^{-r\tau_n} \phi(P_{\tau_n}^n) - \phi(p_n) \right). \tag{6}$$

Since ϕ is twice continuously differentiable and P_t^n evolves as

$$dP_t^n = \lambda_t^n \frac{P_t^n(1-P_t^n)}{\sigma} dB_t$$

where $(B_t)_{t \in \mathbf{R}_+}$ is a standard Brownian motion, we may apply Itô's lemma to the process $(e^{-rt} \phi(P_t^n))_{t \in \mathbf{R}_+}$ to obtain, for each $n \in \mathbf{N}$,

$$e^{-r\tau_n} \phi(P_{\tau_n}^n) = \phi(p_n) - r \int_0^{\tau_n} e^{-rt} \phi(P_t^n) dt + \frac{1}{2} \int_0^{\tau_n} \left(\sqrt{\lambda_t^n} \frac{P_t^n(1-P_t^n)}{\sigma} \right)^2 e^{-rt} \phi''(P_t^n) dt.$$

Substituting in (6) and using (5) yields

$$\begin{aligned} -\gamma_n - \frac{\varepsilon}{2} &\leq \mathbf{E} \left(-e^{-r\tau_n} \varepsilon + r \int_0^{\tau_n} e^{-rt} \left[-\phi(P_t^n) + u(P_t^n) + \lambda_t^n \frac{(P_t^n)^2(1-P_t^n)^2}{2r\sigma^2} \phi''(P_t^n) \right] dt \right) \\ &\leq \mathbf{E} \left(-e^{-r\tau_n} \varepsilon + r \int_0^{\tau_n} e^{-rt} \left[-\phi(P_t^n) + u(P_t^n) \right] dt \right) \end{aligned}$$

$$\begin{aligned}
 & + \bar{\lambda} \frac{(P_t^n)^2 (1 - P_t^n)^2}{2r\sigma^2} \max\{0, \phi''(P_t^n)\} \Big] dt \Big) \\
 & \leq \mathbf{E} \left(-e^{-r\tau_n} \varepsilon + r \int_0^{\tau_n} e^{-rt} (-\varepsilon) dt \right) = -\varepsilon.
 \end{aligned}$$

Since γ_n vanishes as $n \rightarrow \infty$, we have the contradiction $-\varepsilon/2 \leq -\varepsilon$. □

It remains only to show that $v = v_\star = v^\star$, that is, that v is continuous.

THEOREM 1(c). *If u is piecewise continuous, then v is continuous.*

PROOF. We deal separately with $\{0, 1\}$ and $(0, 1)$. Begin with continuity at 0; the argument at 1 is analogous. Take a sequence $(p_n)_{n \in \mathbf{N}}$ in $(0, 1)$ converging to 0; we will show that $v(p_n) \rightarrow u(0) = v(0)$.

At each $n \in \mathbf{N}$, consider the auxiliary problem in which the sender may choose any process $(p_t)_{t \in \mathbf{R}_+}$ satisfying $\mathbf{E}(p_t) = p_n$ for every $t \in \mathbf{R}_+$. The value $V(p_n)$ of this problem must exceed $v(p_n)$ since any belief process the sender can induce in her best-reply problem is available in the auxiliary problem. And we have $V(p_n) \leq (\text{cav } u)(p_n)$ since the auxiliary problem may be broken down into a sequence of independent static persuasion problems, in each of which the optimal value is at most $(\text{cav } u)(p_n)$. Thus we have

$$u(p_n) \leq v(p_n) \leq V(p_n) \leq (\text{cav } u)(p_n) \quad \text{for every } n \in \mathbf{N}.$$

As $n \rightarrow \infty$, $u(p_n) \rightarrow u(0)$ since u is continuous at 0 by piecewise continuity, and $(\text{cav } u)(p_n) \rightarrow u(0)$ since $\text{cav } u$ is continuous and $(\text{cav } u)(0) = u(0)$ because u is continuous at 0. It follows that $u(0) \leq \lim_{n \rightarrow \infty} v(p_n) \leq u(0)$.

To establish that v is continuous on $(0, 1)$, fix a $p \in (0, 1)$. It suffices to show that $\underline{v} \geq \bar{v}$, where

$$\underline{v} := \liminf_{q \rightarrow p} v(q) \quad \text{and} \quad \bar{v} := \limsup_{q \rightarrow p} v(q).$$

By construction, there exist sequences $(\underline{p}_n)_{n \in \mathbf{N}}$ and $(\bar{p}_n)_{n \in \mathbf{N}}$ converging to p along which

$$v(\underline{p}_n) \rightarrow \underline{v} \quad \text{and} \quad v(\bar{p}_n) \rightarrow \bar{v} \quad \text{as } n \rightarrow \infty.$$

Note that v is bounded since u is (being piecewise continuous).

Suppose first that these sequences may both be chosen to lie in $(0, p)$; the case in which they may be chosen to lie in $(p, 1)$ is analogous. Then we may choose them so that $\bar{p}_{n-1} \leq \underline{p}_n \leq \bar{p}_n$ for every $n \in \mathbf{N}$, where $\bar{p}_0 := 0$ by convention. For the sender's best-reply problem with prior $p_0 = \underline{p}_n$, consider a strategy that sets $\lambda = \bar{\lambda}$ while $p_t \in (\bar{p}_{n-1}, \bar{p}_n)$ and $\lambda = 0$ otherwise, and write $(P_t^n)_{t \in \mathbf{R}_+}$ for the induced belief process. Write τ_n for the first time that $(P_t^n)_{t \in \mathbf{R}_+}$ hits $\{\bar{p}_{n-1}, \bar{p}_n\}$. Since this strategy cannot be better than optimal, we have

$$v(\underline{p}_n) \geq \mathbf{E} \left(r \int_0^{\tau_n} e^{-rt} u(P_t^n) dt + e^{-r\tau_n} v(P_{\tau_n}^n) \right) \quad \text{for each } n \in \mathbf{N}.$$

(A standard result ensures that the right-hand-side integrand is measurable, making the expectation well-defined; e.g., [Pham \(2009, Theorem 3.3.1\)](#).) The left-hand side converges to \underline{v} as $n \rightarrow \infty$. The hitting time τ_n vanishes a.s., and $v(P_{\tau_n}^n)$ converges a.s. to \bar{v} . Furthermore, u and v are bounded by piecewise continuity.²⁶ Hence, the right-hand side converges to \bar{v} by the bounded convergence theorem, so that $\underline{v} \geq \bar{v}$.

Suppose instead that the sequences cannot be chosen to lie on the same side of p —without loss of generality, $\underline{p}_n < p < \bar{p}_n$ for every $n \in \mathbf{N}$. For the sender’s problem with $p_0 = \underline{p}_n$, consider a strategy that sets $\lambda = \bar{\lambda}$ while $p_t \in (\underline{p}_{n-1}, \bar{p}_n)$ and $\lambda = 0$ otherwise, and write $(P_t^n)_{t \in \mathbf{R}_+}$ for the induced belief process. Let τ_n be the first time that $(P_t^n)_{t \in \mathbf{R}_+}$ hits $\{\underline{p}_{n-1}, \bar{p}_n\}$. The optimal value must exceed the value from using this strategy:

$$v(\underline{p}_n) \geq \mathbf{E} \left(r \int_0^{\tau_n} e^{-rt} u(P_t^n) dt + e^{-r\tau_n} v(P_{\tau_n}^n) \right) \quad \text{for each } n \in \mathbf{N}. \tag{7}$$

(Again, the right-hand side is well defined.) The left-hand side converges to \underline{v} as $n \rightarrow \infty$. The hitting time τ_n vanishes a.s. since $|\bar{p}_n - \underline{p}_n| \rightarrow 0$. For each $n \in \mathbf{N}$, we have

$$\mathbf{E}(v(P_{\tau_n}^n)) = \gamma_n v(\underline{p}_{n-1}) + (1 - \gamma_n) v(\bar{p}_n)$$

for some $\gamma_n \in (0, 1)$, and the sequences $(\underline{p}_n)_{n \in \mathbf{N}}$ and $(\bar{p}_n)_{n \in \mathbf{N}}$ may be chosen so that $(\gamma_n)_{n \in \mathbf{N}}$ converges to some $\gamma < 1$. Thus, applying the bounded convergence theorem (using the boundedness of u and v) to (7) yields $\underline{v} \geq \gamma \underline{v} + (1 - \gamma) \bar{v}$, which is equivalent to $\underline{v} \geq \bar{v}$ since $\gamma < 1$. □

APPENDIX B: PROOF OF LEMMA 1 (P. 142)

Take $p \in (0, 1) \setminus C$, and suppose toward a contradiction that $v(p) > u(p)$. Since v is continuous and u is upper semicontinuous, we have $v > u$ on an open neighborhood N of p . We will derive a contradiction assuming that $p \notin D$. The result for $p \in D$ then follows from the observation that if $v(p) > u(p)$ for $p \in D$, then since D is discrete, the neighborhood N also contains a $p' \notin D$ at which $v(p') > u(p')$.

We may choose N to not intersect D since the latter is discrete. By [Theorem 1 \(p. 140\)](#) and the fact that $v > u$ on N , v is a viscosity solution of

$$w(p) = u(p) + \bar{\lambda} \frac{p^2(1-p)^2}{2r\sigma^2} w''(p) \tag{8}$$

on N . Observe that u is continuous on N .

We show (constructively) in supplemental Appendix H that (8) has a classical (hence viscosity) solution w^\dagger on N , which satisfies the (Dirichlet) boundary condition $w^\dagger = v$ on ∂N . By the comparison principle (e.g., [Theorem 3.3. in Crandall, Ishii, and Lions \(1992\)](#)), w^\dagger is the unique viscosity solution of (8) on N satisfying this boundary condition. It follows that $v = w^\dagger$.

Since $v > u$ on N , (8) requires that $v'' > 0$ on N . But then $N \subseteq C$, contradicting the supposition that p lies in $(0, 1) \setminus C$. □

²⁶ v is bounded below by u , and is bounded above by $V \leq \text{cav } u$, where V is the value of the auxiliary problem in the first part of the proof.

APPENDIX C: PROOF OF LEMMA 2 (P. 142)

Since a differentiable locally convex function is continuously differentiable (see, e.g., Theorem 24.1 in Rockafellar (1970)), it suffices to show that v is differentiable on C . By local convexity, the left- and right-hand derivatives v'_- and v'_+ of v exist on C and satisfy $v'_- \leq v'_+$ (again, see Theorem 24.1 in Rockafellar (1970)). We must show that $v'_- = v'_+$.

To that end, take a $p \in C$, and suppose toward a contradiction that $v'_-(p) < v'_+(p)$. (That is, there is a convex kink at p .) Then for any $k > 0$, we may find a twice continuously differentiable $\phi : (0, 1) \rightarrow \mathbf{R}$ with $\phi''(p) = k$ such that $v - \phi$ is locally minimized at p .²⁷ Since v is a viscosity supersolution of (HJB) by Theorem 1, it follows that

$$v(p) \geq u_*(p) + \bar{\lambda} \frac{p^2(1-p)^2}{2r\sigma^2} k$$

for any $k > 0$. For large enough k , this contradicts the previously established fact that $v(p) \leq (\text{cav } u)(p)$. \square

APPENDIX D: PROOF OF LEMMA 3 (P. 143)

Since v is locally convex on $C \setminus D$, v'' is nonnegative whenever it exists. Thus by Theorem 1 (p. 140), v is a viscosity solution of the differential equation

$$w(p) = u(p) + \bar{\lambda} \frac{p^2(1-p)^2}{2r\sigma^2} w''(p). \tag{9}$$

on $C \setminus D$. Note that u is continuous on $C \setminus D$.

In supplemental Appendix H, we show constructively that (9) has a classical solution on $C \setminus D$ that can be extended to a continuous function $w^\dagger : C \rightarrow \mathbf{R}$ satisfying the (Dirichlet) boundary condition $w^\dagger = u$ on ∂C . By the comparison principle (e.g., Theorem 3.3. in Crandall, Ishii, and Lions (1992)), w^\dagger is the unique viscosity solution of (9) that satisfies this boundary condition. Since w^\dagger is twice continuously differentiable and satisfies (9) on $C \setminus D$, it suffices to show that $v = w^\dagger$ on $C \setminus D$.

Since v is locally convex on $C \setminus D$, it is twice differentiable a.e. on $C \setminus D$ by the Alexandrov theorem (e.g., Theorem A.2 in Crandall, Ishii, and Lions (1992)), so v'' exists on a dense subset B of $C \setminus D$. Being a derivative, v'' is continuous on a dense subset A of B .²⁸ It follows that v satisfies (9) on A . We have already shown that it satisfies the boundary condition $v = u$ on $\partial C \subseteq [0, 1] \setminus C$.

Since the solution w^\dagger is unique, v coincides with w^\dagger on A . Because A is dense in $C \setminus D$, $v|_A$ admits at most one continuous extension to $C \setminus D$. Since w^\dagger is continuous, it follows that $v = w^\dagger$ on $C \setminus D$. \square

²⁷For example, $\phi(q) := v(p) + \frac{1}{2}[v'_-(p) + v'_+(p)](q-p) + \frac{1}{2}k(q-p)^2$.

²⁸This is a consequence of the Baire category theorem; see Bruckner and Leonard (1966, p. 27).

APPENDIX E: PROOF OF COROLLARY 1 (P. 143)

To emphasize dependence on $\bar{\lambda}$, write the value as $v_{\bar{\lambda}}$.

For the first part, fix an arbitrary $p_0 \in [0, 1]$. Since increasing $\bar{\lambda}$ slackens the constraint $\lambda \leq \bar{\lambda}$ in the best-reply problem (BRP) in every period, it raises the value $v_{\bar{\lambda}}(p_0)$.

Now for the second part. We have established for every $p \in [0, 1]$ that the sequence $(v_{\bar{\lambda}}(p))_{\bar{\lambda} > 0}$ is increasing. Since it lives in the compact set $[u(p), (\text{cav } u)(p)]$ by Proposition 1, it must converge to some $v_{\infty}(p) \in [u(p), (\text{cav } u)(p)]$. In other words, $(v_{\bar{\lambda}})_{\bar{\lambda} > 0}$ converges pointwise to some function $v_{\infty} : [0, 1] \rightarrow \mathbf{R}$ satisfying $u \leq v_{\infty} \leq \text{cav } u$. We claim that $v_{\infty} = \text{cav } u$. Since $\text{cav } u$ is by definition the pointwise smallest concave majorant of u , it suffices to show that v_{∞} is concave.

To that end, take $0 < p' < p < p'' < 1$ in $[0, 1]$, and let $\gamma \in (0, 1)$ be such that $\gamma p' + (1 - \gamma)p'' = p$; we will establish that $\gamma v_{\infty}(p') + (1 - \gamma)v_{\infty}(p'') \leq v_{\infty}(p)$. By changing the units in which time is measured and adjusting discounting accordingly, the sender's best-reply problem (BRP) may be reformulated as

$$v_{\bar{\lambda}}(p) = \sup_{(\lambda'_t)_{t \in \mathbf{R}_+}} \mathbf{E} \left(r' \int_0^{\infty} e^{-r't} u(p_t) dt \right) \quad \text{subject to} \quad dp_t = \sqrt{\lambda'_t} p_t (1 - p_t) d\widehat{B}_t,$$

where $r' := r\sigma^2/\bar{\lambda}$, $\widehat{B}_t := \sqrt{r/r'} B_{tr'/r}$ is a standard Brownian motion, and

$$(\lambda'_t)_{t \in \mathbf{R}_+} := (\lambda_t/\bar{\lambda})_{t \in \mathbf{R}_+}$$

is chosen among all $[0, 1]$ -valued processes adapted to the filtration generated by $(p_t)_{t \in \mathbf{R}_+}$.

Consider the strategy that always sets $\lambda' = 1$, and let $(p_t)_{t \in \mathbf{R}_+}$ be the induced belief process. Write τ for the first time that $(p_t)_{t \in \mathbf{R}_+}$ hits $\{p', p''\}$. Following the proposed strategy until time τ and then behaving optimally cannot be better than optimal, so for every $\bar{\lambda} > 0$ we have

$$\mathbf{E} \left(r' \int_0^{\tau} e^{-r't} u(p_t) dt + e^{-r'\tau} v_{\bar{\lambda}}(p_{\tau}) \right) \leq v_{\bar{\lambda}}(p) \leq v_{\infty}(p),$$

where the second inequality holds since $v_{\bar{\lambda}}$ increases pointwise as $\bar{\lambda}$ increases. As $\bar{\lambda} \rightarrow \infty$, the first term inside the expectation on the left-hand side vanishes a.s., and the second term converges a.s. to $v_{\infty}(p_{\tau})$. Since both terms are bounded, the left-hand side converges to $\mathbf{E}(v_{\infty}(p_{\tau}))$ by the bounded convergence theorem. And we have

$$\mathbf{E}(v_{\infty}(p_{\tau})) = \gamma v_{\infty}(p') + (1 - \gamma)v_{\infty}(p'')$$

by the optional sampling theorem.²⁹

²⁹When $0 < p' < p'' < 1$, we have $\mathbf{E}(\tau) < \infty$, so the optional sampling theorem (Karatzas and Shreve (1991), ch. 1) yields $\mathbf{E}(p_{\tau}) = p$, whence $\mathbf{P}(p_{\tau} = p') = \gamma$ and $\mathbf{P}(p_{\tau} = p'') = 1 - \gamma$ by definition of γ . For the case in which $0 < p' < p'' = 1$ (the other cases are analogous), let τ_n be the first time that $(p_t)_{t \in \mathbf{R}_+}$ hits $\{p', 1 - 1/n\}$, for each $n \in \mathbf{N}$. Then $\mathbf{E}(\tau_n) < \infty$, so $\mathbf{E}(p_{\tau_n}) = p$ by the optional sampling theorem. Since $(p_{\tau_n})_{n \in \mathbf{N}}$ is bounded and converges a.s. to p_{τ} as $n \rightarrow \infty$, the bounded convergence theorem yields $\mathbf{E}(p_{\tau}) = p$.

We have proved that $v_{\bar{\lambda}}$ converges monotonically, pointwise, to $\text{cav } u$. Since $v_{\bar{\lambda}}$ and $\text{cav } u$ are continuous and defined on a compact domain, it follows by Dini's theorem (e.g., Theorem 7.13 in [Rudin \(1976\)](#)) that the convergence of $v_{\bar{\lambda}}$ to $\text{cav } u$ is uniform. \square

APPENDIX F: PROOF OF PROPOSITION 5 (P. 148)

By Proposition 4 (p. 146), it suffices to find an interior belief at which u is discontinuous. Since \mathcal{A} is finite with at least two nonredundant elements, $f(a, \cdot)$ is continuous for each $a \in \mathcal{A}$, and f is strictly single-crossing, the monotone selection theorem of [Milgrom and Shannon \(1994\)](#) yields a finite collection of two or more open intervals of $[0, 1]$, on each of which some action is strictly optimal, with no two intervals associated with the same action; furthermore, the closure of the union of these intervals is $[0, 1]$ itself. It follows that there is at least one interior belief $p \in (0, 1)$ that belongs to the boundary of two such open intervals. At this belief, we have $f(a, p) = f(a', p)$ for two actions $a \neq a'$, with a strictly optimal on a left-neighborhood of p and a' strictly optimal on a right-neighborhood. Note that $a' \succ a$, where \succeq is the total order on \mathcal{A} with respect to which f is strictly single-crossing.

Assume that $p_{D,0} > p_0$; the argument if $p_{D,0} < p_0$ is symmetric. Write $p_D \in (0, 1)$ for the decision-maker's belief when the sender has belief p , that is, $\phi(p_D, p_0, p_{D,0}) = p$. Since $f(a', p) = f(a, p)$, $a' \succ a$ and $p_D > p$, the strict single-crossing property of f yields that $f(a', p_D) > f(a, p_D)$. Since $u = f(a, \cdot)$ on a left-neighborhood of p_D and $u = f(a', \cdot)$ on a right-neighborhood, it follows that u is discontinuous at p_D . \square

APPENDIX G: PROOF OF PROPOSITION 6 (P. 149)

We shall use the following comparative statics lemma.

LEMMA 4. *Let (\mathcal{T}, \succsim) and (\mathcal{X}, \succeq) be totally ordered sets, let $f : \mathcal{T} \times \mathcal{X} \rightarrow \mathbf{R}$ be strictly single-crossing, and let $X : \mathcal{T} \rightarrow \Delta(\mathcal{X})$ satisfy $\text{supp } X(\cdot|t) \subseteq \arg \max_{\mathcal{X}} f(\cdot, t)$ for every $t \in \mathcal{T}$. Then for any $t_1 \succsim t_2 \succsim t_3$ in \mathcal{T} , we have*

$$\int_{\mathcal{X}} f(x, t_1) X(dx|t_2) \geq \int_{\mathcal{X}} f(x, t_1) X(dx|t_3) \quad \text{and}$$

$$\int_{\mathcal{X}} f(x, t_3) X(dx|t_2) \geq \int_{\mathcal{X}} f(x, t_3) X(dx|t_1).$$

PROOF. Fix arbitrary $x_2 \in \text{supp } X(\cdot|t_2)$ and $x_3 \in \text{supp } X(\cdot|t_3)$. Since f is strictly single-crossing, we have $x_2 \succeq x_3$ by the monotone selection theorem of [Milgrom and Shannon \(1994\)](#). As $f(x_2, t_2) \geq f(x_3, t_2)$ and $t_1 \succsim t_2$, single-crossing then yields $f(x_2, t_1) \geq f(x_3, t_1)$. Since $x_2 \in \text{supp } X(\cdot|t_2)$ and $x_3 \in \text{supp } X(\cdot|t_3)$ were arbitrary, it follows that $\int_{\mathcal{X}} f(x, t_1) X(dx|t_2) \geq \int_{\mathcal{X}} f(x, t_1) X(dx|t_3)$. The argument for the second inequality is analogous. \square

To prove the proposition, fix p_0 . Further fix $p_{D,0} < p'_{D,0} < p_0$; the other case is analogous. Write $A : [0, 1] \rightarrow \Delta(A)$ for the decision-maker's myopic and regular Markov strategy, let u, u' be the sender's induced flow payoffs under the two decision-maker priors $p_{D,0}, p'_{D,0}$, and write v, v' for the corresponding value functions of the sender.

Fix an arbitrary belief $p \in [0, 1]$ of the sender. The decision-maker's beliefs $p_D := \phi(p, p_0, p_{D,0})$ and $p'_D := \phi(p, p_0, p'_{D,0})$ satisfy $p_D \leq p'_D \leq p$. We furthermore have $\text{supp } A(\cdot|p') \subseteq \arg \max_{\mathcal{A}} f(\cdot, p')$ at every $p' \in [0, 1]$ because \mathcal{A} is myopic. Hence, since f is strictly single-crossing, Lemma 4 yields

$$u'(p) = \int_{\mathcal{A}} f(a, p) A(da|p'_D) \geq \int_{\mathcal{A}} f(a, p) A(da|p_D) = u(p).$$

Thus $u' \geq u$ since $p \in [0, 1]$ was arbitrary, whence $v'(p_0) \geq v(p_0)$. \square

REFERENCES

- Alonso, Ricardo and Odilon Câmara (2016), “Bayesian persuasion with heterogeneous priors.” *Journal of Economic Theory*, 165, 672–706. [131, 133, 148]
- Au, Pak Hung (2015), “Dynamic information disclosure.” *RAND Journal of Economics*, 46, 791–823. [132]
- Aumann, Robert J. and Michael B. Maschler (1995), *Repeated Games With Incomplete Information*. MIT Press, Cambridge, MA. [132]
- Ball, Ian (2022), “Dynamic information provision.” Working paper. [132]
- Barilla, César and Duarte Gonçalves (2022), “The dynamics of social instability.” Working paper. [133]
- Bergemann, Dirk and Stephen Morris (2019), “Information design: A unified perspective.” *Journal of Economic Literature*, 57, 44–95. [132]
- Billingsley, Patrick (1995), *Probability and Measure*, third edition. Wiley, New York, NY. [144]
- Bizzotto, Jacopo, Jesper Rüdiger, and Adrien Vigier (2021), “Dynamic persuasion with outside information.” *American Economic Journal: Microeconomics*, 13, 179–194. [132]
- Boleslavsky, Raphael and Kyungmin Kim (2021), “Bayesian persuasion and moral hazard.” Working paper. [132]
- Bolton, Patrick and Christopher Harris (1999), “Strategic experimentation.” *Econometrica*, 67, 349–374. [135]
- Brocas, Isabelle and Juan D. Carrillo (2007), “Influence through ignorance.” *RAND Journal of Economics*, 38, 931–947. [132]
- Bruckner, Andrew M. and John L. Leonard (1966) “Derivatives.” *American Mathematical Monthly*, 73, 24–56. [156]
- Calzolari, Giacomo and Alessandro Pavan (2006a), “Monopoly with resale.” *RAND Journal of Economics*, 37, 362–375. [132]
- Calzolari, Giacomo and Alessandro Pavan (2006b), “On the optimality of privacy in sequential contracting.” *Journal of Economic Theory*, 130, 168–204. [132]

- Che, Yeon-Koo and Navin Kartik (2009), “Opinions as incentives.” *Journal of Political Economy*, 117, 815–860. [133]
- Che, Yeon-Koo, Kyungmin Kim, and Konrad Mierendorff (2021). “Keeping the listener engaged.” Working paper. arXiv:2003.07338. [132]
- Crandall, Michael G., Hitoshi Ishii, and Pierre-Louis Lions (1992), “User’s guide to viscosity solutions of second order partial differential equations.” *Bulletin of the American Mathematical Society. New Series*, 27, 31–53. [155, 156]
- Crandall, Michael G. and Pierre-Louis Lions (1983), “Viscosity solutions of Hamilton–Jacobi equations.” *Transactions of the American Mathematical Society*, 277, 1–42. [133]
- Doval, Laura and Vasiliki Skreta (2021), “Mechanism design with limited commitment.” Working paper. arXiv:1811.03579. [132]
- Dworczak, Piotr (2020), “Mechanism design with aftermarket: Cutoff mechanisms.” *Econometrica*, 88, 2629–2661. [132]
- Ely, Jeffrey C. (2017) “Beeps.” *American Economic Review*, 107, 31–53. [132]
- Ely, Jeffrey C. and Martin Szydlowski (2020), “Moving the goalposts.” *Journal of Political Economy*, 128, 468–506. [132]
- Fershtman, Daniel and Alessandro Pavan (2022). “Searching for “arms.”” Working paper. [132]
- Gentzkow, Matthew and Emir Kamenica (2014), “Costly persuasion.” *American Economic Review*, 104, 457–462. [132]
- Georgiadis, George and Balázs Szentes (2020), “Optimal monitoring design.” *Econometrica*, 88, 2075–2107. [132]
- Henry, Emeric and Marco Ottaviani (2019), “Research and the approval process: The organization of persuasion.” *American Economic Review*, 109, 911–955. [132]
- Kamenica, Emir (2019), “Bayesian persuasion and information design.” *Annual Review of Economics*, 11, 249–272. [132]
- Kamenica, Emir and Matthew Gentzkow (2011), “Bayesian persuasion.” *American Economic Review*, 101, 2590–2615. [130, 132, 136, 145]
- Kamenica, Emir, Kyungmin Kim, and Andriy Zapechelnyuk (2021), “Bayesian persuasion and information design: Perspectives and open issues.” *Economic Theory*, 72, 701–704. [132]
- Karatzas, Ioannis and Steven E. Shreve (1991), *Brownian Motion and Stochastic Calculus*, second edition. Graduate Texts in Mathematics. Springer, New York, NY. [135, 144, 157]
- Ke, Tony T. and Miguel J. Villas-Boas (2019), “Optimal learning before choice.” *Journal of Economic Theory*, 180, 383–437. [133]
- Keller, Godfrey and Sven Rady (2020), “Undiscounted bandit games.” *Games and Economic Behavior*, 124, 43–61. [133]

- Kuvalekar, Aditya and Elliot Lipnowski (2020), “Job insecurity.” *American Economic Journal: Microeconomics*, 12, 188–229. [133]
- le Treust, Maël and Tristan Tomala (2019), “Persuasion with limited communication capacity.” *Journal of Economic Theory*, 184, 1–44. [132]
- Milgrom, Paul and Chris Shannon (1994), “Monotone comparative statics.” *Econometrica*, 62, 157–180. [158]
- Nikandrova, Arina and Romans Pancs (2018), “Dynamic project selection.” *Theoretical Economics*, 13, 115–143. [133]
- Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov (2020), “Persuading the principal to wait.” *Journal of Political Economy*, 128, 2542–2578. [132]
- Papanicolaou, Andrew (2016), *Stochastic Analysis Seminar on Filtering Theory*. Lecture Notes. arXiv:1406.1936. [135]
- Pavan, Alessandro and Giacomo Calzolari (2009), “Sequential contracting with multiple principals.” *Journal of Economic Theory*, 144, 503–531. [132]
- Pham, Huyên (2009), *Continuous-Time Stochastic Control and Optimization With Financial Applications*. Stochastic Modelling and Applied Probability. Springer. [150, 152, 153, 155]
- Pham, Huyên (2012), “Errata for “Continuous-time stochastic control and optimization with financial applications”.” <https://sites.google.com/site/phamxuanhuyen/publications/books>. [152]
- Rayo, Luis and Ilya Segal (2010), “Optimal information disclosure.” *Journal of Political Economy*, 118, 949–987. [132]
- Renault, Jérôme, Eilon Solan, and Nicolas Vieille (2017), “Optimal dynamic information provision.” *Games and Economic Behavior*, 104, 329–349. [132]
- Rockafellar, R. Tyrrell (1970), *Convex Analysis*. Princeton University Press, Princeton, NJ. [156]
- Rosar, Frank (2017), “Test design under voluntary participation.” *Games and Economic Behavior*, 104, 632–655. [132]
- Rudin, Walter (1976), *Principles of Mathematical Analysis*, third edition. McGraw-Hill, New York, NY. [146, 158]
- Siegel, Ron and Bruno Strulovici (2020), “The economic case for probability-based sentencing.” Working paper. [132]
- Smolin, Alex (2021), “Dynamic evaluation design.” *American Economic Journal: Microeconomics*, 13, 300–331. [132]
- van den Steen, Eric (2009), “Authority versus persuasion.” *American Economic Review*, 99, 448–453. [133]

Zhong, Weijie (2022), “Optimal dynamic information acquisition.” *Econometrica*, 90, 1537–1582. [133]

Co-editor Marina Halac handled this manuscript.

Manuscript received 12 January, 2022; final version accepted 2 April, 2022; available online 6 April, 2022.