

Career Concerns and Policy Intransigence - A Dynamic Signalling Model

Caroline D. Thomas*

June 27, 2016

Abstract

We investigate why a political leader (manager) may be reluctant to abandon her chosen project, when she is concerned both about social welfare and the electorate's (shareholders') beliefs about her ability. The quality of the project is correlated with her ability, and there is ex ante symmetric information. However, news regarding the project arrives gradually over time and is privately observed by the leader, who may choose whether and when to cancel the project, so that we have a dynamic signalling model. We find that the inefficiency caused by career concerns varies with the information structure. In the "good news" case, there exists an equilibrium that implements the socially efficient decision rule, provided that career concerns are not too large. With larger career concerns, the most efficient equilibrium is characterised by increasing inefficiency over time - a project that should be abandoned early will be done so efficiently, while projects on which there is more initial good news can only be scrapped with increasing delays. No matter how important career concerns, efficiency can be restored when the informational asymmetry between the leader and the electorate is enhanced. In the "bad news" case, the socially efficient policy cannot be implemented, as the leader, here too, has an incentive to delay scrapping the project, no matter how small her career concerns. However inefficiencies are bounded.

JEL Classification Number: *C73, D82, B83, P16.*

Keywords: Career Concerns, Dynamic Signalling, Political Economy, Poisson Bandits, Strategic Experimentation

*Department of Economics, University of Texas at Austin. Email: caroline.thomas@austin.utexas.edu
Thanks to Yiman Sun for research assistance.

1 Introduction

Many unsuccessful political projects arguably went on for longer than necessary, given the information available at the time. In fact, often the governments proposing the policies in question advocated them even more strongly in the face of evidence of their failure. Following the violent 1990 poll tax riots in the UK, Thatcher continued referring to the poll tax as the flagship of her government's political manifesto. Mao's government pushed on with the Great Leap Forward in China despite growing evidence of famine in many parts of the country.

Why are political leaders who undertake projects often reluctant to give them up, despite evidence that the project is failing? One explanation given in the economics literature is that of career concerns. Political leaders care not only about social welfare, but also about reelection. If a policy's success or failure can be informative about a political leader's competence, these two objectives may conflict. Even though the policy-maker may be privately convinced of a policy's worthlessness, she might resist effecting any change in policy that could be interpreted as a sign that her government was not sure of its policy choice in the first place, and might paradoxically press on with the worthless policy to avoid looking incompetent.

I show that the validity of this argument is sensitive to the underlying information structure.

I consider a policy-maker whose competence-level is uncertain and ex-ante unknown both to herself and the electorate. The information is symmetric ex-ante as in the career concerns model of Holmström (1999). The policy-maker is endowed with a policy which can be good or bad, and whose quality is uncertain and also unknown both to the policy-maker and the electorate. However it is common knowledge that the quality of the project and the competence level of the policy-maker are correlated: competent policy-makers are more likely to be endowed with good policies.

Over time the policy-maker receives private, inconclusive information about the project quality. Informational asymmetries arise gradually, inducing a dynamic signalling game. I show that the decision-maker's behaviour – and particularly whether it departs from the socially optimal benchmark – is sensitive to the information structure. In particular, in the good news case, reputational concerns, even when large, do not preclude social efficiency.

In my model, the project quality determines its social payoff. A bad project never yields a payoff. A good project yields no payoff until it succeeds, at an exponentially distributed random time, generating a social payoff with positive discounted present value $g > 0$. The success of the project is a publicly observable event that proves the project to be good. Until then, the project quality remains uncertain for the policy-maker and the electorate. I will say that the policy-maker “experiments” if she keeps active a project that has not yet succeeded.

Experimentation is socially costly. The status-quo policy yields a social payoff with positive discounted present value $s \in (0, g)$. This can be thought of as the cost of “building”

the project (a bridge or a damn), or as the opportunity-cost of not having a functional policy in place.

Time is continuous and the policy-maker chooses at each point in time whether to continue her new policy, or repeal it and revert to the status-quo policy. Repealing the new policy is irreversible. The policy-maker can base her choice on private information which she receives gradually over time. I model the arrival of private news as a standard Poisson process with state-dependent arrival rates, and distinguish two cases. In the “good news” scenario, news events occur more frequently if the project is good, and can be interpreted as “breakthroughs”. Conversely, in the “bad news” scenario, news events occur more frequently if the project is bad, and can be interpreted as “breakdowns”.

Is it possible for a reputationally concerned policy-maker to implement the socially optimal policy? The socially optimal policy is defined as the policy implemented by a policy-maker who does not care about re-election. Such a policy-maker acts as a social planner concerned only with maximising the expected social value of the project, and faces a two-armed bandit problem. The socially optimal decision-rule is characterised by a threshold belief and requires experimenting with the new project and reverting to the status-quo policy if the policy-maker has become too pessimistic about the likelihood that the policy project might ever succeed, so that her posterior belief has fallen below the socially optimal threshold. Clearly, repealing a project that has already succeeded is never optimal.

It may be of benefit to act so as to convince the observer that her private information is supporting the hypothesis that the project is good, since the DM only partially bears the observer’s cost of experimentation.

A policy-maker concerned with re-election might prefer a different policy. While the project’s success, or lack thereof, and the policy-maker’s decisions are observable to the electorate, the policy-maker’s private information is not. Thus, as long as the project has not succeeded, the electorate does not know the policy-maker’s private assessment of the project quality. For example in the good news scenario, the electorate considers it possible that the policy-maker has not repealed an as yet unsuccessful project because she has received encouraging private information about the project quality. Thus, the best assessment of the project quality the electorate can form is to guess the policy-maker’s private assessment based on the publicly available information and, in equilibrium, on the policy-maker’s strategy.

I refer to the electorate’s assessment of the project quality – and hence, indirectly, of the policy-maker’s competence level – as the policy-maker’s *reputation*. When the policy-maker’s objectives are divided between maximising the social welfare, and maximising her reputation, there is scope for her to improve her reputation by departing from the socially optimal decision-rule.

I look for the most efficient PBE (that which maximises the social payoff). I find that whether or not it falls short of the social planner benchmark depends on the preferences of the DM, the information structure, and the restrictions imposed on her actions.

Whether there exists a (perfect Bayesian) equilibrium in which a reputationally con-

cerned policy-maker can implement the socially optimal decision-rule depends on the information structure. The main insights of this paper are as follows:

- In the “good news” case, there exists an equilibrium that implements the socially efficient decision-rule, provided that reputational concerns are not too large. (Proposition 3.)
- When the socially optimal decision-rule cannot be implemented, the most efficient equilibrium is characterised by increasing inefficiency over time. A project that should be abandoned early will be done so efficiently, while projects on which there is more initial good news can only be scrapped with increasing delays. (Proposition 6.)
- Improving the policy-maker’s private information (and thus increasing the informational asymmetry) mitigates her incentive to deviate from the socially efficient decision-rule. In fact, given a policy-maker’s reputational concerns, there exists an information structure under which she implements the socially efficient decision-rule. (Proposition 10.)
- In the “bad news” case, the structure of the equilibrium changes (in particular, beliefs are defined by Bayes’ rule everywhere), and the socially efficient policy cannot be implemented in equilibrium. (Proposition 12.)
- If the decision on whether to continue the project can only be made at given, pre-set dates (for instance party conferences, deadlines, or discrete calendar dates) then reputational concerns always result in excessive continuation of the project, compared to the social optimum. (Proposition 17.)

2 Literature

The closest related work is Majumdar and Mukand (2004) (henceforth MM) which considers a two-period signalling game. In that model, stopping a reform that she has previously initiated perfectly reveals the political leader’s low ability. A political leader’s type is known to her and determines the quality of her private information. A high type learns the quality of a new policy project perfectly while a low type observes a noisy signal about it. In equilibrium, the high ability leader always implements the project in accordance with her information, and thus persist with the project even after a low signal. By separating and stopping the project after a low signal, the low ability leader would perfectly reveal that she is low type. She therefore pools, which is inefficient.

The main difference between my setup and that of MM is the following. My setup is in continuous time and, crucially, I assume that the leader and the electorate have symmetric information ex-ante: both face the same uncertainty about the leader’s ability type, and therefore about the policy project. Informational asymmetries arise gradually over time as the leader privately observes news about the project quality. I obtain very different

results from MM. In particular, there always exists an information structure that makes it possible for a decision-maker to implement the socially optimal decision-rule, even if her concern for social welfare is very small.

Prendergast and Stole (1996) also study a signalling game where an agent, who wants to acquire a reputation for quickly learning the correct course of action, chooses investments on a project over a finite number of periods. In each period the agent may base her choice on her observation of a private signal about the project's profitability. High ability agents receive more precise signals than their less able counterparts and are therefore able to learn about and respond to the economic environment more quickly.

Low ability agents have an incentive to pool and attempt to mimic the behaviour of a high ability agent. They can do this by exaggerating the information of their signal in early periods, and dampening it in later periods. This is because, from an ex-ante point of view, the investment decision of high ability agents is more variable in early periods, and more stable relative to previous investment decisions in later periods.

Earlier work on dynamic signalling starts with Noldeke and Van Damme (1990) and Swinkels (1999), which consider the Spence signalling model, allowing the agent to take actions at various points in time. In these models, the signalling is dynamic, because the agents may act repeatedly, but the information structure does not change over time: the worker has ex-ante private information and her type is persistent. In my model, the agent may act at various points in time, information is ex-ante symmetric, and the informational asymmetry develops over time.

Insofar as there is symmetric information ex-ante and career concerns, my model bears some similarity to Holmström (1999). The main difference is that my decision-maker acquires private information on the equilibrium path, a feature that is absent in Holmström. There, informational asymmetries can only arise off path.

I find that restricting the decision to abandon the project to take place only at given pre-set dates limits the leader's ability to signal her information type. In section 6, I discuss the relation of this result with Abreu, Milgrom, and Pearce (1991), which studies the effect of timing and information structure in a repeated partnership game with imperfect public monitoring.

The formal setup of my model, especially the belief-updating based on the leader's private information, is similar to those used in the strategic experimentation literature, in particular Rosenberg, Salomon, and Vieille (2013), Keller and Rady (2010), and Keller and Rady (2015). The focus of these papers, however, is on information as a public good.

Finally, the economic question addressed in this paper brings to mind some of the literature on reputational cheap talk. Broadly speaking, the problem of an agent who has an incentive to orient her communication strategy away from truth-telling in order to build or maintain a reputation comes in two flavours. The agent's reputation may be to do with her preferences. For instance in Morris (2001) the agent wants to seem ideologically unbiased, since such a reputation confers more credibility to her future policy recommendations. In Benabou and Laroque (1992) it is a reputation for "honesty" that confers credibility to the

report of a market-insider (say a financial journalist or “guru”). If she engages in insider trading, the agent can then exploit her enhanced credibility to manipulate the market to her advantage. Alternatively, as in Ottaviani and Sørensen (2006), the agent’s reputation may be to do with her ability to learn, or the quality of her information.

Contrary to the intuition of the reputational cheap talk literature, in which an expert may want to obfuscate information, here reputational concerns bias my decision-maker towards proceeding and letting the state of the world be revealed.

3 The model

There are two players, a decision-maker – henceforth DM – and an observer¹. Time $t \in [0, \infty)$ is continuous, and the observer’s payoffs are discounted at rate $\rho > 0$. The DM’s competence-level is modeled as a binary state of the world: competent (\mathcal{C}) or incompetent (\mathcal{I}). The state of her competence-level, $\omega \in \{\mathcal{C}, \mathcal{I}\}$, is persistent and selected by nature at the outset of the game according to the distribution $\Pr(\omega = \mathcal{C}) = q$, with $q \in (0, 1)$. Neither the DM nor the observer know the realisation of her competence-level. Thus, the DM has no private information at the beginning of the game: information is symmetric ex-ante.

The DM is endowed with an experiment, the “project”, with unknown quality θ , which can be either good ($\theta = G$) or bad ($\theta = B$). A project is more likely to be good if it is undertaken by a competent DM: $0 \leq \Pr(\theta = G|\omega = \mathcal{I}) < \Pr(\theta = G|\omega = \mathcal{C}) \leq 1$. Let $p_0 := q \Pr(\theta = G|\omega = \mathcal{C}) + (1 - q) \Pr(\theta = G|\omega = \mathcal{I})$ denote the prior probability that the project is good. From now on, to simplify notation, let us set $\Pr(\theta = G|\omega = \mathcal{I}) = 0$ and $\Pr(\theta = G|\omega = \mathcal{C}) = 1$. By this normalisation, the project is good if and only if the DM is competent.

A project’s quality is uncertain and a priori unknown both to the DM and the observer. It determines the observer’s payoff. A bad project never yields a payoff. A good project yields no payoff until a random time $\tau \in [0, \infty)$, and thereafter yields a constant flow payoff with present value $g > 0$. The random time τ follows an exponential distribution with commonly known parameter $\eta > 0$. Let $\gamma := \eta g / (\eta + \rho)$ denote the expected discounted payoff of keeping a good project active until it succeeds. Finally, define the right-continuous process $(S(t))_{t \geq 0}$, with $S(0) = 0$ and $S(t) = \mathbb{1}\{\tau \leq t\}$, that tracks whether at any time $t \geq 0$ the project has as yet succeeded.

We say that the DM “experiments” if she keeps active a project that has not yet succeeded. Repealing the project, thereby reverting to the status-quo policy, yields a payoff with discounted present value $s \in (0, \gamma)$. Accordingly, reverting to the status-quo policy is never socially optimal if the project is known to be good.

¹Depending on the interpretation of the model, the “observer” may be group of individuals (an electorate, the shareholders of a company, etc.).

As long as the project is active, the DM observes private pieces of news. Given the quality of the project $\theta \in \{G, B\}$, one piece of news arrives at each jumping time of a standard Poisson process with state-dependent intensity λ_θ . The constants $\lambda_B > 0$ and $\lambda_G > 0$ are common knowledge. Let $\Delta\lambda := \lambda_G - \lambda_B$. If $\Delta\lambda > 0$, news events can be thought of as “breakthroughs” more likely to occur if the project is good. If $\Delta\lambda < 0$, they are “breakdowns” more likely to occur if the project is bad. This paper analyses both cases.² Let the random variable $N(t)$ be the number of news events observed by the DM up to date t . Observe that this is the DM’s *type* at date t .

The privately observed process $(N(t))_{t \geq 0}$ is a standard Poisson process with intensity λ_θ . For $n \in \{1, 2, \dots\}$, let τ_n , denote the arrival time of the n^{th} jump. Conditional on the state of the world, the processes $(N(t))$ and $(S(t))$, are independent. Let $X(t) = (N(t), S(t))$, and let $\mathcal{F}_t^X = \sigma\{X(s) | 0 \leq s \leq t\}$ reflect the private information available to the DM at date t .³ A strategy for the DM is a stoppingtime, $T \in [0, \infty]$, with respect to the filtration $\{\mathcal{F}_t^X\}_{t \geq 0}$.⁴

Beliefs: Two posterior beliefs about the quality of the project are of importance in this model. The first, $p(t) := \mathbb{P}_{p_0}(\theta = G | \mathcal{F}_t^X)$, is the DM’s posterior belief, and is based on her private information. The second, $\mu(t)$, is the posterior belief of the observer about the state of the project, based only on the publicly available information: the DM’s actions up to date t and whether the project has as yet succeeded.⁵ Under our assumption that $\Pr(\theta = G | \omega = \mathcal{I}) = 0$ and $\Pr(\theta = G | \omega = \mathcal{C}) = 1$, $\mu(t)$ is also the observer’s posterior belief about the DM’s competence level. Therefore, it is also the DM’s “reputation”. Even though the state $N(t)$ of the news process is not observed by the observer, it may be partially inferred from the DM’s strategy and her actions up to date t .

The DM’s private posterior belief satisfies Bayesian updating:

$$(1) \quad p(t) = \frac{\frac{p_0}{1-p_0} \phi(t)}{1 + \frac{p_0}{1-p_0} \phi(t)},$$

²If $\Delta\lambda = 0$, news events provide no information about the state of the world.

³ Formally, let θ be an independent Bernoulli random variable with parameter p . For $\theta \in \{G, B\}$ let $(N_\theta(t))$ be a standard Poisson process with intensity λ_θ , and let $S_\theta(t) = \mathbb{1}\{\tau_\theta < t\}$ where $\tau_B = +\infty$ and τ_G is a random time that follows an exponential distribution with parameter η . Finally let $X_\theta(t) = (N_\theta(t), S_\theta(t))$. The process $(X(t))$ is defined to be $(X_\theta(t))$. Let \mathbb{P}_{p_0} denote the probability measure over the space of realised paths that corresponds to this description. From now on, unless mentioned otherwise, all the expectations are taken under the probability measure \mathbb{P}_{p_0} .

⁴A strategy can equivalently be defined as follows. At the initial moment the DM chooses a deterministic time T_0 with the interpretation that she cancels the project at date T_0 if up until then there has been no news ($\tau_1 \geq T_0$) and no success ($\tau \geq T_0$). If at the random time $\tau_1 < T_0$ a news event occurs, then depending on the value of τ_1 a new time T_1 is chosen, and so on for further news events. If at the random time τ a success occurs, then the project is proven to be good, and the time $T_G = +\infty$ is chosen: the DM adopts the project permanently.

⁵The process $(p(t))_{t \geq 0}$ is adapted to the filtration $\{\mathcal{F}_t^X\}_{t \geq 0}$. The process $(\mu(t))_{t \geq 0}$ is adapted to the natural filtration associated with $(Y(t))_{t \geq 0} := (S(t), a(t))_{t \geq 0}$, where $a(t) := \mathbb{1}\{t \leq T\}$ is a random variable taking the value 0 if at date t the decision-maker has previously repealed the project and reverted to the status-quo policy, and 1 otherwise. The stochastic process $(a(t))_{t \geq 0}$ is the publicly observable *action path*.

where $\phi(t)$ denotes the likelihood ratio $\frac{1-p_0}{p_0} \frac{p(t)}{1-p(t)}$, and is given by

$$\phi(t) = e^{-(\eta+\Delta\lambda)t} \left(\frac{\lambda_G}{\lambda_B} \right)^{N(t)}.$$

Consider the time interval $[t, t+dt)$. If over that interval, there is no news and no success even though the project is active, the DM's posterior belief evolves continuously according to the law of motion

$$dp = -(\eta + \Delta\lambda)p(1-p)dt.$$

If a news event occurs at time t , the DM's posterior belief jumps from $p(t_-)$ (the limit of her posterior beliefs before the news event) to $p(t) = j(p(t_-))$, where

$$j(p) := \frac{p\lambda_G}{\lambda(p)},$$

where $\lambda(p) := p\lambda_G + (1-p)\lambda_B$. If $\Delta\lambda > 0$, $j(p) \in (p, 1)$. If $\Delta\lambda < 0$, $j(p) \in (0, p)$. A success proves that the project is good, and the posterior belief jumps to one.

The paths $p_t^n = \mathbb{P}_p(\theta = G | N(t) = n, S(t) = 0)$ are deterministic functions of t , and are horizontal translations of one another. Following a news event, the magnitude $j(p(t)) - p(t)$ of the jump depends on the value of $p(t)$ and is maximised at $p(t) = 1/2$.

Depending on the parameters η , λ_G and λ_B , the continuous part of the belief may evolve in the opposite direction of the jumps, or in the same direction. When $\Delta\lambda > 0$, then $\eta + \Delta\lambda > 0$ and $p(t)$ drifts down and jumps up. When $\Delta\lambda < 0$ the belief jumps down. It drifts up if $\eta + \Delta\lambda > 0$, and drifts down if $\eta + \Delta\lambda < 0$. The figure below illustrates a sample path of the belief $p(t)$ in each case.

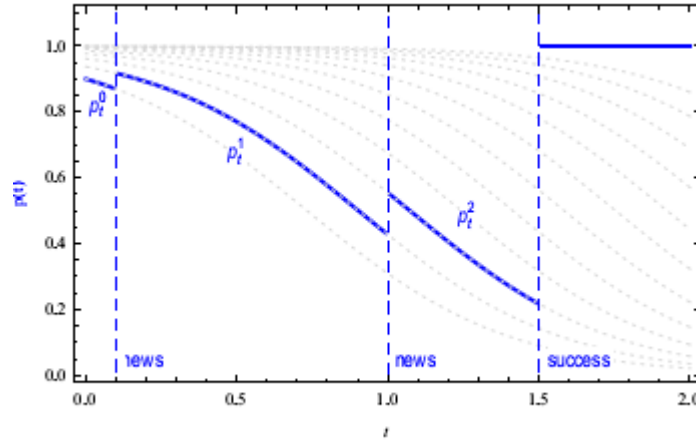


Figure: *Sample path of the belief $p(t)$. (For all figures, unless specified otherwise, we choose $(\lambda_B, \lambda_G, \eta, \rho, g, s) = (3, 5, 1, 1, 4, 1)$.)*

Payoff: A strategy T determines a payoff for the observer and a reputation for the DM. If the project succeeds before the DM repeals it ($\tau < T$), the observer's payoff is

$e^{-\rho\tau}g$ and the DM's reputation is $\mu(\tau) = 1$. If the DM repeals the project first ($\tau \geq T$), the observer's payoff is $e^{-\rho T}s$, while the DM's reputation is $\mu(T)$. The DM cares both about the observer's payoff and about her reputation. Her payoff is defined as a convex combination of the two. Let $\alpha \in [0, 1)$ parameterise the intensity of the DM's reputational concern, and let

$$W_t^T := \mathbb{1}\{\tau < T\} e^{-\rho(\tau-t)}g + \mathbb{1}\{\tau \geq T\} e^{-\rho(T-t)}s$$

be the observer's payoff under the strategy T at date $t < T \wedge \tau$. We'll also refer to it as the “social payoff”, as it accrues to all players. The DM's expected payoff from strategy T is

$$V_t^{\alpha, T} = \mathbb{E} \left[(1 - \alpha) W_t^T + \alpha \mu(T \wedge \tau) \mid \mathcal{F}_t^X \right].$$

The linearity of the payoff function ensures that any bias in favour of experimentation arises solely because of the DM's reputational concern, and not because of her risk-preferences.⁶

When $\alpha = 0$, the DM's has no reputational concerns, her objective coincides with that of the observer. We refer to the DM with preference parameter $\alpha = 0$ as the *social planner*⁷, and to the strategy maximising $V_t^{0, T}$ as the *planner policy*.

Equilibrium: Our solution concept is the perfect Bayesian equilibrium (PBE) in pure strategies. A PBE strategy T maximises the DM's expected payoff, when her and the observer's beliefs satisfy Bayesian updating. In particular, the observer's belief $\mu(T)$ if the project is repealed, whenever it is defined by Bayes' rule, must be consistent with the DM's strategy. Otherwise, it may be arbitrarily chosen from $[p_T^0, 1)$.

The DM's problem can be reduced to an optimal stopping problem for the a-posteriori belief process $(p(t))_{t \geq 0}$.⁸ A strategy T is then summarised by a threshold belief $\hat{p}(t)$ such that $T = \inf\{t : p(t) \leq \hat{p}(t)\}$.

Dynamic signalling game with dynamic types: When the DM is concerned with her reputation ($\alpha > 0$), she is engaged in a dynamic signalling game with the observer. Though the information is symmetric ex-ante, the DM acquires private information over time. By Equation (1), there is a one-to-one mapping between the DM's private information $N(t)$, and her posterior belief $p(t)$ at date t . Without loss of generality, we can therefore let $p(t)$ be the DM's type at date t . The set of possible types at date t is $\Theta(t) := \{p_t^n\}_{t=0}^\infty$.

For type p_t^n , a message is a date $\hat{t} \in [0, \infty]$ at which the DM repeals the project. A strategy for type p_t^n consists of a threshold belief $\hat{p}^n \in [0, 1]$ such that type p_t^n repeals the

⁶See Kamenica and Gentzkow (2011).

⁷The social planner in question puts no weight on the reputational concerns of the DM. For instance, when the “observer” is an entire electorate, this would correspond to the choice of a planner giving equal Pareto-weights to all citizens, including the DM, whose reputational concerns are negligible when the population is sufficiently large.

⁸See Presman and Sonin (1990), Peskir and Shiryaev (2000), Shiryaev (2007), or Cohen and Solan (2013).

project at date t if and only if $p_t^n \leq \hat{p}^n$, and otherwise continues experimenting. Upon observing message $\hat{t} \in [0, \infty]$, the observer forms the posterior belief $\hat{\mu}(p_t^n | \hat{t}) \in \Delta(\Theta(t))$ about the DM's type, and the DM's reputation is $\mu(t) = \mathbb{E}_{\hat{\mu}}[p(\hat{t}) | \hat{t}]$.

Interpretations: It is useful to keep in mind a number of interpretations of the model.

1. The observer is the electorate and the DM a political leader undertaking a policy experiment. The leader cares directly about the social value of the policy experiment, but also about her own future electoral prospects.⁹ Those depend on the electorate's assessment of her competence. At each point in time the political leader may cancel the experiment and revert to a known status-quo policy. She can base her decision on the continuous evaluation of the policy by a government think-tank, though the think-tank's report are not available to the general public. The public does, however, perceive improvements in its standard of living.
2. The decision-maker is a turnaround CEO and the observer is a company's board of directors. While the board can observe a company's recovery once it is reflected in a higher stock price, it is unable to independently assess the effectiveness of the stabilisation and recovery measures instigated by the CEO.
3. The decision-maker is a trader at a financial institution, and the observer the upper-level manager, who has limited knowledge or understanding of the trader's portfolio or investment strategy.

4 Good news

Consider the good news scenario with $\lambda_G > \lambda_B$. Let's begin by deriving the socially optimal stopping time. We'll then show that a reputationally concerned DM may adopt the planner policy in equilibrium if she cares sufficiently about the social payoff, or if her private information is sufficiently good. If the planner policy cannot be implemented, I derive the equilibrium that does best from a social point of view.

4.1 Social planner

When $\alpha = 0$, the DM acts so as to maximise the social payoff. The common value function for the DM and the observer, $U^0(p(t)) := \sup_T V_t^{0,T}$, is convex and continuous. Convexity reflects a nonnegative value of information, and implies continuity in the open unit interval. Continuity at the boundaries follows from the fact that U^0 is bounded above by the full information payoff $p\gamma + (1-p)s$, and bounded below by the payoff $s \vee p\gamma$ of a DM whose strategy T may only take values in $\{0, \infty\}$, with the interpretation that the DM only has

⁹ Rogoff (1990) argues that political leaders derive "ego rents" from being in office, in addition to the utility they derive from public and private consumption goods as ordinary citizen.

a choice between immediately repealing the project, or experimenting forever. Both these payoffs converge to γ as $p \rightarrow 1$ and to s as $p \rightarrow 0$.

Moreover, U^0 solves the Bellman equation:

$$(2) \quad u(p) = \max \left\{ s ; \left(b_S(p, u) + b_N(p, u) - d(p, u) \right) / \rho \right\},$$

where

$$b_S(p, u) = p\eta[g - u(p)]$$

is the expected benefit from a success,

$$b_N(p, u) = \lambda(p)[u(j(p)) - u(p)]$$

is the expected benefit from a piece of news, and

$$d(p, u) = (\eta + \Delta\lambda)p(1 - p)u'(p)$$

measures the deterioration in the DM's outlook when she experiments without observing any event – success or news. As, in the good news case, infinitesimal changes in the belief are always downward, we say that a continuous function u solves the Bellman equation if its left-hand derivative exists on $(0, 1]$ and (1) holds on $(0, 1)$ when this left-hand derivative is used to compute $d(p, u)$. The planner value function U^0 is the unique solution satisfying the boundary conditions $u(0) = s$ and $u(1) = \gamma$. Let $\Omega(p) := (1 - p)/p$ denote the likelihood ratio.

Proposition 1. *The planner's optimal policy is to stop at the first time t such that $p(t) \leq p^*$. The planner threshold belief satisfies*

$$(3) \quad p^* = \frac{\nu(\lambda_G, \lambda_B)s}{(\nu(\lambda_G, \lambda_B) + 1)\gamma - s},$$

where $\nu(\lambda_G, \lambda_B) > 0$ is the positive solution to

$$(4) \quad \lambda_B + \rho - \nu(\eta + \lambda_G - \lambda_B) = \lambda_B \left(\frac{\lambda_B}{\lambda_G} \right)^\nu.$$

The social payoff under the optimal policy is

$$(5) \quad U^0(p) = \begin{cases} u^0(p) & \text{if } p(t) > p^*, \\ s & \text{if } p(t) \leq p^*, \end{cases}$$

where

$$(6) \quad u^0(p) := p\gamma + \frac{1 - p}{1 - p^*} (s - p^*\gamma) \left(\frac{\Omega(p)}{\Omega(p^*)} \right)^{\nu(\lambda_G, \lambda_B)}.$$

Proof. Over values of $p \in (0, 1)$ at which experimentation is optimal, U^0 solves the following ordinary differential difference equation¹⁰:

$$(7) \quad u(p)(p\eta + \rho) + \lambda(p)[u(p) - u(j(p))] + (1-p)p(\eta + \Delta\lambda)u'(p) = p\eta g.$$

Let $p \mapsto u^0(p)$, defined on $(0, 1)$, denote the solution to this differential equation. It is easy to verify that the function $p \mapsto p\gamma$ constitutes a particular solution to (7). The function $p \mapsto (1-p)\Omega(p)^\nu$, $\nu > 0$, captures the option-value of being able to repeal the project¹¹, and constitutes a solution to the homogeneous version of (7) if and only if ν satisfies equation (4).

There are two solutions to (4), one negative, the other positive. We let $\nu(\lambda_G, \lambda_B)$ denote the positive solution. Observe that $\nu(\lambda_G, \lambda_B) \in \left[\frac{\rho}{\eta + \Delta\lambda}, \frac{\rho + \lambda_B}{\eta + \Delta\lambda}\right]$, $\frac{\partial \nu}{\partial \lambda_G}(\lambda_G, \lambda_B) < 0$, and $\frac{\partial \nu}{\partial \lambda_B}(\lambda_G, \lambda_B) > 0$.

The solution to (7) is therefore given by the family of functions

$$u_{C_0}(p) = p\gamma + C_0 (1-p)\Omega(p)^{\nu(\lambda_G, \lambda_B)},$$

where $C_0 \in \mathbb{R}$ is a constant of integration.

Let p^* denote the planner's optimal threshold belief. It satisfies the boundary condition (value-matching): $u^0(p^*) = s$. Solving for the constant C_0 , we obtain the expression in (6) for $u^0(p)$, which is a strictly concave function of p^* on $[0, 1]$. Choosing $p^* \in [0, 1]$ so as to maximise the resulting expression (the solution is interior and smooth-pasting is satisfied), we obtain the expression in (3) for p^* .

In Appendix 4.1.1 we verify that U^0 solves the Bellman equation, (2), with the maximum being achieved under the planner policy. \square

4.1.1 Verification

Proof. For $p \geq p^*$, $(b_S(p, U^0) + b_N(p, U^0) - d(p, U^0))/\rho = U^0(p)$, which is strictly greater than s for every $p > p^*$, and equals s when $p = p^*$.

For $p \leq j^{-1}(p^*)$, $(b_S(p, U^0) + b_N(p, U^0) - d(p, U^0))/\rho = p\eta(g - s)/\rho$. This is strictly less than s for every $p < \bar{p}^*$ (defined in Corollary 2) and consequently for every $p \leq j^{-1}(p^*)$.

For $j^{-1}(p^*) < p < p^*$,

$$\begin{aligned} & (b_S(p, U^0) + b_N(p, U^0) - d(p, U^0))/\rho = p\eta(g - s) + \lambda(p)(U^0(j(p)) - s) \\ & < p\eta(g - s) + \lambda(p)(U^0(j(p^*)) - s) \\ & = p\eta(g - s) + \frac{\lambda(p)}{\lambda(p^*)} \left(p^* \gamma \lambda_G + (s - p^* \gamma) \lambda_B \left(\frac{\lambda_B}{\lambda_G} \right)^{\nu(\lambda_G, \lambda_B)} \right) - \lambda(p)s \\ & < p\eta(g - s) + p^* \gamma \lambda_G + (s - p^* \gamma)(\lambda_B + \rho - \nu(\lambda_G, \lambda_B)(\eta + \Delta\lambda)) - \lambda(p)s \\ & = [\gamma(\eta + \rho) - (\eta + \Delta\lambda)s](p - p^*) + s\rho \\ & < s. \end{aligned}$$

The first inequality follows from the monotonicity of $U^0(j(p))$ in p on $(j^{-1}(p^*), 1)$. The second follow from the monotonicity of $\lambda(p)$ in p on $(0, 1)$ and a substitution from (4). The term in square brackets in the

¹⁰This ODDE bears some resemblance to, but is different from equation (1) in Keller and Rady (2010).

¹¹The assumption of exponential discounting dictates our guess of the form of the option-value $v(p)$.

penultimate expression is bounded below by $\gamma(\rho + \Delta\lambda) > 0$. The last inequality follows, establishing the result. \square

The planner's threshold belief, p^* , is constant over time. Let t_n^* be the date t at which $p_t^n = p^*$. That is the date at which the planner would stop if she had observed n pieces of news prior to her posterior belief falling below the threshold p^* . Observe that since the planner only stops at dates $\{t_n^*\}_{n \geq 0}$, there are unreached information sets under the planner solution. The figure below illustrates the planner solution.

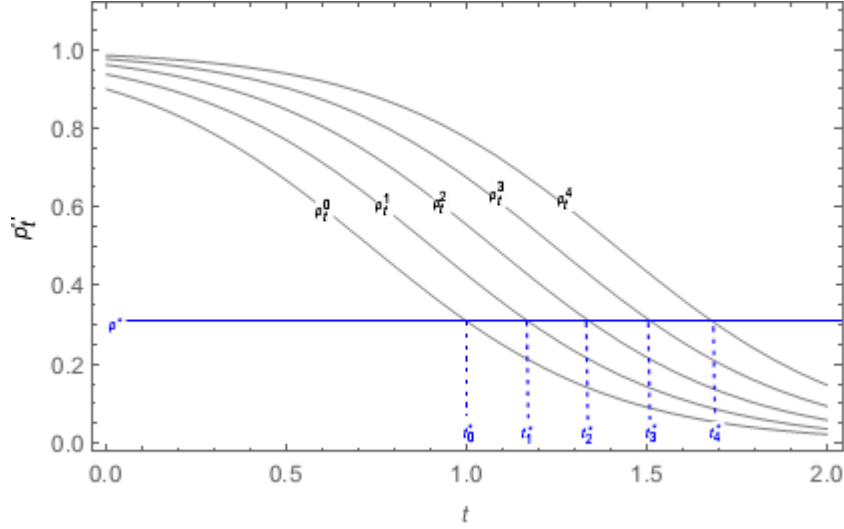


Figure: *The planner solution in the good news case.*

Clearly, p^* and $U^0(p)$ depend on the information structure. Suppose λ_G increases, λ_B decreases, or both, so that the difference $\lambda_G - \lambda_B$ increases, with the interpretation that news events become more informative. Elementary calculations show that p^* decreases, and $U^0(p)$ increases uniformly for all $p \in (p^*, 1)$: the social value of experimentation increases in the presence of more informative news about the project quality.

The next corollary describes the limit cases. Fix $\lambda_G > 0$. When news is uninformative ($\lambda_B = \lambda_G$), then $\nu(\lambda_G, \lambda_B) = \frac{\rho}{\eta}$, and $p^* = \bar{p}^* := \frac{s\rho}{\eta(g-s)}$, as in the exponential bandit decision problem of Keller, Rady, and Cripps (2005)¹². When $\lambda_B = 0$ and news only arrives in the good state, $\nu(\lambda_G, \lambda_B) = \frac{\rho}{\eta + \lambda_G}$, and $p^* = \underline{p}^*$.

In both limit cases $\lambda_B = \lambda_G$ and $\lambda_B = 0$, there is only one possible date at which the planner stops under the optimal policy. When $\lambda_B = \lambda_G$, $p(t)$ drifts down continuously, following the law of motion $dp = -\eta p(1-p)dt$ as long as the project has not succeeded. It is not affected by the arrival of news. The planner stops at the date \bar{t}^* at which $p(\bar{t}^*) = \bar{p}^*$, if the project has not succeeded by then. When $\lambda_B = 0$, $p(t)$ drifts down continuously, following the law of motion $dp = -(\eta + \lambda_G)p(1-p)dt$ as long as no success or no news arrives. The first piece of news conclusively reveals that the project is good, and the

¹² Proposition 3.1 in Keller, Rady, and Cripps (2005).

planner's posterior belief jumps to one, making it optimal to never stop experimenting. The planner stops at the date \underline{t}^* at which $p(\underline{t}^*) = \underline{p}^*$, if the project has not succeeded and no news has arrived before then.

Now fix $\lambda_B \geq 0$. When $\lambda_G \rightarrow \infty$ we obtain the full-information case. In the good state, conclusive news should arrive instantly. In other words, the state is immediately revealed. Therefore, an almost-optimal strategy for the planner is to repeal the project at time ε unless she observes news, where $\varepsilon > 0$ is a small real number. The social payoff under this strategy tends to $p_0\gamma + (1 - p_0)s$.

Corollary 2. (a) Fix $\lambda_G > 0$. If $\lambda_B = \lambda_G$, then $\nu(\lambda_G, \lambda_B) = \frac{\rho}{\eta}$, and $p^* = \bar{p}^* := \frac{s\rho}{\eta(g-s)}$.

(b) Fix $\lambda_G > 0$. If $\lambda_B = 0$, then $\nu(\lambda_G, \lambda_B) = \frac{\rho}{\eta + \lambda_G}$, and $p^* = \underline{p}^* := \frac{\frac{\rho}{\eta + \lambda_G}s}{(\frac{\rho}{\eta + \lambda_G} + 1)\gamma - s}$.

(c) Fix $\lambda_B \geq 0$. If $\lambda_G \rightarrow \infty$, then $\nu(\lambda_G, \lambda_B) \rightarrow 0$, and $p^* \rightarrow 0$.

Proof. (a) Follows from setting $\lambda_B = \lambda_G$ in (4). For (b) and (c), recall from the proof of proposition 1 that $\nu(\lambda_G, \lambda_B) \in \left[\frac{\rho}{\eta + \lambda_G - \lambda_B}, \frac{\rho + \lambda_B}{\eta + \lambda_G - \lambda_B} \right]$, and observe that when $\lambda_B = 0$ this interval is the point $\rho/(\eta + \lambda_G)$, and that when $\lambda_G \rightarrow \infty$ this interval converges to 0. \square

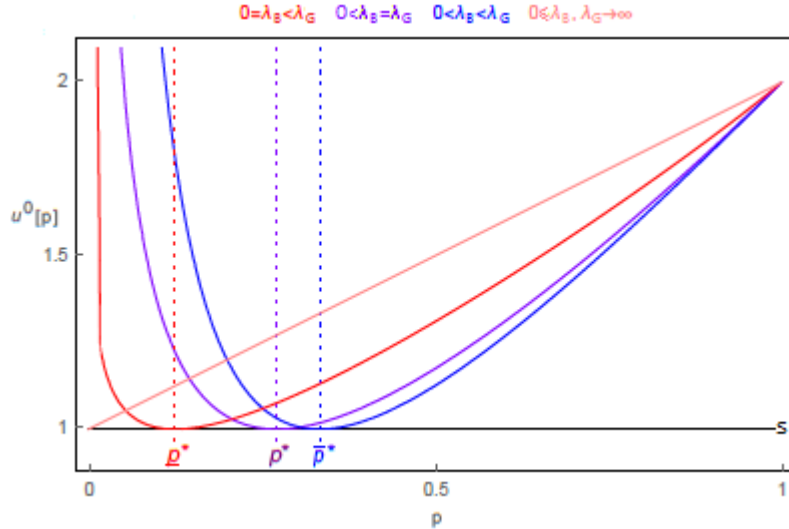


Figure: The function $u^0(p)$ and the planner thresholds for the cases (from lowest to greatest value achieved under the planner solution) $0 < \lambda_B = \lambda_G$; $0 < \lambda_B < \lambda_G$; $0 = \lambda_B < \lambda_G$; $0 = \lambda_B, \lambda_G \rightarrow \infty$.

4.2 Equilibrium

If the planner is not too concerned with her reputation, she can adopt the socially optimal policy in a PBE. If not, I characterise the separating equilibrium that maximises the expected social payoff. I first establish the following result:

Proposition 3. *There exists $\bar{\alpha} \in (0, 1)$ such that the planner policy is a PBE strategy for a DM with preference parameter $\alpha < \bar{\alpha}$.*

Thus, a decision-maker who is sufficiently concerned with the social welfare can implement the planner solution. The threshold $\bar{\alpha}$ need not be small. In fact, I show in section 4.4 that $\bar{\alpha}$ depends on the information structure.

Let's first establish proposition 3. Fix $\alpha < \bar{\alpha}$ and assume that there exists a PBE in which the DM adopts the planner policy. Let $D_n(t, t')$ denote the following deviation: the DM keeps the project active over the time interval $[t, t')$, and resumes the PBE policy from t' on. The subscript n indicates that the deviation is performed by information type $N(t) = n$, i.e. by the DM who has observed n pieces of news up to date t .

The next lemma describes the net payoff from the deviation $D_n(t_n^*, t_{n+1}^*)$. While this deviation always entails a net social loss, it succeeds in convincing the observer that the DM had more good news at t_n^* than is in fact the case, and thus generates a net reputational gain. Therefore this deviation is profitable for a DM with $\alpha > 0$ only if its expected reputational benefit is large enough to outweigh its expected social loss. Lemma 5 then argues that $D_n(t_n^*, t_{n+1}^*)$ is the best possible deviation. Proposition 3 follows.

Lemma 4. *The deviation $D_n(t_n^*, t_{n+1}^*)$ from the planner policy is strictly profitable if and only if*

$$(8) \quad \alpha [p^* - j^{-1}(p^*)] + (1 - \alpha) e^{-\rho(t_{n+1}^* - t_n^*)} [s - u^0(j^{-1}(p^*))] > 0.$$

The first bracket is positive and measures the expected net reputational benefit from deviating. The second bracket is negative and measures the expected net social cost of deviating.

Proof. Let's first evaluate the net reputational benefit entailed by the deviation $D_n(t_n^*, t_{n+1}^*)$. If the DM implements the planner policy and stops at t_n^* , she reveals her private information and the observer learns that her private belief is the planner threshold: $\mu(t_n^*) = p^*$.

Assume that instead the DM follows the deviation $D_n(t_n^*, t_{n+1}^*)$. If in the time interval $[t_n^*, t_{n+1}^*)$ the project succeeds, the DM's reputation jumps to 1 and $\mu(t_{n+1}^*) = p(t_{n+1}^*) = 1$. If it does not, but the DM observes $k \geq 1$ pieces of news, then either the project succeeds (at random time τ) before the DM's belief reaches p^* (at stopping time T), or vice-versa. In both cases, the observer learns the DM's private information and we have $\mu(\tau \wedge T) = p(\tau \wedge T)$. Finally, if in the time interval $[t_n^*, t_{n+1}^*)$ there is no success and the DM observes 0 pieces of news, she stops at t_{n+1}^* . In this case, the observer's belief is defined by Bayes' rule and satisfies $\mu(t_{n+1}^*) = p^*$, whereas the DM's private belief is $p(t_{n+1}^*) = j^{-1}(p^*) < p^*$. By the above argument, we have that

$$(9) \quad \mathbb{E}[\mu(\tau \wedge T) | p(t_n^*) = p^*] - \mathbb{E}[p(\tau \wedge T) | p(t_n^*) = p^*] = \pi(0, p^*, t_{n+1}^* - t_n^*) [p^* - j^{-1}(p^*)]$$

where we let

$$\pi(0, p^*, t_{n+1}^* - t_n^*) := p^* e^{-(\eta + \lambda_G)(t_{n+1}^* - t_n^*)} + (1 - p^*) e^{-\lambda_B(t_{n+1}^* - t_n^*)}$$

denote the probability that no news or success arrives over the time interval $[t_n^*, t_{n+1}^*)$, given the current belief $p(t_n^*) = p^*$.

Since the posterior belief process is an \mathcal{F}^X -martingale, $\mathbb{E}[p(\tau \wedge T) | p(t_n^*) = p^*] = p^*$ and the left-hand side of (9) defines the net reputational benefit from the deviation $D_n(t_n^*, t_{n+1}^*)$. The right-hand side is clearly positive.

Let's now evaluate the net social payoff from the deviation $D_n(t_n^*, t_{n+1}^*)$. We begin with a few observations on $u^0(p)$, defined in (6). It is a convex function of p , reaching its minimum s when $p = p^*$, and with $u^0(1) = \frac{\eta g}{\eta + \rho}$ and $\lim_{p \rightarrow 0} = +\infty$. For all $p_t^n \in (0, 1)$ and for all $\Delta t > 0$, it is easy to verify that

$$(10) \quad u^0(p_t^n) = p_t^n \left(1 - e^{-(\eta + \rho)\Delta t}\right) \frac{\eta g}{\eta + \rho} + \sum_{k=0}^{\infty} \pi(k, p_t^n, \Delta t) e^{-\rho \Delta t} u^0(p_{t+\Delta t}^{n+k}),$$

where

$$\pi(k, p_t^n, \Delta t) := p_t^n \frac{(\lambda_G \Delta t)^k}{k!} e^{-(\eta + \lambda_G)\Delta t} + (1 - p_t^n) \frac{(\lambda_B \Delta t)^k}{k!} e^{-\lambda_B \Delta t}$$

denotes the probability that $S(t + \Delta t) - S(t) = 0$ and $N(t + \Delta t) - N(t) = k$, given $S(t) = 0$ and $N(t) = n$. Equivalently, it is the probability that, given that the DM's belief at t is $p(t) = p_t^n$, there is no success and she observes $k = 0, 1, 2, \dots$ pieces of news over the time interval $[t, t + \Delta t)$. Her resulting posterior belief is then $p_{t+\Delta t}^{n+k}$.

If, given her current belief $p(t_n^*) = p_{t_n^*}^n = p^*$, the DM follows the deviation $D_n(t_n^*, t_n^* + \Delta t)$, for $0 < \Delta t \leq t_{n+1}^* - t_n^*$, the social payoff is as follows. If the project succeeds over the time interval $[t_n^*, t_n^* + \Delta t)$, a payoff with present discounted value g accrues at the success arrival time $\tau \in [t_n^*, t_n^* + \Delta t)$. If there is no success but $k \geq 1$ pieces of news arrive over the time interval $[t_n^*, t_n^* + \Delta t)$, the resulting belief for the DM is $p_{t_n^* + \Delta t}^{n+k}$. It exceeds p^* , so that the DM resumes the planner policy, generating a social payoff of $U^0(p_{t_n^* + \Delta t}^{n+k})$. Finally, if over the time interval $[t_n^*, t_n^* + \Delta t)$ no success and no news occur, the DM's belief, $p_{t_n^* + \Delta t}^n$, is strictly below p^* and the DM repeals the project, thus generating a social payoff of s .

Thus, the expected social payoff when the DM follows the deviation $D_n(t_n^*, t_n^* + \Delta t)$, for $0 < \Delta t \leq t_{n+1}^* - t_n^*$, is

$$(11) \quad (U^0)^{D_n}(t_n^*, t_n^* + \Delta t) := p_{t_n^*}^n \left(1 - e^{-(\eta + \rho)\Delta t}\right) \frac{\eta g}{\eta + \rho} + \sum_{k=1}^{\infty} \pi(k, p_{t_n^*}^n, \Delta t) e^{-\rho \Delta t} U^0(p_{t_n^* + \Delta t}^{n+k}) \\ + \pi(0, p_{t_n^*}^n, \Delta t) e^{-\rho \Delta t} s$$

For $\Delta t = t_{n+1}^* - t_n^*$ we have that, for all $k \geq 1$, $p_{t_n^* + \Delta t}^{n+k} = p_{t_{n+1}^*}^{n+k} \geq p^*$ so that $U^0(p_{t_n^* + \Delta t}^{n+k}) = u^0(p_{t_n^* + \Delta t}^{n+k})$. C Replacing the first line in (11) using (10), we have that

$$(U^0)^{D_n}(t_n^*, t_n^* + \Delta t) = U^0(p_{t_n^*}^n) + \pi(0, p_{t_n^*}^n, t_{n+1}^* - t_n^*) e^{-\rho(t_{n+1}^* - t_n^*)} \left[s - u^0(p_{t_{n+1}^*}^n)\right].$$

Observe that, since $p_{t_{n+1}^*}^n < p^*$, the value $u^0(p_{t_{n+1}^*}^n)$ is strictly greater than s , and the term in square brackets in the expression above is strictly negative. Thus, the expected net social payoff from the deviation $D_n(t_n^*, t_{n+1}^*)$,

$$(U^0)^{D_n}(t_n^*, t_n^* + \Delta t) - U^0(p_{t_n^*}^n) = \pi(0, p_{t_n^*}^n, t_{n+1}^* - t_n^*) e^{-\rho(t_{n+1}^* - t_n^*)} \left[s - u^0(p_{t_{n+1}^*}^n)\right]$$

is strictly negative. This was to be expected: the planner policy maximises the social payoff, therefore any deviation from it generates a net social loss.

Since $\pi(0, p_{t_n^*}^n, t_{n+1}^* - t_n^*) > 0$, the lemma follows. \square

Consequently there exists an $\bar{\alpha} > 0$ such that for all $\alpha < \bar{\alpha}$, the deviation $D_n(t_n^*, t_{n+1}^*)$ is not profitable. For these values of α , longer deviations $D_n(t_n^*, t_{n+k}^*)$ with $k > 1$ are not profitable a fortiori. The following lemma establishes that shorter deviations are not profitable either.

Lemma 5. For all $t' \in (t_n^*, t_{n+1}^*)$ the deviation $D_n(t_n^*, t')$ is less profitable than $D_n(t_n^*, t_{n+1}^*)$.

Proof. Given $t' \in (t_n^*, t_{n+1}^*)$, consider the deviation $D_n(t_n^*, t')$. If no success or news arrives on (t_n^*, t') , the DM stops at t' . Observe that $\mu(t')$ is not determined by Bayes' rule.

Consider the following off-path beliefs: $\mu(t') = p_{t'}^n$. (Motivation: a random event has occurred that prevented the DM from acting for a short duration Δt .) Then, after all histories, $\mu(\tau \wedge T) = p(\tau \wedge T)$. Therefore, by the martingale property, the DM's expected reputational benefit is zero, and cannot outweigh the expected social loss. \square

Thus, Proposition 3 holds. Let us now consider a DM for whom $\alpha \geq \bar{\alpha}$, and establish the following result.

Proposition 6. For $\bar{\alpha} < \alpha < 1$, when the planner policy is not adopted in equilibrium, there exist a separating equilibrium with threshold beliefs $\{\hat{p}_n\}_{n \geq 0}$ satisfying

- $\hat{p}_0 = p_0^*$ (efficient),
- $\hat{p}_n \geq \hat{p}_{n+1}$, $n \geq 0$ (increasing inefficiency over time).

That is, I show that there exists a separating equilibrium characterised by threshold beliefs $\{\hat{p}_n\}_{n \geq 0}$, $\hat{p}_n \geq \hat{p}_{n+1}$ and corresponding stopping dates $\{\hat{t}_n\}_{n \geq 0}$ defined by $\hat{p}_n = p_{\hat{t}_n}^n$.

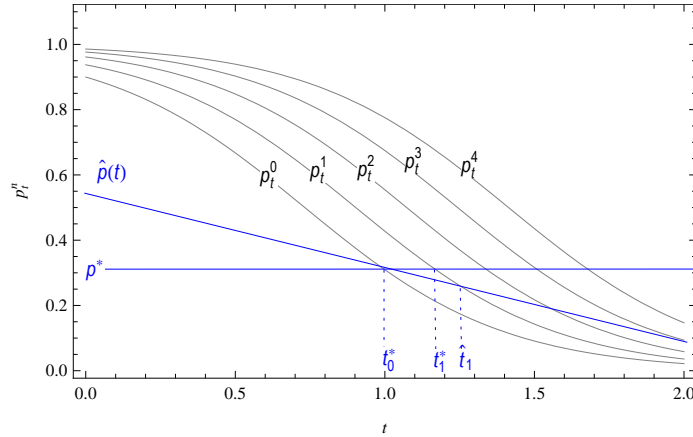


Figure: Sketch of the function $\hat{p}(t)$ representing the sequence of equilibrium threshold beliefs $\{\hat{p}_n\}_{n \geq 0}$.

Proof. Let us first establish that $\hat{p}_0 = p^*$. Assume by way of contradiction that $\hat{p}_0 \neq p^*$. Deviating and using threshold p^* rather than \hat{p}_0 is profitable. The social payoff is maximised by stopping at t_0^* , and the deviation generates a positive expected reputational gain: Though $\mu(t_0^*)$ is not determined by Bayes' rule, it must satisfy $\mu(t_0^*) \geq p_{t_0^*}^0 = p^*$, since $n = 0$ is the worst possible information type of the DM.

Let's now show that $\hat{p}_1 < p^*$. Consider the net payoff at t_0^* to the information type $N(t_0^*) = 0$ from deviating according to $D_0(t_0^*, \hat{t}_1)$, and let us vary \hat{t}_1 . Clearly, the expected net social payoff is negative and decreasing in \hat{t}_1 .

The expected net reputational gain is positive but also decreasing in \hat{t}_1 . Consider the DM's reputation under deviation $D_0(t_0^*, \hat{t}_1)$. If there is no news and no success on $[t_0^*, \hat{t}_1)$, the DM stops at \hat{t}_1 and $p_{\hat{t}_1}^1 =$

$\mu(\hat{t}_1) > p(\hat{t}_1) = p_{\hat{t}_1}^0$. For all other histories, we have $\mu(T \wedge \tau) = p(T \wedge \tau)$. Thus, the expected reputational benefit is

$$\mathbb{P}_p[\text{no news or success on } (t_0^*, \hat{t}_1) | p(t_0^*) = p_{t_0^*}^0] (p_{\hat{t}_1}^1 - p_{\hat{t}_1}^0)$$

Both terms are indeed decreasing in \hat{t}_1 .

By continuity we conclude that for $\alpha > \bar{\alpha}$, there exists a $\hat{t}_1 > t_1^*$ and associated $\hat{p}_1 < p^*$ such that the deviation $D_0(t_0^*, \hat{t}_1)$ is not profitable. \square

4.3 Equilibrium Refinement

Then the efficient PBE is the only PBE satisfying the D1 criterion.

Proposition 7. *Suppose that $\alpha < \bar{\alpha}$. (a) The efficient PBE satisfies the D1 criterion. (b) No inefficient PBE satisfies the D1 criterion.*

Proof. Consider the efficient PBE from Proposition 3. Fix an off-path date $t' \in (t_k^*, t_{k+1}^*)$, and suppose the DM stops at that date. The set of possible types is $\{p_{t'}^n\}_{n=0}^k \cup \{p_{t'}^n\}_{n=k+1}^\infty$. For the types in the first set, a deviation to t' requires delaying stopping when compared with equilibrium. For those in the second set it requires anticipating the equilibrium stopping date. The set of reputations we could offer the DM who stops at t' is $[p_{t'}^0, 1)$.

Proof. (a)

- First, we eliminate all types $p_{t'}^n < p_{t'}^k$.

Suppose the DM stops at date $t' \in (t_k^*, t_{k+1}^*)$. Conditional on $p(t') < p^*$, her type could be $p_{t'}^n \in \{p_{t'}^0, p_{t'}^1, \dots, p_{t'}^k\}$. If $p_{t'}^n = p_{t'}^0$, the DM must have been type $p_{t_0^*}^0$ at date t_0^* and had no success and no news on $[t_0^*, t')$. If $p_{t'}^n = p_{t'}^1$, the DM could either have been type $p_{t_0^*}^0$ at date t_0^* and had no success and one piece of news on $[t_0^*, t')$, or she could have been type $p_{t_1^*}^1$ at date t_1^* and had no success and no news on $[t_1^*, t')$. And so on. If we eliminate deviations up to date t' for types $p_{t_0^*}^0, p_{t_1^*}^1, \dots, p_{t_{k-1}^*}^{k-1}$, in that order, we'll conclude that, conditional on $p_{t'}^n < p^*$, the type stopping at t' must be $p_{t'}^k$.

For $t < t'$, let

$$V_t^T(p(t), \mu(t')) = (1 - \alpha) \mathbb{E} \left[\mathbb{1}_{\{\tau < T\}} e^{-\rho(\tau - t)} g + \mathbb{1}_{\{\tau \geq T\}} e^{-\rho(T - t)} s | p(t) \right] + \alpha \mathbb{E} [\mu(\tau \wedge T) | p(t)]$$

be type $p(t)$'s expected payoff at date t under the arbitrary strategy T , if the reputation from stopping at t' is $\mu(t') \in [0, 1]$, and is as in the efficient PBE for every $t \neq t'$.

Observation 8. *If under T the DM expects to stop at date t' with strictly positive probability, then $V_t^T(p(t), \mu(t'))$ is a strictly increasing function of $\mu(t')$.*

Observation 9. *The payoff $V_t^T(p(t), \mu(t'))$ is a strictly increasing function of $p(t)$.*

Let $\bar{T}_n(\mu(t'))$ be the strategy that maximises $V_{t_n^*}^T(p_{t_n^*}^n, \mu(t'))$. Under that strategy, the DM only stops at dates in $\{t_j^*\}_{j=n}^\infty \cup \{t'\}$, since amongst all strategies that never stop at t' , the planner policy, which we denote T^* , maximises the expected payoff. In other words, the DM can only improve upon the planner policy by adding t' to the support of T^* .

Finally, we let $M_n(t') \subseteq [p_{t'}^0, 1)$ be the set of reputations $\mu(t')$ such that under $\bar{T}_n(\mu(t'))$, the DM who has reputation $p_{t_n^*}^n$ at date t_n^* stops at date $t < t'$, conditional on no success and no news on $[t_n^*, t')$.

We begin by showing that, when $M_n(t') \subset [p_{t'}^0, 1)$, it is the interval $[p_{t'}^0, \underline{\mu}_{t'}^n)$. Let $\underline{\mu}_{t'}^n \in [p_{t'}^0, 1) \setminus M_n(t')$ denote the lowest reputation at t' for which, under $\bar{T}_n(\underline{\mu}_{t'}^n)$, the DM who has reputation $p_{t_n}^n$ at date t_n^* stops at date t' conditional on no success and no news on $[t_n^*, t')$. In this case, the payoff at t_j^* to type $p_{t_j}^n$ under $\bar{T}_n(\underline{\mu}_{t'}^n)$ is

$$V_{t_j^*}^{\bar{T}_n(\underline{\mu}_{t'}^n)}(p_{t_j}^n, \underline{\mu}_{t'}^n) \geq (1 - \alpha)s + \alpha p^*, \quad j = n, \dots, k;$$

where the right-hand side is the payoff at t_j^* to type $p_{t_j}^n$ from stopping at t_j^* . Moreover, for at least one $j \in \{n, \dots, k\}$, the relation above holds with equality. Thus, given $\underline{\mu}_{t'}^n$, at every $t < t'$ type p_t^n has (at least) a weak incentive to continue experimenting until t' , conditional on no event on $[t_n^*, t')$. By Observation 8, for every $\mu(t') \in (\underline{\mu}_{t'}^n, 1)$,

$$V_{t_j^*}^{\bar{T}_n(\mu(t'))}(p_{t_j}^n, \mu(t')) > V_{t_j^*}^{\bar{T}_n(\underline{\mu}_{t'}^n)}(p_{t_j}^n, \underline{\mu}_{t'}^n), \quad j = n, \dots, k.$$

Therefore, under the optimal strategy $\bar{T}_n(\mu(t'))$, we have

$$V_{t_j^*}^{\bar{T}_n(\mu(t'))}(p_{t_j}^n, \mu(t')) > (1 - \alpha)s + \alpha p^*, \quad j = n, \dots, k.$$

Thus, given $\mu(t') \in (\underline{\mu}_{t'}^n, 1)$, at every $t < t'$ type p_t^n has a strict incentive to continue experimenting until t' , conditional on no event on $[t_n^*, t')$. We therefore have $M_n(t') = [p_{t'}^0, \underline{\mu}_{t'}^n)$.

We begin by considering type $p_{t_0}^0$ at date t_0^* . If $M_0(t') = [p_{t'}^0, 1)$, then stopping at date t' can be excluded for type $p_{t'}^0$ by equilibrium dominance.

If $\underline{\mu}_{t'}^n < 1$, then under $\bar{T}_0(\underline{\mu}_{t'}^0)$ we have

$$V_{t_n^*}^{\bar{T}_0(\underline{\mu}_{t'}^0)}(p_{t_n}^0, \underline{\mu}_{t'}^0) \geq (1 - \alpha)s + \alpha p^*, \quad n = 0, \dots, k;$$

For every $n = 1, \dots, k$, the strategy $\bar{T}_0(\underline{\mu}_{t'}^0)$ is available to type $p_{t_n}^1$ at t_n^* . Moreover, by observation 9 we have:

$$V_{t_n^*}^{\bar{T}_0(\underline{\mu}_{t'}^0)}(p_{t_n}^1, \underline{\mu}_{t'}^0) > V_{t_n^*}^{\bar{T}_0(\underline{\mu}_{t'}^0)}(p_{t_n}^0, \underline{\mu}_{t'}^0), \quad n = 1, \dots, k.$$

Therefore, under the optimal strategy $\bar{T}_1(\underline{\mu}_{t'}^0)$, we have

$$V_{t_n^*}^{\bar{T}_1(\underline{\mu}_{t'}^0)}(p_{t_n}^1, \underline{\mu}_{t'}^0) > (1 - \alpha)s + \alpha p^*, \quad n = 1, \dots, k,$$

Thus, given $\underline{\mu}_{t'}^0$, at every $t < t'$ type p_t^1 has a strict incentive to continue experimenting until t' , conditional on no event on $[t_n^*, t')$. Consequently we must have $\underline{\mu}_{t'}^1 < \underline{\mu}_{t'}^0$. Therefore, then stopping at date t' can be excluded for type $p_{t'}^0$ under the D1 criterion¹³.

Finally, we proceed by induction and eliminate all types $p_{t'}^n < p_{t'}^k$.

- Second, we eliminate all types $p_{t'}^n > p_{t'}^{k+1}$.

Consider type $p_{t'}^n > p_{t'}^{k+1}$ at date t' . Her payoff on the equilibrium path is

$$V_{n,t'}^* := (1 - \alpha) U^0(p_{t'}^n) + \alpha p_{t'}^n.$$

¹³ According to the D1 criterion, we can eliminate the deviation t' for type $p_{t'}^n$ if there exists another type $p_{t'}^{n'}$ such that:

$$[\underline{\mu}_t^n, 1) \subset (\underline{\mu}_t^{n'}, 1),$$

or, equivalently, if there exists a type $p_{t'}^{n'}$ for whom $\underline{\mu}_{t'}^{n'} > \underline{\mu}_{t'}^n$. In words: we can eliminate $p_{t'}^n$ if there exists a $p_{t'}^{n'}$ who has a strict incentive to deviate whenever $p_{t'}^n$ has a strict or weak incentive to deviate. By this elimination process, the type who has the strongest incentives to deviate to t' remains.

The payoff from stopping at date t' (off path) is

$$u(p_{t'}^n, t', \mu(t')) = (1 - \alpha) s + \alpha \mu(t').$$

Stopping at date t' is a weakly profitable deviation for $p_{t'}^n$ if

$$(12) \quad \mu(t) \geq \underline{\mu}_{t'}^n := p_{t'}^n + \frac{1 - \alpha}{\alpha} (U^0(p_{t'}^n) - s).$$

If $\underline{\mu}_{t'}^n \geq 1$ then stopping at date t' can be excluded for type $p_{t'}^n$ by equilibrium dominance. If $\underline{\mu}_{t'}^n < 1$, the set of reputations, $\mu(t')$, giving type $p_{t'}^n$ a strict incentive to stop at t' is $(\underline{\mu}_{t'}^n, 1)$, while at $\mu(t') = \underline{\mu}_{t'}^n$ the incentive is weak.

From (12) it is easy to see that $\underline{\mu}_{t'}^n$ is a strictly increasing function of n , so that $\underline{\mu}_{t'}^n < \underline{\mu}_{t'}^{n+1}$ for all $n \geq k$. We can therefore eliminate the deviation t' for all types $p_{t'}^n > p_{t'}^{k+1}$.

- Finally, we show that there exist off path beliefs satisfying the D1 criterion that support the efficient PBE.

For $p_{t_k^*}^k$ to deviate to t' we need

$$(1 - \alpha)e^{-\rho(t' - t_k^*)} [s - u^0(p_{t'}^k)] + \alpha [\mu(t') - p_{t'}^k] > 0.$$

or equivalently:

$$\mu(t') \geq \underline{\mu}_{t'}^k(\alpha) := p_{t'}^k + \frac{1 - \alpha}{\alpha} e^{-\rho(t' - t_k^*)} (u^0(p_{t'}^k) - s)$$

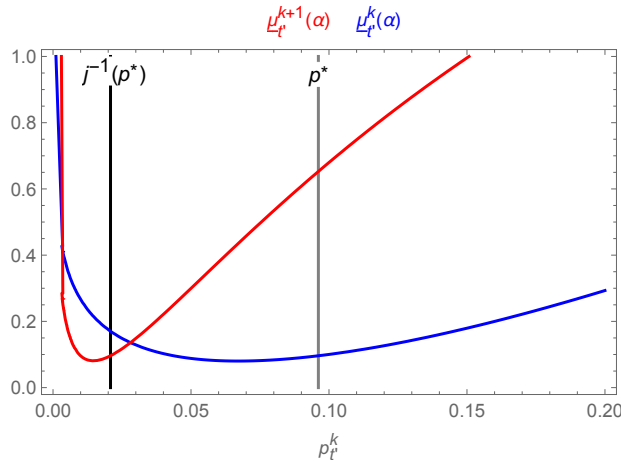
where $e^{-\rho(t' - t_k^*)} = \left(\frac{p^*}{1 - p^*} \frac{1 - p}{p} \right)^{\frac{-\rho}{\eta + \lambda_G - \lambda_B}}$. For $p_{t'}^{k+1}$ to deviate we need

$$(1 - \alpha) [s - U^0(p_{t'}^{k+1})] + \alpha [\mu(t') - p_{t'}^{k+1}] \geq 0.$$

or equivalently:

$$\mu(t') \geq \underline{\mu}_{t'}^{k+1}(\alpha) := j(p_{t'}^k) + \frac{1 - \alpha}{\alpha} (U^0(j(p_{t'}^k)) - s)$$

The figure below illustrates $\underline{\mu}_{t'}^{k+1}(\alpha)$ and $\underline{\mu}_{t'}^k(\alpha)$ for $\alpha \leq \bar{\alpha}$. (For $\alpha > \bar{\alpha}$ the two thresholds intersect at $p < j^{-1}(p^*)$, and we have $\underline{\mu}_{t'}^k(\alpha) < \underline{\mu}_{t'}^{k+1}(\alpha)$ for every $(j^{-1}(p^*), p^*)$.)



Let $t_\mu(\alpha) \in (t_k^*, t_{k+1}^*)$ be the date satisfying:

$$\underline{\mu}_{t_\mu(\alpha)}^k(\alpha) = \underline{\mu}_{t_\mu(\alpha)}^{k+1}(\alpha)$$

Then, under D1, we eliminate $p_{t'}^{k+1}$ for every $t' \in (t_k^*, t_\mu(\alpha))$, and we eliminate $p_{t'}^{k+1}$ for every $t' \in (t_\mu(\alpha), t_{k+1}^*)$. At t_{k+1}^* , both types, $p_{t'}^k$ and $p_{t'}^{k+1}$, are possible.

Thus the off-path reputation satisfying D1 is

$$(13) \quad \mu(t') \in M^{D1}(t') := \begin{cases} \{p_{t'}^k\} & \text{if } t' \in (t_k^*, t_\mu(\alpha)) \\ \{p_{t'}^{k+1}\} & \text{if } t' \in (t_\mu(\alpha), t_{k+1}^*) \\ [p_{t'}^k, p_{t'}^{k+1}] & \text{if } t' = t_\mu(\alpha). \end{cases}$$

Finally, we show that the efficient PBE is supported by off-path beliefs in $M^{D1}(t')$. For every $t' \in (t_k^*, t_\mu(\alpha))$, $\mu t' = p_{t'}^k$, and our equilibrium is supported. It is also supported at date $t' = t_\mu(\alpha)$ if we choose $\mu t_\mu(\alpha) = p_{t_\mu(\alpha)}^k$.

For every $t' \in (t_\mu(\alpha), t_{k+1}^*)$, $\mu t' = p_{t'}^{k+1}$. By definition, on this interval, $\underline{\mu}_{t_\mu(\alpha)}^k(\alpha) > \underline{\mu}_{t_\mu(\alpha)}^{k+1}(\alpha)$. Since $\underline{\mu}_{t_\mu(\alpha)}^{k+1}(\alpha) \geq \mu t' = p_{t'}^{k+1}$, the reputation $\mu t' = p_{t'}^{k+1}$ is not sufficiently high to make a deviation to t' profitable for type $p_{t_k^*}^k$. For type $p_{t'}^{k+1}$, stopping at t' gives the payoff

$$(1 - \alpha)s + \alpha p_{t'}^{k+1} < (1 - \alpha)U^0(p_{t'}^{k+1}) + \alpha p_{t'}^{k+1},$$

where the right-hand side of the inequality is the equilibrium payoff for type $p_{t'}^{k+1}$. Thus, deviating to t' is not profitable for type $p_{t'}^{k+1}$.

We have therefore established Proposition 7(a).

(b) Fix an inefficient PBE.

Consider date $t_1^* < \hat{t}_1$.

Suppose $\mu(t_1^*)$ is such that type $p_{t_1^*}^0$ has a (weak) incentive to engage in deviation $D_0(t_0^*, t_1^*)$. We know that for this to be the case, it is necessary that $\mu(t_1^*) > p^*$. Then the payoff from stopping at t_1^* for type $p_{t_1^*}^1 = p^*$ is

$$(1 - \alpha)s + \alpha \mu(t_1^*)$$

while her expected payoff from abiding by the equilibrium strategy is

$$(1 - \alpha)U^0(p^*) + \alpha p^*$$

Since $U^0(p^*) = s$ and $\mu(t_1^*) > p^*$, the net payoff from deviating is positive and type $p_{t_1^*}^1 = p^*$ has a strict incentive to stop at t_1^* . \square

\square

4.4 Changing the Information Structure

Let the threshold from Proposition 3 explicitly depend on the news technology, that is, on the pair (λ_B, λ_G) . The DM with preference parameter $\bar{\alpha}(\lambda_B, \lambda_G)$ is just indifferent between implementing the planner solution and deviating. A DM with preference parameter $\alpha > \bar{\alpha}(\lambda_B, \lambda_G)$ cannot implement the planner solution in a PBE, as she values the net reputational benefit resulting from a deviation more highly than its expected net social cost.

Proposition 10 describes how the threshold $\bar{\alpha}(\lambda_B, \lambda_G)$ varies with the news technology. In the limit case where news events only occur in the good state ($\lambda_B = 0$), deviating from the planner solution is never profitable for a DM, unless she is only concerned with her reputation, i.e. has preference parameter $\alpha = 1$. In contrast, if news events do occur in the bad state, then a DM may find it profitable to deviate from the planner solution if her preference parameter lies above the threshold $\bar{\alpha}(\lambda_B, \lambda_G)$, which lies in $(0, 1)$. This is the case even if news event become increasingly infrequent in the bad state ($\lim_{\lambda_B \rightarrow 0} \bar{\alpha}(\lambda_B, \lambda_G) < 1$). There is thus a discontinuity between the case $\lambda_B \rightarrow 0$ and the actual limit case $\lambda_B = 0$.

On the other hand, if news events occur with increasing frequency in the good state ($\lambda_G \rightarrow \infty$) we find again that the DM can always implement the planner solution, unless $\alpha = 1$. This proposition indicates that the comparative statics effect of reducing λ_B and augmenting λ_G are not equivalent.

One important implication of Proposition 10 is that, for any given preference parameter $\alpha \in (0, 1)$ of the DM, there exists an information structure $\lambda_G(\alpha), \lambda_B(\alpha)$ such that the planner policy constitutes an equilibrium of the game between the reputationally concerned DM and the observer.

Proposition 10. (a) Fix $\lambda_G > 0$. Then $\bar{\alpha}(0, \lambda_G) = 1$.

(b) Fix $\lambda_G > 0$. Then $\lim_{\lambda_B \rightarrow 0} \bar{\alpha}(\lambda_B, \lambda_G) =: \bar{\bar{\alpha}}(\lambda_G)$, where $\bar{\bar{\alpha}}(\lambda_G) \in (0, 1)$ is the unique value satisfying

$$(14) \quad \frac{\bar{\bar{\alpha}}(\lambda_G)}{1 - \bar{\bar{\alpha}}(\lambda_G)} = \frac{\frac{1}{1-\underline{p}^*} \left(s - \underline{p}^* \frac{\eta g}{\eta + \rho} \right)}{\underline{p}^*},$$

and \underline{p}^* is the planner's threshold belief when $\lambda_B = 0$

(c) Fix $\lambda_B \geq 0$. Then $\lim_{\lambda_G \rightarrow \infty} \bar{\alpha}(\lambda_B, \lambda_G) = 1$. In particular, $\lim_{\lambda_G \rightarrow \infty} \bar{\bar{\alpha}}(\lambda_G) = 1$.

Proof. Consider the DM whose belief at date t_n^* equals the planner threshold p^* . We have argued that the DM implements the planner solution if and only if the deviation $D_n(t_n^*, t_{n+1}^*)$ is not profitable. We wish to describe how the DM's incentive to deviate varies with (λ_B, λ_G) .

By lemma 4, for a DM with preference parameter α , the deviation $D_n(t_n^*, t_{n+1}^*)$ is profitable if and only if

$$(15) \quad \alpha \left[p^* - p_{t_{n+1}^*}^n \right] + (1 - \alpha) e^{-\rho(t_{n+1}^* - t_n^*)} \left[s - u^0(p_{t_{n+1}^*}^n) \right] > 0.$$

Evaluating the terms in the expression above using the fact that $e^{-(\eta + \lambda_G - \lambda_B)(t_{n+1}^* - t_n^*)} = \frac{\lambda_B}{\lambda_G}$, we have

$$(16) \quad p_{t_{n+1}^*}^n = j^{-1}(p^*).$$

Since

$$\lim_{\lambda_B \rightarrow 0} p^* = \underline{p}^*$$

and

$$(17) \quad \lim_{\lambda_B \rightarrow 0} j^{-1}(p^*) = 0$$

we have

$$\lim_{\lambda_B \rightarrow 0} [p^* - p_{t_{n+1}}^n] = \underline{p}^* \in (0, 1).$$

The second term in square brackets in equation (15) is given by

$$s - u^0(j^{-1}(p^*)) = s - \left(j^{-1}(p^*) \frac{\eta g}{\eta + \rho} + \frac{1 - j^{-1}(p^*)}{1 - p^*} \left(s - p^* \frac{\eta g}{\eta + \rho} \right) \left(\frac{1 - j^{-1}(p^*)}{j^{-1}(p^*)} \frac{p^*}{1 - p^*} \right)^{\nu(\lambda_G, \lambda_B)} \right)$$

where we can simplify

$$\left(\frac{1 - j^{-1}(p^*)}{j^{-1}(p^*)} \frac{p^*}{1 - p^*} \right)^{\nu(\lambda_G, \lambda_B)} = \left(\frac{\lambda_G}{\lambda_B} \right)^{\nu(\lambda_G, \lambda_B)}.$$

Consequently,

$$e^{-\rho(t_{n+1}^* - t_n^*)} [s - u^0(j^{-1}(p^*))] = \left(\frac{\lambda_B}{\lambda_G} \right)^{\frac{\rho}{\eta + \lambda_G - \lambda_B}} \left[s - \left(j^{-1}(p^*) \frac{\eta g}{\eta + \rho} + \frac{1 - j^{-1}(p^*)}{1 - p^*} \left(s - p^* \frac{\eta g}{\eta + \rho} \right) \left(\frac{\lambda_G}{\lambda_B} \right)^{\nu(\lambda_G, \lambda_B)} \right) \right].$$

Using (17) and

$$\lim_{\lambda_B \rightarrow 0} \nu(\lambda_G, \lambda_B) = \frac{\rho}{\eta + \lambda_G},$$

we obtain

$$\begin{aligned} \lim_{\lambda_B \rightarrow 0} e^{-\rho(t_{n+1}^* - t_n^*)} [s - u^0(j^{-1}(p^*))] &= \lim_{\lambda_B \rightarrow 0} \left(\frac{\lambda_B}{\lambda_G} \right)^{\frac{\rho}{\eta + \lambda_G - \lambda_B}} \left[s - \frac{1}{1 - p^*} \left(s - p^* \frac{\eta g}{\eta + \rho} \right) \left(\frac{\lambda_G}{\lambda_B} \right)^{\nu(\lambda_G, \lambda_B)} \right] \\ &= -\frac{1}{1 - \underline{p}^*} \left(s - \underline{p}^* \frac{\eta g}{\eta + \rho} \right). \end{aligned}$$

Therefore, $\bar{\alpha}(\lambda_G) := \lim_{\lambda_B \rightarrow 0} \bar{\alpha}(\lambda_B, \lambda_G)$ is the unique value satisfying (14). Finally, since for $\lambda_G < \infty$ the right-hand side of (14) is finite, we have that $\bar{\alpha}(\lambda_G) \in (0, 1)$, establishing (b).

We now prove (c). First, observe that $\lim_{\lambda_G \rightarrow \infty} \underline{p}^* = 0$ so that the right-hand side of equation (14) tends to $+\infty$ and $\lim_{\lambda_G \rightarrow \infty} \bar{\alpha}(\lambda_G) = 1$.

Now consider equation (15). Since $\lim_{\lambda_G \rightarrow \infty} \nu(\lambda_G, \lambda_B) = 0$ we have that $\lim_{\lambda_G \rightarrow \infty} p^* = 0$. Since $p^* > p_{t_{n+1}}^n$, we also have that $\lim_{\lambda_G \rightarrow \infty} p_{t_{n+1}}^n = 0$. Hence, the net reputational benefit from the deviation $D_n(t_n^*, t_{n+1}^*)$ tends to zero:

$$\lim_{\lambda_G \rightarrow \infty} [p^* - p_{t_{n+1}}^n] = 0.$$

However, as $\lambda_G \rightarrow \infty$, the time interval $t_{n+1}^* - t_n^*$ tends to zero and $e^{-\rho(t_{n+1}^* - t_n^*)} [s - u^0(p_{t_{n+1}}^n)]$, the social cost of the deviation $D_n(t_n^*, t_{n+1}^*)$, also tends to zero.

We therefore rewrite equation (15) as

$$(18) \quad \frac{\alpha}{1 - \alpha} > -\frac{e^{-\rho(t_{n+1}^* - t_n^*)} [s - u^0(p_{t_{n+1}}^n)]}{p^* - p_{t_{n+1}}^n}.$$

The right-hand side equals

$$e^{-\rho(t_{n+1}^* - t_n^*)} \left(s \underbrace{\frac{-\left(1 - \frac{1 - p_{t_{n+1}}^n}{1 - p^*} \left(\frac{\lambda_G}{\lambda_B} \right)^{\nu(\lambda_G, \lambda_B)}\right)}{p^* - p_{t_{n+1}}^n}}_{=:A} - \frac{g\eta}{\eta + \rho} \underbrace{\frac{p^* \frac{1 - p_{t_{n+1}}^n}{1 - p^*} \left(\frac{\lambda_G}{\lambda_B} \right)^{\nu(\lambda_G, \lambda_B)} - p_{t_{n+1}}^n}{p^* - p_{t_{n+1}}^n}}_{=:B} \right).$$

Simplifying B , we have

$$B = \frac{1}{1-p^*} \frac{\left(\frac{\lambda_G}{\lambda_B}\right)^{\nu(\lambda_G, \lambda_B)} \frac{\lambda_G}{\lambda_B} - 1}{\frac{\lambda_G}{\lambda_B} - 1}.$$

Since $\lim_{\lambda_G \rightarrow \infty} \left(\frac{\lambda_G}{\lambda_B}\right)^{\nu(\lambda_G, \lambda_B)} = 1$ and $\lim_{\lambda_G \rightarrow \infty} p^* = 0$ we have

$$\lim_{\lambda_G \rightarrow \infty} B = \frac{\frac{\lambda_G}{\lambda_B} - 1}{\frac{\lambda_G}{\lambda_B} - 1} = 1.$$

Since $\lim_{\lambda_G \rightarrow \infty} p_{t_{n+1}^*}^n = 0$ we have

$$\lim_{\lambda_G \rightarrow \infty} A = -\frac{1 - \frac{1}{1-p^*} \left(\frac{\lambda_G}{\lambda_B}\right)^{\nu(\lambda_G, \lambda_B)}}{p^*}.$$

As $\lambda_G \rightarrow \infty$ the interval $\left[\frac{\rho}{\eta + \lambda_G - \lambda_B}, \frac{\rho + \lambda_B}{\eta + \lambda_G - \lambda_B}\right]$ containing $\nu(\lambda_G, \lambda_B)$ converges to the point $1/\lambda_G$. Therefore,

$$\lim_{\lambda_G \rightarrow \infty} A = -\frac{1 - \lambda_G^{\frac{1}{\lambda_G}}}{\frac{1}{\lambda_G}}.$$

Both the numerator and the denominator of the expression above tend to 0 as $\lambda_G \rightarrow \infty$. Applying l'Hospital's rule, we have

$$\lim_{\lambda_G \rightarrow \infty} A = \lim_{\lambda_G \rightarrow \infty} \left(\lambda_G^{\frac{1}{\lambda_G}} (\ln \lambda_G - 1) \right) = +\infty.$$

Hence,

$$\lim_{\lambda_G \rightarrow \infty} e^{-\rho(t_{n+1}^* - t_n^*)} \left(s A - \frac{g\eta}{\eta + \rho} B \right) = +\infty.$$

Therefore, for all $\alpha < 1$ there exists a λ_G such that condition (18) is violated, and the deviation $D_n(t_n^*, t_{n+1}^*)$ is not profitable, establishing (c).

Finally, we establish (a). The planner solution when $\lambda_B = 0$ is described in corollary 2. There is a unique date \underline{t}^* satisfying $p(\underline{t}^*) = \underline{p}^*$ at which the planner might stop on path. The analogue of the deviation $D_0(t_0^*, t_1^*)$ for a DM who has not observed any news by date $\underline{t}^* = t_0^*$ is an infinite deviation under which the DM never repeals the project.

We can argue by continuity that such a deviation cannot be profitable for a DM unless she puts no weight at all on its social cost ($\alpha = 1$). Under an infinite deviation, if no news and no success arrive, the DM's belief $p(t)$ converges to zero. Since $s \in (0, g)$, if $p(t) = 0$, repealling the project strictly dominates continuing for all $\alpha < 1$. By continuity of the posterior belief conditional on no news and no success, and of the function $u^0(p)$, there exists a date $\underline{t}'(\alpha)$ at which repealling the project becomes a dominant strategy for the DM with preference parameter $\alpha < 1$. \square

5 Bad news with downward drift

Let's now turn our attention to the "bad news" setting. Assume that news events are more frequent when the project is bad than when it is good: $\lambda_G < \lambda_B$. News events can be interpreted as "breakdowns". Given her belief $p(t)$, following a privately observable breakdown the DM's private belief jumps down to $j(p(t)) < p(t)$. In this section, maintain

the assumption that $(\eta + \lambda_G - \lambda_B) > 0$, so that as long as no news arrives on the interval $[t, t + dt)$ the DM's private belief drifts *down* according to the law of motion $dp = -p(1 - p)(\eta + \lambda_G - \lambda_B)dt$.

5.1 Planner policy

For the DM with preference parameter $\alpha = 0$, the optimal policy is to stop at the first time at which $p(t) \leq p^b$. Here, p^b denotes the planner threshold belief in the bad news case. Bearing in mind that $\lambda_G < \lambda_B$, the value function V^0 in the planner problem with bad news solves the Bellman equation (2).

In the bad news case, the posterior belief can discontinuously jump into the stopping region, or continuously drift into it. As a result the planner value function V^0 is C^1 everywhere, including at p^b , and “smooth-pasting” is satisfied¹⁴. It admits different expressions on each interval $[p^b, j^{-1}(p^b))$ and $[j^{-n}(p^b), j^{-(n+1)}(p^b))$, $n \geq 1$. The next theorem only gives an explicit expression on the first interval¹⁵. Observe that the optimal threshold belief does not depend on the news process. Here is a heuristic argument explaining it. At the threshold belief, the planner is indifferent between stopping immediately, and committing to keeping the project active for a short time interval dt . If over that time interval the project succeeds, the social payoff g accrues. In all other cases, the planner repeals the project after the time interval dt has elapsed, regardless of how many pieces of news she observed during that time. For an intesimal dt , we loosely obtain the indifference condition: $s = p\eta dt g + (1 - p\eta dt - \rho dt)s$.

Proposition 11. *The planner's optimal policy is to stop at the first time t such that $p(t) \leq p^b$. The planner threshold belief satisfies*

$$(19) \quad p^b = \frac{\rho s}{\eta(g - s)}.$$

On the interval $[0, j^{-1}(p^b))$, the social payoff under the optimal policy is

$$(20) \quad V^0(p) = \begin{cases} s & \text{if } p < p^b, \\ v_1^0(p) & \text{if } p^b \leq p < j^{-1}(p^b), \end{cases}$$

where

$$(21) \quad v_1^0(p) := p \frac{\eta g + \lambda_G s}{\eta + \lambda_G + \rho} + (1 - p) \frac{\lambda_B s}{\lambda_B + \rho} + \left(s - \left(p^b \frac{\eta g + \lambda_G s}{\eta + \lambda_G + \rho} + (1 - p^b) \frac{\lambda_B s}{\lambda_B + \rho} \right) \right) \frac{1 - p}{1 - p^b} \left(\frac{1 - p}{p} \frac{p^b}{1 - p^b} \right)^{\frac{\lambda_B + \rho}{\eta + \lambda_G - \lambda_B}}.$$

¹⁴ In Keller and Rady (2015) the planner's value function does not satisfy the smooth-pasting condition at the optimal stopping threshold, because the belief *drifts away from* and *jumps towards* the stopping region. As a consequence, in their model, the posterior belief can only enter the stopping region following a jump. In my model, the posterior belief would follow a motion similar to that in Keller and Rady (2015) if we had $(\eta + \lambda_G - \lambda_B) < 0$.

¹⁵ See the appendix for the remaining intervals.

On the interval $[j^{-1}(p^b), 1]$, $V^0(p)$ is defined recursively.

Proof. When $\alpha = 0$, the common value function for the DM and the observer, $V^0(p(t)) := \sup_T V_t^{0,T}$, is convex and continuous. The function V^0 solves the Bellman equation (2) — where, in the bad news context, $b_N(p, u)$ now reflects the expected loss from a piece of news — and is the unique solution satisfying the boundary conditions $u(0) = s$ and $u(1) = \gamma$. As will be verified, it is equal to s on $[0, p^b]$ and solves the ordinary differential difference equation (7) in p on $(p^b, 1)^{16}$ for some p^b to be found.

We begin by considering the equation (7) on $[b, 1]$ with some $b \in (0, 1)$ given and fixed. Observe that the function $p \mapsto j(p)$ is decreasing and finds points $\dots > b_2 > b_1 > b_0 := b$ such that $j(b_n) = b_{n-1}$ for $n \geq 1$. Let $I_n := [b_{n-1}, b_n]$ for $n \geq 1$.

Consider the equation (7) on I_1 after setting $u(p) = s$ for $p \in [j(b), b]$:

$$(22) \quad u(p)(p\eta + \rho) + \lambda(p)[u(p) - s] + (1-p)p(\eta + \Delta\lambda)u'(p) = p\eta g.$$

This is then a first-order differential equation which can be solved explicitly. Let $p \mapsto v_1^0(p)$ denote the solution to this differential equation on I_1 . The function $p \mapsto v_0^0(p)$ with

$$(23) \quad v_0^0(p) = \int_0^\infty \left(p e^{-(\eta + \lambda_G + \rho)t} (\eta g + \lambda_G s) + (1-p) e^{-(\lambda_B + \rho)t} \lambda_B s \right) dt = p \frac{\eta g + \lambda_G s}{\eta + \lambda_G + \rho} + (1-p) \frac{\lambda_B s}{\lambda_B + \rho}$$

constitutes a particular solution to (22). It is the payoff to the following policy, given current belief p : repeal the project as soon as a piece of news arrives, provided this occurs before the project succeeds; otherwise keep the project active. The function $p \mapsto (1-p)\Omega(p)^{\frac{\lambda_B + \rho}{\eta + \Delta\lambda}}$ constitutes a solution to the homogeneous version of (22). The solution to (22) on I_1 is therefore given by the family of functions

$$v_{C_1}(p) = v_0^0(p) + C_1 (1-p)\Omega(p)^{\frac{\lambda_B + \rho}{\eta + \Delta\lambda}},$$

where $C_1 \in \mathbb{R}$ is a constant of integration.

Imposing the continuity condition $v_{C_1}(b) = s$ at b (in agreement with $V^0(p^b) = s$), we solve for C_1 and obtain a unique solution $p \mapsto v_1^0(p)$ on I_1 where

$$(24) \quad v_1^0(p, b) := v_0^0(p) + (s - v_0^0(b)) \frac{1-p}{1-b} \left(\frac{\Omega(p)}{\Omega(b)} \right)^{\frac{\lambda_B + \rho}{\eta + \Delta\lambda}},$$

Considering equation (22) on I_2 , using the solution found on I_1 , we then have a first-order differential equation which can be solved explicitly. Imposing a continuity condition over $I_1 \cup I_2$ at b_1 , we obtain a unique solution $p \mapsto v_2^0(p, b)$ on I_2 . And so on by induction.

Choosing $b \in (0, 1)$ so as to maximise $v_1^0(p, b)$ for every $p \in (0, 1)$ (the solution is interior and smooth-pasting is satisfied), we obtain the expression in (19) for p^b .

□

¹⁶As the posterior belief always drifts down, we say that a continuous function solves the following ODDE if its left-hand derivative exists and (7) holds when this left-hand derivative is used to compute $V^{0'}(p)$.

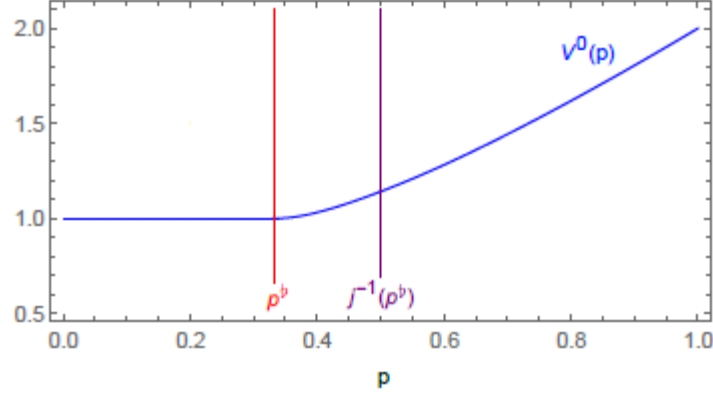


Figure: The planner value, $V^0(p)$, illustrated for $(\eta, \rho, \lambda_G, \lambda_B, G, S) = (10, 10, 3, 6, 4, 1)$

There are two important qualitative differences with the good-news case. First, $t_{n+1}^* < t_n^*$ for all $n \geq 0$, and at any state $t \geq 0$ the DM with $N(t) = 0$ is the most optimistic type. Second, since it is now possible for the DM's private belief to jump into the stopping region, stopping on (t_{n+1}^*, t_n^*) for each $n \geq 0$ occurs on path with strictly positive probability. Stopping after t_0^* does not.

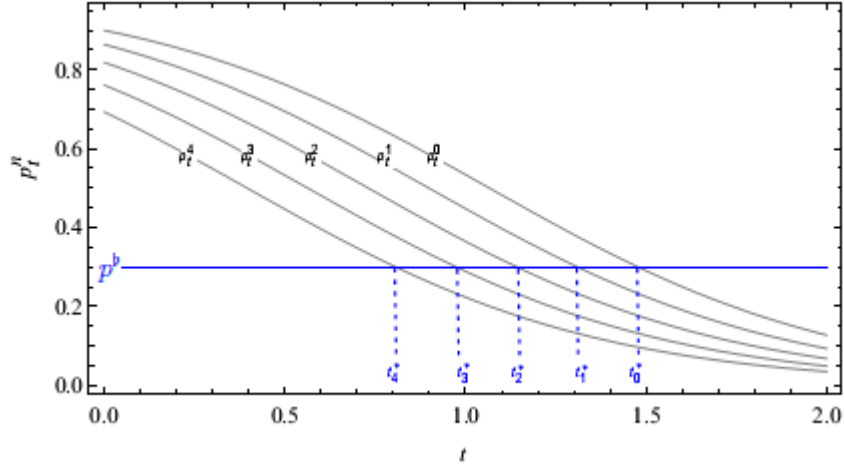


Figure: The planner solution in the bad news case.

5.2 Implementing the Planner Solution

Different from the good news case, with bad news, no preference type finds it optimal to implement the planner solution in a PBE.

Proposition 12. *For every $\alpha > 0$, the social planner solution cannot be implemented as a PBE.*

Proof. I prove the proposition by constructing a profitable deviation from the planner strategy. Consider type $n = 0$ who receives her first piece of news at $t' \in (t_1^*, t_0^*)$ and let us vary t' .

Consider that DM's net payoff from following the deviation $D_0(t', t_0^*)$. Since it departs from the planner solution, her expected net social payoff is negative. However, her expected net reputational gain under

$D_0(t', t_0^*)$ is strictly positive and equal to $p^* - p_{t'}^1$. On path the DM stops at t' and her reputation is $\mu(t') = p_{t'}^1 = p(t')$. Under $D_0(t', t_0^*)$ she stops at t_0^* for all histories. Her reputation is then $\mu(t_0^*) = p^* > p(t_0^*)$.

As $t' \rightarrow t_0^*$, the DM's net reputational gain from deviating, $p^* - p_{t'}^1$, increases; while her expected social loss from deviating tends to zero. By continuity there exists a value $t'' \in (t_1^*, t_0^*)$ such that for all $t' \in (t'', t_0^*)$, the deviation $D_0(t', t_0^*)$ is strictly profitable. \square

This result is strikingly different from the good news case, where a DM not too concerned with her reputation would implement the planner policy in a PBE. A distinguishing feature of the bad news case is that beliefs are pinned down by Bayes' rule at all $t \in (0, t_0^*]$, as under the planner solution, the support of possible stopping dates is connected. In contrast, in the good news case, the support of possible stopping dates under the planner policy, $\{t_n^*\}_{n \geq 0}$ was not connected.

5.3 Equilibrium

A strategy for the DM is characterised by the function $\hat{p}(t)$. The DM stops at the first date at which $p(t) \leq \hat{p}(t)$.

The next lemma states that, in equilibrium, no information type continues experimenting forever.

Lemma 13. *Equilibrium requires that $\exists t_0 \geq 0$ such that $\hat{p}(t_0) \geq p_{t_0}^0$.*

Proof. Suppose, by way of contradiction, that in equilibrium, for every $t \geq 0$, $\hat{p}(t) < p_t^0$. Then, for every t , a DM who sees no news and no success before t has to continue experimenting. In other words, for every $t > 0$, stopping at date t is worse than continuing to experiment for a DM with belief $\hat{p}(t) < p(t) < 1$:

$$\alpha \mu(t) + (1 - \alpha)s \leq \alpha \mathbb{E}[\mu(\tau \wedge T)|p(t)] + (1 - \alpha)\mathbb{E}\left[e^{-\rho(\tau \wedge T)}(g\mathbb{1}_{\{\tau \leq T\}} + s\mathbb{1}_{\{\tau > T\}})|p(t)\right]$$

Observe that $\mu(t) < p_t^0$, so that, as t grows without bounds, the left-hand side above tends to $(1 - \alpha)s$. Moreover, conditional on no success, $p(t) \rightarrow 0$, and the right-hand side above tends to

$$(1 - \alpha)\mathbb{E}\left[e^{-\rho T}s|\theta = B\right] < (1 - \alpha)s;$$

This establishes the contradiction. \square

Suppose that p^0 is high enough that in equilibrium, $\hat{p}(0) < p^0$.

Lemma 14. *In equilibrium, $\hat{p}(t)$ cannot be continuous for every $t \geq 0$.*

Proof. Suppose, by way of contradiction, that in equilibrium, $\hat{p}(t)$ is continuous for every $t \geq 0$. By Lemma 13, there exists t_0 such that $\hat{p}(t_0) \geq p_{t_0}^0$. Moreover we assumed that $\hat{p}(0) < p^0$. Consequently, there exists an integer $m \geq 0$ and a date $t_m \geq 0$ such that $\hat{p}(t_m) = p_{t_m}^m$, and an open interval $(t_m - \varepsilon, t_m + \varepsilon)$ such that

$$\mu(t) = \begin{cases} p_t^{m+1} & \forall t \in (t_m - \varepsilon, t_m) \\ p_t^m & \forall t \in [t_m, t_m + \varepsilon). \end{cases}$$

Consequently, for a DM whose belief enters the stopping region at date $t' \in (t_m - \varepsilon, t_m)$, the deviation $D_{m+1}(t', t_m)$ generates a strictly positive expected net reputational benefit that is increasing in t' . The expected net social cost from such a deviation is strictly decreasing in t' , and tends to zero as $t' \rightarrow t_m$. Thus, there exists a $t'' < t_m$ such that, for every $t'' < t' < t_m$, the deviation $D_{m+1}(t', t_m)$ is strictly profitable. A contradiction. \square

next steps:

$hatp(t)$ must have discontinuities. But let's keep assuming that it has full support, i.e. that every $0 \leq t \leq t_0$ is a possible stopping date.

Can all discontinuities be downward jumps? No, as this would not allow us to satisfy $\hat{p}(0) < p^0$ and $\hat{p}(t_0) \geq p_{t_0}^0$ without having an m and a date t_m at which $\hat{p}(t_m) = p_{t_m}^m$, and the argument from Lemma 14 excluding our strategy profile as an equilibrium

is $hatp(t)$ left- or right-continuous? no restrictions!

Can a discontinuity att' be an upward jump? only if

1) there is a gap in the support (i.e. $\hat{p}(t) = 0$ to the left of the upward jump. otherwise, waiting for the jump is profitable.

2) the function $\hat{p}(t)$ cannot be right-continuous at the jump, otherwise the DM whose belief enters the stopping region at date t' prefers delaying her exit.

So our candidate profiles either have a support which is a finite collection of points at which $\hat{p}(t) > 0$. or (see picture)

=====

Consider the following strategy profile, parameterised by the sequence $\{\hat{t}_k\}_{k=0}^K$ with $\hat{t}_{k+1} < \hat{t}_k$:

$$(25) \quad \begin{aligned} \hat{p}(t) &= 0 & \forall t \in \mathbb{R}^+ \setminus \{\hat{t}_k\}_{k=0}^K \\ \hat{p}(t) &= p^* & \forall t \in \{\hat{t}_k\}_{k=0}^K \end{aligned}$$

Under this strategy profile, the reputation of a DM who repeals the project at date \hat{t}_k is

$$\mu(\hat{t}_k) = \mathbb{E}[p(\hat{t}_k) | p(\hat{t}_k) \leq \hat{p}(\hat{t}_k), p(\hat{t}_{k+1}) > \hat{p}(\hat{t}_{k+1})]$$

for every $k < K$, and $\mu(\hat{t}_K) = \mathbb{E}[p(\hat{t}_K) | p(\hat{t}_K) \leq \hat{p}(\hat{t}_K)]$.

Observe that for every $t \in \mathbb{R}^+ \setminus \{\hat{t}_k\}_{k=0}^K$, $\mu(t) = 0$. Observe further that under this strategy profile, \hat{t}_0 is the last date at which the DM repeals the project.

A strategy profile becomes more socially efficient as the induced support of stopping dates becomes more dense. Consider the strategy $\hat{p}(t)$ defined in (26), and let $\hat{\hat{p}}(t) = f(\hat{p}(t), t')$ be the strategy:

$$(26) \quad \begin{aligned} \hat{\hat{p}}(t) &= 0 & \forall t \in \mathbb{R}^+ \setminus \{\{t'\} \cup \{\hat{t}_k\}_{k=0}^K\} \\ \hat{\hat{p}}(t) &= p^* & \forall t \in \{\{t'\} \cup \{\hat{t}_k\}_{k=0}^K\} \end{aligned}$$

Lemma 15. *For every $t' \in (0, hatt_0)$, the social payoff under the strategy $\hat{\hat{p}}(t)$ is strictly greater than the social payoff under the strategy $\hat{p}(t)$.*

Proof. WRITE □

6 Constrained Stopping Times

Suppose the DM may only make decisions at certain pre-set dates (e.g. discrete calendar dates). Let's consider an extreme case and assume the DM may decide at exogenously given date $t_1 > 0$ whether to repeal the project or keep it active until the exogenously given final date $t_2 > t_1$. If the project doesn't succeed by t_2 it is automatically repealed. The DM optimally uses a threshold strategy: she repeals the project if and only if $S(t_1) = 0$ and $N(t_1) \leq \tilde{n}$, or equivalently if and only if $p(t_1) \leq \tilde{p}$. Index the threshold belief \tilde{p}_α and corresponding number of pieces of news \tilde{n}_α by the DM's preference parameter α . We begin by deriving the planner solution, then show that for all $\alpha > 0$, there exists no socially efficient PBE.

6.1 Planner Solution

Lemma 16. *The social planner's threshold \tilde{p}_0 satisfies*

$$\mathbb{E}[W_{t_1}^{t_2} \mid p(t_1) = \tilde{p}_0] = s.$$

Proof. The term on the left is the expectation at t_1 , conditional on the DM's current belief $p(t_1)$ being at the threshold value \tilde{p}_0 , of the social payoff from keeping the project active at t_1 . \square

6.2 Perfect Bayesian equilibrium

The next proposition says that for all $\alpha > 0$, a PBE strategy cannot maximise the social payoff because the DM's reputational concern biases her in favour of experimentation. Her bias increases with the DM's reputational concern, and a DM who cares only about her reputation never stops at t_1 . The intuition is simple. In a PBE, the marginal type of the DM, $p(t_1) = \tilde{p}_\alpha$, pools with lower types, $p(t_1) < \tilde{p}_0$. She stands to strictly improve her reputation by deviating and keeping the project active at t_1 . Thus, she is willing to incur a strict loss in social welfare.

Proposition 17. *Fix $0 < t_1 < t_2$. (i) $\tilde{p}_0 > \tilde{p}_\alpha > 0$; (ii) \tilde{p}_α decreases with α ; (iii) $\tilde{p}_1 = 0$.*

Proof. (i) The threshold policy with threshold belief $\tilde{p}_\alpha \geq p_{t_1}^0$ is a PBE strategy if the DM with preference parameter α and belief $p(t_1) = \tilde{p}_\alpha$ is indifferent between repealing the project at date t_1 , and keeping it active. Her reputation, conditional on stopping at t_1 , is determined by Bayes rule and satisfies $\mu(t_1) = \mathbb{E}[p(t_1) \mid p(t_1) \leq \tilde{p}_\alpha]$. If she continues at date t_1 , her reputation at t_2 is 1 if the project succeeds on the interval $[t_1, t_2]$, and $\mathbb{E}[p(t_2) \mid S(t_2) = 0, p(t_1) > \tilde{p}_0]$ otherwise.

The threshold belief \tilde{p}_α is therefore pinned down by the indifference condition:

$$(27) \quad \alpha \left[\mathbb{E}[\mu(t_2) \mid p(t_1) = \tilde{p}_\alpha] - \mathbb{E}[p(t_1) \mid p(t_1) \leq \tilde{p}_\alpha] \right] + (1 - \alpha) \left[\mathbb{E}[W_{t_1}^{t_2} \mid p(t_1) = \tilde{p}_\alpha] - s \right] = 0,$$

where the first term is the DM's net reputational payoff from keeping the project active at t_1 , and the second term measures the net social payoff.

We now show that for every $\tilde{p}_\alpha \geq p_{t_1}^0$, the first term is strictly positive. If she stops at date t_1 , the DM's reputation is $\mu(t_1) = \mathbb{E}[p(t_1)|p(t_1) \leq \tilde{p}_\alpha] \leq \tilde{p}_\alpha$. For every $p \in [0, 1]$, the DM's expectation at t_1 of her reputation at t_2 is

$$\mathbb{E}[\mu(t_2)|p(t_1) = p] = p \left(1 - e^{-\eta(t_2-t_1)}\right) + \left(1 - p \left(1 - e^{-\eta(t_2-t_1)}\right)\right) \mathbb{E}[p(t_2)|S(t_2) = 0, p(t_1) > \tilde{p}_\alpha].$$

It is strictly increasing in p . Therefore, for every $p > \tilde{p}_\alpha$,

$$(28) \quad \mathbb{E}[\mu(t_2)|p(t_1) = p] > \mathbb{E}[\mu(t_2)|p(t_1) = \tilde{p}_\alpha].$$

Moreover, the DM's reputation at t_2 is bounded from below:

$$(29) \quad \mathbb{E}[p(t_2)|S(t_2) = 0, p(t_1) > \tilde{p}_\alpha] = \mathbb{E}[p(t_2)|p(t_2) \in (p_{t_2}^{\tilde{n}_\alpha}, 1)] > p_{t_2}^{\tilde{n}_\alpha},$$

where \tilde{n}_α satisfies $p_{t_1}^{\tilde{n}_\alpha} = \tilde{p}_\alpha$, so that

$$\mathbb{E}[\mu(t_2)|p(t_1) = \tilde{p}_\alpha] > \tilde{p}_\alpha \left(1 - e^{-\eta(t_2-t_1)}\right) + \left(1 - \tilde{p}_\alpha \left(1 - e^{-\eta(t_2-t_1)}\right)\right) p_{t_2}^{\tilde{n}_\alpha} = \tilde{p}_\alpha,$$

where the last equality follows from Bayesian updating. It follows that the first term in (27) is strictly positive. For the DM to be indifferent, it must therefore be that

$$\mathbb{E}[W_{t_1}^{t_2}|p(t_1) = \tilde{p}_\alpha] < s.$$

By Lemma 16, we therefore have $\tilde{p}_\alpha < \tilde{p}_0$.

When α is sufficiently close to 0, a solution $\tilde{p}_\alpha \geq p_{t_1}^0$ to (27) exists. (add detail) However, as α increases, we have $\tilde{p}_\alpha \rightarrow 0$. In these cases $\tilde{p}_\alpha < p_{t_1}^0$, and no type $p(t_1)$ of the DM repeals the project at t_1 in equilibrium. □

These results may be contrasted with Abreu, Milgrom, and Pearce (1991) (henceforth AMP). AMP analyse a repeated partnership game with imperfect public monitoring, and examine the effects of changing the period length, i.e. reducing the length of time for which players must keep actions fixed. Information (about the action profile) arrives according to two possible Poisson processes, as in the present paper. As periods become shorter, the quality of information deteriorates, and reducing period length may increase inefficiency. In the good news case, AMP find that any level of cooperation becomes unsustainable as the period length shrinks to zero, while in the bad news case, some cooperation may still be sustained.¹⁷

My substantive finding here is very different from AMP. In my model, when the DM can act in continuous time (i.e. the limit of shorter period length), inefficiencies can be mitigated and may well be eliminated. In contrast, with longer period length (as in this section), inefficiencies obtain. The key difference is that in our case, a shorter period length allows complete separation of DM types, thus mitigating her reputational concerns.

¹⁷Holmstrom and Milgrom (1987) examine an agency model and show that linear incentive schemes become uniquely optimal as periods over which the agent must keep actions fixed shrink.

References

- ABREU, D., P. MILGROM, AND D. PEARCE (1991): “Information and timing in repeated partnerships,” *Econometrica: Journal of the Econometric Society*, pp. 1713–1733.
- BENABOU, R., AND G. LAROQUE (1992): “Using privileged information to manipulate markets: Insiders, gurus, and credibility,” *The Quarterly Journal of Economics*, 107(3), 921–958.
- COHEN, A., AND E. SOLAN (2013): “Bandit problems with Lévy processes,” *Mathematics of Operations Research*, 38(1), 92–107.
- HOLMSTRÖM, B. (1999): “Managerial incentive problems: A dynamic perspective,” *The Review of Economic Studies*, 66(1), 169–182.
- HOLMSTROM, B., AND P. MILGROM (1987): “Aggregation and linearity in the provision of intertemporal incentives,” *Econometrica: Journal of the Econometric Society*, pp. 303–328.
- KAMENICA, E., AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101(6), 2590–2615.
- KELLER, G., AND S. RADY (2010): “Strategic experimentation with Poisson bandits,” *Theoretical Economics*, 5(2), 275–311.
- (2015): “Breakdowns,” *Theoretical Economics*, 10(1), 175–202.
- KELLER, G., S. RADY, AND M. CRIPPS (2005): “Strategic experimentation with exponential bandits,” *Econometrica*, pp. 39–68.
- KRISHNA, V., AND J. MORGAN (2004): “The art of conversation: eliciting information from experts through multi-stage communication,” *Journal of Economic Theory*, 117(2), 147–179.
- MAJUMDAR, S., AND S. W. MUKAND (2004): “Policy gambles,” *American Economic Review*, pp. 1207–1222.
- MILGROM, P., AND J. ROBERTS (1986): “Relying on the information of interested parties,” *The RAND Journal of Economics*, pp. 18–32.
- MILGROM, P. R. (1981): “Good news and bad news: Representation theorems and applications,” *The Bell Journal of Economics*, pp. 380–391.
- MORRIS, S. (2001): “Political correctness,” *Journal of Political Economy*, 109(2), 231–265.
- NOLDEKE, G., AND E. VAN DAMME (1990): “Signalling in a dynamic labour market,” *The Review of Economic Studies*, 57(1), 1–23.

- OTTAVIANI, M., AND P. N. SØRENSEN (2006): “Reputational cheap talk,” *The Rand journal of economics*, 37(1), 155–175.
- PESKIR, G., AND A. N. SHIRYAEV (2000): “Sequential testing problems for Poisson processes,” *Annals of Statistics*, pp. 837–859.
- PRENDERGAST, C., AND L. STOLE (1996): “Impetuous youngsters and jaded old-timers: Acquiring a reputation for learning,” *Journal of Political Economy*, pp. 1105–1134.
- PRESMAN, E. L., AND I. N. SONIN (1990): *Sequential control with incomplete information: the Bayesian Approach to multi-armed bandit problems*. Academic Press.
- ROGOFF, K. (1990): “Equilibrium Political Budget Cycles,” *The American Economic Review*, pp. 21–36.
- ROSENBERG, D., A. SALOMON, AND N. VIEILLE (2013): “On games of strategic experimentation,” *Games and Economic Behavior*, 82, 31–51.
- SCHARFSTEIN, D. S., AND J. C. STEIN (1990): “Herd behavior and investment,” *The American Economic Review*, pp. 465–479.
- SHIRYAEV, A. N. (2007): *Optimal stopping rules*, vol. 8. Springer Science & Business Media.
- SWINKELS, J. M. (1999): “Education signalling with preemptive offers,” *The Review of Economic Studies*, 66(4), 949–970.