Optimal Dynamic Information Provision*

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

February 2016

Abstract

We study a dynamic model of information provision. A state of nature evolves according to a Markov chain. An advisor with commitment power 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, commitment.

JEL Classification: C73, C72

*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 that 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

In this paper, we study a dynamic game of information provision. In each stage, a short-lived "investor" chooses whether or not to choose a risky action, such as a short-run investment. The payoff from the risky action is contingent on some non-persistent state of nature, which is unknown to the investor.

The investor may get information on the current state through an "advisor". The advisor receives a fixed fee whenever investment takes place.

How much information to disclose is the choice of the advisor. We assume that the advisor has commitment power. To be specific, the advisor *ex ante* commits to an information provision rule, which maps histories into distributions over signals.

The problem faced by the advisor reduces to a Markov decision problem (MDP) in which the state space is the set of posterior beliefs of the investors, and the set space is the set of information provision rules. In that MDP, the advisor chooses the provision of information so as to maximize the (expected) discounted frequency of stages 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 of our game. In the first one, the advisor does observe the successive states of nature but does commit to a dynamic information *disclosure* policy, prior to the beginning of the game. An alternative interpretation is to assume instead that the advisor does *not* observe the states of nature. Choosing in a given stage a distribution of signals, contingent on the state, then amounts to choosing a statistical experiment *à la* Blackwell, whose outcome is public. Under this latter interpretation, the advisor chooses which information will be *publicly* obtained.

As in any dynamic optimization problem, the advisor faces a trade-off between maximizing the probability of investment in the current stage, and preserving his continuation value. By disclosing information at a given date, the advisor may increase his current payoff, but then gives up part of his information advantage in later dates (except if successive states are independent).

Writing the dynamic programming equation which characterizes the value of the optimization problem is rather straightforward. *Solving* it, that is, identifying an

optimal policy of the advisor, is not. Our main contribution is to prove that often, but not always, the trade-off faced by the advisor is solved in a very simple, myopic 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 when there are only two states of nature. We then focus on a class of Markov chains, described by a renewal property. Within this class, we prove that the greedy policy is optimal, whenever the distribution of the initial state is close enough to the invariant distribution of the Markov chain of states. We also prove that the greedy policy is also eventually optimal, irrespective of the distribution of the initial state. That is, any optimal policy eventually coincides with the greedy policy. However, as we show by means of a counterexample, the greedy policy may fail to be optimal for some distributions of the initial state.

Our paper belongs to the literature on information design, and on so-called persuasion mechanisms. In the static case, Kamenica and Gentzkow (2011) characterized optimal persuasion mechanisms, using insights from Aumann and Maschler (1995). A dynamic model of persuasion was studied in Ely, Frankel and Kamenica (2014), where the underlying state of nature is fixed. The sequence of posterior beliefs of the uninformed agent then follows a martingale, and the goal of the informed one is to maximize suspense over time, defined as the L^1-variation of the sequence of posteriors. Independently of us, Ely (2015) deals with the same model as ours, fully solves one specific example, and analyzes a number of variants, including the case of two uninformed agents, and of more-than-two actions. Relatedly, Honryo (2011) analyzes a model of dynamic persuasion, where the sender has several pieces of information over the (fixed) quality of a proposal. These pieces, both negative and positive, may be sent sequentially. In Hörner and Skrzypacz's (2012), a seller may transmit gradually (verifiable) information, in exchange for payments.

Che and Hörner (2015) and Halac, Kartic and Lui (2014) address the issue of the optimal design of information disclosure, in the context of social learning/experimentation and of contests respectively.

Our paper also relates to the literature on dynamic Bayesian games, see e.g. Mailath and Samuelson (2001), or Athey and Bagwell (2008), Escobar and Toikka (2013), Renault (2006), Hörner et al. (2010), Hörner, Takahashi and Vieille (2015),

or Renault, Solan and Vieille (2013), who study dynamic sender-receiver games.

2 Model and Preliminary Observations

2.1 Model

We consider a stylized class of two-player dynamic games between an *advisor* and short-lived *investors*. Over time, the advisor privately observes the successive realizations $(\omega_n)_{n\in\mathbf{N}}$ of a changing state of nature, with values in a finite set Ω. In each round n, the advisor chooses which information to disclose to the current investor, who decides whether to invest or not. The game then moves to the next round.[1]

The advisor receives a fee whenever investment takes place and discounts future payoffs according to the discount factor $\delta < 1$. The investor's utility from investing is $r(\omega_n)$, with $r : \Omega \to \mathbf{R}$. When the investor's belief is $p \in \Delta(\Omega)$, the expected payoff from investing (net of the fee to the investor) is therefore given by the scalar product $\langle p, r\rangle = \sum_{\omega\in\Omega} p(\omega)r(\omega)$.

We assume that the sequence $(\omega_n)_{n\in\mathbf{N}}$ is an irreducible Markov chain with transitions $(\pi(\omega' \mid \omega))_{\omega,\omega'\in\Omega}$. While investors know the distribution of the sequence $(\omega_n)_{n\in\mathbf{N}}$, the only additional information received along the play comes from the advisor. That is, previous investment outcomes are not observed by the investor.[2] The investor chooses to invest if and only if the expected (net) payoff from investing is nonnegative.[3]

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.

The game reduces to a stochastic optimization problem, in which the advisor manipulates the posterior beliefs of the investors, so as to maximize the expected discounted frequency of rounds in which investment takes place. An *information disclosure policy* for the advisor specifies for each round, the probability distribution of the message to be sent in that round, as a function of previous messages and of

[1] We are not explicit about the message set. It will be convenient to first assume that it is rich enough, e.g., equal to Ω. We will show that w.l.o.g. two messages suffice for our results.

[2] This assumption is discussed at the end of this section.

[3] In particular, we assume that the investor invests whenever indifferent.

the information privately available to the advisor, that is, past and current states.

Again, an equivalent and alternative interpretation would be to assume that the advisor does *not* observe the sequence $(\omega_n)_{n\in\mathbf{N}}$ and chooses instead in each round a statistical experiment *à la* Blackwell. Such an experiment yields a random outcome, whose distribution is contingent on the current state. Both the experiment choice and the outcome of the experiment are observed by the investor. Under this alternative interpretation, the advisor has no private information, but chooses how much information is being publicly obtained.

As was stressed in the introduction, we assume that the advisor has *commitment power*: the information disclosure policy of the advisor is chosen *ex ante*, and is known to the investors, who therefore know unambiguously how to interpret the successive messages received from the advisor.

The investor uses his knowledge of the distribution of (ω_n) and the successive signals to update his posterior belief. We distinguish the (prior) belief $p_n \in \Delta(\Omega)$ held in round n *before* receiving the message of the advisor, from the (posterior) belief q_n held *after* receiving the message. That is, p_n is the conditional distribution 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. In particular, the investor invests in round n if and only if $q_n \in I$.

The belief p_{n+1} may differ from q_n because states are not fully persistent: for each $\omega' \in \Omega$ one has $p_{n+1}(\omega') = \sum_{\omega\in\Omega} q_n(\omega)\pi(\omega' \mid \omega)$, hence $p_{n+1} = \phi(q_n) := q_n M$, where M is the transition matrix. Note that $\phi(\cdot)$ is a linear map; it describes how beliefs evolve, if the investor discloses no information.

The belief q_n may differ from p_n because of information disclosure. For a given $p \in \Delta(\Omega)$, we denote by $\mathcal{S}(p) \subset \Delta(\Delta(\Omega))$ the set of probability distributions over $\Delta(\Omega)$ with mean p. As a consequence of Bayesian updating, the (conditional) distribution 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 (known as the splitting lemma, see Aumann and Maschler (1995)[4]) states that the converse also holds. That is, given any $p \in \Delta(\Omega)$ 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

[4]Aumann 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.

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 write the decision problem faced by the advisor as a dynamic optimization problem Γ. The state space in Γ is the set $\Delta(\Omega)$ of investors' beliefs and the initial state is the distribution p_1 of ω_1. At each state $p \in \Delta(\Omega)$, the set of available actions is the set $\mathcal{S}(p)$ of splittings, 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 at state p given $\mu \in \mathcal{S}(p)$ is $\mu(q \in I)$.

Throughout the paper, we assume that previous investment outcomes are not observed by investors. Note that this is consistent with each investor observing the outcome of his own investment. For the sake of discussion, assume that investors would get some, possibly noisy, public feedback on earlier outcomes. In such a case, (part of) the private information of the advisor is thereby publicly disclosed. It is natural to expect that this should make the advisor more willing to disclose information than in our benchmark case where no such feedback is available. Whereas we have no direct and formal argument to show that this is indeed the case, it is not difficult, for the results in this paper, to provide a direct (and simpler) proof that the greedy policy remains optimal if the investor learns his payoff whenever he invests. We thus believe it is more striking to prove the optimality of the greedy policy in the setup where we would least expect it to be.

2.2 A two-state, two-stage example

To provide some intuition for our results, we here sketch the analysis of a simple two-stage, two-state version of our model. There are two states, G and B. Investing in state B yields no payoff, whereas investing in state G yields a positive payoff of $\bar{r}$. The investment fee is $c > 0$, hence an investor is willing to invest if and only if the probability p assigned to G satisfies $p \geq p_* := \frac{c}{\bar{r}}$. Below, we normalize c to 1.

There are two stages, hence two consecutive, non-exclusive, investment opportunities. The types of the two consecutive states are correlated, and coincide with probability $\lambda > \frac{1}{2}$. Hence, as a function of the belief q_1 of the investor on the first-

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

stage state, the probability assigned to the second-stage state being G is $p_2 = \phi(q_1) := q_1\lambda + (1-\lambda)(1-q_1) = (2\lambda-1)q_1 + 1 - \lambda$.[6] Here, q_1 stands for the *posterior* belief in stage 1, and p_2 for the *prior* belief prior in stage 2.

We solve this example backwards. Fix the prior belief p_2 of the investor in stage 2. Plainly, the payoff $r(q_2)$ of the advisor in stage 2, as a function of the posterior belief q_2, is equal to 1 if $q_2 \geq p_*$, and to 0 of $p_2 < p_*$, see Figure 1.

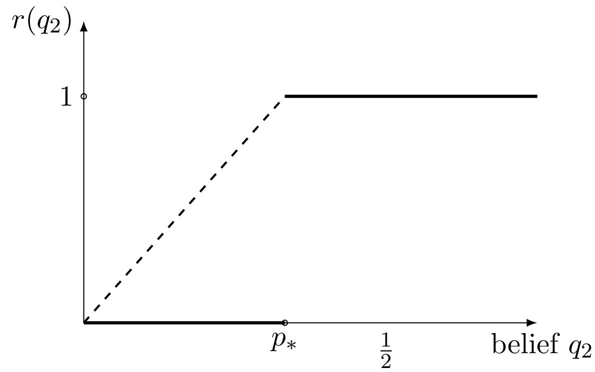

Figure 1: The payoff of the advisor at the second stage.

The goal of the advisor is to pick the distribution μ over $[0,1]$ with mean p_2 so as to maximize $\mathbf{E}_{q_2\sim\mu}[r(q_2)]$, where in the expectation the random variable q_2 has distribution μ. By Aumann and Maschler (1995),[7] it is well-known that the optimal expected payoff of the advisor is equal to $(\text{cav } r)(p_2)$, where cav r is the least concave function above r. It is here equal to $\hat{r}(p_2) = \min\{1, \frac{p_2}{p_*}\}$ (see the dashed line in Figure 1). The optimal policy of the advisor is to disclose nothing if $p_2 \geq p_*$, and to "split" p_2 between 0 and p_* otherwise.

Consider now stage 1, and assume $p_* < \frac{1}{2}$. As a function of the posterior belief q_1 in stage 1, the overall payoff of the investor is

$$U(q_1) = 1_{q_1 \geq p_*} + \hat{r}(\phi(q_1)),$$

which is equal to 2 if $q_1 \geq p_*$, to 1 if $q_1 \in [p_\lambda, p_*)$ and to $\frac{\phi(q_1)}{p_*}$ if $q_1 \in [0, p_\lambda)$, where p_λ is uniquely defined by $\phi(p_\lambda) = p_*$ (see Figure 2).

[6] Note that $|p_2 - \frac{1}{2}| = \lambda|q_1 - \frac{1}{2}|$. Beliefs move towards $\frac{1}{2}$, the invariant probability of state G.

[7] Aumann and Maschler dealt with games with incomplete information. Kamenica and Gentskow (2011) rephrased this in the context of persuasion games.

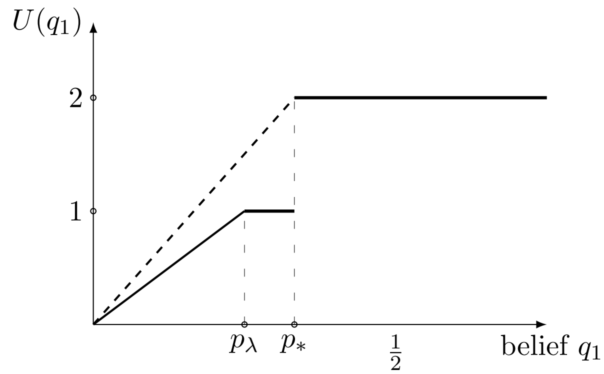

Figure 2: The overall payoff of the investor.

Hence, the objective of the advisor in stage 1 is to find a distribution μ over $[0,1]$ with mean p_1, which maximizes $\mathbf{E}_{q_1 \sim \mu}[U(q_1)]$. As is clear from Figure 2, this goal is uniquely reached with the same policy as in stage 2: disclose nothing if $p_1 \geq p_*$, and split p_1 between 0 and p_* if $p_1 < p_*$.

Thus, the optimal behavior in stage 1 is purely myopic, in that it ignores the impact of stage 1 disclosure on stage 2 payoffs. In particular, it is easily checked that the expected payoff in stage 2 is strictly lower that it would be, should the advisor choose not to disclose any information in stage 1.

One can verify that this conclusion is not sensitive to the choice of parameters. It also holds when stage 2 payoffs are discounted, because it makes the advisor all the more willing to ignore stage 2. By contrast, the extension to an infinite horizon is unclear, as in this case the weight of future stages increases, so that the negative future implications of disclosing much information early become more severe.

The main question we are addressing is the extent to which the above observation on the optimality of myopic play is valid.

2.3 The (static) value of information

We denote by $V_\delta(p_1)$ the value of the dynamic optimization problem 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 equation[8]

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

We first argue that the value of the information held by the advisor is positive. This property is expressed by the concavity of the value function V_δ.

Lemma 1 *The value 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 $a' + a'' = 1$. To guarantee $a'V_\delta(p') + a''V_\delta(p'')$ at p, the advisor sends signals to the investor in such a way that the investor's belief becomes p' with probability a' and p'' with probability a''; the advisor then follows an optimal strategy for the resulting belief.[9] ■

This result has implications on the structure of the advisor's optimal strategy. Two of these are discussed here. Proofs are in Appendix A. The first one is that the advisor will not disclose information when the investor's current 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 splitting $\mu_p \in \mathcal{S}(p)$ which assigns probability 1 to p achieves the maximum in (1). The intuition for this result is as follows. When $p \in I$, the disclosure of information cannot possibly increase the current payoff and would therefore make

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

[9]That is, the advisor picks the element $\mu \in \mathcal{S}(p)$ that assigns probabilities a' 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 compound lottery obtained as the result of first using μ, and then the first choice of an optimal strategy in either $\Gamma(p')$ or $\Gamma(p'')$.

sense only as a means to increase continuation payoffs. Yet, this disclosure may as well be postponed to the next round and combined with any information that is revealed in that round, so that there is no urgency to disclose information when $p \in I$.

Note that Corollary 2 does not rule out the existence of additional optimal policies that disclose information in I.

A second consequence of Lemma 1 is that the advisor can limit himself to splittings that involve only two posterior 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 $\Delta(\Omega)$, which belong to I and J respectively. 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.[10]

We now explain the intuition of this result. For the sake of discussion, consider a splitting $\mu \in \mathcal{S}(p)$, which assigns positive probabilities to posterior beliefs in I and in J, and let q_I (resp. q_J) be the expected posterior belief, conditional on being in I (resp. in J). The posterior beliefs q_I and q_J are thus obtained by pooling together some of the posterior beliefs. The outcome of the splitting μ may be viewed as the result of a two-step splitting: The advisor first splits p into q_I and q_J, and next possibly further splits q_I and q_J. This second step does not affect the current payoff of the investor and, thanks to the concavity of V_δ, may be dispensed with.

Note that Corollary 3 does not imply that it is optimal to disclose information whenever $p \in J$.

3 The greedy policy

The tradeoff faced by the advisor is simple. When $p \notin I$, disclosing information is needed in order to get a positive payoff in the current round, but this is potentially

[10]More generally, if the stage payoff assumes k values $x_1, \cdots, x_k$, and the level sets $\{p \colon V_\delta(p) \geq x_l\}$ are convex for every $l = 1, 2, \ldots, k$, then the optimal choice is carried by at most k points.

harmful in terms of continuation payoffs, because of the concavity of the value function. Corollaries 2 and 3 above show qualitatively – but not explicitly– what the nature of an optimal compromise is. Our main conclusion is that the optimal compromise is often to minimize the amount of information released, subject to current payoffs being maximized.

In this section, we formally define this information disclosure policy, coined the *greedy policy*. We then show how to effectively compute the corresponding splitting at any $p \in \Delta(\Omega)$, and discuss the evolution of beliefs.

3.1 Greedy policy: the definition

Whenever there is no ambiguity, it will be convenient to identify a convex decomposition $p = a_1 q_1 + a_2 q_2$ (where $q_1, q_2 \in \Delta(\Omega)$, and $a_1, a_2 \geq 0$, $a_1 + a_2 = 1$) with the splitting μ which selects q_1 and q_2 with probabilities a_1 and a_2 respectively.

Definition 4 *The* greedy policy *is the Markov stationary policy* σ_* *in* Γ *defined as follows:*

G1 *For* $p \in I$, $\sigma_*(p) = \mu_p$*:* σ_* *discloses no information.*

G2 *For* $p \notin I$, $\sigma_*(p) \in \mathcal{S}(p)$ *is a solution to the problem* $\max a_I$*, under the constraints* $p = a_I q_I + a_J q_J$, $q_I \in I$, $a_I + a_J = 1$, $a_I, a_J \geq 0$.

We will refer to the optimal decomposition $p = a_I q_I + a_J q_J$ in **G2** as *the greedy splitting at p*.

The value $\widehat{r}(p)$ of the optimization problem in **G2** is equal to the highest expected stage payoff of the advisor when the investor's belief is p. Notice that $\widehat{r}$ coincides with the value function V_0 of a fully myopic advisor. Hence, it follows from Lemma 1 that $\widehat{r}$ is concave over $\Delta(\Omega)$.

Note also that the greedy policy σ_* is independent of the discount factor δ, and of the transitions π; transitions affect the information disclosure policy only through the calculation of the posterior belief.

As an illustration, let Ω be the three-point set $\{A, B, C\}$, and consider Figure 3 below.

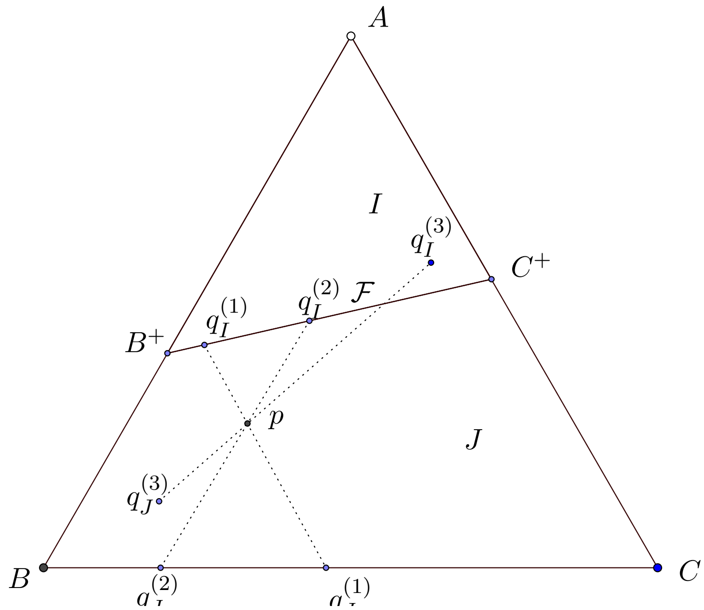

Figure 3: Three splittings at p.

In this example, $r(A)$ is positive, and both $r(B)$ and $r(C)$ are negative. The investment frontier is the segment $[B^+, C^+]$. The fact that B^+ is closer to B than C^+ is to C shows that $r(B) > r(C)$.

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 $a_I^{(i)} = \dfrac{\|p - q_J^{(i)}\|_2}{\|q_I^{(i)} - q_J^{(i)}\|_2}$, 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. It is plain here that σ_* splits p between B^+ and a point on the lower edge $[B, C]$.

The noninvestment region J is divided into two triangles by the segment $[B^+, C]$, see Figure 4. 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 (B^+, C^+, C) are split by σ_* between C and a point on the line segment $[B^+, C^+]$.

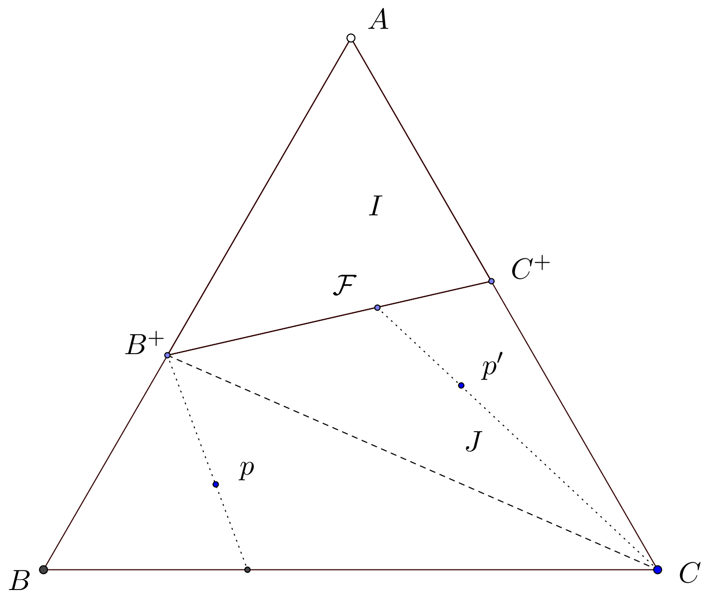

Figure 4: The decomposition of the noninvestment region.

The greedy splitting is uniquely defined unless when the line (B^+, C^+) is parallel to the line (B, C). This is the (nongeneric) case[11] where $r(B) = r(C)$.

3.2 Overview of the main results

Fix a belief p in the interior of J and a line ℓ that passes through p and crosses the investment frontier $\mathcal{F}$. Denote by p_1 the intersection point of ℓ with $\mathcal{F}$, and by p_2 the second intersection point of ℓ and the boundary of J. In particular, $p_1 \in I$ and $p_2 \in J$. Since p lies on ℓ between p_1 and p_2, there are positive real numbers a_1, a_2 that sum up to 1 such that $p = a_1 p_1 + a_2 p_2$. As explained above, this implies that when the investor's belief is p the adviser can reveal information in such a way that the investor's belief is p_1 with probability a_1 and p_2 with probability a_2.

Since $p_1 \in I$, one optimal strategy at p_1 is not to reveal any information, and have the investor invest. Since p_2 is on the boundary of J and not in I, one optimal strategy at p_2 is not to reveal information, and have the investor not invest.

[11]See Lemma 6 below.

Suppose that the value function V_δ is linear on the line segment $\ell(p_1, p_2)$ that connects p_1 and p_2 (which is the intersection of ℓ with $J \cup \mathcal{F}$), and in this case say that V_δ is *quasi-linear* at p. Then, at p, any strategy of the investor that reveals some information and continues by following an optimal strategy at the resulting belief is optimal, as long as the belief after the information revelation remains in $\ell(p_1, p_2)$ with probability 1. The greedy strategy is one such strategy. In particular, the greedy strategy is optimal whenever the value function is linear on J. It turns out that this is the case whenever there are two states, so that when $|\Omega| = 2$ the greedy policy is optimal (Theorem 9 below).

In general the value function is not linear on J. However, in one important special case the value function is piecewise linear, and we can derive positive conclusions. The evolution ϕ is an *homothety* if the state remains constant until a random shock occurs, in which case the next state is chosen according to a fixed probability distribution m. It turns out that when ϕ is an homothety, the value function is piecewise linear, and moreover for $p = m$, the points p_1 and p_2 lie in the set $\mathcal{O}$, which is the largest set that contains m and on which V_δ is quasi-linear. It turns out that in this case the greedy policy is optimal as well, provided the initial belief is in $\mathcal{O}$ (Theorem 11 below).

When ϕ is an homothety, the invariant distribution is m, and one can show that the belief of the investor eventually enters $\mathcal{O}$, that is, it reaches a belief at which the greedy strategy is optimal. We thus deduce that there is an optimal strategy that eventually coincides with the greedy strategy (Theorem 12 below).

3.3 Greedy policy: the computation

We here discuss how to compute the greedy splitting at an arbitrary $p \notin J$. We 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. We also let $\mathcal{E}$ be the set of extreme points of the investment frontier $\mathcal{F}$. Given a finite set $Y \subset \mathbf{R}^\Omega$ we denote[12] by cone(Y) the closed convex hull of $Y \cup \{0\}$.

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 $\omega^- \in \Omega^-$, $\omega^+ \in \Omega^+$.

[12] Elements of cone(Ω) are best seen as "sub"-probability measures.

Intuitively, the greedy splitting at p writes p as a convex combination of two points, one in $\mathcal{F}$ and the other in $\Delta(\Omega^-)$, in a way so as to maximize the weight assigned to the point in $\mathcal{F}$, see **G2** in Definition 4. This optimization problem can be recast as the following program

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

where the maximum is over pairs $(\pi_1, \pi_2) \in \text{cone}(\mathcal{E}) \times \text{cone}(\Omega^-)$ such that $\pi_1 + \pi_2 = p$. The program (LP) is in turn equivalent to a simpler linear program (see Appendix B).

Lemma 5 *The value of the program (LP) is equal to the value of the following problem (LP').*

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

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

According to this lemma, finding the greedy splitting is equivalent to solving a (continuous) knapsack problem. Formally, order the elements of Ω^- by decreasing payoff: $0 > r(\omega_1) \geq \cdots \geq r(\omega_{|\Omega^-|})$. Define

$$k_* := \min\{k \geq 1 : \sum_{\omega \in \Omega^+} p(\omega) r(\omega) + \sum_{i \leq k} p(\omega_i) r(\omega_i) \leq 0\}.$$

The optimal solution π^* of (LP') assigns weight $\pi^*(\omega) = p(\omega)$ for any state $\omega \in \Omega^+ \cup \{\omega_1, \ldots, \omega_{k_*-1}\}$, and it assigns weight $\pi^*(\omega_i) = 0$ whenever $i > k_*$. The weight $\pi^*(\omega_{k_*})$ assigned to ω_{k_*} is then uniquely determined by the condition $\langle \pi^*, r \rangle = 0$.

The vector π^* is the unique solution of (LP') as long as no two states in Ω^- yield the same payoff. [13]. To sum up, we have proven the lemma below.

Lemma 6 *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.

[13] Otherwise, the ordering of Ω^- is not uniquely defined

For $k \leq |\Omega^-|$, we denote by $\mathcal{O}(k) \subset J$ the set of beliefs p such that the index k_* is equal to k, and by $\bar{\mathcal{O}}(k)$ the closure of $\mathcal{O}(k)$. The set $\bar{\mathcal{O}}(k)$ is a polytope for each k. In Figure 4, the sets $\bar{\mathcal{O}}(1)$ and $\bar{\mathcal{O}}(2)$ coincide with the triangles (B, B^+, C) and (B^+, C^+, C) respectively.

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

Lemma 7 *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 $\bar{\mathcal{O}}(k)$.*

3.4 Greedy policy: the beliefs dynamics

We now comment on the beliefs dynamic, using the simple 3-state example of Figure 4. To fix ideas, we assume that $\phi : \Delta(\Omega) \to \Delta(\Omega)$ is an homothety with ratio between zero and one, and that the invariant measure m belongs to the triangle $J_0 := (B^+, C^+, C)$, see Figure **n** below.

If the current belief p_n of the investor lies in the triangle J_0, the greedy policy splits p_n between C and some point on the edge $[B^+, C^+]$. Then, irrespective of the outcome q_n of this splitting, and given the homothetic assumption, the next belief p_{n+1} still lies in J_0 – so that J_0 is stable under the splitting policy. Combined with the property that the current expected payoff of the advisor is affine on J_0 – a consequence of Thales theorem – this will imply that the overall payoff of the advisor is affine on J_0.

Whenever the current belief p_n lies instead in (B, B^+, C) the greedy policy splits p_n between B^+ and some point on the edge $[B, C]$. We next introduce the point $O_2 \in [B, C]$ defined by the condition that $P_2 := \phi(O_2)$ lies on the edge $[B^+, C]$, and denote the triangle (B^+, O_2, C) by J_1. If $p_n \in J_1$, then q_n is either B^+ or some point between O_2 and C. In either case, the next belief $p_{n+1} = \phi(q_n)$ lies in J_1.

Plainly, this generalizes into the pattern which is suggested by Figure **n**. If e.g. $p_n \in J_4$, then the greedy policy splits p_n between B^+ and some point between O_5 and O_4. In the former case, the belief p_{n+1} enters J_0. In the latter case, the belief p_{n+1} will belong to the triangle J_3, and will be split between B^+ and some point between O_4 and O_3, etc.

It is noteworthy that the sequence (p_n) of beliefs will eventually enter the triangle J_0, and stay there. The actual number of rounds until p_n first enters J_0 is random and depends on the messages sent by the advisor. The maximal number of rounds also depends on the location of the initial distribution p_1. It is equal to five in the case of Figure 5, but can be made arbitrarily large by playing with the degree of state persistency: when the states are quite persistent, the ratio of the homothety is close to one, and the points O_2, O_3, etc. are all very close.

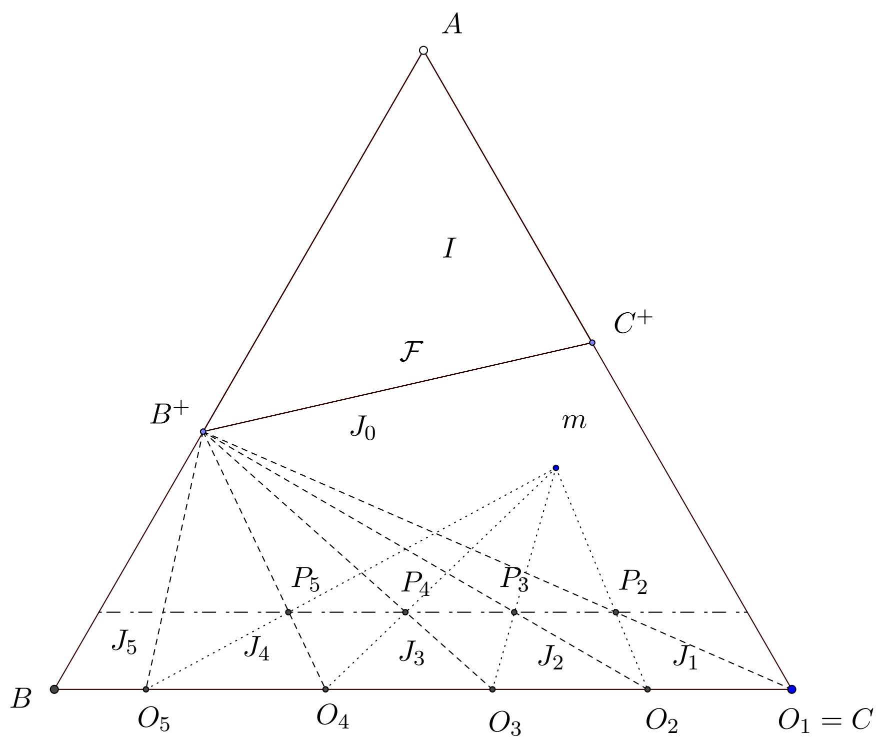

F

4 Results

When successive states are *independent*, the greedy policy is clearly optimal. When successive states are *fully persistent*, the greedy policy is also optimal.[14] We focus on intermediate cases.

[14]This follows for instance from Theorem 11 below.

It turns out that a simple observation provides a necessary and sufficient condition for the optimality of σ_*. For $\delta < 1$, we denote by $\gamma(p)$ the payoff induced by σ_* as a function of the initial belief p. For $p \in J$, the overall payoff of the advisor if he discloses no information in the first round and then switches to σ_* in round 2 is equal to $\delta\gamma(\phi(p))$. If σ_* is optimal for *all* initial distributions, then γ coincides with V_δ, and is therefore concave. In addition, the optimality of σ_* and the dynamic programming principle imply that $\gamma(p) \geq \delta\gamma(\phi(p))$ for all $p \in J$. Somewhat surprisingly, the converse implication also holds. For the proof, see Appendix C.

Lemma 8 *The greedy policy σ_* is optimal for all initial distributions if and only if γ is concave over $\Delta(\Omega)$ and $\gamma(p) \geq \delta\gamma(\phi(p))$ for all $p \in J$.*

4.1 Main results

Theorem 9 *If $|\Omega| = 2$, the greedy strategy is optimal, for all initial distributions p_1.*

The proof of Theorem 9, which relies on Lemma 8, appears in Appendix D. We here provide a sketch. Assume $|\Omega^+| = |\Omega^-| = 1$ (otherwise the result is trivial), and identify a belief with the probability $p \in [0, 1]$ assigned to the good state. Then the investment region is equal to the subinterval $[p_*, 1]$ for some cutoff value p_*. Because the greedy policy always splits $p \notin I$ into the beliefs 0 and p_*, the payoff function $\gamma(\cdot)$ is affine on $[0, p_*]$. Since ϕ is linear and contracting, the function γ is piecewise affine on the entire interval $[0, 1]$, and consequently γ is concave. Proving that the difference $\gamma(p) - \delta\gamma(\phi(p))$ is always nonnegative on $[0, p_*]$ requires some work. As a function of p, this is the difference between an affine map and a concave function, hence the difference is convex, continuous, and piecewise affine on $[0, p_*]$, equals zero at $p = 0$ and equals $1 - \delta$ for $p = p_*$. Consequently, it is enough to prove that the slope of the difference is positive at $p = 0$. The computation of the slope depends on the relative location of m and p_*, and on the transition matrix.

Although interesting in its own sake, the two-state case is quite specific, and we next investigate whether conclusion of Theorem 9 extends to an arbitrary finite set Ω of states. We will restrict ourselves to a class of Markov chains in which the state may change only when an exogenous shock occurs, at a random time. The durations between successive shocks are i.i.d. random variables with a geometric distribution.

When a shock occurs, the next state is drawn according to a fixed distribution, and may thus coincide with the previous state. Equivalently, these are the chains where in each round, a shock occurs with a fixed probability λ, in which case a new state is drawn according to a fixed distribution m; if no shock occurs, the state is kept fixed. Accordingly, the transition function is given by

$$\pi(\omega \mid \omega) = (1-\lambda)m(\omega) + \lambda, \tag{2}$$

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

for some $\lambda \in [0, 1)$. Note that the invariant distribution is equal to m and the map $\phi : \Delta(\Omega) \to \Delta(\Omega)$ that describes the evolution of the investors' 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 nonnegative ratio λ. It turns out that even in this restricted class of chains, and with as few as three states, Theorem 9 does not extend without qualifications.

Proposition 10 *Let $|\Omega| = 3$. The greedy strategy need not be optimal for all initial distributions.*

Indeed, we exhibit in Section 5 an example in which, for some initial distributions, it is strictly optimal not to disclose any information in early stages. This negative result hinges on fairly extreme choices of the invariant measure and of the initial distribution.

We provide below a few positive results. We first prove that when the initial distribution is (close to) m, the greedy policy is optimal.

Theorem 11 *Let the cardinality of Ω be arbitrary and suppose that ϕ is an homothety. If $p_1 = m$, then the greedy policy is optimal.*

In contrast with the other results, the proof does not use Lemma 8. Rather, we show that when the initial belief is m,[15] the greedy policy achieves the first-best payoff at every period, and is therefore optimal. Let $n \geq 1$, and recall that $\widehat{r}(p_n)$ is an upper

[15] Or in a polytope $\bar{\mathcal{O}}(k)$ that contains m.

bound on the conditional expected payoff of the advisor in round n, given the belief p_n of the investor. Define $\bar{p}_n := \phi^{(n-1)}(p_1)$ to be the *unconditional* distribution of the state in round n. Then $\mathbf{E}_\sigma[p_n] = \bar{p}_n$ for every policy σ of the advisor. In particular, by concavity of $\widehat{r}$ and by Jensen's inequality, the expected payoff of the advisor in round n cannot exceed $\widehat{r}(\bar{p}_n)$, so that

$$\gamma_*(p_1) := (1-\delta)\sum_{n=1}^{\infty}\delta^{n-1}\widehat{r}(\bar{p}_n) \tag{4}$$

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

We now provide a quick sketch. If $m \in I$, and if the initial distribution p_1 is also in I, then no information is ever disclosed, and all successive beliefs are in I. The overall payoff $\gamma(p_1)$ is thus equal to one, and σ_* is trivially optimal.

Assume now that $m \notin I$. Recall from Section 3.3 that the greedy policy induces a decomposition of the noninvestment region J into finitely many polytopes $\bar{\mathcal{O}}(k)$, on which the nature of the greedy splitting is qualitatively constant. Any point in $\bar{\mathcal{O}}(k)$ is split under σ_* into boundary points of $\bar{\mathcal{O}}(k)$. Let $\bar{\mathcal{O}}$ be one[16] of these polytopes containing m. Because the polytope $\bar{\mathcal{O}}$ is stable under both ϕ and under the greedy splitting, all successive beliefs belong to $\bar{\mathcal{O}}$, as long as $p_1 \in \bar{\mathcal{O}}$ and moreover the payoff γ is linear on $\bar{\mathcal{O}}$. This property will imply that the payoff $\gamma(p)$ induced by σ_* is equal to the upper bound $\gamma_*(p)$, for all $p \in \bar{\mathcal{O}}$. This in turn implies that σ_* is optimal for all initial distributions in the polytope $\bar{\mathcal{O}}$.

Building on Theorem 11, we prove that, irrespective of the initial distribution, it is eventually optimal to use the greedy policy.

Theorem 12 *Let the cardinality of Ω and the initial distribution be arbitrary, and suppose that ϕ is an homothety. There is an optimal policy σ and an a.s. finite stopping time after which σ coincides with the greedy policy.*

The driving force behind Theorem 12 is that when ϕ is an homothety, with probability 1 the sequence of beliefs $(p_n)_{n\in\mathbf{N}}$ eventually enters some polytope $\mathcal{O}(k)$ that contains m, and the result follows from Theorem 11. Under the assumption that no two states yield the same payoff, the conclusion of Theorem 12 holds for *every*

[16] $\bar{\mathcal{O}}$ is uniquely defined unless when m lies on the common boundary between two polytopes.

optimal policy σ. That is, the suboptimality identified in Proposition 10 is typically transitory. On almost every history, the advisor will at some point switch to the greedy policy. Whether or not it is possible to put a *deterministic* upper bound on this stopping time is unknown to us.

We provide a sketch of the proof for the case where m belongs to the *interior* of the polytope $\bar{\mathcal{O}}$. Let σ be some optimal policy. Let the initial distribution p_1 belong[17] to the complement $J \setminus \bar{\mathcal{O}}$ of $\bar{\mathcal{O}}$ in J, and let $\mathcal{C}$ be the (closure of the) connected component of $J \setminus \bar{\mathcal{O}}$ which contains p_1. A crucial observation is that, for each $p \in \mathcal{C}$, the splitting $\sigma(p)$ at p under the optimal policy σ necessarily assigns[18] probability 1 to $\mathcal{C} \cup \bar{\mathcal{O}}$. Since m is in the interior of $\bar{\mathcal{O}}$, the sequence $\bar{p}_n = \phi^{(n)}(p_1)$ of unconditional beliefs eventually enters $\bar{\mathcal{O}}$. Combining these two facts implies that the probability that the actual belief p_n reaches $\bar{\mathcal{O}}$ in finite time is bounded away from zero. Repeating this argument shows that the sequence (p_n) of beliefs reaches $\bar{\mathcal{O}}$ in finite time, with probability 1.

4.2 The three-state case

We fconclude with more detailed results when $|\Omega| = 3$. When $|\Omega^-| = 2$, we use the notations of Figure 3: $\Omega^- = \{B, C\}$ with $r(B) \geq r(C)$, and the vertices of $\mathcal{F}$ are denoted by B^+ and C^+.

Theorem 13 *Assume $|\Omega| = 3$ and suppose that ϕ is an homothety. The policy σ_* is optimal in the following cases:*

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

The conditions in Theorem 13 are by no means necessary. More generally, one can prove that for a given function $r : \Omega \to \mathbf{R}$, the policy σ_* is optimal whenever the advisor is not very patient (δ small enough), or when states are not very persistent (λ small enough). The exact joint conditions on δ and λ are omitted.

[17] The case where $p_1 \in \bar{\mathcal{O}}$ is dealt with in Theorem 11. If now $p_1 \in I$, then we know it is optimal not to disclose information until p_n enters J.

[18] The proof of this observation is not difficult, but a bit technical.

5 A counterexample: proof of Proposition 10

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 6). Assume that the invariant distribution is $m = \omega_2$, and that $\lambda = \frac{1}{2}$. Let the initial belief be $p_1 = 2\varepsilon\omega_1 + (1-2\varepsilon)\omega_3$.

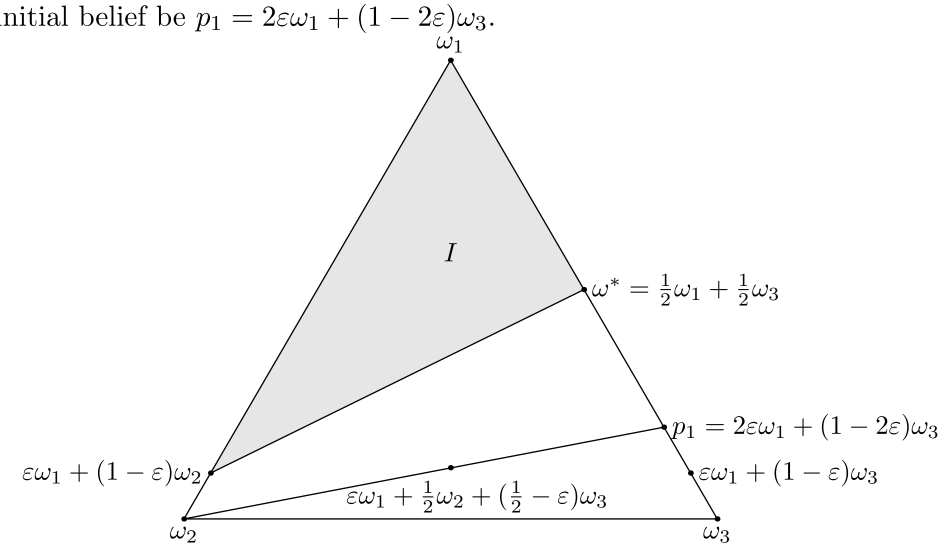

Figure 6: The counterexample.

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 information in the first stage, so that

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

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 $q_2 = \varepsilon\omega_1 + (1-\varepsilon)\omega_3$ (with the complementary probability). The expected payoff in the second stage is therefore $\frac{1}{2(1-\varepsilon)}$. 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. A similar result therefore holds for some invariant distribution m in the interior of $\Delta(\Omega)$.

The example shows that the greedy strategy is not always optimal. 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.

Assume to the contrary that there is an optimal strategy σ that, at every belief p, either does not provide information or reveals information according to the greedy splitting. By Lemma 15 below, either the greedy splitting is optimal throughout the entire line segment $[\omega^*, \omega^3]$, or at no belief in $[\omega^*, \omega^3]$. Given that the greedy splitting is not optimal at $2\varepsilon\omega_1 + (1-2\varepsilon)\omega_3$, the former holds and $\sigma(p) = \mu_p$ for each $p \in [\omega^*, \omega^3]$. This implies that

$$V_\delta(p) = (1-\delta)\mathbf{1}_{\{p\in\mathcal{F}\}} + \delta V_\delta(\phi(p)),\ p \in [\omega^*, \omega^3].$$

Whereas the functions V_δ and $V_\delta \circ \phi$ are easily seen to be continuous, the function $\mathbf{1}_{\{p\in\mathcal{F}\}}$ is not continuous on $[\omega^*, \omega^3]$. This is the desired contradiction.

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] Che, Y.K. and J. Hörner (2015) Optimal Design for Social Learning, Cowles Foundation Discussion Paper #2000.

[4] Ely, J. (2015), Beeps, *mimeo*.

[5] Ely, J., A. Frankel, and K. Kamenica (2015), Suspense and Surprise, *Journal of Political Economics*, **123**, 215–260.

[6] Escobar, J.F. and J. Toikka (2013), Efficiency in Games with Markovian Private Information, *Econometrica*, **81**, 1887–1934.

[7] Forges, F. and F. Koessler (2008), Long Persuasion Games, *Journal of Economic Theory*, **143**, 1–35.

[8] Halac, M., N. Kartik, and Q. Liu (2014), Contests for Experimentation, Preprint.

[9] Honryo, T. (2011), Dynamic Persuasion. Preprint.

[10] Hörner, J. and A. Skrzypacz (2012), Selling Information. Preprint.

[11] Hörner, J., D. Rosenberg, E. Solan and N. Vieille (2010), On a Markov Game with One-Sided Incomplete Information, *Operations Research*, **58**, 1107–1115.

[12] Hörner, J., S. Takahashi and N. Vieille (2015) Truthful Equilibria in Dynamic Bayesian Games, *Econometrica*, **83**, 1795–1848.

[13] Kamenica, E. and M. Gentzkow (2011), Bayesian Persuasion, *American Economic Review*, **101**, 2590–2615.

[14] Mailath, G. and L. Samuelson (2001), Who Wants a Good Reputation?, *Review of Economic Studies,* **68**, 415–441.

[15] Renault, J. (2006), The Value of Markov Chain Games with Lack of Information on One Side, *Mathematics of Operations Research*, **31**, 490–512.

[16] Renault, J. , E. Solan and N. Vieille (2013), *Journal of Economic Theory*,**148**, 502–534

[17] Sorin, S. (2002). A first course on zero-sum repeated games, Springer.

A Appendix: Proofs for Section 2

A.1 Proof of Corollary 2

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

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

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 Eq. (1) is at most $(1-\delta) + \delta V_\delta(\phi(p))$, and this upper bound is achieved for $\mu = \mu_p$.

A.2 Proof of Corollary 3

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, by Jensen's inequality

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

hence μ_p yields a (weakly) higher continuation payoff. 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 $q_I := \mathbf{E}_\mu\left[q \mid q \in I\right]$ (resp. $q_J := \mathbf{E}_\mu\left[q \mid q \in J\right]$) 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

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

while

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

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 Eq. (1) than μ. This proves the result.

B Appendix: Proofs for Section 3

B.1 Proof of Lemma 5

Recall that $p \notin I$ in **G2**. If $(\pi_1, \pi_2) \in \mathrm{cone}(\mathcal{E}) \times \mathrm{cone}(\Omega^-)$ 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 $\pi \in \mathrm{cone}(\mathcal{E})$. Set $\pi' := p - \pi \in \mathrm{cone}(\Omega)$. It is readily checked that $\pi'(\Omega^+) = 0$, for otherwise the corresponding probability could be transferred to π. Hence $\pi' \in \mathrm{cone}(\Omega^+)$ and (π, π') is a feasible solution of (LP). This implies that the value of (LP) is at least the value of (LP').

B.2 Proof of Lemma 7

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

$$\pi^*(\omega^+) = p(\omega^+), \quad \omega^+ \in \Omega^+, \tag{5}$$

$$\pi^*(\omega_i) = p(\omega_i), \quad 1 \leq i \leq k-1, \tag{6}$$

$$0 \leq \pi^*(\omega_k) \leq p(\omega_k). \tag{7}$$

By Lemma 6, 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 6, 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 $q_J \in \bar{\mathcal{O}}(k)$.

C Proofs for Section 4

C.0.1 Proof of Lemma 8

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

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

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.

D The two-state case: proof of Theorem 9

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 long 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 $p_* r(\omega^+) + (1-p_*)r(\omega^-) = 0$.

The invariant measure m assigns probability $\dfrac{\pi(\omega^+ \mid \omega^-)}{\pi(\omega^+ \mid \omega^-) + \pi(\omega^- \mid \omega^+)}$ to ω^+. With our notations, for $q \in [0,1] (= \Delta(\Omega))$ one has

$$\phi(q) = m + \left(1 - \pi(\omega^+ \mid \omega^-) - \pi(\omega^- \mid \omega^+)\right)(q - m),$$

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

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

$$\gamma(p) = \frac{p}{p_*}\gamma(p_*) + \left(1 - \frac{p}{p_*}\right)\gamma(0) \text{ for } p \in [0, p_*], \tag{8}$$

and

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

Eq. (8) shows that $\gamma(\cdot)$ is affine over $[0, p_*]$ (but not necessarily on $[p_*, 1]$). Note that $\gamma(0) = \delta(\gamma \circ \phi)(0)$. Note also that $\widehat{r}(p) = \dfrac{p}{p_*}$ for $p \in [0, p_*]$.

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.

Case 1: $p_* \geq m$ and $\lambda \geq 0$ (see Figure 7).

Figure 7: the various beliefs in Case 1.

We here argue that $\gamma(p) = \gamma_*(p)$ for each p, where $\gamma_*(p)$ is the upper bound on payoffs identified in Eq. 4.

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 $\mathbf{E}_{\sigma_*}[\widehat{r}(p_n)]$. Since $\widehat{r}$ is affine on $[0, p_*]$, the expected payoff in stage n is also equal to $\widehat{r}(\mathbf{E}_{\sigma_*}[p_n]) = \widehat{r}(\bar{p}_n)$, 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 $n_* := \inf\{n \geq 1 : \bar{p}_n < p_*\}$ 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 $q_n = \bar{p}_n$ and $r(q_n) = 1 = \widehat{r}(\bar{p}_n)$ 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.

All other cases below follow the same pattern. We first argue that γ is piecewise affine, increasing on $[0, p_*]$ and concave on $[0, 1]$. Combined with the fact that γ is affine on $J = [0, p_*)$, this implies that the difference $d(p) := \gamma(p) - \delta\gamma(\phi(p))$ is piecewise affine, and convex on $[0, p_*]$. We then compute the slope of d at zero. In each case, we obtain that $d'(0) = (1 - (\delta\lambda)^i)\gamma'$ for some $i \in \mathbf{N}$, where γ' is the slope of γ on $[0, p_*]$.

Case 2: $p_* \leq m$ and $\lambda \geq 0$ (see Figure 8).

Figure 8: the various beliefs in Case 2.

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 $p \geq p_*$, and therefore γ is piecewise affine. Using Eq. (8), this implies

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

Since $\gamma(0) < 1$ it follows that γ is increasing on $[0, p_*]$, and therefore that γ is concave on $[0, 1]$.

We next compute $d'(0)$. If $\phi(0) \geq p_*$, one has $\phi([0, p_*]) \subset I$, so that $d(p) = \gamma(p) - \delta$ for $p \leq p_*$, and $d'(0) = \gamma'$. If instead $\phi(0) < p_*$, there exists a unique q_* such that $\phi(q_*) = p_*$, and $q_* \in [0, p_*]$. Therefore, $\phi([0, q_*]) \subset [0, p_*]$, and $d'(0) = \gamma' - \delta\phi'(0)\gamma'(\phi(0)) = (1 - \delta\lambda)\gamma'$.

Case 3: $p_* \geq m$ and $\lambda \leq 0$ (see Figure 9).

Figure 9: the various beliefs in Case 3.

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

$$\gamma(p) = (1 - \delta) + \delta\gamma(\phi(p)) \text{ for } p \in I, \tag{9}$$

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 (9) yields $\gamma'(p) = \delta\lambda\gamma'(\phi(p))$ for $p > p_*$, hence the two slopes have opposite signs. Since $\gamma(p_*) = (1 - \delta) + \delta\gamma(\phi(p_*))$, we have $\gamma(p_*) > \gamma(\phi(p_*))$, so that γ is increasing on $[0, p_*]$ (and then decreasing on $[p_*, 1]$).

We next compute $d'(0)$. If $\phi(0) \leq p_*$, one has $\phi([0, p_*]) \subset [0, p_*]$, hence d is affine on $[0, p_*]$ and $d'(0) = (1 - \delta\lambda)\gamma'$. If instead $\phi(0) > p_*$, there is a unique q_*, such that $\phi(q_*) = p_*$ and $q_* \in [0, p_*]$. On the interval $[0, q_*]$, d is therefore affine with $d'(0) = \gamma' - \delta\lambda\gamma'(\phi(0)) = (1 - (\delta\lambda)^2)\gamma'$.

Case 4: $p_* \leq m$ and $\lambda \leq 0$ (see Figure 10).

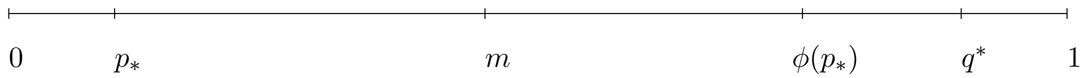

Figure 10: the various beliefs in Case 4.

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 $\phi(1) < p_*$, there is a unique $q_* \in [m, 1]$ with $\phi(q_*) = p_*$. 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 $p \geq q_*$, $\gamma(p) = (1-\delta) + \delta\gamma(\phi(p))$. 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 $[q_*, 1]$.

We now compute $d'(0)$. If $\phi(0) \leq q_*$, one has $\phi([0, p_*]) \subset [p_*, q_*]$, hence $d(p) = \gamma(p) - \delta$ on $[0, p_*]$, so that $d'(0) = \gamma'$. If instead $\phi(0) > q_*$, one has $d'(0) = \gamma' - \delta\lambda\gamma'(\phi(0))$. Since $\phi(0) \in (q_*, 1] \subset I$, $\gamma'(\phi(0)) = \delta\lambda\gamma'(\phi^{(2)}(0)) = (\delta\lambda)^2\gamma'$, since $\phi^{(2)}(0) \in [0, p_*)$.

E Proof of Theorem 11

As indicated in the text, we will prove a strengthened version of Theorem 11, to be used in the proof of Theorem 12. We recall from Section 3.3 that $L_k(\cdot)$ is the linear map defined by

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

With the notations of Section 3.3, the optimal solution π^* to (LP) is affine on the set $\bar{\mathcal{O}}(k) = \{L_k(\cdot) \geq 0 \geq L_{k+1}(\cdot)\}$ and, therefore, so is $\widehat{r}(\cdot)$.

Theorem 14 *Let k be such that $m \in \bar{\mathcal{O}}(k)$. Then $\gamma(p) = \gamma_*(p)$ for $p \in \bar{\mathcal{O}}(k)$. Therefore, the greedy policy σ_* is optimal whenever the initial distribution belongs to $\bar{\mathcal{O}}(k)$.*

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. By Lemma 7, 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 σ_*. That is, under σ_* and for $p_1 \in \bar{\mathcal{O}}(k)$, one has $p_n \in \bar{\mathcal{O}}(k)$ a.s. for every n. Since $\widehat{r}$ is affine on $\bar{\mathcal{O}}(k)$, this implies

$$\mathbf{E}_{\sigma_*}\left[\widehat{r}(p_n)\right] = \widehat{r}\left(\mathbf{E}_{\sigma_*}\left[p_n\right]\right) = \widehat{r}(\bar{p}_n)$$

for each round n. The result follows. ■

E.1 Proof of Theorem 12

We will first 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 15 *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 that

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

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 14, $\gamma(p) = \gamma_*(p)$ 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 16 *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 $p' \in \bar{\mathcal{O}}(k) \cap [q_I, q_J]$. Since $p \in \bar{J} \setminus \bar{\mathcal{O}}(k)$, one has $a_I, a_J > 0$. By Lemma 15, 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 policy, the expected payoff in each round is equal to the first best payoff in that round. In particular, any optimal splitting at p' must be the greedy one. By Lemma 7 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

$$L_{k-1}(m) > 0 = L_k(m) = \cdots = L_l(m) > L_{l+1}(m) \text{ for some } k \leq l. \tag{10}$$

We then set $P := \{p \in \bar{J}, L_{k-1}(p) \geq 0 \geq L_{l+1}(p)\} = \bar{\mathcal{O}}(k-1) \cup \cdots \cup \bar{\mathcal{O}}(l)$. By construction, m belongs to the interior of P. By (10), one has $m \in \bar{\mathcal{O}}(i)$ for $i = k-1, \ldots, l$, hence the set P is stable under ϕ. This implies that σ_* is optimal whenever $p_1 \in P$.

Lemma 17 *Assume that all connected components of $\bar{J} \setminus P$ in $\Delta(\Omega)$ are convex. Then the conclusion of Theorem 12 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. Let N be a compact neighborhood of m contained in Q.

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 $\tilde{p} \in N$ and all $\mu \in \mathcal{S}(\tilde{p})$, one has $\mu(q \in \bar{Q}) \geq c$.

Since $\Delta(\Omega)$ is compact and ϕ is contracting, there exists $\bar{n} \in \mathbf{N}$ such that $\phi^{(\bar{n})}(p) \in N$ 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 16, 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 12 follows from Lemma 17.

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 $\theta := \inf\{n \geq 1 \colon q_n \in I\}$ 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 0, 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 $\mathbf{P}_\tau(\theta \leq \bar{n}) \geq c$ for all initial distributions p_1. This implies the result, as in the first part of the proof.

E.2 Proof of Theorem 13

The analysis relies on a detailed study of the belief dynamics under σ_*. We will organize the discussion according to the cardinality of Ω^-.

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

We prove the optimality of σ_* in two steps. We first argue that γ is concave and that d is nonnegative on the straight line joining C and m, see Figure 11. 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. Hence[19] it follows from Section D that γ is concave and that d nonnegative on

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

$\mathcal{L}$.

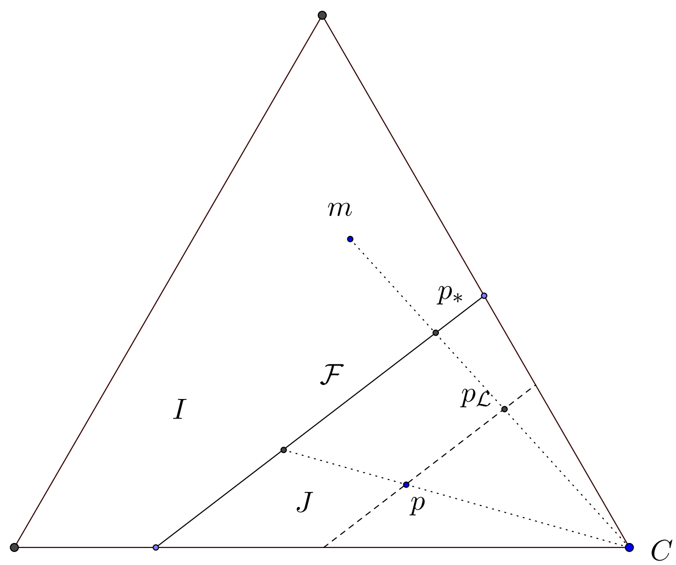

Figure 11: The case $|\Omega^-| = 1$.

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 $\widehat{r}(p) = \widehat{r}(p_{\mathcal{L}})$, where $p_{\mathcal{L}} \in \mathcal{L}$ is such that $(pp_{\mathcal{L}})$ 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 d.

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

We will exhibit a 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 denote by J_0 the triangle (C^+, B^+, C). 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), see Figure 6. 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 12.

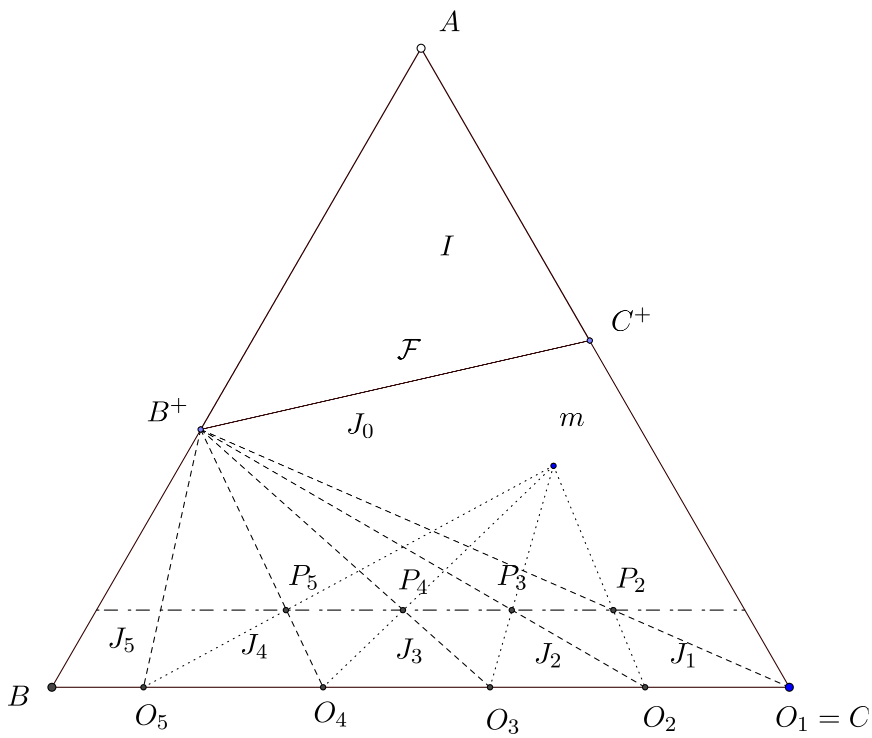

Figure 12: The points $(O_k)_k$ and $(P_k)_k$.

We first argue that this construction indeed ends in a finite number of steps.

Claim 18 $K < +\infty$.

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 13.

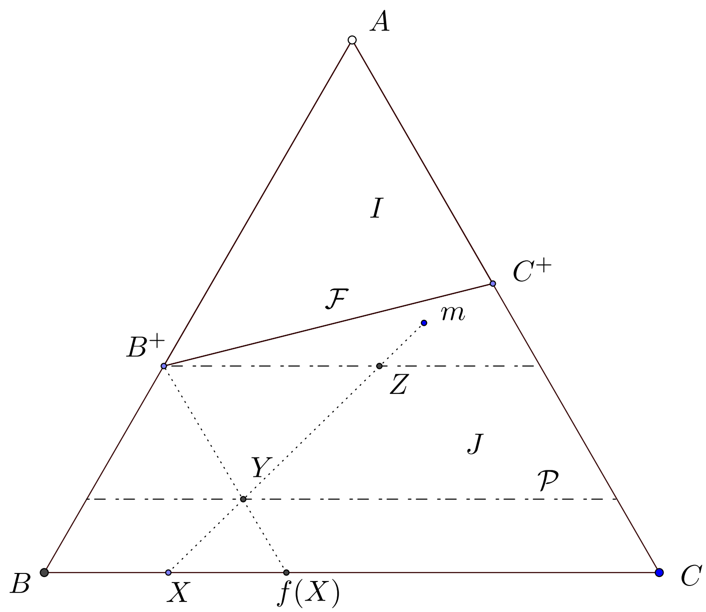

Figure 13: The definition of f.

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. 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_kO_{k+1} increases with k. In both cases, $K < +\infty$. ■

For $k = 1, \ldots, K-1$, 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 19 *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 20 *The function γ is concave on $J_k \cup J_{k+1}$ for $k = 1, \ldots, K-2$.*

Proof. We will use the following elementary observation. Let $g_1, g_2 : \mathbf{R}^2 \to \mathbf{R}$ 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 concave[20] 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

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

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 $[p_*, \tilde{p}_0]$.[21] Since

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

[21] We 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 $[0, 1]$.

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 21 *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 $J_{k_1}, J_{k_1+1}, \ldots, J_{k_2}$ 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 22 *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 $k \geq 1$.[22] 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 21 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 23 *$\gamma(p) \leq \gamma(\mathcal{F})$ for all $p \in J$.*

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

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 $\gamma(C) \leq \gamma(\mathcal{F})$.

[22] For other cases, the concavity of γ on $[\tilde{p}_1, \tilde{p}_2]$ follows from either Step 1 or Claim 21.

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

$$d(p_\varepsilon) = \gamma(p_\varepsilon) - \delta\gamma(\phi(p_\varepsilon)) = \gamma(p_\varepsilon) - \delta\gamma(q_\varepsilon) + \delta\left(\gamma(q_\varepsilon) - \gamma(\phi(p_\varepsilon))\right). \tag{11}$$

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

$$\gamma(p_\varepsilon) - \delta\gamma(q_\varepsilon) = \varepsilon\gamma(B^+)(1-\delta). \tag{12}$$

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

$$\gamma(q_\varepsilon) - \gamma(\phi(p_\varepsilon)) = \gamma_k(q_\varepsilon) - \gamma_k(\phi(p_\varepsilon)) = \gamma_k(q_\varepsilon - \phi(p_\varepsilon)) = \varepsilon\gamma_k(B^+ - \phi(B^+)). \tag{13}$$

Substituting (12) and (13) into (11) one gets

$$d(p_\varepsilon) = \varepsilon\left(\gamma(B^+)(1-\delta) + \delta\gamma_k(B^+ - \phi(B^+))\right). \tag{14}$$

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 (14), one finally gets

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

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 $[O_{k+1}, p]$.