

Optimal Dynamic Information Provision*

Jérôme Renault,[†] Eilon Solan,[‡] and Nicolas Vieille[§]

August 20, 2018

Abstract

We study a dynamic model of information provision. A state of nature evolves according to a Markov chain. An informed advisor decides how much information to provide to an uninformed decision maker, so as to influence his short-term decisions. We deal with a stylized class of situations, in which the decision maker has a risky action and a safe action, and the payoff to the advisor only depends on the action chosen by the decision maker. The greedy disclosure policy is the policy which, at each round, minimizes the amount of information being disclosed in that round, under the constraint that it maximizes the current payoff of the advisor. We prove that the greedy policy is optimal in many cases – but not always.

Keywords: Dynamic information provision, optimal strategy, greedy algorithm.

*The research of Renault and Vieille was supported by Agence Nationale de la Recherche (grant ANR-10-BLAN 0112). Solan acknowledges the support of the Israel Science Foundation, Grants #212/09 and #323/13. The authors thank Omri Solan for numerics which have led to the counterexample presented in Section 5.

[†]TSE (GREMAQ, Université Toulouse 1), 21 allée de Brienne, 31000 Toulouse, France. E-mail: jerome.renault@tse-fr.eu.

[‡]School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: eilons@post.tau.ac.il.

[§]Departement Economics and Decision Sciences, HEC Paris, 1, rue de la Libération, 78 351 Jouy-en-Josas, France. E-mail: vieille@hec.fr.

1 Introduction

Market conditions evolve over time, and information that is privately available to a market participant is a valuable asset. In this paper we study the optimal provision of information by an informed “expert” with no decision power, to an uninformed agent in a dynamic setup. We develop a stylized model in which an “investor” chooses at each date whether or not to choose a risky action, such as a short-run investment. The payoff from investing depends on some underlying state of nature, which is unknown to the investor. This state accounts for all relevant external factors and evolves exogenously according to a Markov chain.

At each date, the investor may get information through the advisor. How much information is being disclosed is the choice variable of the advisor. To be specific, the advisor publicly chooses an information provision rule, which maps each history into a distribution over signals. The investor observes both the rule chosen by the advisor and the realized signal. We assume that the advisor receives a fixed fee whenever investment takes place, and that the investor invests whenever the expected net payoff given his current posterior belief is nonnegative.

This allows us to recast the problem faced by the advisor as a Markov decision problem (MDP) in which the state space is the compact set of posterior beliefs of the investor, and the action space is the set of information provision rules. In that MDP, the advisor chooses dynamically the provision of information so as to maximize the (expected) discounted frequency of dates in which investment takes place. Advising is thus both honest, in that realized signals cannot be manipulated, and strategic, in that the information content of the signal is strategic.

There are two (mutually exclusive) interpretations that befit this description. In the first one, the advisor does not observe the underlying state, and chooses how much information will be publicly obtained. In other words, he chooses a statistical experiment *à la* Blackwell, whose outcome is public. In the second interpretation, the advisor does observe the successive states of nature but commits *ex ante* to a dynamic information provision policy.

The basic trade-off faced by the advisor is the following. By disclosing information at a given date, the advisor may increase his payoff at that date, but then gives up part of his information advantage for later dates, as soon as

successive states are autocorrelated. Our focus is on structural properties of the model. Characterizing optimal information provision in general is out-of-reach, and we instead focus on the optimality of rules in which the above trade-off is solved in a very simple way. We define the greedy policy as the one that, at any given date, minimizes the amount of information being disclosed, subject to the current payoff of the advisor being maximized. We prove that this policy is optimal in the case of two states of nature. We then exhibit a class of Markov chains, described by a renewal property, for which this policy is optimal for a large range of initial distributions of the state (including most natural ones), and is eventually optimal, for any initial distribution of the state. Our main message is thus that this policy is likely to perform very well in a large number of cases, but not always, as we show by means of a counterexample.

Our modelling of information acquisition/disclosure is a dynamic version of the persuasion mechanisms of Kamenica and Gentskow (2011) who study optimal signals in a broader, yet static, setup. It also parallels the independent paper by Ely (2014). Our paper joins the growing literature on dynamic models in which uncertainty evolves, see, e.g., Mailath and Samuelson (2001), Phelan (2006), Wiseman (2008), or Athey and Bagwell (2008), and Escobar and Toikka (2013) for economic applications. These references focus on game models, whose mathematical analysis is in general quite challenging, see Renault (2006) and Hörner et al. (2010). Although our basic model is a game-theoretic one, its reduced form, and the commitment assumption makes it more comparable to contract theory ones, see e.g. Battaglini (2005), Zhang and Zenios (2008) or Zhang, Nagarajan and Sosis (2008).

2 Model and Main Results

2.1 Model

We consider the following stylized class of two-player games between an “advisor” (Player 1) and an “investor” (Player 2). The advisor observes a stochastic process $(\omega_n)_{n \in \mathbf{N}}$ with values in a finite set of states Ω , and may provide the investor with information regarding the current or past values of the process. In each round, the investor chooses whether to invest or not. The investor’s payoff from investing in round n is $r(\omega_n)$, where $r : \Omega \rightarrow \mathbf{R}$. The advisor receives a

fee whenever investment takes place (this fee is already accounted for in r) and discounts future payoffs according to a discount factor δ .

While the investor knows the law of the sequence $(\omega_n)_{n \in \mathbf{N}}$, he receives no information on the realized states, except through the advisor. It is then natural to assume that he chooses to invest whenever his expected (net) payoff from investing is nonnegative, where the expectation is computed using the information released by the advisor.^{1,2} Thus, the game reduces to a stochastic optimization problem, in which the advisor chooses whether and how to reveal information to the investor, so as to maximize the expected discounted frequency of rounds in which investment takes place.

We assume that the process $(\omega_n)_{n \in \mathbf{N}}$ follows an irreducible Markov chain with transition matrix $M = (\pi(\omega' | \omega))_{\omega, \omega' \in \Omega}$ and invariant measure $m \in \Delta(\Omega)$. The set $\Delta(\Omega)$ is the set of probability distributions over Ω , whose elements are potential “beliefs” of the investor. Throughout, we identify each $\omega \in \Omega$ with a unit basis vector in \mathbf{R}^Ω , and $\Delta(\Omega)$ with the $(|\Omega| - 1)$ -dimensional unit simplex in \mathbf{R}^Ω , endowed with the induced topology.

The game is played as follows. In each round n , the state ω_n is drawn according to $\pi(\cdot | \omega_{n-1})$, the advisor observes ω_n and chooses which message to send to the investor; The investor next chooses whether to invest, and the game moves to the next round.³ When the investor’s belief is $p \in \Delta(\Omega)$, his expected net payoff from investing is given by the scalar product $\langle p, r \rangle = \sum_{\omega \in \Omega} p(\omega)r(\omega)$. Accordingly, the *investment region* is $I := \{p \in \Delta(\Omega), \langle p, r \rangle \geq 0\}$ and the *investment frontier* is $\mathcal{F} := \{p \in \Delta(\Omega), \langle p, r \rangle = 0\}$. We also denote by $J := \Delta(\Omega) \setminus I$ the *noninvestment region*.

Throughout, we will denote by $\Omega^+ := \{\omega \in \Omega, r(\omega) \geq 0\}$ and $\Omega^- := \{\omega \in \Omega, r(\omega) < 0\}$ the states with nonnegative and negative payoff respectively, so that Ω^+ and Ω^- form a partition of Ω .

¹From the literature on dynamic games we know that more sophisticated equilibria may possibly be designed. Besides being natural, our assumption allows to cover the case of short-lived investors or of a large number of investors.

²To simplify the analysis we will assume that the investor also invests on the investment frontier, that is, when his expected profit is 0. Indeed, otherwise, whenever the investor’s belief is on the investment frontier, the advisor would reveal a small amount of additional information, so as to push the investor’s belief to the region where the investor strictly prefers investing to not investing.

³We are not explicit about the message set. It will be convenient to first assume that it is rich enough, e.g., equal to $\Delta(\Omega)$. We will show that w.l.o.g. two messages suffice.

An *information disclosure policy* for the advisor specifies for each round, the probability law of the message being sent in that round, as a function of previous messages and the information privately available to the advisor, that is, past and current states.

We will assume that the advisor has *commitment power*. To be specific, we assume that in any given round, the investor knows which disclosure policy was used in that round, and therefore knows unambiguously how to interpret the message received from the advisor.

An equivalent and alternative interpretation is to assume that the advisor does *not* observe the process $(\omega_n)_{n \in \mathbb{N}}$ and chooses in each round a statistical experiment *à la* Blackwell. Such an experiment yields a random outcome, whose distribution is contingent on the current state. Under this alternative interpretation, the advisor has no private information, but by choosing the experiment, he effectively determines how much information is being publicly obtained, and the investor observes both the experiment choice and the outcome of the experiment.

2.2 A Reformulation

Given an information disclosure policy, the investor uses the successive messages received from the advisor to update his belief on the current state. We find it convenient to distinguish the beliefs p_n and q_n held in round n , respectively *before* and *after* receiving the message of the advisor. Formally, p_n is the conditional law of ω_n given the messages received prior to round n , while q_n is the updated belief, once the round n message has been received, so that the investor invests in round n if and only if $q_n \in I$.

The beliefs q_n and p_{n+1} differ because the state evolves: ω_n and ω_{n+1} need not be equal, and one has $p_{n+1} = \phi(q_n) := q_n M$. The difference between p_n and q_n is the result of the information provided by the advisor.

For a given $p \in \Delta(\Omega)$, denote by $\mathcal{S}(p) \subset \Delta(\Delta(\Omega))$ the set of probability distributions over $\Delta(\Omega)$ with mean p . We denote by $\mu_p \in \mathcal{S}(p)$ the distribution over $\Delta(\Omega)$ that assigns probability 1 to p .

As a consequence of Bayesian updating, the (conditional) law μ of q_n belongs to $\mathcal{S}(p_n)$, for every information disclosure policy. Conversely, a classical result from the literature of repeated games with incomplete information (see Aumann

and Maschler (1995)⁴) states that the converse also holds. That is, given any distribution $p \in \Delta(\Omega)$ and any distribution $\mu \in \mathcal{S}(p)$ of beliefs with mean p , the advisor can correlate the message with the state in such a way that the investor's updated belief is distributed according to μ . Elements of $\mathcal{S}(p)$ will be called *splittings at p* , as is common in the literature.⁵

These observations allow us to reformulate the decision problem faced by the advisor as a dynamic optimization problem Γ . The state space in Γ is the set $\Delta(\Omega)$ of investor's beliefs and the initial state is p_1 , the law of ω_1 . At each state $p \in \Delta(\Omega)$, the set of available actions is the set $\mathcal{S}(p)$, so that the advisor chooses a distribution μ of posterior beliefs that is consistent with p . Given the posterior belief q , the current payoff is 1 if $q \in I$ and 0 if $q \notin I$, and the next state in Γ is $\phi(q)$. Thus, the (expected) stage payoff given μ is $\mu(q \in I)$.

We denote by $V_\delta(p_1)$ the value of Γ as a function of the initial distribution p_1 . The value function V_δ is characterized as the unique solution of the dynamic programming equation⁶

$$V_\delta(p) = \max_{\mu \in \mathcal{S}(p)} \{(1 - \delta)\mu(q \in I) + \delta \mathbf{E}_\mu [V_\delta(\phi(q))]\}, \quad \forall p \in \Delta(\Omega). \quad (1)$$

2.3 The (static) value of information

We first argue that the value function V_δ is concave. This result has a number of implications on the structure of the advisor's optimal strategy. We will point at two such implications which are especially useful in the sequel.

Lemma 1 *The function V_δ is concave on $\Delta(\Omega)$.*

Proof. This is a standard result in the literature on zero-sum games with incomplete information, see, e.g., Sorin (2002, Proposition 2.2). While the setup here is different, the proof follows the same logic, and we only sketch it. We need to prove that $V_\delta(p) \geq a'V_\delta(p') + a''V_\delta(p'')$ whenever $p = a'p' + a''p''$, with $a', a'' \geq 0$ and $a' + a'' = 1$. Starting from p , consider the following strategy σ for the advisor. Pick first the element $\mu \in \mathcal{S}(p)$ that assigns probabilities a'

⁴Aumann and Maschler (1995) contains a proof when the distribution μ has a finite support. Their proof readily extends to the case in which the support of μ is general.

⁵Or simply *splitting*, if p is clear from the context.

⁶We write max on the right-hand side because it is readily checked that V_δ is Lipschitz over $\Delta(\Omega)$, the expression between braces is upper hemi-continuous w.r.t. μ in the weak-* topology on $\Delta(\Delta(\Omega))$, and $\mathcal{S}(p)$ is compact in that topology. Details are standard and omitted.

and a'' to p' and p'' respectively, and next follow an optimal strategy in $\Gamma(p')$ or $\Gamma(p'')$, depending on the outcome of μ . Thus, the advisor's behavior at p is a so-called *compound lottery* obtained as the result of first using μ , and then the first choice of an optimal strategy in either $\Gamma(p')$ or $\Gamma(p'')$.

The strategy σ yields at p a payoff equal to $a'V_\delta(p') + a''V_\delta(p'')$, hence the result. ■

The first consequence of Lemma 1 is that the advisor does not reveal information when the investor's belief is in the investment region.

Corollary 2 *At any $p \in I$, it is optimal for the advisor not to provide information to the investor.*

That is, the distribution $\mu_p \in \mathcal{S}(p)$ that assigns probability one to p achieves the maximum in (1).

The intuition is as follows. When $p \in I$, revealing information cannot increase the current payoff, and therefore, such revelation may only possibly be beneficial in subsequent stages. However, every information that is disclosed today could instead be revealed tomorrow, so that there is no reason to provide information to the investor when $p \in I$. Note that we do not rule out the possibility that there are *additional* optimal strategies that do reveal information in I .

Proof. Fix $\mu \in \mathcal{S}(p)$. By the concavity of the function $q \mapsto V_\delta(\phi(q))$ and Jensen's inequality, one has

$$\mathbf{E}_\mu [V_\delta(\phi(q))] \leq V_\delta(\phi(\mathbf{E}_\mu[q])) = V_\delta(\phi(p)),$$

with equality for $\mu = \mu_p$. Moreover, $\mu(q \in I)$ cannot exceed 1, and is equal to 1 for $\mu = \mu_p$. Therefore the right-hand side in (1) is at most $(1 - \delta) + \delta V_\delta(\phi(p))$, and this upper bound is achieved for $\mu = \mu_p$. ■

A second corollary of Lemma 1 states that in the investment region, the advisor can restrict himself to splitting the investor's belief among at most two beliefs.

Corollary 3 *At any $p \notin I$, there is an optimal choice $\mu \in \mathcal{S}(p)$, which is carried by at most two points.*

That is, at each $p \notin I$ it is either optimal not to disclose information, or to disclose information in a coarse way so that the posterior belief of the investor takes only two well-chosen values in $\Delta(\Omega)$. This result hinges on the fact that (i) the advisor's stage payoff assumes two values only, and (ii) the investment region I is convex.

Proof. Let $p \notin I$ and $\mu \in \mathcal{S}(p)$ be arbitrary. Assume first that $\mu(q \in I) = 0$ and compare the distribution μ to the distribution μ_p in which no information is revealed. The two distributions yield the same current payoff, because $\mu(q \in I) = \mu_p(q \in I) = 0$. However, μ_p yields a (weakly) higher continuation payoff, because by Jensen's inequality

$$\mathbf{E}_{\mu_p} [V_\delta(\phi(q))] = V_\delta(\phi(p)) \geq \mathbf{E}_\mu [V_\delta(\phi(q))].$$

Assume now that $\mu(q \in I) > 0$. Since $p \in J$ and I is convex, one also has $\mu(q \in J) > 0$.

Denote by $q_I := \mathbf{E}_\mu [q \mid q \in I]$ (resp. $q_J := \mathbf{E}_\mu [q \mid q \in J]$) the expected posterior belief conditional on it being in (resp. not in) the investment region. Then

$$p = \mu(q \in I)q_I + \mu(q \in J)q_J.$$

Denote by $\tilde{\mu} \in \mathcal{S}(p)$ the two-point distribution that assigns probabilities $\mu(q \in I)$ and $\mu(q \in J)$ to q_I and q_J respectively. Plainly, $\tilde{\mu}(q \in I) = \mu(q \in I)$ and

$$\mathbf{E}_{\tilde{\mu}} [V_\delta(\phi(q))] = \mu(q \in I)V_\delta(\phi(q_I)) + \mu(q \in J)V_\delta(\phi(q_J)),$$

while

$$\begin{aligned} \mathbf{E}_\mu [V_\delta(\phi(q))] &= \mu(q \in I)\mathbf{E}_\mu [V_\delta(\phi(q)) \mid q \in I] + \mu(q \in J)\mathbf{E}_\mu [V_\delta(\phi(q)) \mid q \in J] \\ &\leq \mu(q \in I)V_\delta(\phi(\mathbf{E}_\mu [q \mid q \in I])) + \mu(q \in J)V_\delta(\phi(\mathbf{E}_\mu [q \mid q \in J])) \\ &\leq \mathbf{E}_{\tilde{\mu}} [V_\delta(\phi(q))]. \end{aligned}$$

To sum up, for any given μ , we have shown that either the no disclosure policy μ_p , or some two-point distribution $\tilde{\mu}$ yields a weakly higher right-hand side in (1) than μ . This proves the result. ■

Note that it may still be optimal not to disclose information at $p \in J$. This is in particular the case whenever $p(\Omega^+) = 0$.

2.4 Main Results

The intuition behind Corollaries 2 and 3 above is clear. When $p \in I$, no information should be revealed, because it cannot help to increase the current payoff, and can only hurt continuation values. When $p \notin I$, there are two conflicting effects at play. For the sake of maximizing payoffs, the advisor should release information. But releasing information may only hurt continuation payoffs, because of concavity.

Corollary 3 shows qualitatively (but not explicitly) how to compromise between the two effects. The main message of our results is that in many cases but not all, the explicit compromise is simple: the advisor should minimize the amount of information released, subject to current payoffs being maximal. We define accordingly the *greedy strategy* σ_* as follows.

Definition 4 *The greedy strategy for the advisor is the strategy σ_* that depends on the investor's current belief p , and plays as follows:*

G1 *At any $p \in I$, the strategy σ_* discloses no information.*

G2 *At any $p \notin I$, the strategy σ_* chooses a solution $\tilde{\mu} \in \mathcal{S}(p)$ to the problem $\max a_I$, under the constraints $p = a_I q_I + a_J q_J$, $q_I \in I$, $a_I + a_J = 1$, $a_I, a_J \geq 0$.*

Thus, the greedy strategy is stationary in the investor's belief, which can be computed by the advisor using the investor's initial belief p_1 and the messages sent to the investor in earlier rounds.

An important point is that σ_* does not depend on the discount factor, nor on the transition matrix. So it can be implemented without knowing δ nor π .

It will be convenient to identify, whenever there is no ambiguity, a decomposition $p = a_I q_I + a_J q_J$ with the splitting μ which selects q_I and q_J with probabilities a_I and a_J respectively. We will call the decomposition in **G2** *the greedy splitting at p* .

As an illustration, consider Figure 1 below, with $\Omega = \{A, B, C\}$.




Figure 1: Three splittings at p .

Three different splittings at p have been drawn: $p = a_I^{(i)} q_I^{(i)} + a_J^{(i)} q_J^{(i)}$, $i \in \{1, 2, 3\}$, with $a_I^{(i)} = \frac{\|p - q_J^{(i)}\|_2}{\|q_I^{(i)} - q_J^{(i)}\|_2}$, so that $a_I^{(1)} > a_I^{(2)} > a_I^{(3)}$. Since $a_I^{(i)}$ is the current payoff under splitting i , the first of the three splittings yields a higher payoff.

For every two points $p_1, p_2 \in \Delta(\Omega)$ denote by (p_1, p_2) the line that passes through p_1 and p_2 , and by $[p_1, p_2]$ the line segment that connects p_1 and p_2 . The noninvestment region J is divided into two triangles by the segment $[B^+, C^+]$, see Figure 2 below.




Figure 2: The decomposition of the noninvestment region.

Because the line (B^+, C^+) has a positive slope, every point p in the lower triangle (B^+, B, C) is split by the greedy strategy σ_* between B^+ and a point on the line segment $[B, C]$, and points p' in the upper triangle are split by σ_* between C and a point on the line segment $[B^+, C^+]$.

Only in the case where the line (B^+, C^+) is parallel to the line (B, C) are there several optimal splittings. This is a nongeneric situation⁷ where $r(B) = r(C)$.

Theorem 5 *If $|\Omega| = 2$, the greedy strategy is optimal, irrespective of the initial distribution p_1 .*

⁷see lemma 11 later.

Although interesting in its own sake, the two-state problem is specific in many respects, and we next investigate the robustness of the conclusion of Theorem 5.

From now on, we restrict ourselves to a class of Markov chains, in which shocks occur at random times, and the state remains unchanged between two consecutive shocks. When a shock occurs, the next state is drawn according to a fixed distribution (and may thus coincide with the previous state). The durations between successive shocks are i.i.d. random variables with a geometric distribution. Note that the invariant distribution m is then equal to the fixed distribution according to which new states are drawn. Equivalently, these are the chains with a transition function given by

$$\pi(\omega | \omega) = (1 - \lambda)m(\omega) + \lambda, \quad (2)$$

$$\pi(\omega' | \omega) = (1 - \lambda)m(\omega') \text{ if } \omega' \neq \omega, \quad (3)$$

for some $\lambda \in [0, 1)$. Note that the drift map $\phi : \Delta(\Omega) \rightarrow \Delta(\Omega)$ that describes the evolution of the investor's belief when no new information is provided is given by

$$\phi(p) - m = \lambda(p - m),$$

so that ϕ is an homothety on the simplex with center m and ratio λ . Notice that we only consider homotheties with non negative ratio.

It turns out that even in this restricted class of chains, and with as few as three states, Theorem 5 does not extend without qualifications.

Proposition 6 *Let $|\Omega| = 3$. The greedy strategy need not be optimal for all initial distributions.*

Indeed, we exhibit in Section 5 a simple counterexample in which, for some initial distributions, it is strictly optimal not to disclose any information in early stages.

Yet, this counterexample hinges on fairly extreme choices of the invariant measure and the initial distribution. In many cases the greedy strategy is a very relevant strategy. We substantiate this claim by means of three results.

First, it may be natural to assume that the initial distribution and the invariant measure coincide.⁸ In that case, the conclusion of Theorem 5 extends to an arbitrary number of states.

⁸Or are very close. This is in particular relevant when the interaction between the advisor and the investor starts at a given date, long after the Markov chain has started evolving.

Theorem 7 *Let the cardinality of Ω be arbitrary and suppose that ϕ is an homothety. If $p_1 = m$, then the greedy strategy is optimal.*

In fact, we will identify a polytope of initial distributions in $\Delta(\Omega)$ of full dimension that contains m in its interior, for which the greedy strategy is optimal. This allows us to prove that, irrespective of the initial distribution, it is eventually optimal to use the greedy strategy.

Theorem 8 *Let the cardinality of Ω and the initial distribution be arbitrary, and suppose that ϕ is an homothety. There is an optimal strategy σ and an a.s. finite stopping time after which σ coincides with the greedy strategy.*

Under the assumption that no two states yield the same payoff, the conclusion of Theorem 8 holds for every optimal strategy σ . That is, the suboptimality identified in Proposition 6 is typically transitory. On almost every history, the advisor will at some point switch to the greedy strategy. Whether or not it is possible to put a deterministic upper bound on this stopping time is unknown to us.

We finally provide an in-depth analysis of the three-state case. As it turns out, σ_* is optimal in most circumstances.

When $|\Omega^-| = 2$, we use the notations of Figure 1: $\Omega^- = \{B, C\}$ with $r(B) \geq r(C)$, and the vertices of \mathcal{F} are denoted by B^+ and C^+ .

Theorem 9 *Assume $|\Omega| = 3$ and suppose that ϕ is an homothety. The strategy σ_* is optimal in the following cases:*

- $|\Omega^-| = 1$;
- $|\Omega^-| = 2$ and m belongs to either I or to the triangle (C^+, B^+, C) .

When instead m belongs to the triangle (B, B^+, C) , the greedy strategy may fail to be optimal only when three conditions are met simultaneously: (i) the advisor is very patient, that is, δ is close to one; (ii) the state is very persistent, that is, λ is close to one; and (iii) the line segment \mathcal{F} is close to parallel to the line (B, C) , that is, $r(B)$ and $r(C)$ do not differ by much. While the first two conditions are natural, we have no intuition to offer for the last condition.

3 Preparations

3.1 The greedy strategy

In this section we provide more details and results on the greedy strategy σ_* . We let \mathcal{E} be the set of extreme points of \mathcal{F} . One can verify that for each $\omega^- \in \Omega^-$ and $\omega^+ \in \Omega^+$, the line segment $[\omega^-, \omega^+]$ contains a unique point in \mathcal{E} . Conversely, any $e \in \mathcal{E}$ lies on a line segment $[\omega^-, \omega^+]$ for some $\omega^- \in \Omega^-$, $\omega^+ \in \Omega^+$.

It is convenient to reformulate the optimization problem **G2** in Definition 4 as a linear program. Given a finite set $A \subset \mathbf{R}^\Omega$ we denote⁹ by $\text{cone}(A)$ the closed convex hull of $A \cup \{0\}$. The optimization program in **G2** is equivalent to the following linear program

$$(LP) : \max \pi_1(\Omega),$$

where the maximum is over pairs $(\pi_1, \pi_2) \in \text{cone}(\mathcal{E}) \times \text{cone}(\Omega^-)$ such that $\pi_1 + \pi_2 = p$.

Lemma 10 *The value of the program (LP) is equal to the value of the following problem (LP').*

$$(LP') : \max \pi(\Omega),$$

where the supremum is over all $\pi \in \text{cone}(\Omega)$ such that $\pi \leq p$ and $\sum_{\omega \in \Omega} \pi(\omega)r(\omega) \geq 0$.

Proof. Recall that in **G2** $p \notin I$. If $(\pi_1, \pi_2) \in \text{cone}(\mathcal{E}) \times \text{cone}(\Omega^-)$ is an optimal solution of (LP) then π_1 is a feasible solution of (LP'), and therefore the value of (LP') is at least the value of (LP).

Fix now an optimal solution π of (LP'). If $\sum_{\omega \in \Omega} \pi(\omega)r(\omega) > 0$, then by increasing the weight of states in Ω^- we can increase $\pi(\Omega)$, which would contradict the fact that π is an optimal solution of (LP'). The weight of some states in Ω^- can be increased because $\pi \leq p$ and $\langle p, r \rangle < 0$. It follows that $\pi \in \text{cone}(\mathcal{E})$. Set $\pi' := p - \pi \in \text{cone}(\Omega)$. It is readily checked that $\pi'(\Omega^+) = 0$, for otherwise the corresponding probability could be transferred to π . Hence $\pi' \in \text{cone}(\Omega^+)$ and (π, π') is a feasible solution of (LP). This implies that the value of (LP) is at least the value of (LP'). ■

⁹Elements of $\text{cone}(\Omega)$ are best seen as “sub”-probability measures.

This reformulation allows for a straightforward description of the greedy strategy at $p \in J$. Intuitively, the weight $p(\omega^-)$ of each $\omega^- \in \Omega^-$ should be “allocated” between \mathcal{F} and $\Delta(\Omega^-)$ so as to maximize the total weight assigned to \mathcal{F} . Since \mathcal{F} is defined by the equality $\langle \pi, r \rangle = 0$, it is optimal to allocate to \mathcal{F} the states ω^- in which the payoff $r(\omega^-)$ is the least negative.

Formally, we order the elements of Ω^- into $\omega_1, \dots, \omega_{|\Omega^-|}$ by decreasing payoff: $0 > r(\omega_1) \geq \dots \geq r(\omega_{|\Omega^-|})$. Next, we define linear maps $L_1, \dots, L_{|\Omega^-|}$ over $\Delta(\Omega)$ by

$$L_k(p) := \sum_{\omega \in \Omega^+} p(\omega)r(\omega) + \sum_{i \leq k} p(\omega_i)r(\omega_i).$$

$L_k(p)$ is a linear combination of the payoff of all states whose payoff is positive or whose index is at most k ; that is, this linear combination assumes only states which are “better” than state k .

Observe that $L_1(\cdot) \geq \dots \geq L_{|\Omega^-|}(\cdot)$. We set $k^* := \inf\{k : L_k(p) \leq 0\}$ to be the minimal index for which the linear combination $L_k(p)$ is nonpositive. With these notations the optimal solution π^* of (LP') is given by

- $\pi^*(\omega^+) = p(\omega^+)$ for $\omega^+ \in \Omega^+$;
- $\pi^*(\omega_i) = p(\omega_i)$ for $i < k^*$;
- $\pi^*(\omega_i) = 0$ for $i > k^*$;
- $\pi^*(\omega_{k^*}) = -\frac{L_{k^*-1}(p)}{r(\omega_{k^*})}$.

The vector π^* is the unique solution of (LP') as soon as no two states in Ω^- yield the same payoff. If different states yield the same negative payoff, the ordering of Ω^- is nonunique. To sum up, we have proven the lemma below.

Lemma 11 *Assume that no two states in Ω^- yield the same payoff: $r(\omega) \neq r(\omega')$ for every $\omega \neq \omega' \in \Omega^-$. Then the greedy splitting $p = a_I q_I + a_J q_J$ is uniquely defined at each $p \in J$. In addition, $q_I \in \mathcal{F}$ and $q_J \in \Delta(\Omega^-)$.*

The distributions q_I and q_J are obtained by renormalizing π^* and $p - \pi^*$, respectively. Note that for $p \in \Delta(\Omega^-)$ one has $a_I = 0$, so that formally speaking, q_I is indeterminate. Yet, the solution to (LP) is unique.

For $k \in \{1, \dots, |\Omega^-|\}$, we let $\bar{\mathcal{O}}(k) := \{p \in J : L_{k-1}(p) \geq 0 \geq L_k(p)\}$ (with $L_0 = 1$). The following figure depicts the sets $\bar{\mathcal{O}}(1)$ and $\bar{\mathcal{O}}(2)$ when $\Omega = \{\omega^+, \omega_1, \omega_2\}$, $r(\omega^+) = 2$, $r(\omega_1) = -1$ and $r(\omega_2) = -4$.




Figure 3: The sets $\bar{\mathcal{O}}(1)$ and $\bar{\mathcal{O}}(2)$.

A useful consequence of the solution of Problem (LP') is that each set $\bar{\mathcal{O}}(k)$ is stable under the greedy splitting.

Lemma 12 *If $p \in \bar{\mathcal{O}}(k)$ and if $p = a_I q_I + a_J q_J$ is the greedy splitting at p , then q_I and q_J are in $\bar{\mathcal{O}}(k)$.*

Proof. Fix $p \in \bar{\mathcal{O}}(k)$. Then the optimal solution π^* to (LP') satisfies

$$\pi^*(\omega^+) = p(\omega^+), \quad \omega^+ \in \Omega^+, \quad (4)$$

$$\pi^*(\omega_i) = p(\omega_i), \quad 1 \leq i \leq k-1, \quad (5)$$

$$0 \leq \pi^*(\omega_k) \leq p(\omega_k). \quad (6)$$

By Lemma 11, q_I is the normalization of π^* . However, $L_{k-1}(\pi^*) > 0$ and $L_k(\pi^*) = 0$, so that $L_{k-1}(q_I) > 0$ and $L_k(q_I) = 0$, and therefore $q_I \in \bar{\mathcal{O}}(k)$.

By Lemma 11, q_J is the normalization of $p - \pi^*$. This implies that $q_J(\omega^+) = 0$ for every $\omega^+ \in \Omega^+$ and $q_J(\omega_i) = 0$ for every $1 \leq i \leq k-1$, so that $L_{k-1}(q_J) = 0$ and $L_k(q_J) \leq 0$, and therefore $q_J \in \bar{\mathcal{O}}(k)$. ■

3.2 Preparatory results

For later use we collect in this section a number of simple, yet general and useful observations. None of the results here uses the specific structure of the Markov chain.

For $p \in \Delta(\Omega)$ we let $\hat{r}(p) := \max_{\mu \in \mathcal{S}(p)} \mu(q \in I)$ be the highest stage payoff of the advisor when the investor's belief is p . Notice that \hat{r} coincides with the value function V_0 with null discount factor, so as an immediate corollary of lemma 1 we obtain:

Lemma 13 *The map \hat{r} is concave.*

Lemma 13 has the following noteworthy implication. Fix $n \geq 1$ and let $\bar{p}_n := \phi^{(n-1)}(p_1)$ be the (unconditional) distribution of the state in round n . Then $\mathbf{E}_\sigma[p_n] = \bar{p}_n$ for every strategy σ of the advisor. In particular, by concavity of the function \hat{r} and Jensen's inequality, the expected payoff of the advisor in round n cannot exceed $\hat{r}(\bar{p}_n)$, so that

$$\gamma_*(p_1) := (1 - \delta) \sum_{n=1}^{\infty} \delta^{n-1} \hat{r}(\bar{p}_n)$$

is an upper bound on the total discounted payoff to the advisor.

Fix $\delta < 1$. We denote by $\gamma(p)$ the payoff induced by the greedy strategy as a function of the initial belief p . We also set

$$d(p) := \gamma(p) - \delta\gamma(\phi(p)).$$

For $p \in J$, the quantity $d(p)$ is the payoff difference when playing greedy, compared to disclosing *no* information in the first round and then switching to the greedy strategy in round 2.

If the greedy strategy is optimal for *all* initial distributions, then γ coincides with V_δ , and therefore γ is concave and $d(\cdot) \geq 0$ over $\Delta(\Omega)$. Somewhat surprisingly, the converse implication also holds.

Lemma 14 *Assume that γ is concave and that $d \geq 0$ over J . Then σ_* is optimal for all initial distributions.*

Proof. It suffices to show that $\gamma(\cdot)$ solves the dynamic programming equation, that is,

$$\gamma(p) = \max_{\mu \in \mathcal{S}(p)} \{(1 - \delta)\mu(q \in I) + \delta \mathbf{E}_\mu[(\gamma \circ \phi)(q)]\}.$$

Denoting by μ_p^* the greedy splitting at p , we have

$$\gamma(p) = (1 - \delta)\mu_p^*(q \in I) + \delta \mathbf{E}_{\mu_p^*}[\gamma \circ \phi(p)]$$

and therefore

$$\gamma(p) \leq \max_{\mu \in \mathcal{S}(p)} \{(1 - \delta)\mu(q \in I) + \delta \mathbf{E}_\mu [(\gamma \circ \phi)(q)]\}.$$

We now show the reverse inequality. Let $\mu \in \mathcal{S}(p)$ be arbitrary. Because $d(\cdot) \geq 0$ on J , one has for each $q \in \Delta(\Omega)$,

$$(1 - \delta)1_{\{q \in I\}} + \delta(\gamma \circ \phi)(q) \leq \gamma(q).$$

Taking expectations w.r.t. μ and using the concavity of γ , one gets

$$(1 - \delta)\mu(q \in I) + \delta \mathbf{E}_\mu [(\gamma \circ \phi)(q)] \leq \mathbf{E}_\mu [\gamma(q)] \leq \gamma(\mathbf{E}_\mu[q]) = \gamma(p).$$

This concludes the proof. ■

4 The 2-state case: proof of Theorem 5

We here assume that $\Omega = \{\omega^-, \omega^+\}$ is a two-point set. W.l.o.g. we assume that $r(\omega^+) > 0 > r(\omega^-)$, and we identify a belief over Ω with the probability assigned to state ω^+ . Here, the investor is willing to invest as soon as the probability assigned to ω^+ is high enough, and the investment region is the interval $I = [p_*, 1]$, where $p_* \in (0, 1)$ solves $p_*r(\omega^+) + (1 - p_*)r(\omega^-) = 0$.

At any $p < p_*$, σ_* chooses the distribution $\mu \in \mathcal{S}(p)$ which assigns probabilities $\frac{p}{p_*}$ and $1 - \frac{p}{p_*}$ to p_* and 0, respectively, and does not disclose information if $p \geq p_*$. In particular,

$$\gamma(p) = \frac{p}{p_*}\gamma(p_*) + \left(1 - \frac{p}{p_*}\right)\gamma(0) \text{ for } p \in [0, p_*], \quad (7)$$

and

$$\gamma(p) = (1 - \delta) + \delta(\gamma \circ \phi)(p) \text{ for } p \in [p_*, 1].$$

Eq. (7) shows that $\gamma(\cdot)$ is affine over $[0, p_*]$ (but need not be affine on $[p_*, 1]$). Note that $\gamma(0) = \delta(\gamma \circ \phi)(0)$.

In this setup,¹⁰ concavity of $\gamma(\cdot)$ alone is equivalent to the optimality of σ_* . Indeed, assume $\gamma(\cdot)$ is concave. Recall that $d(p) = 0$ for $p \in I$. In addition, $\gamma(\cdot)$ is affine on $[0, p_*]$, and since ϕ is affine and γ is concave, the composition $(\gamma \circ \phi)(\cdot)$ is concave on $[0, p_*]$. Thus, $d(\cdot)$ is convex on $[0, p_*]$. Observe now that

¹⁰This observation does not generalize to $|\Omega| \geq 3$.

$d(0) = d(p_*) = 0$, hence $d(\cdot) \geq 0$ on $\Delta(\Omega)$, and the optimality of σ_* then follows from Lemma 14.


It is left to prove that γ is concave. The invariant measure m assigns probability $\frac{\pi(\omega^+ | \omega^-)}{\pi(\omega^+ | \omega^-) + \pi(\omega^- | \omega^+)}$ to ω^+ . With our notations, for $q \in [0, 1](= \Delta(\Omega))$ one has

$$\phi(q) = m + (1 - \pi(\omega^+ | \omega^-) - \pi(\omega^- | \omega^+))(q - m),$$

hence ϕ is a homothety on $[0, 1]$ centered at m with ratio $\lambda := 1 - \pi(\omega^+ | \omega^-) - \pi(\omega^- | \omega^+) \in (-1, 1)$.

It is convenient to organize the proof below according to the relative values of p_* and m , and to the sign of the ratio λ . In the first case we provide a direct argument. In the following cases we prove the concavity of γ .

Case 1: $p^* \geq m$ and $\lambda \geq 0$.



Establishing directly the concavity of $\gamma(\cdot)$ is possible, yet involved, as $\gamma(\cdot)$ fails to be affine on I . We instead argue that $\gamma(p) = \gamma_*(p)$ for each p , where $\gamma_*(p)$ is the upper bound on payoffs identified earlier.

Assume first that $p_1 \in [0, p_*]$. Since the interval $[0, p_*]$ is stable under ϕ under σ_* , one has $q_n \in \{0, p_*\}$ for each $n \geq 1$, and $p_n \in \{\phi(0), \phi(p_*)\}$ for each $n > 1$. In each stage $n \geq 1$, conditional on the previous history, the strategy σ_* maximizes the expected payoff in stage n , so that the expected payoff in stage n is given by $\mathbf{E}_{\sigma_*} [\hat{r}(p_n)]$. Since \hat{r} is affine on $[0, p_*]$, the expected payoff in stage n is also equal to $\hat{r}(\mathbf{E}_{\sigma_*} [p_n]) = \hat{r}(\bar{p}_n)$, so that $\gamma(p_1) = \gamma_*(p_1)$.

Assume now that $p_1 \in I$. Then the sequence $(\bar{p}_n)_{n \geq 1}$ is decreasing (towards m). Let $n_* := \inf\{n \geq 1 : \hat{p}_n < p_*\}$ be the stage in which the unconditional distribution of the state leaves I . Under σ_* , the advisor discloses no information up to stage n_* , so that $r(q_n) = 1 = \hat{r}(\bar{p}_n)$ for all $n < n_*$. That is, σ_* achieves the upper bound on the payoff in *each* stage $n < n_*$, and, by the previous argument, in each stage $n \geq n_*$ as well.

Case 2: $p^* \leq m$ and $\lambda \geq 0$.



Since $m \geq p_*$, one has $\phi([p_*, 1]) \subseteq [p_*, 1]$: the investment region is stable under ϕ . Thus, once in I , σ_* yields a payoff of 1 in each stage: $\gamma(p) = 1$ for $p \geq p_*$. Using (7), one thus has

$$\gamma(p) = \frac{p}{p_*} + \left(1 - \frac{p}{p_*}\right) \gamma(0) \text{ for } p > p_*.$$

Since $\gamma(0) < 1$ it follows that γ is increasing (and affine) on $[0, p_*]$. Hence the concavity of γ on $[0, 1]$.

Case 3: $p^* \geq m$ and $\lambda \leq 0$.



Recall that γ is affine on $[0, p_*]$. From the formula $\phi(p) = m + \lambda(p - m)$, one has $\phi(p) \leq \phi(p_*) \leq p_*$ for all $p \geq p_*$, that is, I is mapped into $[0, p_*]$ under ϕ . Since

$$\gamma(p) = (1 - \delta) + \delta(\gamma \circ \phi)(p) \text{ for } p \in I, \quad (8)$$

this implies that γ is also affine on $[p_*, 1]$. To establish the concavity of γ we need to compare the slopes of γ on I and $J = [0, p^*]$. Differentiating (8) yields $\gamma'(p) = \delta\lambda(\gamma' \circ \phi)(p)$ for $p > p_*$, hence the two slopes are of opposite signs. Note finally that $\gamma(p_*) = (1 - \delta) + \delta(\gamma \circ \phi)(p_*)$, hence $\gamma(p_*) > (\gamma \circ \phi)(p_*)$, so that γ is increasing on $[0, p_*]$ (and then decreasing on $[p_*, 1]$).

Case 4: $p^* \leq m$ and $\lambda \leq 0$.



The dynamics of the belief under σ_* is here slightly more complex. If $\phi(1) \geq p_*$, the investment region I is stable under ϕ , hence $\gamma(p) = 1$ for all $p \in I$, and the concavity of γ follows as in Case 2. If instead $\phi(1) < p_*$, we introduce the cutoff $q_* \in [m, 1]$ defined by $\phi(q_*) = p_*$. Since ϕ is contracting, the length of the interval $[\phi(q_*), \phi(p_*)]$ is smaller than that of $[p_*, q_*]$, which implies that the

interval $[p_*, q_*]$ is stable under ϕ . Therefore, $\gamma(p) = 1$ for all $p \in [p_*, q_*]$. As in Case 2, this implies that γ is increasing (and affine) on $[0, p_*]$.

For $p \geq q_*$, $\gamma(p) = (1 - \delta) + (\gamma \circ \phi)(p)$. Since the interval $[q_*, 1]$ is mapped into $[0, p_*]$ under ϕ , this implies in turn that γ is affine on $[q_*, 1]$, with slope given by $\gamma'(p) = \lambda\delta(\gamma' \circ \phi)(p) < 0$. That is, γ is piecewise affine, increasing on $[0, p_*]$, constant on $[p_*, q_*]$ and decreasing on $[q_*, 1]$.

5 A counterexample: proof of Proposition 6

We here provide an example in which σ_* fails to be optimal for some initial distribution p_1 .

There are three states, $\Omega = \{\omega_1, \omega_2, \omega_3\}$, and the investment region is the triangle with vertices ω_1 , $\varepsilon\omega_1 + (1 - \varepsilon)\omega_2$, and $\frac{1}{2}\omega_1 + \frac{1}{2}\omega_3$ where $\varepsilon > 0$ is sufficiently small (see Figure 1). Assume first that the invariant distribution is $m = \omega_2$, and that $\lambda = \frac{1}{2}$. Let the initial belief be $p_1 = 2\varepsilon\omega_1 + (1 - 2\varepsilon)\omega_3$.




Figure 4: The counterexample.

According to σ_* , at the first stage p_1 is split between ω_3 (with probability $1 - 4\varepsilon$) and $\frac{1}{2}\omega_1 + \frac{1}{2}\omega_3$ (with probability 4ε). Because the line segment $[\omega_2, \omega_3]$ is contained in J and $m = \omega_2$, the payoff to the investor once the belief reaches ω_3 is 0. It follows that the payoff under σ_* is $\gamma(p_1) = 4\varepsilon\gamma(\frac{1}{2}\omega_1 + \frac{1}{2}\omega_3) \leq 4\varepsilon$.

Consider the alternative strategy, in which the advisor discloses no informa-

tion in the first stage, so that

$$p_2 = \phi(p_1) = \frac{1}{2}\omega_2 + \frac{1}{2}p_1 = \varepsilon\omega_1 + \frac{1}{2}\omega_2 + \left(\frac{1}{2} - \varepsilon\right)\omega_3,$$

and then at the second stage splits p_2 between $q_2 = \varepsilon\omega_1 + (1 - \varepsilon)\omega_2$ (with probability $\frac{1}{2(1-\varepsilon)}$) and $q_3 = \varepsilon\omega_1 + (1 - \varepsilon)\omega_3$ (with the complementary probability). The expected payoff in the second stage is therefore $\frac{1}{2(1-\varepsilon)}$.

Hence, the alternative strategy is better than σ_* as soon as $4\varepsilon < \frac{1}{2(1-\varepsilon)} \times \delta(1 - \delta)$. For fixed $\delta \in (0, 1)$, this is the case for small ε .

In this example, the invariant distribution m is on the boundary of $\Delta(\Omega)$. However, for fixed δ , the above argument is robust to a perturbation of the transition probabilities. In particular we obtain a similar result for an invariant distribution m that is in the interior of $\Delta(\Omega)$.

The example shows that the greedy strategy is not always optimal. Then, a natural question is whether there is always an optimal strategy that satisfies the following property: whenever the strategy provides information to the investor, it does so according to the greedy splitting. The answer is negative, and we end this section by showing that the optimal strategy in this example sometimes splits the investor's belief in a way that is not the greedy splitting.

Assume then to the contrary that in this example there is an optimal strategy σ that, at every belief p , either does not provide information or reveals information according to the greedy splitting. Consider the line segment $[\omega^*, \omega^3]$. If there is a belief p on this line segment for which the greedy splitting is optimal, then by Lemma 16 below the greedy splitting is optimal for every belief on this line segment, which contradicts the fact that the greedy splitting is not optimal at $2\varepsilon\omega_1 + (1 - 2\varepsilon)\omega_3$. Thus, σ does not provide information for any p on this line segment. In particular,

$$V_\delta(p) = (1 - \delta)\mathbf{1}_{\{p \in \mathcal{F}\}} + \delta V_\delta(\phi(p)).$$

Whereas the functions V_δ and $V_\delta \circ \phi$ are continuous, the function $\mathbf{1}_{\{p \in \mathcal{F}\}}$ is not continuous on the line segment $[\omega^*, \omega^3]$, a contradiction.

6 Invariant initial distributions: proof of Theorem 7

We will prove a strengthened version of Theorem 7, which will be used in the proof of Theorem 8. We recall from Section 3.1 that $L_k(\cdot)$ is the linear map defined by

$$L_k(p) = \sum_{\omega \in \Omega^+} p(\omega)r(\omega) + \sum_{i \leq k} r(\omega_i)p(\omega_i).$$

Note that, with the notations of Section 3.1, the map $p \mapsto \pi_1^*(p)$ is affine on the set $\bar{\mathcal{O}}(k) = \{L_k(\cdot) \geq 0 \geq L_{k+1}(\cdot)\}$ and, therefore, so is $\hat{r}(\cdot)$.

Theorem 15 *Let k be such that $m \in \bar{\mathcal{O}}(k)$. Then $\gamma(p_1) = \gamma_*(p_1)$ for every initial distribution $p_1 \in \bar{\mathcal{O}}(k)$. In particular, the greedy strategy σ_* is optimal whenever $p_1 \in \bar{\mathcal{O}}(k)$.*

Proof. Let $p \in \bar{\mathcal{O}}(k)$ be arbitrary, and denote by $p = a_I q_I + a_J q_J$ the greedy splitting at p . Lemma 12 implies that both q_I and q_J belong to $\bar{\mathcal{O}}(k)$. Since $m \in \bar{\mathcal{O}}(k)$, the set $\bar{\mathcal{O}}(k)$ is stable under the greedy strategy σ_* . That is, if $p_1 \in \bar{\mathcal{O}}(k)$ then under σ^* we have $p_n \in \bar{\mathcal{O}}(k)$ for every n .

Since \hat{r} is affine on $\bar{\mathcal{O}}(k)$, one has for each stage n ,

$$\mathbf{E}_{\sigma_*} [\hat{r}(p_n)] = \hat{r}(\mathbf{E}_{\sigma_*} [p_n]) = \hat{r}(\bar{p}_n).$$

Hence the result. ■

7 Eventually greedy strategies: proof of Theorem 8

We will assume that $m \in J$, which is the more difficult case. The case where $m \in I$ is dealt with at the end of the proof. We start with an additional, simple, observation on the shape of the value function.

Lemma 16 *Let $p \in J$ be given, and let $p = a_I q_I + a_J q_J$ be an optimal splitting at p . If $a_I, a_J > 0$, then*

1. V_δ is affine on $[q_I, q_J]$;

2. at each $p' \in [q_I, q_J]$, it is optimal to split between q_I and q_J .

We stress that $p = a_I q_I + a_J q_J$ need not be the greedy splitting at p .

Proof. By assumption, $V_\delta(p) = a_I V_\delta(q_I) + a_J V_\delta(q_J)$, hence the first statement follows from the concavity of V_δ on $[q_I, q_J]$. Given a point $p' = a'_I q_I + a'_J q_J \in [q_I, q_J]$, this affine property implies

$$V_\delta(p') = a'_I V_\delta(q_I) + a'_J V_\delta(q_J).$$

On the other hand, splitting p' into q_I and q_J yields $a'_I V_\delta(q_I) + a'_J V_\delta(q_J)$, hence the second statement. ■

In the sequel, we let k be such that $m \in \bar{\mathcal{O}}(k)$. By Theorem 15, $\gamma(p) = \gamma_*(p)$ for every $p \in \bar{\mathcal{O}}(k)$. We denote by $\bar{J} := J \cup \mathcal{F} = \{p \in \Delta(\Omega), \langle p, r \rangle \leq 0\}$ the closure of J .

Lemma 17 *Let $p \in \bar{J} \setminus \bar{\mathcal{O}}(k)$ be given, and let $p = a_I q_I + a_J q_J$ be an optimal splitting at p . Then $[q_I, q_J] \cap \bar{\mathcal{O}}(k) = \emptyset$.*

Proof. We argue by contradiction and assume that there exists $p' \in \bar{\mathcal{O}}(k) \cap [q_I, q_J]$. By Lemma 16, the splitting $p' = a'_I q_I + a'_J q_J$ is optimal at p' . Since $p' \in \bar{\mathcal{O}}(k)$, one has $V_\delta(p') = \gamma_*(p')$. This implies that under the optimal strategy, the expected payoff in each stage is equal to the first best payoff in that stage. In particular, any optimal splitting at p' must be the greedy one. By Lemma 12 this implies that both q_I and q_J belong to $\bar{\mathcal{O}}(k)$, hence by convexity $p \in \bar{\mathcal{O}}(k)$ – a contradiction. ■

We will need to make use of a set P of the same type as $\bar{\mathcal{O}}(k)$, which contains m in its interior, and starting from which σ_* is optimal.

If m belongs to the interior of $\bar{\mathcal{O}}(k)$ for some k , we simply set $P := \bar{\mathcal{O}}(k)$. Otherwise, one has

$$L_{k-1}(m) > 0 = L_k(m) = \dots = L_l(m) > L_{l+1}(m) \text{ for some } k \leq l. \quad (9)$$

We then set $P := \{p \in \bar{J}, L_{k-1}(p) \geq 0 \geq L_{l+1}(p)\} = \bar{\mathcal{O}}(k-1) \cup \dots \cup \bar{\mathcal{O}}(l)$. By construction, m belongs to the interior of P . By (9), one has $m \in \bar{\mathcal{O}}(i)$ for $i = k-1, \dots, l$, hence the set P is stable under the Markov chain. This implies that σ_* is optimal whenever $p_1 \in P$.

Lemma 18 *Assume that all connected components of $\bar{J} \setminus P$ in $\Delta(\Omega)$ are convex. Then the conclusion of Theorem 8 holds.*

Proof. Let \mathcal{C} be an arbitrary connected component of $\bar{J} \setminus P$. Since \mathcal{C} is convex, there is an hyperplane H (in $\Delta(\Omega)$) that (weakly) separates \mathcal{C} from P , and we denote by Q the open half-space of $\Delta(\Omega)$ that contains m .

We will make use of the following observation. Since $\bar{Q} \cap \Delta(\Omega)$ is compact, there is a constant $c > 0$ such that the following holds: for all $\tilde{p} \in \bar{Q} \cap \Delta(\Omega)$ and all $\mu \in \mathcal{S}(\tilde{p})$, one has $\mu(q \in \bar{Q}) \geq c$.

Since $m \in Q$, the distance from m to H is positive. Since $\Delta(\Omega)$ is compact and ϕ is contracting, there exists $\bar{n} \in \mathbf{N}$ such that $\phi^{(\bar{n})}(p) \in Q$ for all $p \in \Delta(\Omega)$.

Fix $p \in \mathcal{C}$ and let τ be any optimal policy when $p_1 = p$. We let $\theta := \inf\{n \geq 1, q_n \in P\}$ be the stage at which the investor's belief reaches P . We prove below that $\theta < +\infty$ with probability 1 under τ . This proves the result, since θ is an upper bound on the actual stage at which the advisor can switch to σ_* .

Since $m \in J$, under τ one has $q_n \in \bar{J}$ with probability 1 for all n . By Lemma 17, one has $q_n \in \mathcal{C}$ on the event $n < \theta$. On the other hand, the (unconditional) law of q_n belongs to $\mathcal{S}(\bar{p}_n)$ for each n : $\mathbf{E}[q_n] = \bar{p}_n$. This implies that $\mathbf{P}_\tau(q_{\bar{n}} \in Q) \geq c$, so that $\mathbf{P}_\tau(\theta \leq \bar{n}) \geq c$.

The same argument, applied more generally, yields $\mathbf{P}_\tau(\theta \leq (j+1)\bar{n} \mid \theta > j\bar{n}) \geq c$ for all $j \in \mathbf{N}$. Therefore, $\mathbf{P}(\theta < +\infty) = 1$, as desired. ■

The complement of P in \bar{J} is the *disjoint* union of $\{p \in J : L_k(p) < 0\}$ and $\{p \in J : L_l(p) > 0\}$. Both sets are convex, hence Theorem 8 follows from Lemma 18.

For completeness, we now provide a proof for the case $m \in I$. In that case, the entire investment region I is stable under σ_* . Hence, it is enough to prove that the stopping time $\theta := \inf\{n \geq 1 : q_n \in I\}$ is a.s. finite, for any initial distribution $p \in J$ and any optimal policy τ . Observe first that the payoff $\gamma(p)$ under σ_* is bounded away from zero and therefore so is $V_\delta(p) \geq \gamma(p)$. For a fixed δ , this implies the existence of a constant $c > 0$ and of a stage $\bar{n} \in \mathbf{N}$, such that $\mathbf{P}_\tau(\theta \leq \bar{n}) \geq c$. This implies the result, as in the first part of the proof.

8 The case of 3 states: proof of Theorem 9

The analysis relies on a detailed study of the belief dynamics under σ_* . We will exhibit a simplicial decomposition of $\Delta(\Omega)$ with respect to which γ is affine. This partition will be used to prove that $\gamma(\cdot)$ is concave and $d(\cdot)$ nonnegative on $\Delta(\Omega)$. We will organize the discussion in two cases, depending on the size of Ω^- .

Case 1: $\Omega^- = \{C\}$.

We prove the optimality of σ_* in two steps. We first argue that γ is concave and d nonnegative on the straight line joining C and m . We next check that both γ and d are constant on each line parallel to \mathcal{F} . These two steps together readily imply that γ is concave and d nonnegative throughout $\Delta(\Omega)$, as desired.

Step 1. Denote by \mathcal{L} the line (C, m) , and by p_* the intersection of \mathcal{L} and \mathcal{F} . The line \mathcal{L} is stable under ϕ , and σ_* splits any $p \in \mathcal{L} \cap J$ between C and p_* . The dynamics of beliefs and of payoffs thus follows the same pattern as in the two-state case. Hence¹¹ it follows from Section 4 that γ is concave and $d(\cdot)$ nonnegative on \mathcal{L} .

¹¹We emphasize however that this is not sufficient to conclude the optimality of σ_* on \mathcal{L} .




Figure 5: The case $|\Omega^-| = 1$.

Step 2. With the notations of Figure 5, σ_* splits any $p \in J$ between C and a point in the investment frontier \mathcal{F} , and $\hat{r}(p) = \hat{r}(p_{\mathcal{L}})$, where $p_{\mathcal{L}}$ is a point for which $(pp_{\mathcal{L}})$ is parallel to \mathcal{F} . Note that any line parallel to \mathcal{F} is mapped by ϕ into some line parallel to \mathcal{F} . This implies that γ and $\gamma \circ \phi$ are constant on each line parallel to \mathcal{F} , and so is $\Delta(\cdot)$.

Denote by J_0 the triangle (C^+, B^+, C) .

Case 2: $\Omega^- = \{B, C\}$ and $m \in I \cup J_0$.

Again, we proceed in several steps. We first prove that γ is concave and d nonnegative on $I \cup J_0$. We next explicit the dynamics of beliefs under σ_* . This in turn leads to the concavity of γ in Step 3. In Step 4, we prove that $d \geq 0$ on $\Delta(\Omega)$.

Step 1. The function γ is concave and $d \geq 0$ on $I \cup J_0$.

The analysis is identical to that in **Case 1**. First, it follows from the two-state case that the conclusion holds on the line (C, m) . Next, as before, both γ and $\gamma \circ \phi$ are constant on each line segment contained in $I \cup J_0$ and parallel to \mathcal{F} .

Step 2. The belief dynamics under σ_* .

We construct recursively a finite sequence O_1, \dots, O_K of points in the line segment $[B, C]$ as follows. Set first $O_1 = C$ and let $k \geq 1$. If ϕ maps B into the triangle (C^+, O_k, O_{k-1}) (or J_0 , if $k = 1$), we set $K = k$. Otherwise, O_{k+1} is the unique point in the line segment $[C, O_k]$ such that $P_{k+1} := \phi(O_{k+1}) \in [C, O_k]$.

Since ϕ is an homothety, all points $(P_k)_{k \leq K}$ lie on some line \mathcal{P} parallel to (B, C) , see Figure 6.




Figure 6: The points $(O_k)_k$ and $(P_k)_k$.

The next claim states that this algorithm ends in a finite number of steps.

Claim 19 $K < +\infty$.

Proof. We introduce the map f from the line segment $[B, C]$ to the line (B, C) as follows. Given $X \in [B, C]$, we let $f(X)$ be the intersection of (B^+, Y) with (B, C) , where Y is the intersection of (X, m) with \mathcal{P} , see Figure 7.




Figure 7: The definition of f .

Since m belongs to $I \cup J_0$ and to the relative interior of $\Delta(\Omega)$, $f(X)$ is well-defined and $f(B)$ lies strictly “to the right” of B .

Observe that (by Thales Theorem), the Euclidian distance $Xf(X)$ is proportional to the distance B^+Z . Hence, as m moves away from B towards C , $Xf(X)$ increases if $m \in I$, and decreases if $m \in J_0$. In the former case, this implies that $O_{k+1}f(O_{k+1}) = O_{k+1}O_k \geq Bf(B)$ for each k . In the latter one, this implies that $O_k O_{k+1}$ increases with k . In both cases, $K < +\infty$. ■

For $k = 1, \dots, K - 1$, denote by J_k the triangle (B^+, O_k, O_{k+1}) (see Figure 6), and observe that $\phi([O_k, O_{k+1}]) = [P_k, P_{k+1}]$. The belief dynamics is similar for any initial belief in J_k . Any $p_1 \in J_k$ is first split between B^+ and some $q_1 \in [O_k, O_{k+1}]$. In the latter case, q_1 is mapped to $p_2 := \phi(q_1) \in [P_k, P_{k+1}]$.

The belief p_2 is then split between B^+ and $q_2 \in [O_{k-1}, O_k]$, etc. The (random) belief p_{k+1} in stage $k+1$ lies in $I \cup J_0$.

Step 3. The function γ is concave on $\Delta(\Omega)$.

We proceed with a series of claims.

Claim 20 *The function γ is affine on J_k , for every $k \leq K$.*

Proof. We argue by induction and start with $k = 0$. We denote by $\gamma(\mathcal{F})$ the constant value of γ on \mathcal{F} . Given $p = xC + (1-x)q \in J_0$ with $q \in [B, C]$, one has $\gamma(p) = x\gamma(C) + (1-x)\gamma(\mathcal{F})$, hence the affine property. For later use, note also that, as p moves towards $[A, C^+]$ on a line parallel to $[B, C]$, the weight x decreases, hence $\gamma(\cdot)$ is decreasing on such a line.

Assume now that γ is affine on J_{k-1} for some $k \geq 1$. For $p \in [O_k, O_{k+1}]$, $\gamma(p) = \delta\gamma \circ \phi(p)$. Since $\phi(p) \in J_{k-1}$, γ is affine on $[O_k, O_{k+1}]$. Next, for $p = x_I B^+ + x_k O_k + x_{k+1} O_{k+1} \in J_k$,

$$\begin{aligned} \gamma(p) &= x_I \gamma(B^+) + (x_k + x_{k+1}) \gamma\left(\frac{x_k O_k + x_{k+1} O_{k+1}}{x_k + x_{k+1}}\right) \\ &= x_I \gamma(B^+) + x_k \gamma(O_k) + x_{k+1} \gamma(O_{k+1}). \end{aligned}$$

That is, γ is affine on J_k . ■

Claim 21 *The function γ is concave on $J_k \cup J_{k+1}$ for $k = 1, \dots, K-2$.*

Proof. We will use the following elementary observation. Let $g_1, g_2 : \mathbf{R}^2 \rightarrow \mathbf{R}$ be affine maps. Let \mathcal{L} be a line in \mathbf{R}^2 , such that $g_1 = g_2$ on \mathcal{L} . Let H_1 and H_2 be the two half-spaces defined by \mathcal{L} , and let h be the map that coincides with g_i on H_i . Assume that for $i = 1, 2$, there is a point A_i in the relative interior of H_i such that h is concave on $[A_1, A_2]$. Then h is concave¹² on \mathbf{R}^2 .

We prove the claim by induction. Pick first $\tilde{p}_0 \in J_0 \cap \mathcal{P}$ and $\tilde{p}_1 \in J_1 \cap \mathcal{P}$, and let p_* be the point of intersection of \mathcal{P} with the line (B^+, C) . Under σ_* , any point $p \in [\tilde{p}_1, p_*]$ is split as $p = (1-x)B^+ + xq_J$, where $q_J \in (B, C)$. Note that x does not depend on p , and

$$\begin{aligned} \gamma(p) &= (1-x)\gamma(B^+) + x\gamma\left(\frac{p - (1-x)B^+}{x}\right) \\ &= (1-x)\gamma(B^+) + x\delta\gamma \circ \phi\left(\frac{p - (1-x)B^+}{x}\right). \end{aligned}$$

¹²If $g_1 = g_2$ everywhere the conclusion holds trivially. Otherwise, g_1 and g_2 coincide only on \mathcal{L} , and then $h = \min\{g_1, g_2\}$.

As p moves from \tilde{p}_1 towards p_* , $\phi\left(\frac{p - (1-x)B^+}{x}\right)$ moves from p_* towards \tilde{p}_0 . Hence, the derivative of γ on $[\tilde{p}_1, p_*]$ is equal to $\delta\lambda$ times the derivative of γ on $[p_*, \tilde{p}_0]$.¹³ Since the latter derivative is negative, and $\delta\lambda < 1$, $\gamma(\cdot)$ is concave on $[\tilde{p}_1, \tilde{p}_0]$. The concavity of γ on $J_1 \cup J_0$ then follows from the preliminary observation.

Assume now that γ is concave on $J_k \cup J_{k-1}$ for some $k \geq 1$. For $p \in [O_{k+1}, O_{k-1}]$, we have $\gamma(p) = \delta\gamma(\phi(p))$. Since $\phi(p) \in [P_{k+1}, P_{k-1}] \subset J_k \cup J_{k-1}$, the function γ is concave on $[O_{k+1}, O_k]$ hence by the preliminary observation it is also concave on $J_{k+1} \cup J_k$. ■

Claim 22 *The function γ is concave on J .*

Proof. Let \tilde{p}_1 and \tilde{p}_2 be given in the relative interior of J_{k_1} and J_{k_2} respectively, with $k_1 \leq k_2$. Since the intersection of the line segment $[\tilde{p}_1, \tilde{p}_2]$ with each of the sets $J_{k_1}, J_{k_1+1}, \dots, J_{k_2}$ is a line segment with a nonempty interior, the concavity of the function γ on each $J_k \cup J_{k+1}$ implies its concavity on $[\tilde{p}_1, \tilde{p}_2]$.

The concavity of the function γ on J follows by continuity. ■

Claim 23 *The function γ is concave on $\Delta(\Omega)$.*

Proof. As above, it suffices to prove that γ is concave on the relative interior $\overset{\circ}{\Delta}(\Omega)$ of $\Delta(\Omega)$. Pick $\tilde{p}_1, \tilde{p}_2 \in \overset{\circ}{\Delta}(\Omega)$, with $\tilde{p}_1 \in I$ and $\tilde{p}_2 \in J_k$ for some $k \geq 1$.¹⁴ Since $[\tilde{p}_1, \tilde{p}_2] \subset \overset{\circ}{\Delta}(\Omega)$, there is a line segment $[p_*, p_{**}] \subseteq [\tilde{p}_1, \tilde{p}_2]$ with $p_*, p_{**} \in J_0$ and $p_* \neq p_{**}$. By Step 1 the function γ is concave on $[\tilde{p}_1, p_{**}]$ and by Claim 22 it is concave on $[p_*, \tilde{p}_2]$. Therefore it is concave on $[\tilde{p}_1, \tilde{p}_2]$. ■

Step 4. $d \geq 0$ on $\Delta(\Omega)$.

We start with the intuitive observation that the payoff under σ_* is higher when starting from \mathcal{F} than from J .

Claim 24 $\gamma(p) \leq \gamma(\mathcal{F})$ for all $p \in J$.

Proof. This is trivial if $m \in I$, since $\gamma(B^+)$ is then equal to 1. Assume then that $m \in J_0$.

¹³We are here identifying any point $p = y\tilde{p}_0 + (1-y)\tilde{p}_1$ of $[\tilde{p}_1, \tilde{p}_0]$ with the real number y , and we view γ as defined over $[0, 1]$.

¹⁴For other cases, the concavity of γ on $[\tilde{p}_1, \tilde{p}_2]$ follows from either Step 1 or Claim 22.

We prove inductively that $\gamma(p) \leq \gamma(\mathcal{F})$ for all $p \in J_k$. Note first that $\gamma(C) = \delta\gamma(\phi(C))$, so that $\gamma(C) \leq \gamma(\phi(C))$. Since $m, C \in J_0$, we have $\phi(C) \in J_0$, hence $\gamma(\phi(C))$ is a convex combination of $\gamma(C)$ and $\gamma(\mathcal{F})$. This implies that $\gamma(C) \leq \gamma(\mathcal{F})$. Note next that, for $p \in J_0$, the quantity $\gamma(p)$ is a convex combination of $\gamma(C)$ and $\gamma(\mathcal{F})$, hence $\gamma(p) \leq \gamma(\mathcal{F})$.

Assume that the conclusion holds on J_{k-1} for some $k \geq 1$. For $p \in [O_{k+1}, O_k]$, since $\phi(p) \in J_{k-1}$, we have $\gamma(p) = \delta\gamma(\phi(p)) \leq \gamma(\mathcal{F})$. Observe finally that for some $p \in J_k$, the quantity $\gamma(p)$ is a convex combination of $\gamma(\mathcal{F})$ and of $\gamma(q)$ for some $q \in [O_{k+1}, O_k]$, hence $\gamma(p) \leq \gamma(\mathcal{F})$ and the conclusion holds on J_k as well. ■

We conclude with the tricky part of the proof.

Claim 25 *For $k \geq 1$, we have $d \geq 0$ on some neighborhood of O_{k+1} in J_k .*

Proof. Given $\varepsilon > 0$, let $p_\varepsilon := \varepsilon B^+ + (1 - \varepsilon)O_{k+1} \in J_k$. Fix $\varepsilon > 0$ small enough so that $\phi(p_\varepsilon) \in J_{k-1}$. Observe that both γ and $\gamma \circ \phi$ are affine on the triangle $(p_\varepsilon, O_{k+1}, O_k)$, hence d is affine on this triangle as well. Since $d = 0$ on $[O_{k+1}, O_k]$ it thus suffices to prove that $d(p_\varepsilon) \geq 0$.

We denote by $\gamma_k : \Delta(\Omega) \rightarrow \mathbf{R}$ the affine map which coincides with γ on J_k . Set $q_\varepsilon := \varepsilon B^+ + (1 - \varepsilon)P_{k+1}$ and observe that

$$d(p_\varepsilon) = \gamma(p_\varepsilon) - \delta\gamma(\phi(p_\varepsilon)) = \gamma(p_\varepsilon) - \delta\gamma(q_\varepsilon) + \delta(\gamma(q_\varepsilon) - \gamma(\phi(p_\varepsilon))). \quad (10)$$

Since $\gamma(p_\varepsilon) = \varepsilon\gamma(B^+) + (1 - \varepsilon)\delta\gamma(P_{k+1})$ and $\gamma(q_\varepsilon) = \varepsilon\gamma(B^+) + (1 - \varepsilon)\gamma(P_{k+1})$, one has

$$\gamma(p_\varepsilon) - \delta\gamma(q_\varepsilon) = \varepsilon\gamma(B^+)(1 - \delta). \quad (11)$$

On the other hand, since q_ε and $\phi(p_\varepsilon)$ belong to J_{k-1} , one has

$$\gamma(q_\varepsilon) - \gamma(\phi(p_\varepsilon)) = \gamma_k(q_\varepsilon) - \gamma_k(\phi(p_\varepsilon)) = \gamma_k(q_\varepsilon - \phi(p_\varepsilon)) = \varepsilon\gamma_k(B^+ - \phi(B^+)). \quad (12)$$

Substituting (11) and (12) into (10) one gets

$$d(p_\varepsilon) = \varepsilon(\gamma(B^+)(1 - \delta) + \delta\gamma_k(B^+ - \phi(B^+))). \quad (13)$$

Now rewrite $B^+ - \phi(B^+)$ as

$$\begin{aligned} B^+ - \phi(B^+) &= B^+ - O_k + O_k - P_k + P_k - \phi(B^+) \\ &= O_k - P_k + (1 - \lambda)(B^+ - O_k) \end{aligned}$$

(recall that $P_k = \phi(O_k)$).

Since all three points O_k, P_k and B^+ belong to J_{k-1} , one has

$$\begin{aligned}\gamma_k(B^+ - \phi(B^+)) &= \lambda\gamma_k(O_k) - \gamma_k(P_k) + (1 - \lambda)\gamma_k(B^+) \\ &= \lambda\gamma(O_k) - \gamma(P_k) + (1 - \lambda)\gamma(B^+) \\ &= (1 - \lambda)\gamma(B^+) - (1 - \lambda\delta)\gamma(P_k).\end{aligned}$$

Plugging into (13), one finally gets

$$d(p_\varepsilon) = \varepsilon(1 - \lambda\delta) (\gamma(B^+) - \delta\gamma(P_k)),$$

which is nonnegative by Claim 1. ■

We now conclude the proof of Step 4. Let $p \in J_k$ be given. Since γ is affine on J_k and concave on $\Delta(\Omega)$, the function d is convex on J_k . Since $d(O_{k+1}) = 0$ and $d \geq 0$ in a neighborhood of O_{k+1} (in J_k), d is nonnegative on the entire line segment $[O_{k+1}, p]$.

References

- [1] Athey, S. and K. Bagwell (2008), Collusion with Persistent Cost Shocks, *Econometrica*, **76**(3), 493–540.
- [2] Aumann, R.J. and M. Maschler (2005), Repeated Games with Incomplete Information, *MIT Press*.
- [3] Battaglini, M. (2005), Long-term contracting with Markovian consumers, *American Economic Review*, **95**, 637–658.
- [4] Ely, J. (2014), Beeps, *mimeo*.
- [5] Escobar, J.F. and J. Toikka (2013), Efficiency in Games with Markovian Private Information, *Econometrica*, **81**, 1887–1934.
- [6] Hörner, J., D. Rosenberg, E. Solan and N. Vieille (2010), On a Markov Game with One-Sided Incomplete Information,” *Operations Research*, **58**, 1107–1115.
- [7] Kamenica, E and M. Gentskow (2011), Bayesian Persuasion *American Economic Review*, **101**, 2590–2615.

- [8] Mailath, G. and L. Samuelson (2001), Who Wants a Good Reputation?, *Review of Economic Studies*, **68**, 415–441.
- [9] Phelan, C. (2006), Public Trust and Government Betrayal, *Journal of Economic Theory*, **127**(1), 27–43.
- [10] Renault, J. (2006), The Value of Markov Chain Games with Lack of Information on One Side, *Mathematics of Operations Research*, **31**, 490–512.
- [11] Sorin, S. (2002). A first course on zero-sum repeated games, Springer.
- [12] Wiseman, T. (2008), Reputation and Impermanent Types, *Games and Economic Behavior*, **62**, 190–210.
- [13] Zhang, H., M. Nagarajan and G. Sodic (2010), Dynamic Supplier Contracts Under Asymmetric Inventory Information, *Operations Research*, **58**(5), 1380–1397.
- [14] Zhang, H. and S. Zenios (2008), A Dynamic Principal-Agent Model with Hidden Information: Sequential Optimality Through Truthful State Revelation, *Operations Research*, **58**, 214–228.