Strategic research funding*

Matteo Escudé
European University Institute

Ludvig Sinander
Northwestern University

21 March 2019

Abstract

We study a dynamic game in which information arrives gradually as long as a principal funds research, and an agent takes an action in each period. In equilibrium, the principal's patience is the key determinant of her information provision: the lower her discount rate, the more eagerly she funds. When she is sufficiently patient, her information provision and value function are well-approximated by the 'Bayesian persuasion' model. If the conflict of interest is purely belief-based and information is valuable, then she provides full information if she is patient. We also obtain a sharp characterisation of the principal's value function. Our proofs rely on a novel dynamic programming principle rooted in the theory of viscosity solutions of differential equations.

1 Introduction

When is research in society's best interest? Consider the scientific investigation of the extent and implications of man-made climate change. Such research has the clear benefit of informing policy-making: mitigation is socially desirable if, and only if, climate change is a severe threat. But policy-makers do not always take actions that are socially optimal given the available information. When, then, is it socially beneficial to fund research?

In this paper, we study optimal research funding over time. In our model, a planner ('principal') decides how much information to provide by funding research over time, and the public ('agent') takes an action in each period.

*We are grateful to Eddie Dekel, Piero Gottardi, Alessandro Pavan and Bruno Strulovici for guidance and comments, and to Gregorio Curello, Jeff Ely, George Georgiadis, Yingni Guo, Konstantin Milbradt, Francisco Poggi, Marciano Siniscalchi, Pierpaolo Soravia, Quitzé Valenzuela-Stookey and audiences at Northwestern and EUI for helpful comments and suggestions.

We characterise all Markov perfect equilibria (with the common belief as state variable) in which tie-breaking is well-behaved. We characterise how the planner's equilibrium information provision depends on her discount rate, and under what conditions it is close to information provision in Bayesian persuasion, the natural static benchmark model. When the conflict of interest is rooted in a difference in prior beliefs, we show that a patient planner optimally provides full information. In addition, we characterise her value function, and relate it to the value function in the persuasion benchmark. Our proofs rely on a novel extension of the dynamic programming principle to a class of singular and discontinuous stochastic optimisation problems.

Our first result (Proposition 2) shows that socially optimal research funding (and thus information provision) is decreasing in the discount rate and in the noisiness of research findings. Thus, the debate about what discount rate should be used to assess the costs of climate change is also of key importance for optimal research funding policy. In particular, Stern's (2007) proposal that the social discount rate should be low implies that it is optimal to continue research for longer (and thus to provide more information) than if the higher rate favoured by Nordhaus (2007) is used.

Our second result (Proposition 3) relates equilibrium information provision to that in Bayesian persuasion, a natural static benchmark for our model. It is always socially optimal to provide less information than in the persuasion benchmark. But if a sufficiently low social discount rate is employed, then optimal information provision is close to that in the benchmark. Likewise if research findings are sufficiently precise (not noisy).

Thirdly, we consider the salient special case in which the conflict of interest is purely belief-based: the planner and public have common preferences, but different priors. We show in Proposition 4 that provided the common preference values information, it is optimal to provide (exactly) full information whenever the social discount rate (or noise) is low enough, regardless of the priors. In the context of climate change research, this suggests that if the only friction in policy-making is that the public's prior belief differs from the planner's, then research is always socially beneficial under Stern's (2007) low social discount rate.

Fourth, in Proposition 1, we provide a full characterisation of the planner's value function. We show that the value may be viewed as a generalisation to the impatient case of the *concave envelope,* which is the value function in the static persuasion benchmark. If the social discount rate and noisiness of research findings are sufficiently low, then the value function is well-approximated by the concave envelope (Corollary 1).

Our proof strategy for the aforementioned propositions is to first characterise the planner's value function using the HJB (Hamilton–Jacobi–Bellman) equation (obtaining Proposition 1), then to derive implications for behaviour (Propositions 2, 3 and 4). However, our model is sufficiently general that the value function need not satisfy the HJB equation in the classical sense.

To deal with this, we prove a novel dynamic programming principle (Lemma 1) according to which the value function is a *viscosity solution* of the HJB equation, extending existing results from the mathematics literature. This permits us to use the powerful theory of viscosity solutions of differential equations to characterise the value function in Proposition 1. We view this as a methodological contribution, and believe that our viscosity approach is likely to prove useful in other stochastic models in continuous time.

1.1 An example with two actions

There are two characters, the public (or their government) and a social planner. The state of the world is θ, equal to 1 if anthropogenic climate change is real and 0 otherwise. The players have a common prior belief $p_0 \in (0, 2/3)$ that $\theta = 1$ at time t. Write p_t for their (common) belief that $\theta = 1$ at time $t = 0$.

At each instant, the public chooses whether to take action to mitigate and/or adapt to climate change. No action ($a = 0$) has no cost or benefit. Taking action costs 1/2, yields a benefit of 3/4 for the public if man-made climate change is happening ($\theta = 1$), and has no benefit otherwise.[1] Given the belief p, the public's expected flow payoff is

$$f_{\mathrm{a}}(a,p) = \begin{cases} 0 & \text{for } a = 0 \\ \frac{3}{4}p - \frac{1}{2} & \text{for } a = 1. \end{cases}$$

The social benefit of taking action against climate change when $\theta = 1$ is 3/2, exceeding the private benefit of 3/4. The planner's payoff is therefore

$$f_P(a,p) = \begin{cases} 0 & \text{for } a = 0 \\ \frac{3}{2}p - \frac{1}{2} & \text{for } a = 1. \end{cases}$$

The planner discounts flow payoffs at the social discount rate $r > 0$.[2]

[1] The public does not observe flow payoffs in real time, so cannot infer θ.

[2] The conflict of interest can be motivated by impatience on the part of the public. Suppose that the cost of action is paid up-front, and that the benefit is $(3/2)r$ in *every* future period. If the public's discount rate is $2r$, then it values this stream of benefits at $\int_0^\infty e^{-2rt}(3/2)r\mathrm{d}t = (1/2r)(3/2)r = 3/4$. The planner values it at 3/2.

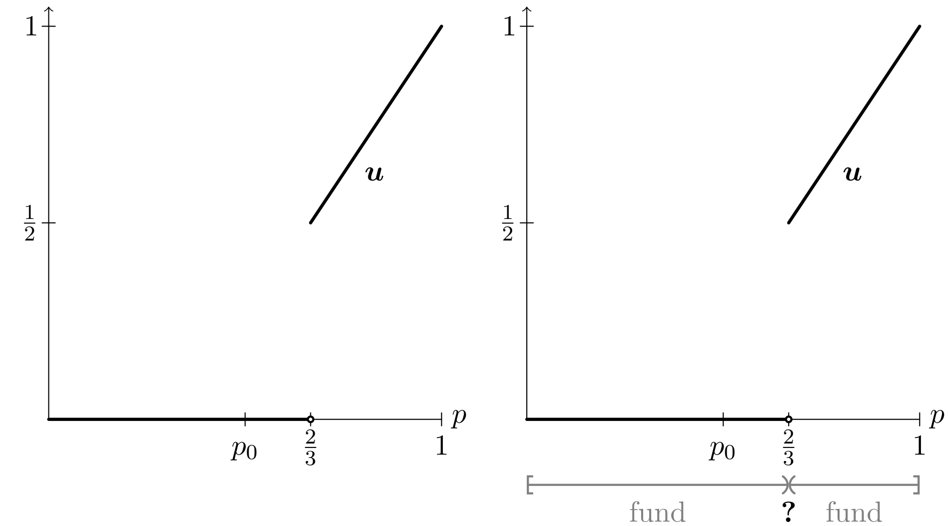

(a) Flow payoff u. (b) (Part of) an optimal policy.

Figure 1 – Planner's payoffs and optimal policy.

At each moment, the planner decides whether or not to fund climate change research.[3] Without funding, no research occurs, so the belief p_t stays put. When research is funded, the belief p_t evolves as a continuous (Brownian) martingale: it never jumps, and its changes $\mathrm{d}p_t$ have mean zero.

Suppose that the planner uses a *Markov* strategy, meaning that her funding policy at time t depends only on the current common belief p_t. Then the public cannot influence the planner's future behaviour, so finds it optimal to behave *myopically,* maximising its expected flow payoff period by period. In particular, it makes policy as follows:

$$A(p) = \begin{cases} 0 & \text{for } p \in [0, 2/3) \\ 1 & \text{for } p \in [2/3, 1]. \end{cases}$$

The planner's flow payoff as a function of the current common belief is then $u(p) \coloneqq f_P(A(p), p)$, drawn in Figure 1a.

Most of the planner's best-reply problem is easily solved. When $p \in (0, 2/3)$, it is strictly optimal for her to fund research since her flow payoff can only improve. When $p \in (2/3, 1)$, it is weakly optimal to fund research. To see this, observe that she could fund research only as long as $p > 2/3$. Then

[3]We neglect the cost of funding research. We consider this reasonable because the costs and benefits of mitigating climate change are of much larger magnitude in practice.

at each instant that she funds, the belief changes by a mean-zero random increment $\mathrm{d}p$. Since u is affine on the interval $(2/3, 1]$, the expected payoff is $\mathbf{E}(u(p + \mathrm{d}p)) = u(p)$, the same as from not funding. This policy is depicted in Figure 1b.

The planner faces a non-trivial trade-off at $p = 2/3$, however. By stopping research, she can lock in a moderate payoff of $1/2$ forever. If she funds research, then she may increase her payoff toward 1 (if p rises). But if the belief initially declines, then the planner suffers a zero flow payoff in the near future. The optimal resolution of this trade-off depends on the planner's discount rate: a patient planner will fund research at $p = 2/3$, while an impatient one will not.

This leads to our first result: the socially optimal funding of climate change research hinges on the choice of discount rate. If we follow Stern (2007) in using a very low social discount rate, then research probably ought to be funded at $p = 2/3$. But if we adopt Nordhaus's (2007) proposal that the planner should discount using a (much higher) market rate, then it may be socially optimal to stop research at $p = 2/3$. This insight generalises: we show in Proposition 2 that no matter what the payoffs of the planner and the public, the planner funds more eagerly (and thus provides more information) the more patient she is. Similarly, the planner funds more keenly the more informative research is about the state θ.

Our second result relates optimal funding to information provision in Bayesian persuasion (Kamenica & Gentzkow, 2011). In this static benchmark model, the planner designs an arbitrary *information structure*, which induces a mean-p_0 distribution of posterior beliefs p. Kamenica and Gentzkow (2011) have shown that the planner's value function in this problem is the concave envelope of u, which in this case is $(\operatorname{cav} u)(p) = p$, depicted in Figure 2. The planner's expected value is therefore p_0. She achieves this by providing a fully informative signal about θ, inducing the beliefs 0 and 1 (with probabilities $1 - p_0$ and p_0, respectively).

These are the same beliefs induced with positive probability in the long run by a patient planner: she funds research until the belief hits either 0 or 1,[4] then stops. Thus, when the principal is patient, the static Bayesian persuasion benchmark provides a good approximation to optimal funding policy. We generalise this insight in Proposition 3: no matter what the planner's and public's payoffs, the behaviour of a sufficiently patient planner is well-approximated by the static persuasion benchmark. The same is not true of an impatient planner, who allows research only while $p \neq 2/3$, and

[4]Technically, the belief may *converge* to 0 or 1, but (a.s.) does not *hit* either.

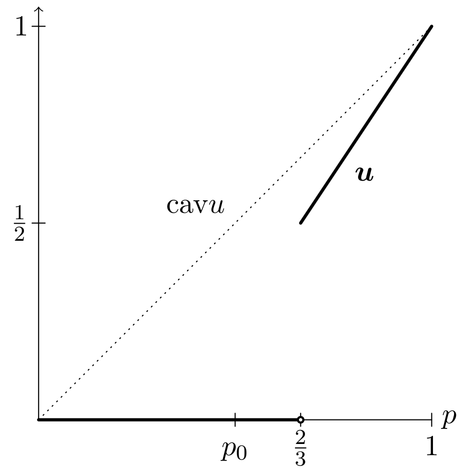

Figure 2 – Value cav u in the static persuasion benchmark.

thus induces the less informative long-run beliefs $\{0, 2/3\}$.

The long-run beliefs $\{0, 1\}$ induced by a patient planner are *fully* informative. The root of this feature is that the conflict of interest can be re-interpreted as purely-belief based: the public's behaviour A can be rationalised by assuming that it shares the planner's preferences f_P, but has a different prior. In particular, suppose that the public's prior $p_{a,0}$ is lower than the planner's. Then no matter what information arrives, its posterior belief is lower than the principal's. If $p_{a,0}$ is chosen appropriately, then the public's belief exceeds the threshold $1/3$ above which it is optimal to take action precisely when the principal's belief exceeds $2/3$, so that the strategy A is (myopically) optimal for the public.

Our third result provides conditions under which the planner optimally provides full information when the conflict of interest is purely belief-based. In particular, Proposition 4 asserts that for any common preference that values information (a convexity condition), a sufficiently patient planner provides (exactly) full information, whatever the priors. The convexity condition is satisfied by the preferences in our two-action example, and more generally by expected-utility preferences. This result suggests that according to Stern's (2007) low social discount rate, climate change research is always socially advantageous.

Our fourth result characterises the planner's value function as a generalisation of the concave envelope. In the example, the value function v is depicted in Figures 3a and 3b for low and high values of the discount rate.

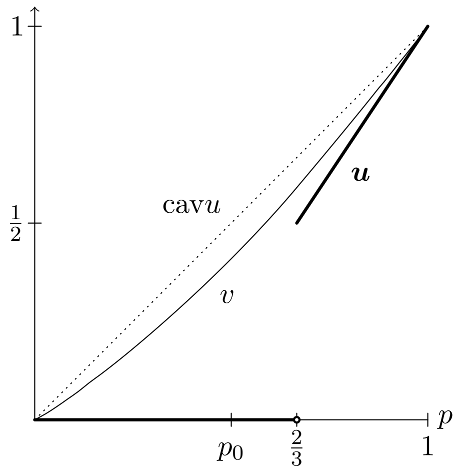

(a) Patient planner.

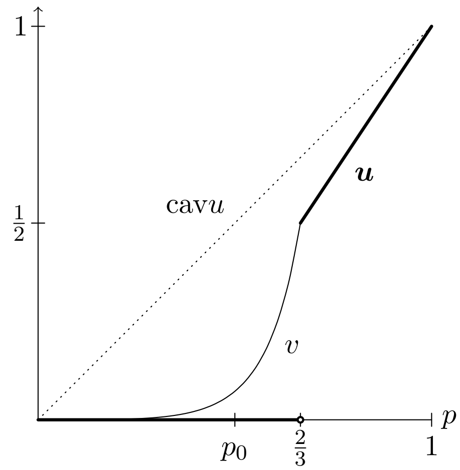

(b) Impatient planner.

Figure 3 – Planner's value function v.

Both v and $\operatorname{cav} u$ are upper envelopes of u that exceed u when convex and coincide with u when concave. But whereas $\operatorname{cav} u$ is affine whenever it exceeds u, v is strictly convex when it exceeds u due to the planner's impatience. Moreover, when the planner is patient (Figure 3a), her value v is (uniformly) close to $\operatorname{cav} u$. We show that these conclusions hold for arbitrary payoffs: the value function is an impatience-adjusted cousin of the concave envelope (Proposition 1), and is uniformly close to the latter when the planner is patient (Corollary 1).

It is intuitive that the dynamic programming principle should hold, meaning that the planner's value function v should satisfy the HJB equation:

$$v(p) = u(p) + \frac{1}{r} \max \left\{ 0, \left(\frac{p(1-p)}{\sigma} \right)^2 \frac{v''(p)}{2} \right\}.$$

In words, this says that the value at state p is the flow payoff $u(p)$ plus the discounted option value of funding, which is proportional to the value of information $v''(p)/2$.[5] But the HJB equation simply cannot hold in the classical sense in Figure 3b: the right-hand side is ill-defined wherever v has a kink, since v'' does not exist there. These singularities stem from the planner's ability to 'freeze' the state variable p_t by stopping research.

[5]The term $p(1-p)/\sigma$ is the rate at which information arrives while the planner funds research. We will see in §2.2 why it has this form.

To deal with this issue, we prove a dynamic programming principle for problems with singularities and discontinuous flow payoffs. We build on a result from the mathematics literature: allowing for singularities, if u is continuous, then v is a *viscosity solution* of the HJB equation. In Lemma 1, we extend this theorem to allow the flow payoff u to be discontinuous. This extension is essential for any economic application in which payoffs depend on the behaviour of other players, and may have uses beyond this paper. We use our dynamic programming principle to characterise (for arbitrary payoffs) the value function (Proposition 1), and thence information provision (Propositions 2, 3 and 4).

1.2 Related literature

As our principal and agent are symmetrically informed throughout, there is no communication friction. This distinguishes our work from e.g. Pei (2015), Argenziano, Severinov and Squintani (2016) and Frug (2017).

Our paper is more closely related to the persuasion literature initiated by Kamenica and Gentzkow (2011), Rayo and Segal (2010) and Aumann and Maschler (1995), and in particular to dynamic persuasion. In contrast to Ely (2017) and Renault, Solan and Vieille (2017), the state of the world is constant in our environment.

Orlov, Skrzypacz and Zryumov (2018), Ely and Szydlowski (2017) and Bizzotto, Rüdiger and Vigier (2017) study dynamic persuasion problems in which the state is fixed and the agent chooses at each instant whether to quit irrevocably. By contrast, our agent chooses freely among her actions at each instant. This is an important difference: whereas the principal would not benefit from being able to commit ex ante to an information policy in our environment, she does benefit strictly in these papers.

Henry and Ottaviani (2019) and Siegel and Strulovici (2018, §4) study dynamic persuasion models with a fixed state in which the *principal* chooses when to stop irreversibly, whereupon the agent takes an action and payoffs are realised.[6] Their agent acts only once, and their principal earns a payoff at that time; by contrast, our agent acts in every period, and our principal earns a flow payoff throughout. Until she stops, their principal incurs a flow cost $c > 0$, absent in our model.[7] Both papers focus on how institutional design can improve welfare, whereas we study equilibrium behaviour in a fixed game. Siegel and Strulovici (2018, §4) do not have results analogous to

[6] Brocas and Carrillo (2007) study a similar model, but with no discounting ($r = 0$) and in discrete time.

[7] Another difference is that payoffs are assumed linear in the belief p in these papers.

any of ours, but their model exhibits similar properties to the special case of our model with only two or three actions.

Henry and Ottaviani (2019) further assume that there are only two actions and that the principal's payoff is independent of the state. In this special case, the dependence of information provision on patience that we emphasise (Proposition 2) is entirely absent when (as in our model) $c = 0$. It becomes important as soon as the principal's payoff is allowed to be state-dependent (the two-action example in §1.1) or there are more than two actions (we give an example on p. 20). The authors characterise the principal's value function, and show that information provision is close to that in Bayesian persuasion when the cost c is small. These two results are analogous to the specialisation of our Propositions 1 and 3 to two actions and state-independent preferences.[8] The special case is simple enough to permit elementary proofs, with no need for viscosity methods. Since their assumptions rule out belief-based disagreement, these authors have no analogue of our Proposition 4.

A question distinct from, but related to, ours is which static models of information acquisition are equivalent to a model of sequential information acquisition. In probability theory, Skorokhod-type embedding theorems can be used to characterise the set of distributions of posterior beliefs that arise from *some* sequential sampling rule. Morris and Strack (2017) ask what static models of *costly* information acquisition are equivalent to a sequential model of costly information acquisition.

On the technical side, our model is a stochastic differential game, meaning a stochastic game in continuous time in which the state evolves according to an Itô diffusion. Stochastic differential games were introduced by Isaacs (1954, 1965), and have made occasional appearances in economics (see Fudenberg and Tirole (1991, §13.3) for some examples). The central difficulties in the study of general stochastic differential games are absent in our model because only one player (the principal) is able directly to influence the evolution of the state. Like much of this literature, we avoid the technical issues associated with defining strategies in continuous time (e.g. Simon and Stinchcombe (1989)) by focussing on Markov perfect equilibria.

Viscosity solutions of differential equations were introduced by Crandall and Lions (1983). We give a brief exposition and some references in appendix J. Viscosity solutions have begun to be used in macroeconomics; see in particular Achdou, Han, Lasry, Lions and Moll (2017). The best-reply problem in these models is simpler than in ours because the flow payoff is continuous,

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

singularities cannot arise, and the HJB equation is a first- as opposed to a second-order differential equation.

2 Model

Our model is a stochastic game in continuous time. There is a principal and an agent. At each moment, the agent and principal take respective actions a_t and λ_t. Flow payoffs depend on the agent's action a_t and on the state p_t. The principal's action λ_t affects the stochastic evolution of the state p_t. In particular, λ_t is the rate at which public information is permitted to arrive, 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 principal and the agent have a common prior belief p_0 that the state is $\theta = 1$.[9] Time $t \in \mathbf{R}_+$ is continuous. The agent takes an action $a \in \mathcal{A}$ at each moment, where $\mathcal{A}$ is a finite set.

When the agent takes action a and the common belief is p, the principal's and agent's respective flow payoffs are $f_P(a, p)$ and $f_\mathrm{a}(a, p)$, both continuous in p. Expected utility ($f_P(a, \cdot)$ and $f_\mathrm{a}(a, \cdot)$ affine) is a natural special case. The principal and agent discount flow payoffs at rates $r > 0$ and $r_\mathrm{a} > 0$, respectively.

2.2 The principal's information provision

At each instant, the principal can costlessly permit a small amount of public information to arrive by funding research. In particular, she chooses $\lambda_t \in [0, 1]$, and everyone observes the process

$$\mathrm{d}X_t = \theta\lambda_t \mathrm{d}t + \sigma\sqrt{\lambda_t}\, \mathrm{d}\widetilde{B}_t,$$

where $\widetilde{B}$ is a standard Brownian motion and $\sigma > 0$.[10] The constraint $\lambda_t \leq 1$ bounds the rate at which the principal can information. Our qualitative conclusions would remain intact if this constraint were replaced with a cost of funding research, but some tractability would be lost.

The assumption that the noise in the signal X is Brownian rules out information arriving in discrete lumps.[11] The volatility σ of the noise con-

[9] We will drop the common-prior assumption in §6.

[10] In the background, there is a suitable probability space that carries the process $\widetilde{B}$.

[11] Allowing for lumps complicates the analysis without adding (much) insight.

strains how rapidly information can arrive. The limit $\sigma \to 0$ corresponds to the case in which information can arrive arbitrarily fast.

The signal process admits a natural micro-foundation. For concreteness, consider the planner–public example: λ_t is the planner's funding of research, and $\theta = 1$ iff anthropogenic climate change is happening. Research produces a body of evidence which may be summarized by a 'score'. Write Y_Λ for the (random) cumulative score following cumulative research funding Λ, so that $\dot{Y}_\Lambda$ is today's score. The score $\dot{Y}_\Lambda$ has mean θ, but is subject to noise with variance $\sigma^2 > 0$. Since a white noise is the rate of change of a random walk, which in continuous time means a Brownian motion, we may write $\mathrm{d}Y_\Lambda = \theta \mathrm{d}\Lambda + \sigma \mathrm{d}\overline{B}_\Lambda$, where $\overline{B}$ is a standard Brownian motion.

The 'time-changed' white noise $t \mapsto \mathrm{d}\overline{B}_{\Lambda_t}$ has the same law as the *scaled* white noise $t \mapsto \sqrt{\mathrm{d}\Lambda_t/\mathrm{d}t}\,\mathrm{d}\widetilde{B}_t$, where $\widetilde{B}$ is a(nother) standard Brownian motion.[12] Cumulative funding evolves over time according to $\Lambda_t = \int_0^t \lambda_s \mathrm{d}s$. The change over time of the cumulative score $X_t \coloneqq Y_{\Lambda_t}$ may be therefore be expressed as

$$\mathrm{d}X_t = \theta \mathrm{d}\Lambda_t + \sigma \mathrm{d}\overline{B}_{\Lambda_t} = \theta \frac{\mathrm{d}\Lambda_t}{\mathrm{d}t}\mathrm{d}t + \sigma \sqrt{\frac{\mathrm{d}\Lambda_t}{\mathrm{d}t}}\,\mathrm{d}\widetilde{B}_t = \theta \lambda_t \mathrm{d}t + \sigma \sqrt{\lambda_t}\,\mathrm{d}\widetilde{B}_t.$$

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

$$\mathrm{d}p_t = \sqrt{\lambda_t}\,\frac{p_t(1-p_t)}{\sigma}\mathrm{d}B_t,$$

where B is a standard Brownian motion according to the common belief.[13] See e.g. Bolton and Harris (1999, Lemma 1) for a heuristic derivation. By inspection, 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' behaviour 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.

[12] This is standard—see e.g. Lowther (2010).

[13] This is (a special case of) the Kushner–Stratonovich equation (e.g. Papanicolaou (2016, §4.2.2)). The process B is given by $\mathrm{d}B_t = \mathrm{d}X_t - p_t \mathrm{d}t$ 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 seen from the Girsanov theorem (e.g. Karatzas and Shreve (1991, §3.5)).

We remark, however, that non-Markov equilibria (suitably defined) can differ qualitatively from Markov perfect ones.

Definition 1. A function $[0,1] \to \mathbf{R}^n$ is *piecewise continuous* iff its discontinuities form a discrete subset of $(0,1)$.[14]

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 principal (agent) is a $[0,1]$-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 principal* is a measurable map $\Lambda : [0,1] \to [0,1]$. A *Markov strategy of the agent* is a piecewise continuous map $A : [0,1] \to \Delta(\mathcal{A})$.[15] We identify a Markov strategy Λ (A) with the strategy (stochastic process) that it induces via $\lambda_t \coloneqq \Lambda(p_t)$ (a_t distributed according to $A(p_t)$, independently over time).

The restriction to piecewise continuous strategies is a mild assumption on the agent's tie-breaking that has no payoff consequences for her, provided her preferences are non-degenerate in a weak sense.[16] The role of piecewise continuity is to ensure that the principal's best-reply problem satisfies a dynamic programming principle (Lemma 1 in §4.1).

Definition 3. A strategy $(\lambda_t)_{t\in\mathbf{R}_+}$ of the principal is a *best reply* at $\bar{p} \in [0,1]$ to a Markov strategy $A : [0,1] \to \Delta(\mathcal{A})$ of the agent iff it maximises

$$\mathbf{E}\left(\int_0^\infty e^{-rt} \left[\int_{\mathcal{A}} f_P(a, p_t) A(\mathrm{d}a|p_t) \right] \mathrm{d}t \right),$$
$$\text{where} \quad \mathrm{d}p_t = \sqrt{\lambda_t}\, \frac{p_t(1-p_t)}{\sigma} \mathrm{d}B_t \quad \text{and} \quad p_0 = \bar{p},$$

over all $[0,1]$-valued processes $(\lambda_t)_{t\in\mathbf{R}_+}$ adapted to the filtration generated by $(p_t)_{t\in\mathbf{R}_+}$.[17] (Note that a best reply must do better than all other strategies, not only Markov ones.)

A strategy $(a_t)_{t\in\mathbf{R}_+}$ of the agent is *undominated* iff there is no other strategy $(a'_t)_{t\in\mathbf{R}_+}$ that yields the same expected payoff and has $f_{\mathrm{a}}(a'_t, p) >$

[14] A subset of $(0,1)$ is *discrete* iff each of its members is an isolated point: it lives in a neighbourhood that contains no other members.

[15] $\Delta(\mathcal{A})$ denotes the set of probability distributions over the (finite) set $\mathcal{A}$.

[16] We explain why in supplemental appendix I. It is obvious for expected-utility preferences.

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

$f_{\mathrm{a}}(a_t, p)$ a.s. for some $p \in [0,1]$. A strategy $(a_t)_{t\in\mathbf{R}_+}$ of the agent is a *best reply* at $\bar{p} \in [0,1]$ to a strategy $(\lambda_t)_{t\in\mathbf{R}_+}$ of the principal iff it is undominated and maximises

$$\mathbf{E}\left(\int_0^\infty e^{-r_{\mathrm{a}}t} f_{\mathrm{a}}(a_t, p_t)\mathrm{d}t\right),$$
$$\text{where} \quad \mathrm{d}p_t = \sqrt{\lambda_t}\,\frac{p_t(1-p_t)}{\sigma}\mathrm{d}B_t \quad \text{and} \quad p_0 = \bar{p},$$

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 agent as uninteresting. Accommodating them is not hard, but 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 Other interpretations

In §1, we gave the following interpretation of the model: θ is whether man-made climate change is real, the principal is a social planner who allocates research funding, and the agent is the public. Two other interpretations are as follows.

First, suppose that the state θ is whether the president has misbehaved. The attorney general (principal) can start an independent investigation into the president's alleged wrongdoing, and stop it at any time. While it is ongoing, the investigation releases any findings in real time. The public (agent) give more support to the president's party the less likely they think the president is to be a crook. The attorney general has no private information about the state, but favours the president's party. How will the attorney general behave?

Alternatively: in 1996, pro-gun legislators in the US Congress passed the Dickey Amendment, which prohibits the Centers for Disease Control and Prevention from conducting research into the causal link between gun ownership and violence. To model this, let the state θ be whether gun ownership causes violence. Pro-gun legislators (principal) have a majority in the legislature, so can decide whether to allow or to ban gun violence research. The public (agent) agitate for or against gun rights depending on their belief about the link between guns and violence, and the legislators want them to support gun rights.

2.5 Extensions

In §6, we extend the model to allow the principal and agent to have different priors. Other natural extensions of our model include multiple agents, an external source of public news (so that p_t cannot be totally frozen by setting $\lambda_t = 0$), a cost of information provision, a stochastically evolving state θ, more than two states θ, and lumpy information. All of these extensions modify our results in obvious ways without substantially affecting the qualitative conclusions. Some of them, such as lumpy information, make the model less tractable.

A natural alternative model would have the principal provide information until she decides to stop irreversibly, whereupon the agent acts once and for all and payoffs are realised. Such a model may describe a prosecutor gathering evidence before going to trial, or a pharmaceutical firm running clinical trials before requesting regulatory approval. Models along these lines are studied by Henry and Ottaviani (2019) and Siegel and Strulovici (2018), and turn out to have similar properties to ours.

2.6 Benchmark: Static Bayesian persuasion

In static persuasion, the principal flexibly provides information once and for all—she is not restricted to do it gradually. If the principal induces belief p, then the agent takes an action $A(p)$ that maximises $f_a(\cdot, p)$, giving the principal a payoff of $u(p) \coloneqq f_P(A(p), p)$. Assume that the agent breaks ties such that u is upper semi-continuous.[18]

Kamenica and Gentzkow (2011) studied this model, and showed the following. The principal is able to induce all and only distributions of beliefs whose mean is p_0 ('splits' of the prior). The principal's value at prior p_0 is $(\operatorname{cav} u)(p_0)$, where $\operatorname{cav} u$ is the concave envelope of u (the smallest concave function that majorises u). The principal has an optimal policy that induces either two beliefs (if $(\operatorname{cav} u)(p_0) > u(p_0)$) or one belief (if $(\operatorname{cav} u)(p_0) = u(p_0)$).

We shall see (Proposition 3, Corollary 1) that when the principal is patient or can provide information rapidly (in particular, when $r\sigma^2$ is low), the principal's value and information provision are close to those in the static persuasion benchmark.

[18]Without this assumption, an optimal policy need not exist.

3 Myopic behaviour by the agent

To avoid uninteresting technicalities, we will focus on MPEs in which the agent's tie-breaking is well-behaved. Say that a Markov strategy $A : [0,1] \to \Delta(\mathcal{A})$ of the agent is *regular* iff it breaks ties such that the principal's induced payoff $u(p) \coloneqq \int_{\mathcal{A}} f_P(a,p) A(\mathrm{d}a|p)$ is upper semi-continuous. Regular Markov strategies exist—indeed, *any* Markov strategy of the agent need be modified only on a discrete (hence finite or countable) set of beliefs p to be made regular, and this modification leaves the agent's flow payoff unchanged.[19]

Call a Markov strategy $A : [0,1] \to \Delta(\mathcal{A})$ of the agent *myopic* iff at each belief $p \in [0,1]$, every action in the support of $A(\cdot|p)$ maximises $f_a(\cdot, p)$.

Observation 1. A regular strategy of the agent is part of a MPE iff it is myopic.

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

Proof. If the principal uses a Markov strategy, then since the agent cannot affect the evolution of the state, a strategy of hers is a best reply iff it is myopic. The 'only if' part follows.

For the 'if' part, fix a regular and myopic Markov strategy $A : [0,1] \to \Delta(\mathcal{A})$ of the agent. We show in §5.1 the principal has a best reply, in fact a Markov one, to any regular A.[20] Since A is myopic, it is a best reply. We have found a MPE. ■

Observation 1 implies that Markov perfect equilibria exist, and further that a principal-preferred MPE exists (by a simple continuity argument). It also follows that the agent's behaviour 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.[21]

In light of Observation 1, it remains only to characterise the principal's best reply to a given regular and myopic strategy of the agent. We will in

[19] Recall that a Markov strategy is piecewise continuous by definition, and that the agent's flow payoff $f_a(a, \cdot)$ is continuous.

[20] There is no circularity here: the subsequent arguments concerning the principal's best reply do not rely on Observation 1.

[21] In particular, for Lebesgue almost all expected-utility f_a (viewed as vectors in $\mathbf{R}^{2|\mathcal{A}|}$), the agent's strategy differs across MPEs only on a countable set of beliefs. We will see in §5.1 that the principal's best reply is generically partially unique.

fact characterise her best reply to an arbitrary regular Markov strategy. We proceed in two steps, studying the principal's value function §4, and then her best reply in §5.

4 The principal's value function

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

$$u(p) := \int_{\mathcal{A}} f_P(a,p) A(\mathrm{d}a|p).$$

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

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

$$\begin{aligned} v(p_0) = \sup_{(\lambda_t)_{t\in\mathbf{R}_+}} &\mathbf{E}\left(r \int_0^\infty e^{-rt} u(p_t) \mathrm{d}t \right) \\ \text{s.t.} \quad \mathrm{d}p_t &= \sqrt{\lambda_t}\, \frac{p_t(1-p_t)}{\sigma} \mathrm{d}B_t, \end{aligned} \tag{BRP}$$

where $(\lambda_t)_{t\in\mathbf{R}_+}$ is chosen among all $[0,1]$-valued processes adapted to the filtration generated by $(p_t)_{t\in\mathbf{R}_+}$, and p_0 is given.

4.1 The HJB equation and viscosity solutions

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

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

or equivalently

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

In well-behaved problems, the principal's best reply problem satisfies a *dynamic programming principle:* her 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 value of information.

Our problem is not well-behaved, however: we saw in Figure 3b that the value function v may have a kink. Since v'' does not exist at kinks, v is not a classical solution of (HJB) (the right-hand side is ill-defined).[22]

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 $u^\star$ ($u_\star$) denote the upper (lower) semi-continuous envelope of u.[23] The envelopes $u^\star$ and $u_\star$ differ only on a discrete set since u is piecewise continuous, and we have $u_\star \leq u = u^\star$ since u is upper semi-continuous.

Definition 5. $w : [0,1] \to \mathbf{R}$ is a *viscosity sub-solution (super-solution)* of (HJB) iff it is upper (lower) semi-continuous, and for any twice continuously differentiable $\phi : (0,1) \to \mathbf{R}$ and local minimum $p \in (0,1)$ of $\phi - w$ (of $w - \phi$),

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

w is a *viscosity solution* of (HJB) iff it is both a viscosity sub-solution and super-solution.

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 J. Observe that if w is a viscosity solution of (HJB) and is twice continuously differentiable on a neighbourhood of $p \in (0,1)$, then it satisfies (HJB) in the classical sense at p.[24]

[22] Even in the absence of a kink (as in Figure 3a), 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.

[23] That is, the pointwise smallest (largest) upper (lower) semi-continuous function that majorises (minorises) u.

[24] This is because we may choose a twice continuously differentiable ϕ that coincides with w on a neighbourhood of p, so that $\phi - w$ and $w - \phi$ are locally minimised at p and $\phi''(p) = w''(p)$.

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

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

The proof is in appendix A.

We view Lemma 1 as a technical contribution. It extends a well-known theorem from the stochastic 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 behaviour of other players.[25] Lemma 1 may prove useful for studying other models of strategic interaction in continuous-time stochastic environments.

4.2 Characterisation of the principal's value function

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

Definition 6. $w : [0,1] \to \mathbf{R}$ is *locally (strictly) convex* at $p \in (0,1)$ iff

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

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 iff the reverse (strict) inequality holds.[26]

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 $u \leq v \leq \operatorname{cav} u$. On C, we have $v < \operatorname{cav} u$, and v is once continuously differentiable. On $C \setminus D$, we have further that v is twice continuously differentiable and satisfies

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

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

Proposition 1 is summarised in Table 1, where $\mathcal{C}^k$ means 'continuous and k times continuously differentiable'. The 'smooth pasting' property in

[25] By contrast, we have derived our hypothesis of piecewise continuity of u from piecewise continuity of A, a tie-breaking assumption with no payoff consequences for the agent.

[26] Note well that (strict) local convexity/concavity at p is weaker than (strict) convexity/concavity on a neighbourhood of p. The function v depicted in Figure 3b is locally strictly concave at 2/3, but not concave on any neighbourhood of 2/3.

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

Table 1 – Summary of Proposition 1. $\mathcal{C}^k$ means 'continuous and k times continuously differentiable'.

the entry for the region $C \cap D$ 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 $p \in C \cap D$,[27] we have

$$\lim_{n \to \infty} v'(p_n) = v'(p).$$

Remark that since v is only $\mathcal{C}^0$ on $(0,1) \setminus C$, it may have (locally concave) kinks in this region—we saw an example of this in Figure 3b.

The characterisation of v in Proposition 1 is a generalisation of the concave envelope $\operatorname{cav} u$. Both are upper envelopes of u that exceed u when convex and coincide with u when concave. But whereas $\operatorname{cav} u$ 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, and if there is one then $C' = C$. We give some details in supplemental appendix F.

In the two-action example of §1.1, Proposition 1 implies that the value function must have either the strictly convex shape in Figure 3a or the convex-affine shape in Figure 3b. Proposition 1 further rules out one of the two; for example, for $r\sigma^2$ large, the convex candidate violates $u \leq v$ at $2/3$.[28]

As another example, consider the flow payoff u depicted in Figure 4a, corresponding to three actions.[29] The concave envelope and value function for high and low values of $r\sigma^2$ are depicted in Figures 4b–4d.[30]

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

[28] See supplemental appendix F.1 for the details.

[29] The underlying model has actions $\mathcal{A} = \{0, 1, 3\}$, flow payoff $f_P(a,p) = a$ for the principal, and payoffs $f_a(0,p) = 0$, $f_a(1,p) = 2p - 1$ and $f_a(3,p) = \frac{14}{3}p - 3$ for the agent.

[30] See supplemental appendix F.2 for the details.

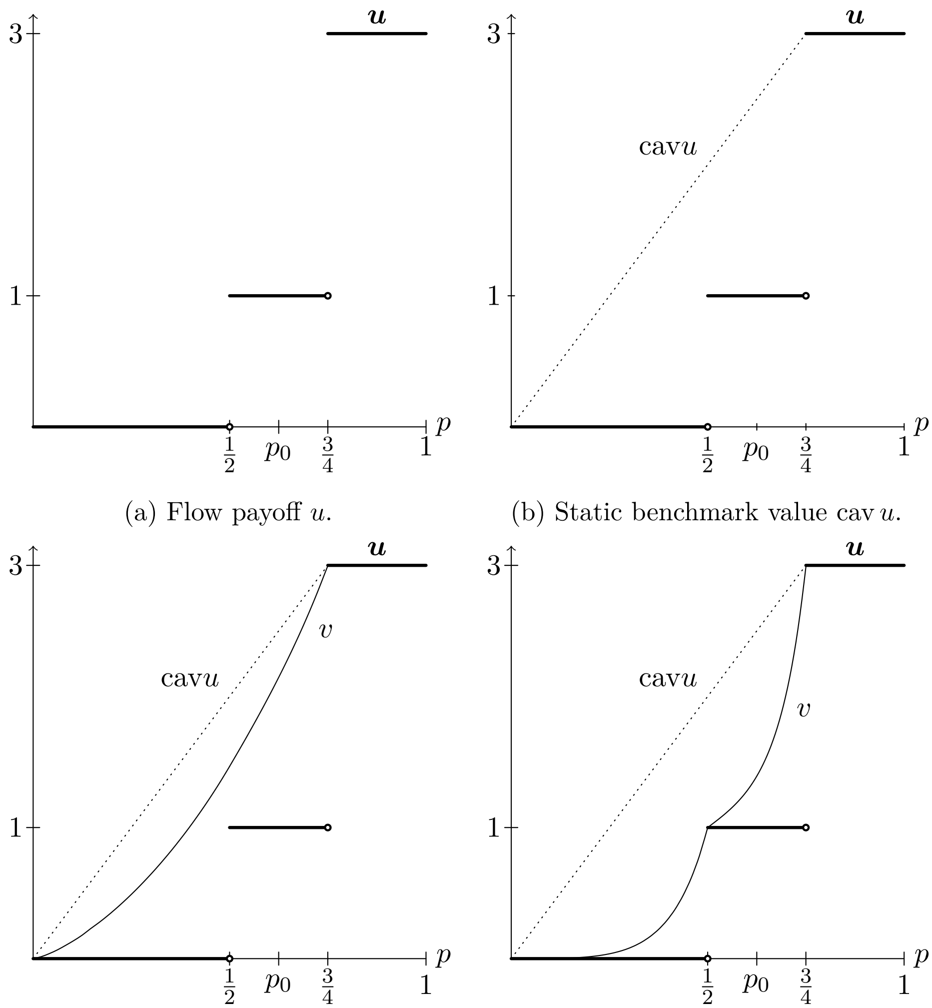

(a) Flow payoff u.

(b) Static benchmark value cav u.

(c) Value v for low $r\sigma^2$.

(d) Value v for high $r\sigma^2$.

Figure 4 – Example with three actions.

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

Proof of Proposition 1. By Lemma 1, v is a viscosity solution of (HJB) satisfying $v = u$ on $\{0, 1\}$. It follows that v is continuous. The fact that $u = \operatorname{cav} u$ on $\{0, 1\}$ follows from upper semi-continuity 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 $v \leq \operatorname{cav} u$, take any $p \in [0, 1]$, and consider the auxiliary problem in which the principal 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 principal can induce in her best-reply problem is available in the auxiliary problem. And we have $V(p) = (\operatorname{cav} u)(p)$ since the auxiliary problem consists of a sequence of independent static persuasion problems (one for each instant t), in each of which the optimal value is $(\operatorname{cav} u)(p)$ since u is upper semi-continuous.

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

Lemma 2. On $(0, 1) \setminus C$, we have $v = u$.

Now for C. In appendix C, we prove that

Lemma 3. v is continuously differentiable on C.

To show that $v < \operatorname{cav} u$ 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

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

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

Lemma 4. On $C \setminus D$, v is twice continuously differentiable and satisfies (∂).

This lemma is proved in appendix D. ■

Taking the patient limit $r\sigma^2 \to 0$ in Proposition 1 yields

Corollary 1. As $r\sigma^2$ decreases, v increases pointwise. As $r\sigma^2 \to 0$, v converges uniformly to $\operatorname{cav} u$.

Corollary 1 shows that when the principal is patient or able to provide information rapidly, the static persuasion model provides a good approximation to the principal's equilibrium value. We saw by example in Figures 3b and 4d that $\operatorname{cav} u$ may approximate v poorly when the principal is impatient or highly restricted.

Proof. To emphasise dependence on parameters, write the value as $v_{r\sigma^2}$.

For the first part, fix an arbitrary $p_0 \in [0,1]$. Using the the change of variable $s \coloneqq rt$, we may rewrite the best-reply problem (BRP) at prior $p_0 \in [0,1]$ as

$$v_{r\sigma^2}(p_0) = \sup_{(\lambda'_s)_{s\in\mathbf{R}_+}} \mathbf{E}\left(\int_0^\infty e^{-s} u(p_s)\mathrm{d}s\right) \quad \text{s.t.} \quad \mathrm{d}p_s = \sqrt{\lambda'_s}\, p_s(1-p_s)\mathrm{d}\widehat{B}_s,$$

where $\widehat{B}_s \coloneqq \sqrt{r} B_t$ is a standard Brownian motion, and

$$(\lambda'_s)_{s\in\mathbf{R}_+} = \left(\lambda_s / r\sigma^2\right)_{s\in\mathbf{R}_+}$$

is chosen among all $[0, 1/r\sigma^2]$-valued processes adapted to the filtration generated by $(p_s)_{s\in\mathbf{R}_+}$. Since lowering $r\sigma^2$ slackens the constraint $\lambda'_s \leq 1/r\sigma^2$ in every period, it raises the value $v_{r\sigma^2}(p_0)$.

Now for the second part. We have established for every $p \in [0,1]$ that the sequence $(v_{r\sigma^2}(p))_{r\sigma^2>0}$ is increasing. Since it lives in the compact set $[u(p), (\operatorname{cav} u)(p)]$ by Proposition 1, it must converge to some $v_0(p) \in [u(p), (\operatorname{cav} u)(p)]$. In other words, $(v_{r\sigma^2})_{r\sigma^2>0}$ converges pointwise to some function $v_0 : [0,1] \to \mathbf{R}$ satisfying $u \leq v_0 \leq \operatorname{cav} u$. We claim that $v_0 = \operatorname{cav} u$. Since $\operatorname{cav} u$ is by definition the pointwise smallest concave majorant of u, it suffices to show that v_0 is concave.

To that end, take $p' < p < p''$ in $[0,1]$, and let $\gamma \in (0,1)$ be such that $\gamma p' + (1-\gamma)p'' = p$; we will establish that $\gamma v_0(p') + (1-\gamma)v_0(p'') \leq v_0(p)$. Consider the principal's best-reply problem (BRP) with prior $p_0 = p$ and parameters $r > 0$ and $\sigma^2 > 0$. By a change of variable akin to the one in the first part of the proof, we may consider the payoff-equivalent problem with unit signal noise and discount rate $r\sigma^2$ (and λ_t chosen from $[0,1]$). 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 $r\sigma^2 > 0$ we have

$$\mathbf{E}\left(r\sigma^2 \int_0^\tau e^{-r\sigma^2 t} u(p_t)\mathrm{d}t + e^{-r\sigma^2\tau} v_{r\sigma^2}(p_\tau)\right) \leq v_{r\sigma^2}(p) \leq v_0(p),$$

where the second inequality holds since $v_{r\sigma^2}$ increases pointwise as $r\sigma^2$ declines. As $r\sigma^2 \to 0$, the first term inside the expectation on the left-hand side vanishes a.s., and the second term converges a.s. to $v_0(p_\tau)$. Since both terms are bounded, the left-hand side converges to $\mathbf{E}(v_0(p_\tau))$ by the bounded convergence theorem (e.g. Theorem 16.5 in Billingsley (1995)). And we have

$$\mathbf{E}(v_0(p_\tau)) = \gamma v_0(p') + (1-\gamma) v_0(p'')$$

by the optional sampling theorem (e.g. Theorem 3.22 in Karatzas and Shreve (1991, ch. 1)).[31]

We have proved that $v_{r\sigma^2}$ converges monotonically, pointwise, to $\operatorname{cav} u$. Since $v_{r\sigma^2}$ and $\operatorname{cav} u$ 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 $v_{r\sigma^2}$ to $\operatorname{cav} u$ is uniform. ■

5 Equilibrium information provision

Having characterised the principal's value function (Proposition 1), we are ready to study her information provision in equilibrium. We shall first establish that she provides more information the more patient she is (Proposition 2, §5.2). Secondly, in Proposition 3 (§5.3), we will show that her information provision is less generous than in the static persuasion benchmark, but close to the latter if she is sufficiently patient.

5.1 Induced beliefs in the long run

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

Corollary 2. The following Markov strategy is a best reply:

$$\Lambda^\star(p) = \begin{cases} 1 & \text{if } v(p) > u(p) \\ 0 & \text{otherwise.} \end{cases}$$

Proof. When $v = u$, setting $\lambda = 0$ clearly attains the value v in the principal's problem (BRP). When $v > u$, $\lambda > 0$ must be optimal. By inspection of the principal's problem, $\lambda = 1$ is optimal whenever $\lambda > 0$ is. ■

[31] When $0 < p' < p'' < 1$, we have $\mathbf{E}(\tau) < \infty$, so the optional sampling theorem 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 $\mathbf{E}\left(p_{\tau_n}\right) = p$ by the optional sampling theorem. Since $\left(p_{\tau_n}\right)_{n \in \mathbf{N}}$ is bounded and converges a.s. to p_τ as $n \to \infty$, the bounded convergence theorem yields $\mathbf{E}\left(p_\tau\right) = p$.

The strategy $\Lambda^\star$ provides information at full tilt when it is strictly valuable, and provides none otherwise. It is partially unique.[32]

Under this strategy, the belief p_t evolves according to $\mathrm{d}p_t = \sigma^{-1} p_t(1 - p_t)\mathrm{d}B_t$ until it hits the (closed) set $\{v = u\}$, then stays constant forever. Since $(p_t)_{t \in \mathbf{R}_+}$ is a bounded martingale, it converges a.s. by the martingale convergence theorem.[33] Write F for the distribution of the limiting random variable. The support of F is the set of beliefs that the principal induces (with positive probability) in the long run. Note that F has mean p_0 since each p_t does.[34]

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

$$\begin{aligned} p^- &\coloneqq \sup\{p \in [0, p_0] : v(p) = u(p)\} \\ p^+ &\coloneqq \inf\{p \in [p_0, 1] : v(p) = u(p)\}. \end{aligned}$$

Proof. The best reply $\Lambda^\star$ from Corollary 2 obviously induces $\{p^-, p^+\}$. ■

Provided $v(p_0) > u(p_0)$ (the interesting case), the long-run induced beliefs $\{p^-, p^+\}$ that we study are generically the unique ones consistent with a best reply.[35] In general, they are the least extreme beliefs induced by some best reply.[36]

5.2 Comparative statics

The principal provides more information the more patient she is:

Proposition 2 (comparative statics)**.** Fix a prior p_0. As $r\sigma^2$ decreases, p^- decreases and p^+ increases.

Proof. Write $\{v > u\}$ for the set of beliefs at which the strategy $\Lambda^\star$ from Corollary 2 provides information. As $r\sigma^2$ decreases, v increases pointwise by Corollary 1, so $\{v > u\}$ increases in the sense of set inclusion, and thus

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

decreases and p^+ similarly increases. ■

[32]Precisely: any best reply must have $\Lambda > 0$ on $\{v > u\}$, $\Lambda = 1$ 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 3b (p. 7), we have $\{v = u\} \cap K = (2/3, 1)$.)

[33]E.g. Theorem 3.15 in Karatzas and Shreve (1991, ch. 1).

[34]By the bounded convergence theorem (e.g. Theorem 16.5 in Billingsley (1995)).

[35]See supplemental appendix G for a discussion.

[36]Because any best reply must set $\Lambda > 0$ on $\{v > u\}$.

By way of illustration, consider the two-action example of §1.1: whereas a patient planner induces the long-run beliefs $\{0, 1\}$, an impatient one induces the less informative beliefs $\{0, 2/3\}$. Similarly, in the three-action example of Figure 4 (p. 20), the long-run induced beliefs are more extreme when the principal is patient ($\{0, 3/4\}$ rather than $\{1/2, 3/4\}$).

Recall from §1 the controversy between Stern (2007) and Nordhaus (2007) about what discount rate society should use to value how climate change will affect future generations. According to Proposition 2, the choice social discount rate also bears crucially on optimal funding of climate change research. In particular, when funding research is socially optimal given Stern's low proposed discount rate (as at $p = 1/2$ in Figure 3a), it may be that stopping research is socially optimal if Nordhaus's higher rate is used instead (as in Figure 3b). The underlying trade-off is illustrated by the two-action example of §1.1: while stopping at $p = 2/3$ is safe, continuing yields possible losses in the near future and potential gains in the further future, and thus what is optimal depends on how severely future payoffs are discounted.

Proposition 2 also states that more information is provided the lower the noisiness σ^2 of the signal process. Intuitively, this is because lowering the noise accelerates the rate at which information arrives, effectively 'speeding up time', which is equivalent in payoff terms to lowering the discount rate.[37]

It is worth noting that Proposition 2 generalises: in particular, it continues to hold if information can arrive in discrete lumps. Indeed, Quah and Strulovici (2013) have shown that in a large class of Markov stopping problems, a less patient decision-maker stops more eagerly. While the economic force is the same, our result is not a special case of theirs because their assumptions are not satisfied by our model.[38]

5.3 Comparison with the static persuasion benchmark

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

$$P^- := \sup\{p \in [0, p_0] : (\operatorname{cav} u)(p) = u(p)\}$$
$$P^+ := \inf\{p \in [p_0, 1] : (\operatorname{cav} u)(p) = u(p)\}.$$

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

[37] This can be seen formally in the proof of Corollary 1 (p. 21).

[38] In particular, Quah and Strulovici (2013) assume that the stochastic process $U_t := u(p_t)$ is right-continuous, which only true in our model if u is continuous.

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 $r = 0$ of perfect patience. The following result shows in addition that information provision by a patient principal is close to that in the persuasion benchmark:

Proposition 3 (persuasion connection)**.** Fix a prior p_0. For any $r\sigma^2 > 0$, we have $P^- \leq p^- \leq p^+ \leq P^+$. As $r\sigma^2 \to 0$, $p^- \to P^-$ and $p^+ \to P^+$.

Proposition 3 also implies that holding the impatience $r > 0$ of the principal fixed, her information provision is close to the persuasion benchmark if the signal process is precise (low noisiness σ^2).

To illustrate, consider the two-action example of §1.1. In the patient case depicted in Figure 3a, the same information is provided as in persuasion: $P^- = p^- = 0 < 1 = p^+ = P^+$. But in the impatient case in Figure 3b, strictly less information is provided than in the benchmark: $P^- = p^- = 0$ and $p^+ = 2/3 < 1 = P^+$. Similarly, in the three-action example in Figure 4, the benchmark and patient long-run induced beliefs are $\{0, 3/4\}$, but an impatient planner induces the less extreme beliefs $\{1/2, 3/4\}$. The examples show that if the principal is impatient, then the persuasion benchmark may provide a poor approximation to equilibrium information provision.

In the two- and three-action examples, the long-run induced beliefs of a sufficient patient planner *coincide* with the induced beliefs in the persuasion benchmark rather than merely converging as per Proposition 3. This need not happen—we give an example in supplemental appendix H in which $P^- < p^- < p^+ < P^+$ for every $r\sigma^2 > 0$.

Although it is intuitive that long-run induced beliefs should converge to the persuasion ones as $r\sigma^2 \to 0$, the result is not obvious. To see this, observe that the analogous result for the impatient limit $r\sigma^2 \to \infty$ is false! For this limit, the natural static benchmark 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 $r\sigma^2 \to \infty$: indeed, in the two-action example of §1.1, the (uniquely optimal) long-run induced beliefs are $\{p^-, p^+\} = \{0, 2/3\}$ for *every* low value of $r\sigma^2 > 0$.[39]

Proof. To emphasise dependence on parameters, write $v_{r\sigma^2}$, $p^-_{r\sigma^2}$ and $p^+_{r\sigma^2}$ for the value and long-run beliefs. Let $\{v_{r\sigma^2} > u\}$ be the set of beliefs at which the strategy $\Lambda^\star$ in Corollary 2 funds research.

[39] The key formal difference between the two limits is that the convergence of v to $\operatorname{cav} u$ as $r\sigma^2 \to 0$ is uniform by Corollary 1 (p. 21), whereas the convergence of v to u as $r\sigma^2 \to \infty$ is merely pointwise (unless u is continuous).

For the first part, fix a $r\sigma^2 > 0$. Clearly $P^- \leq P^+$. By Proposition 1, $u(p) < v_{r\sigma^2}(p)$ implies $u(p) < (\operatorname{cav} u)(p)$, so that $\{v_{r\sigma^2} > u\} \subseteq \{\operatorname{cav} u > u\}$. Therefore

$$P^- = \sup\left([0, p_0] \setminus \{\operatorname{cav} u > u\}\right) \leq \sup\left([0, p_0] \setminus \{v_{r\sigma^2} > u\}\right) = p^-_{r\sigma^2},$$

and similarly $p^+_{r\sigma^2} \leq P^+$.

Now for the second part. Since $p^-_{r\sigma^2}$ decreases monotonically as $r\sigma^2$ does by Proposition 2, and lives in the compact set $[P^-, p_0]$, it converges to some limit $p^-_0 \in [P^-, p_0]$. We wish to show that $p^-_0 = P^-$, so suppose toward a contradiction that $p^-_0 > P^-$. On the one hand,

$$(\operatorname{cav} u)\left(p^-_{r\sigma^2}\right) \to (\operatorname{cav} u)\left(p^-_0\right) > u\left(p^-_0\right) \tag{1}$$

by continuity of $\operatorname{cav} u$ and $p^-_0 > P^-$ (recall the definition of P^-). On the other hand, (recalling the definition of $p^-_{r\sigma^2}$,)

$$\left|(\operatorname{cav} u)\left(p^-_{r\sigma^2}\right) - u\left(p^-_{r\sigma^2}\right)\right| = \left|(\operatorname{cav} u)\left(p^-_{r\sigma^2}\right) - v_{r\sigma^2}\left(p^-_{r\sigma^2}\right)\right| \to 0$$

since $v_{r\sigma^2}$ converges *uniformly* to $\operatorname{cav} u$ by Corollary 1 (p. 21).[40] It follows by upper semi-continuity of u that

$$(\operatorname{cav} u)\left(p^-_{r\sigma^2}\right) \to \lim_{r\sigma^2 \to 0} u\left(p^-_{r\sigma^2}\right) \leq u\left(p^-_0\right)$$

a contradiction with (1). A similar argument shows that $p^+_{r\sigma^2} \to P^+$. ■

6 Heterogeneous priors

In this section, we drop the common-prior assumption. We first show (§6.1) that our preceding characterisation of equilibrium welfare and information provision remains valid.

We then focus on the special case in which the conflict of interest between the principal and agent derives from their differing beliefs alone. We show in Proposition 4 (§6.3) that when information is valuable, a sufficiently patient principal provides full information, regardless of the priors. If disagreement about climate change policy is rooted in differing prior beliefs, then this result suggests that climate change research is always socially valuable according to Stern's (2007) social discount rate.

In supplemental appendix E, we give conditions under which the principal is better-off the smaller the prior disagreement.

[40] See e.g. Theorem 7.11 in Rudin (1976).

6.1 Equilibrium characterisation

The model is as in §2, except that the priors $p_0, p_{a,0} \in (0,1)$ of the principal and agent may differ. They agree to disagree—the difference in priors is not rooted in private information. Write p_t and $p_{a,t}$ for the principal's and agent's beliefs at time t.

The extended model is tractable because we need not keep track of the agent's belief, as we can back it out from the principal's belief and the priors:

Observation 2. The agent's time-t belief is $p_{a,t} = \phi(p_t, p_0, p_{a,0})$, where

$$\phi(p_t, p_0, p_{a,0}) := \frac{p_t}{p_t + (1-p_t)D(p_0, p_{a,0})}$$

$$\text{and} \qquad D(p_0, p_{a,0}) := \frac{p_0}{1-p_0} \Big/ \frac{p_{a,0}}{1-p_{a,0}} .$$

Proof. Write ℓ_t^1/ℓ_t^0 for the likelihood ratio of the (random) observation $(X_s)_{s\in[0,t]}$.[41] By Bayes's rule,

$$\frac{p_t}{1-p_t} = \frac{p_0}{1-p_0}\frac{\ell_t^1}{\ell_t^0}$$

and

$$\frac{p_{a,t}}{1-p_{a,t}} = \frac{p_{a,0}}{1-p_{a,0}}\frac{\ell_t^1}{\ell_t^0} = \frac{1}{D(p_0, p_{a,0})}\frac{p_0}{1-p_0}\frac{\ell_t^1}{\ell_t^0}.$$

Combining and rearranging yields the result. ■

To illustrate, consider the two-action example from §1.1. Let the planner have prior $p_0 = 1/2$. Suppose that the public shares the planner's preference f_P, but that its prior is $p_{a,0} = 1/5$. Then by Observation 2, the public's payoffs from the two actions in terms of the *planner's* belief p are $f_a(a,p) = 0$ and

$$f_a(1,p) = \frac{3}{4}\phi(p, 1/2, 1/5) - \frac{1}{2} = \frac{3}{4}\frac{p}{4-3p} - \frac{1}{2},$$

depicted in Figure 5b. The public finds it myopically optimal to take action $a = 1$ when $p_{a,0} \geq 1/3$, which occurs precisely when $p \geq 2/3$. This is the agent's (Markov) strategy from §1.1.

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 (§3), a regular

[41] Formally, 'ℓ_t^1/ℓ_t^0' is the Radon–Nikodým derivative of the law of $(X_s)_{s\in[0,t]}$ under $\theta = 1$ with respect to its law under $\theta = 0$.

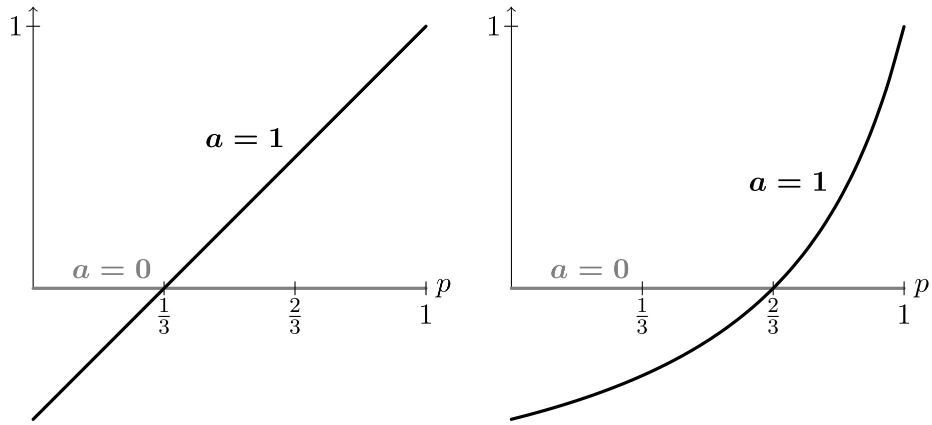

(a) Planner's flow payoff. (b) Public's flow payoff.

Figure 5 – Flow payoffs of actions as a function of the *planner's* belief p.

Markov strategy $A : [0,1] \to \Delta([0,1])$ is part of a MPE iff it is myopic. Given a regular Markov strategy A, the principal's induced flow payoff is now

$$u(p) := \int_{\mathcal{A}} f_P(a,p) A(\mathrm{d}a|\phi(p, p_0, p_{\mathrm{a},0}))$$

since the agent's belief is $\phi(p, p_0, p_{\mathrm{a},0})$ when the principal'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 principal's best-reply problem is unchanged, noting again that p_t is the *principal's* belief.[42] All preceding results therefore remain valid: the principal's value function is a generalised concave envelope (Proposition 1), she provides more information the more patient she is (Proposition 2), and her information provision is well-approximated by static persuasion when she is sufficiently patient (Proposition 3).

6.2 Belief-based conflict of interest

With heterogeneous priors, there is typically a conflict of interest between the principal and agent even if they have the same objectives. To focus on this issue, we assume henceforth that preferences at each belief $p \in [0,1]$ are the same, so that $f_P = f_{\mathrm{a}} = f$.[43] We ask how the extent of disagreement

[42] It is important that we use the principal's belief. The agent's belief is not a martingale from the principal's perspective, as she expects it to drift toward her own.

[43] In the special case of expected utility, this is equivalent to players having the same preferences conditional on the state θ.

between the priors p_0 and $p_{\mathrm{a},0}$ impacts information provision (§6.3) and the principal's welfare (supplemental appendix E).

We saw in the previous section that the agent's behaviour A in the two-action example from §1.1 is consistent with a belief-based conflict of interest. By contrast, the agent's behaviour in the three-action example in Figure 4 cannot arise from a purely belief-based conflict of interest: if the agent had the principal's payoff $f_P(a,p) = a$, then she would strictly prefer action 3 at every belief (no matter what the priors).

6.3 Full information provision

When the conflict of interest is purely belief-based, there is an additional force that favours more information provision by the principal. When beliefs are close, the principal can rely on the agent to make good use of information by taking actions that the principal likes, suggesting that the principal improves her payoff by providing information. And providing information tends to make the beliefs move closer: indeed, if research is funded forever, then *both* beliefs converge a.s. to either 0 or 1.

There are two gaps in this reasoning. First, for arbitrary preferences f, information need not be valuable: the principal could well prefer less-than-full information even if she were to choose actions herself.[44] To ensure that information is valuable, we assume that $f(a,\cdot)$ is convex for every action $a \in \mathcal{A}$. This allows for expected utility, where $f(a,\cdot)$ is affine.

Secondly, although information provision tends to bring the beliefs closer, it may take a long time. In the interim, the agent may take actions that the principal dislikes, and this may cause the principal to provide little or no information. To make sure that this short-term effect does not dominate, we consider the patient case ($r\sigma^2$ small).

To avoid trivialities, we also assume that f is *non-degenerate,* by which we mean that (i) there are actions $a, a' \in \mathcal{A}$ and beliefs $p, p' \in [0,1]$ such that $f(a,p) > f(a',p)$ and $f(a,p') < f(a',p')$, and (ii) there is an $\varepsilon > 0$ such that $\max_{a\in\mathcal{A}} f(a,\cdot) = f(a_0,\cdot)$ on $[0,\varepsilon)$ ($= f(a_1,\cdot)$ on $(1-\varepsilon,1]$) for some $a_0, a_1 \in \mathcal{A}$. Part (ii) says that the agent doesn't want to switch actions infinitely frequently near 0 or 1, and holds for expected-utility preferences. For similarly technical reasons, we focus on pure-strategy MPEs.

Proposition 4 (full information)**.** Suppose that $f_P = f_{\mathrm{a}} = f$, where $f(a,\cdot)$ is convex for each $a \in \mathcal{A}$ and f is non-degenerate. Let the agent use a pure

[44] Consider $\mathcal{A} = \{0\}$ and $f(0,\cdot) = p(1-p)$. Then u is strictly concave regardless of the priors, so the principal strictly prefers to stop immediately.

and regular Markov strategy. Then for any fixed priors p_0 and $p_{a,0}$, we have $\{p^-, p^+\} = \{0, 1\}$ for all sufficiently small $r\sigma^2 > 0$.

In the context of climate change research, the conflict of interest may be purely belief-based. Proposition 4 then suggests that on Stern's (2007) proposal of a very low social discount rate, climate change research is always socially beneficial, so should be funded. By contrast, climate change research may sometimes be socially detrimental if we adopt the higher discount rate advocated by Nordhaus (2007).

We showed by example that convexity and high patience are both needed for the result. The three-action example in Figure 4 (p. 20) shows that it is also essential that the conflict of interest be purely belief-based: by inspection, the principal never induces more informative long-run beliefs than $\{0, 3/4\}$.

Proof. Fix priors p_0 and $p_{a,0}$. We proceed in three steps. We first show that $\operatorname{cav} u > u$ on $(0, 1)$, by convexity and part (i) of non-degeneracy. Next, using Corollary 1 and part (ii) of non-degeneracy, we find an $\eta > 0$ such that $r\sigma^2 < \eta$ implies $v > u$ on the closure of the set D' of discontinuities of the agent's strategy A. Finally, we use convexity to show that $v > u$ on all of $(0, 1)$ whenever $r\sigma^2 < \eta$, which by Corollary 2 suffices to prove the proposition.

Step 1: $\operatorname{cav} u > u$ on $(0, 1)$: Define $\bar{u}(p) := \max_{a \in \mathcal{A}} f(a, p)$, the principal's flow payoff if she were to choose actions herself. This function is convex since each $f(a, \cdot)$ is. It follows that $\operatorname{cav} \bar{u}$ is affine, and (since $\bar{u}$ is upper semi-continuous) that $\operatorname{cav} \bar{u} = \bar{u}$ on $\{0, 1\}$. By part (i) of non-degeneracy of f, $\bar{u}$ cannot be affine.[45] Since $\bar{u}$ is convex and non-affine, it must satisfy $\operatorname{cav} \bar{u} > \bar{u}$ on $(0, 1)$.

Let u be the flow payoff corresponding to the fixed priors p_0 and $p_{a,0}$. Clearly $\bar{u} \geq u$, with equality on $\{0, 1\}$. Thus $\operatorname{cav} \bar{u}$ is a concave function majorising u. There is no pointwise smaller concave function majorising u since $\operatorname{cav} \bar{u}$ is affine and coincides with u on $\{0, 1\}$. This shows that $\operatorname{cav} u = \operatorname{cav} \bar{u}$, and thus $\operatorname{cav} u = \operatorname{cav} \bar{u} > \bar{u} \geq u$ on $(0, 1)$, as desired.

Step 2: $v > u$ on $\operatorname{cl} D'$: Let A be the agent's (pure) strategy. Write D' for the set of of discontinuities of A, and $\operatorname{cl} D'$ its closure.[46] D' is a discrete subset of $(0, 1)$ since a strategy is piecewise continuous by definition.

By part (ii) of non-degeneracy, we have $D' \subseteq [\varepsilon, 1 - \varepsilon]$ for some $\varepsilon > 0$, and thus likewise for $\operatorname{cl} D'$. Observe that $\operatorname{cav} u - u$ is a lower semi-continuous

[45] In detail: suppose to the contrary that $\bar{u}$ is affine. Then by definition of $\bar{u}$ and the convexity of $f(a, \cdot)$ for each $a \in \mathcal{A}$, it must be that $\bar{u}(\cdot) = f(a, \cdot)$ for some single $a \in \mathcal{A}$, which contradicts non-degeneracy of f.

[46] Clearly D' contains the set D of discontinuities of u.

function since $\operatorname{cav} u$ is continuous and u is upper semi-continuous. Since $\operatorname{cav} u - u > 0$ on the compact set $[\varepsilon, 1-\varepsilon] \subseteq (0,1)$, it follows that

$$k \coloneqq \inf_{p \in \operatorname{cl} D'} [(\operatorname{cav} u)(p) - u(p)] \geq \inf_{p \in [\varepsilon, 1-\varepsilon]} [(\operatorname{cav} u)(p) - u(p)] > 0.$$

Now choose $\delta \in (0, k)$. Since v converges uniformly to $\operatorname{cav} u$ as $r\sigma^2 \to 0$ by Corollary 1 (p. 21), we may find an $\eta > 0$ such that $r\sigma^2 < \eta$ implies

$$\sup_{p \in \operatorname{cl} D'} [(\operatorname{cav} u)(p) - v(p)] < \delta.$$

(We will use this η throughout the remainder of the proof.) We then have

$$\begin{aligned}
\inf_{p \in \operatorname{cl} D'} [v(p) - u(p)] &= \inf_{p \in \operatorname{cl} D'} \{[(\operatorname{cav} u)(p) - u(p)] - [(\operatorname{cav} u)(p) - v(p)]\} \\
&\geq \inf_{p \in \operatorname{cl} D'} [(\operatorname{cav} u)(p) - u(p)] - \sup_{p \in \operatorname{cl} D'} [(\operatorname{cav} u)(p) - v(p)] \geq k - \delta > 0.
\end{aligned}$$

Step 3: $v > u$ on $(0,1)$*:* To complete the proof, it suffices by Corollary 2 (p. 23) to show that $v > u$ on $(0,1)$, for then $\{p^-, p^+\} = \{0,1\}$.

Suppose first that D' is empty. Then since agent's strategy A is pure, there is an action $a \in \mathcal{A}$ such that $u(\cdot) = f(a, \cdot)$ on $(0,1)$. Thus u is convex on $(0,1)$ since $f(a,\cdot)$ is. Since u is continuous at 0 and at 1 by piecewise continuity, it follows that u is convex on all of $[0,1]$.

Take any $p \in (0,1)$, and consider the principal's best-reply problem (BRP) with prior $p_0 = p$. The strategy that sets $\lambda = 1$ always cannot be better than optimal, so

$$v(p) \geq \mathbf{E}\left(\int_0^\infty e^{-rt} u(p_t) \mathrm{d}t\right) > \int_0^\infty e^{-rt} u(p) \mathrm{d}t = u(p),$$

where the strict inequality used the fact that u is convex and non-affine.

For the remainder, we assume that D' is non-empty. Take any $p' \in D'$. We claim that p' has a successor in $\operatorname{cl} D' \cup \{1\}$.[47] Suppose toward a contradiction that it does not. Since it clearly has upper bounds (viz. 1), it must be that for every $\varepsilon > 0$, there is $p_+ \in (p', p'+\varepsilon) \cap (\operatorname{cl} D' \cup \{1\})$ such that $p' < p_+ < p' + \varepsilon$. It follows that p' is not an isolated point of $\operatorname{cl} D'$. But then p' cannot lie in D' since the latter is discrete—a contradiction.

Furthermore, $p' = 0$ has a successor in $\operatorname{cl} D' \cup \{1\}$. The least upper bound of 0 in $\operatorname{cl} D' \cup \{1\}$ is evidently $\inf \operatorname{cl} D'$, which lies in $\operatorname{cl} D'$. To see that $0 \neq \inf \operatorname{cl} D'$, simply recall that $D' \subseteq [\varepsilon, 1-\varepsilon]$ for some $\varepsilon > 0$.

[47] The *successor* of $p' \in (0,1)$ in a set $A \subseteq [0,1]$ is the least upper bound of p' in $A \setminus \{p'\}$.

It follows that $(0,1)$ may be partitioned into $\operatorname{cl} D'$ along with (at most countably many) intervals (p',p''), where either (a) $p' \in D' \cup \{0\}$ and p'' is the successor of p' in $\operatorname{cl} D' \cup \{1\}$, or (b) p' is an element of $\operatorname{cl} D' \setminus D'$ that has a successor in $\operatorname{cl} D' \cup \{1\}$ and p'' is that successor. Since $v > u$ on $\operatorname{cl} D'$ whenever $r\sigma^2 < \eta$ by step 2, it remains only to show that $v > u$ on every (p',p'') when $r\sigma^2 < \eta$.

To that end, fix an interval (p',p'') of the sort just described. Since we supposed D' to be non-empty, it must be either that $0 < p'$ or that $p'' < 1$. We assume the former, omitting the analogous argument for the latter case. Define $u^\circ : [p',p''] \to \mathbf{R}$ by

$$u^\circ(p) \coloneqq \begin{cases} u(p'+) & \text{for } p = p' \\ u(p) & \text{for } p \in (p',p'') \\ u(p''-) & \text{for } p = p''. \end{cases}$$

u° is well-defined and continuous since u is continuous on $(p',p'') \subseteq (0,1) \setminus D'$. It clearly satisfies $u^\circ \leq u^\star = u$.

We claim that u° is convex. To see this, remark first that since (p',p'') does not intersect D' and the agent's strategy is pure, there must be an action $a \in \mathcal{A}$ such that $u(\cdot) = f(a,\cdot)$ on (p',p''). Since $f(a,\cdot)$ is convex, it follows that u° is convex on (p',p''). To conclude, observe that u° is continuous at p' and at p'' by construction, so is convex on its entire domain $[p',p'']$.

Take $p \in (p',p'')$; we will show that $v(p) > u(p)$. Since $r\sigma^2 < \eta$, we have $v > u$ on $\operatorname{cl} D'$, so in particular $v(p') > u(p')$ and $v(p'') \geq u(p'')$.[48] Consider the strategy that lets $\lambda = 1$ until the first time τ that the belief process $(p_t)_{t \in \mathbf{R}_+}$ with initial condition $p_0 = p$ hits $\{p',p''\}$, and subseqently behaves optimally. By the optional sampling theorem (e.g. Theorem 3.22 in Karatzas and Shreve (1991, ch. 1)), the probability that $p_\tau = p'$ is γ given by $p = \gamma p' + (1-\gamma)p''$.[49] Since behaving optimally is weakly better than this

[48] The latter inequality strict iff $p'' \neq 1$.

[49] The optional sampling theorem applies directly when $p'' < 1$, as then clearly $\mathbf{E}(\tau) < \infty$. In the case $p'' = 1$, let τ_n be the first time that $(p_t)_{t \in \mathbf{R}_+}$ hits $\{p', 1 - 1/n\}$ for $n \in \mathbf{N}$. Clearly $\mathbf{E}(\tau_n) < \infty$, 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 $n \to \infty$, we may apply the bounded convergence theorem (e.g. Theorem 16.5 in Billingsley (1995)) to obtain $\mathbf{E}(p_\tau) = p$, as desired.

strategy, and $u \geq u^{\circ}$, we have

$$\begin{aligned}
v(p) &\geq \mathbf{E}\left(r\int_0^{\tau} e^{-rt}u^{\circ}(p_t)\mathrm{d}t + \gamma e^{-r\tau}v(p') + (1-\gamma)e^{-r\tau}v(p'')\right) \\
&\geq \mathbf{E}\left(r\int_0^{\tau} e^{-rt}u^{\circ}(p)\mathrm{d}t + \gamma e^{-r\tau}v(p') + (1-\gamma)e^{-r\tau}v(p'')\right) \\
&= \mathbf{E}\left(\left(1-e^{-r\tau}\right)u^{\circ}(p) + e^{-r\tau}\left[\gamma v(p') + (1-\gamma)v(p'')\right]\right) \\
&> \mathbf{E}\left(\left(1-e^{-r\tau}\right)u^{\circ}(p) + e^{-r\tau}\left[\gamma u^{\circ}(p') + (1-\gamma)u^{\circ}(p'')\right]\right) \\
&\geq \mathbf{E}\left(\left(1-e^{-r\tau}\right)u^{\circ}(p) + e^{-r\tau}u^{\circ}(p)\right) \\
&= u(p),
\end{aligned}$$

where we used convexity of u° (the second inequality), $v(p') > u(p') \geq u^{\circ}(p')$ and $v(p'') \geq u(p'') \geq u^{\circ}(p'')$ (the strict inequality), convexity of u° again (the final inequality), and $u^{\circ} = u$ on $(0,1)$ (the last line). ■

Appendix

A Proof of Lemma 1 (p. 18)

The boundary condition $v = u$ on $\{0,1\}$ holds by inspection of the principal's best-reply problem (BRP), since these are absorbing states. For the remainder, let $v_{\star}$ and $v^{\star}$ be the lower and upper semi-continuous envelopes of v. In Lemma 1(a), we show that piecewise continuity of u suffices for $v_{\star}$ to be a viscosity super-solution. In Lemma 1(b), we prove that $v^{\star}$ is a viscosity sub-solution, without requiring piecewise continuity. In both cases, we adapt the standard proof. Finally, we establish in Lemma 1(c) that v is continuous, so that $v_{\star} = v^{\star} = v$.[50] We will make occasional use of the fact that $v \geq u$, which holds since the value u is attainable (by never funding research).

First, the super-solution property of $v_{\star}$, which relies on piecewise continuity of u:

Lemma 1(a). If u is piecewise continuous, then $v_{\star}$ is a viscosity super-solution 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 emphasise the juncture at which a new argument is needed to accommodate merely piecewise continuity of u.

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

Take any $p \in (0,1)$ and any twice continuously differentiable $\phi : (0,1) \to \mathbf{R}$ 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

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

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

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

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

$$\gamma_n \coloneqq v(p_n) - \phi(p_n)$$

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

Consider the 'full tilt forever' strategy which sets $\lambda_t = 1$ 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 behaviour:

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

Using $v - \phi \geq v_\star - \phi \geq 0$ and the definition of γ_n yields

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

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

$$\mathrm{d}P_t^n = \frac{P_t^n\left(1 - P_t^n\right)}{\sigma} \mathrm{d}B_t$$

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

$$\begin{aligned} e^{-rh_n}\phi\left(P_{h_n}^n\right) = \phi\left(p_n\right) - r \int_0^{h_n} e^{-rt}\phi\left(P_t^n\right) \mathrm{d}t& \\ + \frac{1}{2}\int_0^{h_n} \left(\frac{P_t^n\left(1 - P_t^n\right)}{\sigma}\right)^2 e^{-rt}\phi''\left(P_t^n\right) \mathrm{d}t&. \end{aligned}$$

Substituting in (3) and rearranging slightly yields

$$\frac{\gamma_n}{h_n} \geq \mathbf{E}\left(r\frac{1}{h_n}\int_0^{h_n} e^{-rt}u\left(P_t^n\right)\mathrm{d}t - r\frac{1}{h_n}\int_0^{h_n} e^{-rt}\phi\left(P_t^n\right)\mathrm{d}t \right.$$
$$\left. + r\frac{1}{h_n}\int_0^{h_n} e^{-rt}\frac{\left(P_t^n\right)^2\left(1-P_t^n\right)^2}{2r\sigma^2}\phi''\left(P_t^n\right)\mathrm{d}t \right) \quad \forall n \in \mathbf{N}. \quad (4)$$

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

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

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

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

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

so that

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

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

$$\lim_{\varepsilon\to 0}\frac{1}{2\varepsilon}\int_{p-\varepsilon}^{p+\varepsilon} u \geq \min\left\{u\left(p-\right), u\left(p+\right)\right\} \geq u_\star(p).$$

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 exeeds $u_\star(p)$.

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 according to the occupancy measure of $(P_t^n)_{t\in\mathbf{R}_+}$, for every $n \in \mathbf{N}$.[51] It follows by the bounded convergence theorem that the right-hand side of (4) converges to a limit exceeding

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

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

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

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

Unlike the super-solution property of $v_\star$, the sub-solution property of $v^\star$ holds for general u:

Lemma 1(b). $v^\star$ is a viscosity sub-solution 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 $u^\star$.

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

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

We shall derive a contradiction.

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

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

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

$$k_\eta \coloneqq \min_{q\in\partial B_\eta} |\phi(q) - v^\star(q)|,$$

[51]This 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 $\phi - v^\star$ at p is strict. Since ϕ and ϕ'' are continuous and $u^\star$ is upper semi-continuous, we may find an $\eta > 0$ and an $\varepsilon \in (0, k_\eta]$ small enough that

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

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

$$\gamma_n \coloneqq \phi(p_n) - v(p_n)$$

vanishes. Let $(\lambda_t^n)_{t\in\mathbf{R}_+}$ be an $\varepsilon/2$-best reply in the principal'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 behaviour is even better, so certainly attains value at least $v(p_n) - \varepsilon/2$:

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

Subtracting $\phi(p_n)$ from both sides and using the fact that

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

yields

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

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

$$\mathrm{d}P_t^n = \lambda_t^n \frac{P_t^n\left(1 - P_t^n\right)}{\sigma} \mathrm{d}B_t$$

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

$$\begin{aligned} e^{-r\tau_n}\phi\left(P_{\tau_n}^n\right) = \phi\left(p_n\right) - r \int_0^{\tau_n} e^{-rt} \phi\left(P_t^n\right) \mathrm{d}t \\ + \frac{1}{2} \int_0^{\tau_n} \left(\sqrt{\lambda_t^n} \frac{P_t^n\left(1 - P_t^n\right)}{\sigma} \right)^2 e^{-rt} \phi''\left(P_t^n\right) \mathrm{d}t. \end{aligned}$$

Substituting in (6) and using (5) yields

$$
\begin{aligned}
& -\gamma_n - \frac{\varepsilon}{2} \leq \mathbf{E}\Bigg(-e^{-r\tau_n}\varepsilon \\
& \quad + r\int_0^{\tau_n} e^{-rt}\left[-\phi\left(P_t^n\right) + u\left(P_t^n\right) + \frac{\left(P_t^n\right)^2\left(1-P_t^n\right)^2}{2r\sigma^2}\lambda_t^n \phi''\left(P_t^n\right)\right]\mathrm{d}t\Bigg) \\
& \leq \mathbf{E}\Bigg(-e^{-r\tau_n}\varepsilon \\
& \quad + r\int_0^{\tau_n} e^{-rt}\left[-\phi\left(P_t^n\right) + u\left(P_t^n\right) + \frac{\left(P_t^n\right)^2\left(1-P_t^n\right)^2}{2r\sigma^2}\max\left\{0, \phi''\left(P_t^n\right)\right\}\right]\mathrm{d}t\Bigg) \\
& \leq \mathbf{E}\left(-e^{-r\tau_n}\varepsilon + r\int_0^{\tau_n} e^{-rt}(-\varepsilon)\mathrm{d}t\right) = -\varepsilon.
\end{aligned}
$$

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

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

Lemma 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 $v(p_n) \to u(0) = v(0)$.

At each $n \in \mathbf{N}$, consider the auxiliary problem in which the principal 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 principal can induce in her best-reply problem is available in the auxiliary problem. And we have $V(p_n) \leq (\operatorname{cav} u)(p_n)$ since the auxiliary problem may broken down into a sequence of independent static persuasion problems, in each of which the optimal value is at most $(\operatorname{cav} u)(p_n)$.[52] Thus we have

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

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

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

$$\underline{v} \coloneqq \liminf_{q\to p} v(q) \quad \text{and} \quad \overline{v} \coloneqq \limsup_{q\to p} v(q).$$

[52] If u is upper semi-continuous, then $V = \operatorname{cav} u$.

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

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

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 principal's best-reply problem with prior $p_0 = \underline{p}_n$, consider a strategy that sets $\lambda = 1$ 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

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

The left-hand side converges to $\underline{v}$ as $n \to \infty$. The hitting time τ_n vanishes a.s., and $v\left(P_{\tau_n}^n\right)$ converges a.s. to $\bar{v}$. Furthermore, u and v are bounded by piecewise continuity.[53] 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 principal's problem with $p_0 = \underline{p}_n$, consider a strategy that sets $\lambda = 1$ while $p_t \in \left(\underline{p}_{n-1}, \bar{p}_n\right)$ 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 $\left\{\underline{p}_{n-1}, \bar{p}_n\right\}$. The optimal value must exceed the value from using this strategy:

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

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

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

for some $\gamma_n \in (0, 1)$, and the sequences $\left(\underline{p}_n\right)_{n\in\mathbf{N}}$ 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 $\gamma < 1$. ■

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

B Proof of Lemma 2 (p. 21)

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 semi-continuous, we have $v > u$ on an open neighbourhood 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 neighbourhood 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 Lemma 1 (p. 18) and the fact that $v > u$ on N, v is a viscosity solution of

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

on N. Observe that u is continuous on N.

We show (constructively) in supplemental appendix F 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 $(0,1) \setminus C$. ■

C Proof of Lemma 3 (p. 21)

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 $\phi : (0,1) \to \mathbf{R}$ with $\phi''(p) = k$ such that $v - \phi$ is locally minimised at p.[54] Since v is a viscosity super-solution of (HJB) by Lemma 1, it follows that

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

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

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

D Proof of Lemma 4 (p. 21)

Since v is locally convex on $C \setminus D$, v'' is non-negative whenever it exists. Thus by Lemma 1 (p. 18), v is a viscosity solution of the differential equation

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

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

In supplemental appendix F, we show constructively that (9) has a classical solution on $C \setminus D$ that can be extended to a continuous function $w^\dagger : C \to \mathbf{R}$ satisfying the (Dirichlet) boundary condition $w^\dagger = u$ on ∂C. By the comparison principle (e.g. Theorem 3.3. in Crandall et al. (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 Aleksandrov theorem (e.g. Theorem A.2 in Crandall et al. (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.[55] 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 $C \setminus D$. ■

Supplemental appendix

E Welfare with heterogeneous priors (§6)

In this appendix, we give a condition on preferences under which when the conflict of interest is purely belief-based, the principal is better-off the closer the agent's prior is to her own. For simplicity, we focus on equilibria in which the agent uses a pure strategy $A : [0,1] \to \mathcal{A}$.

Since the agent's preferences differ from the principal's only due to their belief disagreement, it is natural to conjecture that the principal is better-off the smaller the prior disagreement. This intuition is incorrect in general because for arbitrary preferences f, the principal need not always

[55] This is a consequence of the Baire category theorem—see Bruckner and Leonard (1966, p. 27).

prefer actions taken by an agent with a belief closer to her own.[56] But the assertion is true for expected-utility preferences, and more generally under a single-crossing assumption:

Proposition 5. Suppose that $f_P = f_a = f$, and that the actions $\mathcal{A}$ can be ordered so that f is strictly single-crossing.[57] Then for a fixed p_0 and $p_{a,0} < p_0$ ($p_{a,0} > p_0$), $v(p_0)$ increases pointwise as $p_{a,0}$ increases (decreases).

Strict single-crossing is satisfied whenever f has the expected-utility form.[58] Its role is to ensure that the principal prefers actions taken by an agent with a belief closer to her own. This is a general property of single-crossing, embodied in the following comparative statics lemma.

Lemma 5. Let $(\mathcal{T}, \gtrsim)$ and $(\mathcal{X}, \succeq)$ be partially ordered sets, let $f : \mathcal{T} \times \mathcal{X} \to \mathbf{R}$ be strictly single-crossing, and let $x : \mathcal{T} \to \mathcal{X}$ be a selection from $t \mapsto \arg\max_{x \in \mathcal{X}} f(x,t)$. Then for $t_1 \gtrsim t_2 \gtrsim t_3$ in $\mathcal{T}$, we have $f(x(t_2), t_1) \geq f(x(t_3), t_1)$ and $f(x(t_2), t_3) \geq f(x(t_1), t_3)$.

Proof. Since f is strictly single-crossing, x must be increasing by the monotone selection theorem of Milgrom and Shannon (1994), so in particular $x(t_2) \geq x(t_3)$. Since $x(t_2)$ is optimal at parameter t_2, we have $f(x(t_2), t_2) \geq f(x(t_3), t_2)$. Since $x(t_2) \succeq x(t_3)$ and $t_1 \gtrsim t_2$, and f is single-crossing, it follows that $f(x(t_2), t_1) \geq f(x(t_3), t_1)$. The argument for $f(x(t_2), t_3) \geq f(x(t_1), t_3)$ is analogous. ■

Proof of Proposition 5. Consider $p_{a,0} < p'_{a,0} < p_0$; the other case is analogous. Write $A : [0,1] \to \mathcal{A}$ for the agent's pure, myopic and regular Markov strategy. Write u and u' for the corresponding induced flow payoff for the principal, and v and v' for her value function.

[56] Consider $\mathcal{A} = \{0,1\}$, $f(0,p) = 1/2$ and $f(1,p) = 1 - \mathbf{1}_{(1/3,2/3)}(p)$, and $p_0 = 5/6$. If the agent's prior is $p_{a,0} = 1/6$, then the principal can (and does) earn 1 forever by never funding, so $v(p_0) = 1$. If instead $p_{a,0} = 1/2$, then $v(p_0) < 1$ since she earns $1/2$ early on.

[57] That is: there is a partial order $\succeq$ on $\mathcal{A}$ such that for $a' \succ a$ and $p' > p$, $f(a',p) \geq f(a,p)$ implies $f(a',p') > f(a,p')$.

[58] An expected utility f satisfies $f(a,p) = (1-p)u_0(a) + pu_1(a)$ for some maps $u_0, u_1 : \mathcal{A} \to \mathbf{R}$. Define an order $\succeq$ on $\mathcal{A}$ by $a' \succeq a$ iff $u_1(a') - u_0(a') \geq u_1(a) - u_0(a)$. Then f is strictly single-crossing: for given $a' \succ a$ and $p' > p$, $f(a',p) \geq f(a,p)$ is equivalent to

$$p([u_1(a') - u_0(a')] - [u_1(a) - u_0(a)]) \geq u_0(a) - u_0(a'),$$

which since $p' > p$ and $a' \succ a$ implies that

$$p'([u_1(a') - u_0(a')] - [u_1(a) - u_0(a)]) > u_0(a) - u_0(a'),$$

which is equivalent to $f(a',p') > f(a,p)$.

Since A is myopic, it is a selection from

$$p \mapsto \arg\max_{a \in \mathcal{A}} f(a, p).$$

Fix any belief $p \in [0,1]$ of the principal. Write $p_{\mathrm{a}} \coloneqq \phi(p, p_0, p_{\mathrm{a},0})$ and $p'_{\mathrm{a}} \coloneqq \phi(p, p_0, p'_{\mathrm{a},0})$ for the agent's corresponding beliefs given her prior, and note that $p_{\mathrm{a}} \leq p'_{\mathrm{a}} \leq p$. Since f is strictly single-crossing, it follows by Lemma 5 that

$$\begin{aligned} u'(p) = f(A(\phi(p, p_0, p_{\mathrm{a}',0})), p) &= f(A(p'_{\mathrm{a}}), p) \\ &\geq f(A(p_{\mathrm{a}}), p) = f(A(\phi(p, p_0, p_{\mathrm{a},0})), p) = u(p). \end{aligned}$$

Since $p \in [0,1]$ was arbitrary, this shows that $u' \geq u$, which obviously implies that $v'(p_0) \geq v(p_0)$. ■

F Solving for the value function

In this appendix, we explain how to solve for the principal's value function using Proposition 1. We detail in particular how the value is computed in the two- and three-action examples, allowing us to draw Figures 3 and 4.

Partition $C \setminus D$ into maximal intervals $(R_k)_{k=1}^K$.[59] Fix a continuity interval R_k. The homogeneous part (without the u) of the differential equation (∂) in Proposition 1 has general solution $AH_1 - BH_2$ for constants $A, B \in \mathbf{R}$, where

$$H_1(p) \coloneqq p^{\xi}(1-p)^{1-\xi} \quad \text{and} \quad \xi \coloneqq 1/2 + \sqrt{1/4 + 2r\sigma^2},$$

and $H_2(p) \coloneqq H_1(1-p)$.

A particular solution may be obtained from formula (6.2) in Coddington (1961, ch. 3). Things are easier when the principal has expected-utility preferences, so that u is affine, as u itself is then a particular solution. This is the case in the two- and three-action examples. In the expected-utility case, the value function is given on each maximal interval R_k of $C \setminus D$ as

$$v(p) = u(p) + A_{R_k} H_1(p) - B_{R_k} H_2(p) \quad \text{for all } p \in R_k,$$

where the constants $(A_{R_k}, B_{R_k})_{k=1}^K$ are the unique ones that ensure that the properties in Proposition 1 are satisfied: the boundary condition $v = u$ on $\{0, 1\}$, the continuity of v on D, and smooth pasting on $C \cup D$.

[59] In the two-action example, the maximal intervals are $[0, 2/3)$ and $(2/3, 1]$. In the three-action example, they are $[0, 1/2)$, $(1/2, 3/4)$ and $(3/4, 1]$.

F.1 The two-action example (§1.1)

Here $D = \{2/3\}$, and C contains $[0, 2/3)$ and may or may not contain $[2/3, 1]$. In either case,

$$v(p) = \begin{cases} A_{[0,2/3)}H_1(p) - B_{[0,2/3)}H_2(p) & \text{for } p \in [0, 1/2) \\ \alpha p - \beta + A_{(2/3,1]}H_1(p) - B_{(2/3,1]}H_2(p) & \text{for } p \in (2/3, 1], \end{cases}$$

where $\alpha = 3/2$ and $\beta = 1/2$.[60] The boundary conditions require that $B_{[0,2/3)} = 0$ and $A_{(2/3,1]} = 0$. Continuity of v at $2/3$ requires that

$$A_{[0,2/3)}H_1(2/3) = \alpha(2/3) - \beta - B_{(2/3,1]}H_2(2/3).$$

If $r\sigma^2$ is sufficiently low, then $2/3 \in C$, in which case smooth pasting must hold at $2/3$:

$$A_{[0,2/3)}H_1'(2/3) = \alpha - B_{(2/3,1]}H_2'(2/3).$$

Thus the constants are uniquely pinned down.

If $r\sigma^2$ is high, then $2/3 \notin C$, in which case $v = u$ on $[2/3, 1]$. Thus $B_{(2/3,1]} = 0$, whence $A_{[0,2/3)}$ is uniquely pinned down by the continuity condition.

To determine which case applies for a given value of $r\sigma^2$, calculate $A_{[0,2/3)}$ assuming that the first case applies. If

$$A_{[0,2/3)}H_1(2/3) \geq u(1/2) = 1/2,$$

then the first case does indeed apply; if not, then not.

F.2 The three-action example (Figure 4)

Clearly C contains $[0, 1/2)$ and $(1/2, 3/4)$, and does not contain $[3/4, 1]$. Thus the value function off D is

$$v(p) = \begin{cases} A_{[0,1/2)}H_1(p) - B_{[0,1/2)}H_2(p) & \text{for } p \in [0, 1/2) \\ \ell + A_{(1/2,3/4)}H_1(p) - B_{(1/2,3/4)}H_2(p) & \text{for } p \in (1/2, 3/4) \\ h, & \end{cases}$$

[60] If C contains $[2/3, 1]$ then the expression for $p \in (2/3, 1]$ holds since (HJB) must be satisfied in the classical sense by Proposition 1. If not, then Proposition 1 requires that $v = u$, which amounts to setting $A_{(2/3,1]} = B_{(2/3,1]} = 0$.

where $\ell = 1$ and $h = 3$. The boundary condition at $p = 0$ again requires that $B_{[0,1/2)} = 0$. Continuity of v at $1/2$ and at $3/4$ requires that

$$A_{[0,1/2)}H_1(1/2) = \ell + A_{(1/2,3/4)}H_1(1/2) - B_{(1/2,3/4)}H_2(1/2)$$

and

$$\ell + A_{(1/2,3/4)}H_1(3/4) - B_{(1/2,3/4)}H_2(3/4) = h.$$

These are two equations in three unknowns.

If $r\sigma^2$ is sufficiently low that $1/2 \in C$, then smooth pasting must hold at $1/2$, giving us the third equation

$$A_{[0,1/2)}H_1'(1/2) = A_{(1/2,3/4)}H_1'(1/2) - B_{(1/2,3/4)}H_2'(1/2).$$

If $r\sigma^2$ is sufficiently high that $1/2 \notin C$, then $v(1/2) = u(1/2) = \ell$. We thus obtain a third equation from the requirement that v be continuous at $1/2$:

$$A_{[0,1/2)}H_1(1/2) = \ell.$$

To discern which case applies, compute $A_{[0,1/2)}$ assuming that the first (patient) case applies. If

$$A_{[0,1/2)}H_1(1/2) \geq \ell,$$

then the patient case does indeed apply; otherwise, it does not.

G Generic uniqueness of long-run beliefs

We claimed in §5.1 that provided $v(p_0) > u(p_0)$, generically, all best replies of the principal induce the same long-run beliefs (viz. the beliefs $\{p^-, p^+\}$ defined in Corollary 3 (p. 24)).

To see how uniqueness can fail, consider the three-action example from Figure 4 (p. 20). Figure 6a depicts the knife-edge case in which $r\sigma^2$ is such that the patient-case value function with the convex-flat shape in Figure 4c touches u at $1/2$.[61] In this case, the principal is indifferent between funding and not funding at $1/2$, and strictly prefers to fund on $(0, 1/2)$ and $(0, 3/4)$. The best reply $\Lambda^\star$ from Corollary 2 (p. 23) stops at $1/2$, inducing the long-run beliefs $\{p^-, p^+\} = \{1/2, 3/4\}$ from Corollary 3. But since the principal is indifferent at $1/2$, she also has a best reply that funds at $1/2$, and this induces the long-run beliefs $\{0, 3/4\}$.

[61] We thank Jeff Ely for pointing out this scenario.

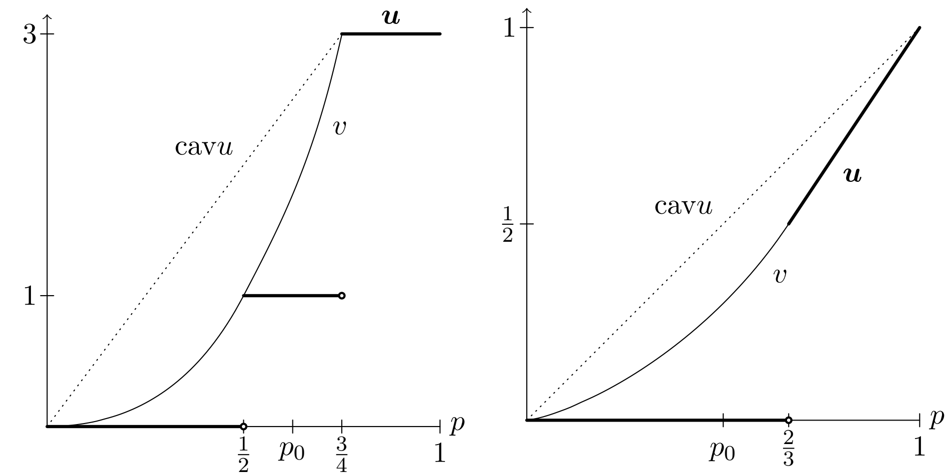

(a) Three-action example from Figure 4.

(b) Two-action example from §1.1.

Figure 6 – Knife-edge cases in which long-run beliefs are not unique.

This scenario is non-generic in the sense that slightly decreasing $r\sigma^2$ puts us back in Figure 4c, where the principal strictly prefers to fund at $1/2$, whereas slightly increasing $r\sigma^2$ puts us in Figure 4d, where she strictly prefers to stop at $1/2$.

Similarly, Figure 6b depicts the case in the two-action example of §1.1 in which $r\sigma^2$ has exactly the value needed for the patient-case value function with the convex shape in Figure 3a to just touch u at $2/3$. In this example, there is more multiplicity: the principal is indifferent on $[1/2, 1]$, so has best replies that induce any mean-p_0 distribution of long-run beliefs supported on $\{0\} \cup [2/3, 1]$. (The best reply $\Lambda^\star$ induces the beliefs $\{0, 2/3\}$.) Again, perturbing $r\sigma^2$ makes the principal's preference strict at $2/3$, so that long-run induced beliefs are unique (either $\{0, 2/3\}$ or $\{0, 1\}$).

The non-genericity of multiplicity in these examples is a general phenomenon. Multiplicity occurs for some prior p_0 with $v(p_0) > u(p_0)$ precisely if the principal is indifferent between stopping and continuing at some $p \in (0, 1)$ and weakly prefers to continue on a neighbourhood of p. In such cases, her preference becomes strict when $r\sigma^2$ is perturbed slightly.

H An example with gradual convergence

In the two- and three-action examples, the long-run induced beliefs of a sufficiently patient planner coincide exactly with those in the persuasion benchmark. Figure 7 shows that this need not happen: here the long-run induced beliefs converge to the persuasion beliefs P^- and P^+ as $r\sigma^2 \to 0$ as per Proposition 3, but remain strictly less extreme for every $r\sigma^2 > 0$.

I Piecewise continuity is merely tie-breaking

We asserted in §2.3 that provided the flow payoff f_a is non-degenerate in a mild sense, it is without loss of optimality for her to restrict attention to piecewise continuous Markov strategies $A : [0,1] \to \Delta(\mathcal{A})$.

To justify this claim, begin by recalling from §3 that the agent best-replies to a Markov strategy of the principal by myopically maximising $f_a(a,p)$ at each p.[62] Fix two actions $a, a' \in \mathcal{A}$, and write

$$\psi(p) \coloneqq f_a(a,p) - f_a(a',p)$$

for their payoff difference. Say that ψ *strictly up-crosses* at $p \in (0,1)$ iff $\psi(p) = 0$ and for any $\varepsilon > 0$, there are $p' \in (p-\varepsilon, p)$ and $p'' \in (p, p+\varepsilon)$ such that $\psi(p') < 0 < \psi(p'')$, *strictly down-crosses* if the reverse inequalities hold, and simply *strictly crosses* if either is the case. Write $K \subseteq (0,1)$ for the set on which ψ strictly crosses. We claim that given some weak non-degeneracy condition on f_a, the crossing set K is discrete, so that the agent strictly prefers to switch actions only on a discrete set. (It suffices to consider only two arbitrary actions $a, a' \in \mathcal{A}$ because $\mathcal{A}$ is finite.)

To see what can go wrong, suppose that $f_a(a,p) = 0$ and that $p \mapsto f_a(a',p)$ is a typical path of a standard Brownian motion. Then ψ is continuous, but the strict crossing set K is non-empty with no isolated points (see e.g. Theorem 9.6 in Karatzas and Shreve (1991, ch. 2)). This preference dithers maniacally, wishing to switch actions back and forth extremely frequently.

As a first pass, observe that if ψ is monotone, or more generally if ψ or $-\psi$ has the single-crossing property ($\psi(p) \geq (>)\ 0$ implies $\psi(p') \geq (>)\ 0$ for $p < p'$), then K is empty or a singleton, so certainly discrete. These assumptions are satisfied by expected-utility preferences.

A weak non-degeneracy condition that suffices is *local single-crossing:* for each $p \in K$, we have either $\psi \geq 0$ or $\psi \leq 0$ on a left-neighbourhood of p, and similarly on a right-neighbourhood. Then each $p \in K$ is manifestly the unique strict crossing of ψ on a neighbourhood, hence isolated. A sufficient condition for this is *local monotonicity:* for each $p \in K$, we have $\psi(p-\varepsilon) \leq 0 \leq \psi(p+\varepsilon)$ for all sufficiently small $\varepsilon > 0$, or the reverse inequality.

[62] This remains true if the agent is allowed to use any map $A : [0,1] \to \Delta(\mathcal{A})$.

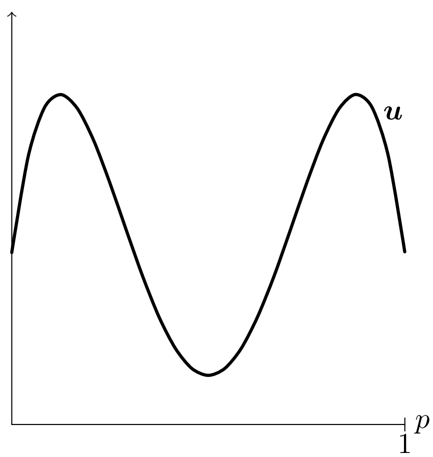

(a) Flow payoff u.

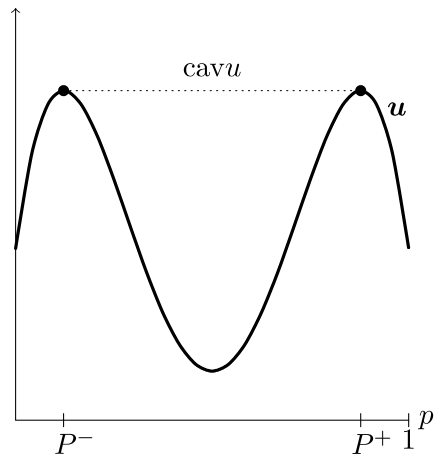

(b) Static benchmark value cav u.

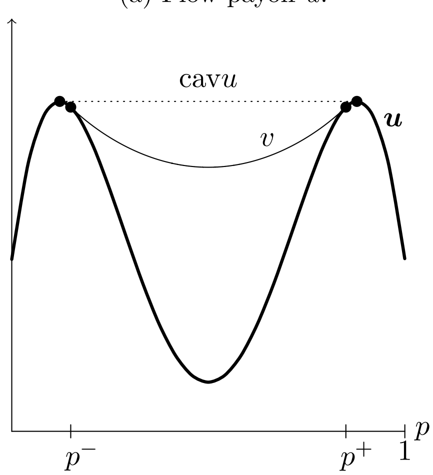

(c) Value v for low $r\sigma^2$.

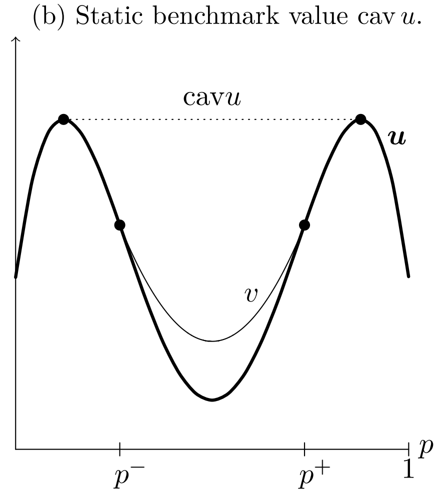

(d) Value v for high $r\sigma^2$.

Figure 7 – Example in which the long-run beliefs $\{p^-, p^+\}$ are strictly less extreme than the static persuasion beliefs $\{P^-, P^+\}$ for every $r\sigma^2 > 0$.

J A very brief introduction to viscosity solutions

Crandall (1997), Katzourakis (2015) and Crandall et al. (1992) provide overviews of the theory of viscosity solutions of second-order differential equations. Moll (2017), Evans (2010, ch. 10), Calder (2018) and Bressan (2011) give easier treatments that deal mostly with first-order equations.

The general idea of viscosity solutions is as follows. If w is a viscosity solution of (HJB), then it must satisfy (HJB) in the classical sense on any neighbourhood on which w'' exists and is continuous. If w'' does not exist at $p \in [0, 1]$, we require instead that (HJB) hold with the appropriate inequality when $w''(p)$ is replaced by $\phi''(p)$ for some twice continuously differentiable local approximation ϕ to w at p. (The formal definition was given on p. 17.)

J.1 Illustration of the definition

Consider the three-action example from Figure 4a (p. 20). Write $\mathcal{C}^2$ for the set of twice continuously differentiable functions $(0, 1) \to \mathbf{R}$. Begin by observing that v is continuous, hence upper and lower semi-continuous.

Consider a p in whose vicinity v is twice continuously differentiable, e.g. $p = 2/5$. We may easily find $\phi_1, \phi_2 \in \mathcal{C}^2$ such that $\phi_1 - v$ and $v - \phi_2$ are locally minimised at p, as in Figure 8a. But in particular, we may choose $\phi \in \mathcal{C}^2$ to coincide with v on a neighbourhood of p. Then $\phi - v$ and $v - \phi$ are *both* locally minimised at p, and $\phi''(p) = v''(p)$. Since v is a viscosity sub-solution (super-solution) by Lemma 1 (p. 18), and $u(p) = u^\star(p) = u_\star(p)$, it follows that

$$v(p) \leq (\geq)\ u(p) + \frac{p^2(1-p)^2}{2r\sigma^2} \max\left\{0, v''(p)\right\}.$$

So (HJB) must be satisfied in the classical sense at p.

Next consider a point at which v'' is undefined, e.g. $p = 1/2$. There are many $\phi \in \mathcal{C}^2$ such that $\phi - v$ has a local minimum at p; an example is depicted in Figure 8b. Since v is a viscosity sub-solution of (HJB) and $u^\star(p) = u(p)$, we must have

$$v(p) \leq u(p) + \frac{p^2(1-p)^2}{2r\sigma^2} \max\left\{0, \phi''(p)\right\}$$

for any such ϕ. In fact, ϕ can be chosen so that $\phi''(p) \leq 0$: the ϕ depicted in Figure 8c is affine, so has $\phi''(p) = 0$. The sub-solution condition therefore requires precisely that

$$v(p) \leq \inf_{\substack{\phi \in \mathcal{C}^2: \\ \phi - v \text{ loc. min. at } p}} \left\{ u(p) + \frac{p^2(1-p)^2}{2r\sigma^2} \max\left\{0, \phi''(p)\right\} \right\} = u(p),$$

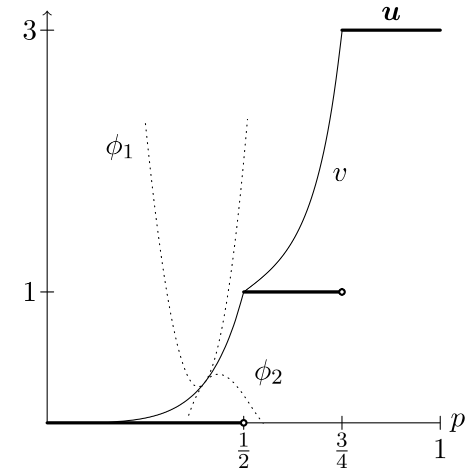

(a) $\phi_1, \phi_2 \in \mathcal{C}^2$ such that $\phi_1 - v$ and $v - \phi_2$ have local minima at 2/5.

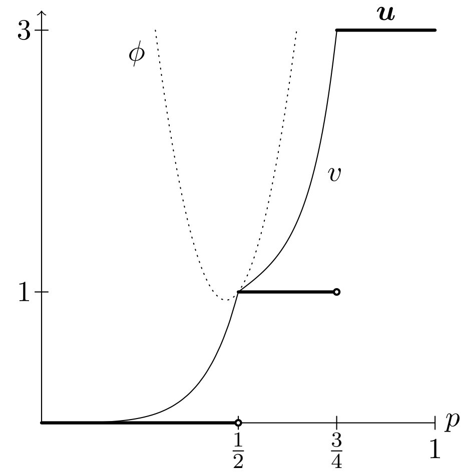

(b) $\phi \in \mathcal{C}^2$ such that $\phi - v$ has a local minimum at 1/2.

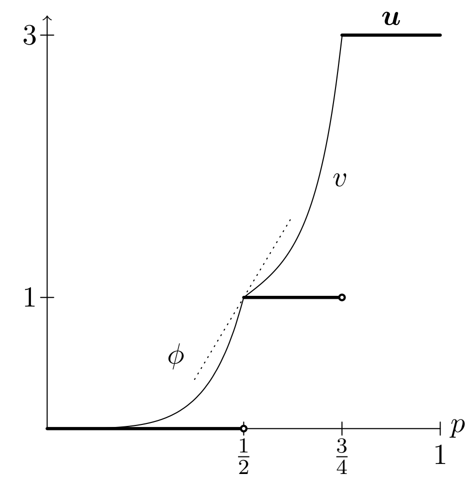

(c) $\phi \in \mathcal{C}^2$ such that $\phi - v$ has a local minimum at 1/2 and $\phi''(1/2) = 0$.

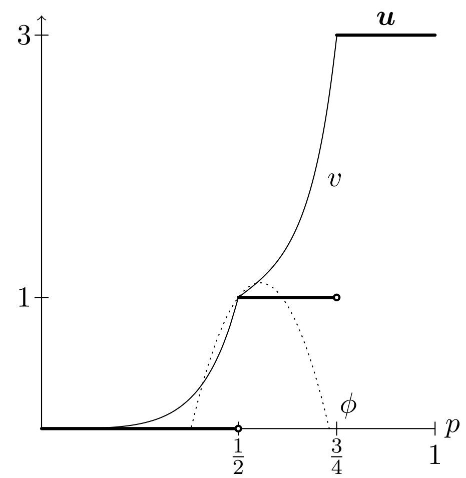

(d) $\phi \in \mathcal{C}^2$ for which $v - \phi$ does not have a local minimum at 1/2.

Figure 8 – Functions $\phi \in \mathcal{C}^2$ that approximate v locally.

which holds (with equality, in fact).

By contrast, there are no $\phi \in \mathcal{C}^2$ such that $v - \phi$ has a local minimum at p; a (failed) attempt to find such a ϕ is drawn in Figure 8d. The fact that v is a viscosity super-solution of (HJB) therefore has no bite at $p = 1/2$.

J.2 Some properties of viscosity solutions

There are other non-classical notions of 'solution' of a differential equation, most importantly distributional solutions (e.g. Evans (2010, chs. 5–9)). But for many differential equations, including HJB equations, viscosity solutions are the appropriate notion. The chief reasons are twofold: viscosity solutions exist, and they satisfy a comparison principle.

Begin with existence. Many HJB equations, including ours, fail to have a classical solution. Many also fail to have non-classical solutions of e.g. the distributional variety. By contrast, HJB equations always have a viscosity solution.

The other principal virtue of viscosity solutions is that they satisfy a comparison principle (also called a 'maximum principle') of the following kind: if $\underline{w}$ is a sub-solution on (a, b), $\overline{w}$ is a super-solution on (a, b), and $\underline{w} \leq \overline{w}$ on $\{a, b\}$, then $\underline{w} \leq \overline{w}$ on (a, b). (See Crandall et al. (1992, Theorem 3.3).) Classical sub- and super-solutions also satisfy a comparison principle, but other non-classical notions of 'solution' do not.

The comparison principle may be used to obtain uniqueness results; a standard one is that the HJB equation has at most one viscosity solution with the right boundary conditions satisfying a linear-growth condition. It follows that the value function is the unique solution with the right boundary conditions and linear growth. (See Fleming and Soner (2006, ch. V).) We use the comparision principle in this manner in the proofs of Lemmata 2 and 4 (appendices B and D).

The comparison principle may also be used to establish the continuity of solutions, and thus of the value function. In particular, suppose that we have shown that the upper (lower) semi-continuous envelope $v^\star$ ($v_\star$) of the value v is a sub-solution (super-solution) of the HJB equation, and that $v_\star = v^\star$ on $\{0, 1\}$. (We do precisely this in the proof of Lemma 1 in appendix J.) A comparison principle then yields $v^\star \leq v_\star$, which since $v_\star \leq v \leq v^\star$ implies that v is itself a viscosity solution, hence continuous.

In our proof of Lemma 1 (appendix A), we eschew this approach in favour of a direct proof that v is continuous. We do this because we are not aware of a comparision principle that applies assuming only piecewise continuity of u. The closest result that we know of is Theorem 3.3 in Soravia (2006), which would be applicable under the additional hypotheses that u has only *finitely* many discontinuities and satisfies $u(p) \in [u(p-) \wedge u(p+), u(p-) \vee u(p+)]$ at every $p \in (0,1)$.

References

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., & Moll, B. (2017). *Income and wealth distribution in macroeconomics: A continuous-time approach*. NBER working paper 23732, Aug 2017. doi:10.3386/w23732

Argenziano, R., Severinov, S., & Squintani, F. (2016). Strategic information acquisition and transmission. *American Economic Journal: Microeconomics*, *8*(3), 119–155. doi:10.1257/mic.20140135

Aumann, R. J. & Maschler, M. B. (1995). *Repeated games with incomplete information*. Cambridge, MA: MIT Press.

Billingsley, P. (1995). *Probability and measure* (3rd). New York, NY: Wiley.

Bizzotto, J., Rüdiger, J., & Vigier, A. (2017). *How to persuade a long-run decision maker*. Working paper, Apr 2017. doi:10.2139/ssrn.2706683

Bolton, P. & Harris, C. (1999). Strategic experimentation. *Econometrica*, *67*(2), 349–374. doi:10.1111/1468-0262.00022

Bressan, A. (2011). *Viscosity solutions of Hamilton–Jacobi equations and optimal control problems*. Lecture notes. Retrieved from https://www.math.psu.edu/bressan/PSPDF/HJ-lnotes.pdf

Brocas, I. & Carrillo, J. D. (2007). Influence through ignorance. *RAND Journal of Economics*, *38*(4), 931–947. doi:10.1111/j.0741-6261.2007.00119.x

Bruckner, A. M. & Leonard, J. L. (1966). Derivatives. *American Mathematical Monthly*, *73*(4), 24–56. doi:10.2307/2313749

Calder, J. (2018). *Lecture notes on viscosity solutions*. Lecture notes. Retrieved from http://www-users.math.umn.edu/~jwcalder/8590F18/viscosity_solutions.pdf

Coddington, E. A. (1961). *An introduction to ordinary differential equations*. New York, NY: Dover.

Crandall, M. G. (1997). Viscosity solutions: A primer. In I. Capuzzo Dolcetta & P.-L. Lions (Eds.), *Viscosity solutions and applications* (pp. 1–43). Berlin: Springer.

Crandall, M. G., Ishii, H., & Lions, P.-L. (1992). User's guide to viscosity solutions of second order partial differential equations. *Bulletin of the American Mathematical Society*. New Series, *27*(1), 31–53. doi:10.1090/S0273-0979-1992-00266-5

Crandall, M. G. & Lions, P.-L. (1983). Viscosity solutions of Hamilton–Jacobi equations. *Transactions of the American Mathematical Society*, *277*, 1–42. doi:10.1090/S0002-9947-1983-0690039-8

Ely, J. C. (2017). Beeps. *American Economic Review*, *107*(1), 31–53. doi:10.1257/aer.20150218

Ely, J. C. & Szydlowski, M. (2017). *Moving the goalposts*. Working paper, 11 Aug 2017.

Evans, L. C. (2010). *Partial differential equations* (2nd). Graduate Studies in Mathematics. Providence, RI: American Mathematical Society.

Fleming, W. H. & Soner, H. M. (2006). *Controlled Markov processes and viscosity solutions* (2nd). New York, NY: Springer.

Frug, A. (2017). Strategic gradual learning and information transmission. *Journal of Economic Theory*, *177*, 594–615. doi:10.1016/j.jet.2018.07.006

Fudenberg, D. & Tirole, J. (1991). *Game theory*. Cambridge, MA: MIT Press.

Henry, E. & Ottaviani, M. (2019). Research and the approval process: The organization of persuasion. *American Economic Review*, *109*(3), 911–955. doi:10.1257/aer.20171919

Isaacs, R. (1954). *Differential games, I–IV* (tech. rep. No. RM–1391, RM–1399, RM–1411 and RM–1486). RAND Corporation.

Isaacs, R. (1965). *Differential games: A mathematical theory with applications to warfare and pursuit, control and optimization*. Mineola, NY: Dover.

Kamenica, E. & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, *101*(6), 2590–2615. doi:10.1257/aer.101.6.2590

Karatzas, I. & Shreve, S. E. (1991). *Brownian motion and stochastic calculus* (2nd). Graduate Texts in Mathematics. New York, NY: Springer.

Katzourakis, N. (2015). *An introduction to viscosity solutions for fully nonlinear PDE with applications to calculus of variations in L^∞*. SpringerBriefs in Mathematics. Cham, CH: Springer.

Lowther, G. (2010). *Time-changed Brownian motion*. Retrieved from https://almostsure.wordpress.com/2010/04/20/time-changed-brownian-motion/

Milgrom, P. & Shannon, C. (1994). Monotone comparative statics. *Econometrica*, *62*(1), 157–180. doi:10.2307/2951479

Moll, B. (2017). *Viscosity solutions for dummies (including economists)*. Slides. Retrieved from http://www.princeton.edu/~moll/viscosity_slides.pdf

Morris, S. & Strack, P. (2017). *The Wald problem and the equivalence of sequential sampling and static information costs*. Working paper, 6 Sep 2017. doi:10.2139/ssrn.2991567

Nordhaus, W. D. (2007). A review of the *Stern Review on the Economics of Climate Change*. *Journal of Economic Literature*, *45*(3), 686–702. doi:10.1257/jel.45.3.686

Orlov, D., Skrzypacz, A., & Zryumov, P. (2018). *Persuading the principal to wait*. Working paper, 11 Jan 2018. doi:10.2139/ssrn.2767234

Papanicolaou, A. (2016). *Stochastic analysis seminar on filtering theory*. Lecture notes. arXiv: 1406.1936

Pei, H. D. (2015). Communication with endogenous information acquisition. *Journal of Economic Theory*, *160*, 132–149. doi:10.1016/j.jet.2015.08.011

Pham, H. (2009). *Continuous-time stochastic control and optimization with financial applications*. Stochastic Modelling and Applied Probability. Springer-Verlag.

Pham, H. (2012). *Errata for 'Continuous-time stochastic control and optimization with financial applications'*. Retrieved from https://docs.google.com/file/d/0BwiI-nyqtVlRVk9DSVpNSEl2Y3c/view

Quah, J. K.-H. & Strulovici, B. (2013). Discounting, values, and decisions. *Journal of Political Economy*, *121*(5), 896–939. doi:10.1086/673867

Rayo, L. & Segal, I. (2010). Optimal information disclosure. *Journal of Political Economy*, *118*(5), 949–987. doi:10.1086/657922

Renault, J., Solan, E., & Vieille, N. (2017). Optimal dynamic information provision. *Journal of Political Economy*, *104*, 329–349. doi:10.1016/j.geb.2017.04.010

Rockafellar, R. T. (1970). *Convex analysis*. Princeton, NJ: Princeton University Press.

Rudin, W. (1976). *Principles of mathematical analysis* (3rd). New York, NY: McGraw-Hill.

Siegel, R. & Strulovici, B. (2018). *Multiverdict systems*. Working paper, 2 Apr 2018.

Simon, L. K. & Stinchcombe, M. B. (1989). Extensive form games in continuous time: Pure strategies. *Econometrica*, *57*(5), 1171–1214. doi:10.2307/1913627

Soravia, P. (2006). Uniqueness results for fully nonlinear degenerate elliptic equations with discontinuous coefficients. *Communications on Pure & Applied Analysis*, *5*(1), 213–240. doi:10.3934/cpaa.2006.5.213

Stern, N. (2007). *The economics of climate change: The Stern Review*. Cambridge, UK: Cambridge University Press.