MARKOVIAN PERSUASION

EHUD LEHRER AND DIMITRY SHAIDERMAN

Abstract. In the classical Bayesian persuasion model an informed player and an uninformed one engage in a static interaction. The informed player, the sender, knows the state of nature, while the uninformed one, the receiver, does not. The informed player partially shares his private information with the receiver and the latter then, based on her belief about the state, takes action. This action, together with the state of nature, determines the utility of both players. We consider a dynamic Bayesian persuasion situation where the state of nature evolves according to a Markovian law. In this repeated persuasion model an optimal disclosure strategy of the sender should, at any period, balance between obtaining high a stage payoff and disclosing information which may have negative implications on future payoffs. We discuss optimal strategies under different discount factors and characterize when the asymptotic value achieves the maximal value possible.

Date: March 6, 2022
JEL Classification: D72, D82, D83, K40, M31
Keywords: Markovian persuasion, Dynamic Bayesian persuasion, Markov chain, Asymptotic value, Absorbing set, Homothety.

The first author acknowledges the support of grants ISF 591/21 and DFG KA 5609/1-1. Andrei Iacob and Eilon Solan and greatly acknowledged for their comments.

1. Introduction

The literature devoted to Bayesian persuasion studies optimal policies by which an informed agent, the sender, discloses information to an uninformed agent, the receiver. Kamenika and Gentzkow [21] present a case where a prosecutor, who is fully informed about the state of nature, i.e., whether the suspect is guilty or innocent, wishes to persuade a judge, e.g., to convict the suspect of a crime. This is a static scenario: upon the prosecutor's disclosure, the judge takes a decision and the game is over.

In this paper we study a dynamic model where the interaction between the sender, who has a commitment power, and the receiver, evolves over time. At each stage the sender is informed about a randomly chosen state of nature. The decision of how much information to disclose to the receiver at each stage is at the sender's discretion. The latter publicly announces an information provision policy, to which he is committed throughout. The receiver knows the details of this policy and thus, when she receives a signal from the sender, she updates her belief about the true state accordingly. She then takes action based solely on her posterior belief. This action, together with the realized state, determines not only her own payoff, but also that of the sender. In short, in this dynamic signaling interaction, the sender is in charge of the informational aspects of the interaction (talking), whereas the receiver is in charge of acting. Neither may affect the evolution of the realized states, which is entirely exogenous.

We assume that the evolution of states follows a Markov chain. The dynamic of the receiver's belief is governed by both the disclosure policy of the sender and a Markov transition matrix. After observing the signal sent by the sender, the receiver updates her belief according to Bayes' law. Due to the Markovian transition law, this updated belief shifts to another belief. Since all terms are known to both players, the nature of the shift from the prior belief to a posterior one is also known to both.

When committing to an information provision policy, the sender takes into account the resulting posterior belief of the receiver, which has two effects. The first is on the action taken by the receiver and directly on his own stage payoff. In a dynamic setting, as discussed here, there is also an effect on the belief over future states, and consequently on future payoffs. The posterior resulting from a Bayesian updating in one period is then shifted by the Markov transition matrix and becomes the initial belief in the next one. Balancing between these two, potentially contradicting effects, is the dynamic programming problem faced by the sender.

In this paper we study the long-run optimal values that the sender may achieve in the case of an irreducible and aperiodic Markov chain.[1] This value depends on his discount factor. As in Aumann and Maschler [7] and in Kamenika and Gentzkow [21], the optimal signaling strategy follows a splitting scheme of beliefs in accordance with the concavification of a certain function[2]. While in a static model this function is the payoff function of the sender, here, this function takes into account the underlying dynamics and combines present and future payoffs.

The first set of results deal with a patient sender. In this case the stationary distribution of the Markov chain plays a central role. The reason is that in case the sender keeps mute, the receiver's beliefs quickly[3] become convergent to the stationary distribution.

We show that as the sender becomes increasingly patient,[4] the optimal values converge uniformly to a certain level, called the asymptotic value. Moreover, this value cannot exceed the optimal level obtained in the static (one-shot) case when the prior belief of the Markov chain is its unique stationary distribution. One may ask under what conditions this static optimal value can be obtained in the dynamic model. Our main theorem provides a full characterization. In order to describe it we introduce the notion of an M-absorbing set, where M is the irreducible, aperiodic stochastic matrix governing the transitions of states in the Markov chain.

A set C of beliefs is said to be M-absorbing, if whenever $p \in C$, the belief pM can be decomposed as a convex combination of beliefs in C. To better understand this notion, suppose that C is convex. This set is M-absorbing if any point p in C is shifted by M to another point in C, and this belief is further shifted to another point within C, etc. Thus, the image of C (under the transformation M) remains in C.

Kamenika and Gentzkow [21] proved that when the prior distribution is the stationary distribution π_M of M, the static value is equal to the concavification of the one-shot payoff function, evaluated at π_M. This

[1]In a companion paper we examine, based on the current study, the general case.

[2]That is, the lowest concave function above a given function.

[3]At an exponential rate (see e.g., Theorem 4.9 in [28]).

[4]The method of treating patient players in not new to the literature of dynamic interactions. Fudenberg and Maskin [15] introduced a Folk Theorem characterizing the set of equilibria in a repeated game played by patient players. Typically in these kind of games, it is impossible to provide a description of the value or of the set of equilibria that correspond to a specific discount factor that is more informative than that given by Bellman equation (see Abreu et. al. [1]).

value is either the sender's payoff function evaluated at π_M, or obtained by splitting π_M. The latter means that (i) π_M is expressed as a convex combination of other beliefs, which are the posteriors induced by the information the sender provides to the receiver, and (ii) these posteriors guarantee the highest expected payoff the sender may obtain this way. An optimal splitting method is effective when posteriors induce an expected payoff that surpasses the sender's payoff function evaluated at the π_M. From a geometric point of view, the combination of posteriors that corresponds to the best split around π_M generates a hyperplane that touches the graph of the sender's payoff function from above. This is a tangent hyperplane that we call the *central tangent*, because it is the supporting hyperplane to the graph of the concavification of the sender's payoff function at the stationary distribution π_M.

The main theorem characterizes, in terms of the primitives of the model, the conditions under which the asymptotic value is as high as it can get (i.e., equals the static value at π_M). The theorem asserts that this occurs precisely when there is an M-absorbing set of beliefs over which the central tangent supports (or in less formal words, touches) the sender's payoff function. The fact that this set is M-absorbing is expressed in terms of the parameters of the dynamics, while the condition that on this set the central tangent supports the sender's payoff function is expressed in terms of the sender's preferences.

The importance of such M-absorbing sets of beliefs stems from two factors. First, it consists of beliefs on which the sender's payoff function coincides with their values on the hyper plane, namely, the corresponding payoffs are relatively high. Second, the shift of these belief by the Markov chain keeps them within the M-absorbing set itself, enabling a perpetual use of this set in order to obtain high payoffs.

The intuition of this result is as follows. Suppose first that the initial belief is π_M and furthermore, that C is an M-absorbing set of beliefs on which the central tangent supports the sender's payoff function. At the first stage the sender splits π_M to a set of posteriors in C in order to obtain (in expectation) the optimal static value at π_M. At this point in time, the Markov chain comes into play. The shift of those beliefs by M constitutes the set of receiver's possible beliefs at the start of the second stage. Due to the fact that C is M-absorbing, these beliefs are within C, enabling the sender to split each one to beliefs that lie within C. In other words, the posterior beliefs of the receiver at the second stage also lie within C. This feature applies to all subsequent stages as well.

It turns out that this scheme guarantees that the sender will obtain the optimal static value at π_M in the entire future. Indeed, the posteriors at any stage, generated by the sender's messages and by the Markov motion, actually define a split of π_M that remains within C. This is true because the expectation of the posterior beliefs at any stage is equal to π_M and moreover, any such split, being confined to C, guarantees the static value at π_M. Note that this splitting scheme may be applied not only to π_M, but also to any initial belief within the boundaries of C.

Now suppose that the initial belief is not π_M, nor within C. In case the sender provides no information to the receiver, the receiver's beliefs are determined solely by the shifts of M. These beliefs become closer and closer to π_M, and this process is very fast. Now the splitting strategy described above can be employed. The effect is twofold. First, all the posteriors remain within C, and second, as time goes by, the posteriors converge, in expectation, to π_M. Both effects guarantee that when the sender is patient enough, the time it takes for the initial belief to get into C does not affect significantly the sender's discounted payoff in the entire interaction. Moreover, the payoffs in subsequent stages are rather close to the static value at π_M. We may therefore conclude that when an M-absorbing set of beliefs on which the central tangent supports the sender's payoff function exists, the asymptotic value achieves the highest level possible.

The other direction of the main theorem states that when such a set does not exist, the asymptotic value is strictly below the optimal static value at π_M. In order to show this direction, we develop and use tools and techniques adopted from the literature devoted to repeated games with incomplete information (see Renault [33]).

The second type of results is non-asymptotic in nature. For a certain region around the stationary distribution we provide a closed-form expression for the values corresponding to any level of patience. In the case where this region includes a neighborhood of the stationary distribution, we provide asymptotically effective[5] two-sided bounds for the values corresponding to any level of patience. Moreover, the effectiveness of those bounds depends on the geometry and size of the described neighborhood.

The closest paper to the present one is that of Renault et. al. [32]. It deals with a specific type of Markov chains, homothetic ones, and with a utility function which gives a fixed amount to the sender in a certain region and zero in its complement. Renault et. al. [32] show that in

[5]That is, bounds whose difference is arbitrarily small for a patient enough sender.

this model the optimal information provision policy is a greedy one (starting from a certain random time period on). Namely, the sender's greedy policy instructs him to maximize his current stage payoff at any time period, ignoring future effects of this policy. Peşki and Toikka [29] refer to homothetic chains in the context of zero-sum stochastic games when the underlying state evolves according to a Markov operator and is being observed only by one player. Also closely related to our work are the works of Ely [11] and Farhadi and Teneketzis [13], which deal with a case where there are two states, whose evolution is described by a Markov chain with one absorbing state. In these two works, the receiver is interested in detecting the jump to the absorbing state, whereas the sender seeks to prolong the duration of time until detection.

In a broader context, our paper can be ascribed to a vast growing literature concerned with dynamic information design problems, specifically with the analysis of situations in which the sender is required to convey signals sequentially to a receiver, which he may base on his additional private information. Without elaborating on the exact details of their models, whose frameworks employ diverse methodologies, we refer the reader to Mailath and Samuelson [25], Honryo [18], Orlov et al. [27], Phelan [30], Lingenbrink and Iyer [24], Wiseman [37], Kolotilin et al. [22], Athey and Bagwell [4], Arieli and Babichenko [2], Arieli et al. [3], Au [5], Dziuda and Gradwohl [10], Escobar and Toikka [12], Guo and Shmaya [35], Ganglmair and Tarantino [16], Guo and Shmaya [17], and Augenblick and Bodoh-Creed [6] for an acquaintance with the literature.

The mathematical innovation of the paper is the incorporation of tools and techniques from the field of repeated games with incomplete information [33]. We find it appropriate to refer the reader to works on repeated games with incomplete information involving Markov chains, particularly Renault [31], Hörner et al. [20], and Bressaud and Quas [9]. Forges [14] surveys the close relationships between repeated games with incomplete information and modern economic cornerstones such as cheap talk and persuasion.

The paper is organized as follows. Section 2 presents the model and an example by which we explain the notations, new concepts and the main results. The asymptotic results as well as the main theorem are provided in Section 3. Results related to non-asymptotic values are given in Section 4. Section 5 provides a characterization of homothetic matrices in terms of the asymptotic value. The proofs are given in Section 6.

2. The Model

Let $K = \{1, ..., k\}$ be a finite set of states. Assume that $(X_n)_{n\geq 1}$ is an irreducible and aperiodic Markov chain over K with prior probability $p \in \Delta(K)$ and a transition rule given by the stochastic matrix M. We assume that $(X_n)_{n\geq 1}$ are defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$.

A *sender* is an agent who is informed at each time period n of the realized value x_n of X_n. Upon obtaining this information a sender is prescribed to send a signal s_n, from a finite set of signals S with cardinality at least k.[6]

A *receiver* is an agent who, at any time period n, is instructed to make a decision b_n from a set of possible set of decisions B, assumed to be a compact metric space. This decision may take into account the first n signals $s_1, ..., s_n$ of the sender.

The payoffs of the sender and the receiver at time period n are given by the utilities $v(x_n, b_n)$ and $w(x_n, b_n)$, respectively, so that they depend solely on the realized state x_n and the decision b_n. Both the sender and the receiver discount their payoffs by a factor $\lambda \in [0, 1)$. We denote this game by $\Gamma_\lambda(p)$. As in the models of Renault et al. [32], Ely [11], and Farhadi and Teneketzis [13], the receiver obtains information only through the sender.

A *signaling strategy* σ of the sender in $\Gamma_\lambda(p)$ is described by a sequence of stage strategies (σ_n), where each σ_n is a mapping $\sigma_n : (K \times S)^{n-1} \times K \to \Delta(S)$. Thus, the signal s_n sent by the sender at time n is distributed by the lottery σ_n, which may depend on all past states $x_1, ..., x_{n-1}$ and past signals $s_1, ..., s_{n-1}$ together with the current state x_n. Let Σ be the space of all signaling strategies.

A standard assumption in many Bayesian persuasion models is that of *commitment* by the sender. That is, we assume that the sender commits to a signaling strategy σ at the start of the game $\Gamma_\lambda(p)$, and makes it known to the receiver. The commitment assumption enables the receiver to update her beliefs on the distribution of states (X_n) based on the signals (s_n) she receives from the sender. Formally, by Kolmogorov's Extension Theorem, each signaling strategy σ together with $(X_n)_{n\geq 1}$ induces a unique probability measure $\mathbb{P}_{p,\sigma}$ on the space

[6] This assumption is in place to make sure the sender can disclose $(x_n)_n$.

$\mathcal{Y} = (K \times S)^{\mathbb{N}}$, determined by the laws

$$\mathbb{P}_{p,\sigma}(x_1, s_1, ..., x_n, s_n) = \left(p(x_1) \prod_{i=1}^{n-1} M_{x_i, x_{i+1}} \right) \times \left(\prod_{i=1}^{n} \sigma_i(x_1, s_1, ..., x_{i-1}, s_{i-1}, x_i)(s_i) \right). \tag{1}$$

Thus, the posterior probability p_n^ℓ the receiver assigns to the event $\{X_n = \ell\}$, given the signals $s_1, ..., s_n$ and the strategy σ, is given by the formula

$$p_n^\ell = \mathbb{P}_{p,\sigma} \left(X_n = \ell \mid s_1, ..., s_n \right). \tag{2}$$

Set $p_n = (p_n^\ell)_{\ell \in K}$. A second key assumption of our model is that the receiver's decision at any time period n depends only on p_n. Such an assumption includes the natural situation in which the receiver seeks to maximize her expected payoff based on her current belief (e.g., Renault et al. (2015)). Denote by $\theta : \Delta(K) \to B$ the decision policy of the receiver, that is, the mapping which depicts the decision of the receiver as a function of her belief. As in previous related models, we assume that the decision policy of the receiver is known to the sender. The last assumption of our model is that the function $u : \Delta(K) \to \mathbb{R}$ defined by $u(q) = \sum_{\ell \in K} q^\ell v(\ell, \theta(q))$ is continuous. To summarize, our assumptions imply that the signaling strategy σ of the sender determines his payoff at any time period n. Moreover, the total expected payoff to the sender in $\Gamma_\lambda(p)$ under the signaling strategy σ can now be written as

$$\gamma_\lambda(p, \sigma) := \mathbb{E}_{p,\sigma} \left[(1 - \lambda) \sum_{n=1}^{\infty} \lambda^{n-1} u(p_n) \right], \tag{3}$$

where $\mathbb{E}_{p,\sigma}$ is the expectation w.r.t. $\mathbb{P}_{p,\sigma}$. The value of the game $\Gamma_\lambda(p)$ is $v_\lambda(p) = \sup_{\sigma \in \Sigma} \gamma_\lambda(p, \sigma)$.

Example 1. *Each day Anna is producing one unit of a divisible good and sells it to Bob. Anna is committed to sell to Bob any portion of this good he desires. The quality of the good can be either high (state H) or low (state L). When Bob purchases the proportion $x \in [0, 1]$, at state L his utility $w(L, x)$ equals $2 - x^2$, while at state H his utility $w(H, x)$ is equal to $2 - (1 - x)^2$.*

Suppose that Bob's belief is $(p, 1 - p)$, where p is the probability he assigns to the state H. When deciding to purchase a proportion x of Anna's good, his expected utility is $p(2 - (1 - x)^2) + (1 - p)(2 - x^2) = -x^2 + 2px + 2 - p$. The maximum of this function is attained at p.

Thus, when Bob's belief is $(p, 1-p)$*, he purchases the portion* p *of Anna's product and leaves a* $1-p$ *to her. In particular, when Bob assigns a probability of* 1 *to the quality of the good being high, he will purchase the entire unit of it, as opposed to the case where he purchases nothing in the case he assigns probability* 1 *to the quality being low.*

As for Anna, when the state is H*, her utility derives both from selling the product to Bob and from self-consumption. Her utility from selling is increasing linearly with the proportion* p *she sells: it is* $0.5p$*. Her utility from self-consumption is increasing linearly as long as* $1-p$*, the proportion she is left with, is below* 0.5*. In this range her utility is* $3.5(1-p)$*. The point is that Anna cannot consume more than* 0.5*. At this level she gets saturated, she obtains the maximal utility* 1.75*, but she has to dump the leftover. The cost of dumping is increasing linearly with the quantity being dumped. In order to dump* $(1-p)-0.5 = 0.5-p$ *she has to pay* $2.5(0.5-p)$*. To sum up, her total utility becomes*

$$v(H,p) = \begin{cases} 0.5p + 1.75 - 2.5(0.5-p), & \text{if } p \le 0.5, \\ 0.5p + 3.5(1-p), & \text{if } p > 0.5, \end{cases}$$

which is equal to $2 - 3|p - 0.5|$*.*

In case the state is L*, Anna does not want to consume her product and her total utility* $v(L,p)$*, when Bob decides to purchase the proportion* p*, combining sells and dumping equals* $p/10$*. In other words, the net payoff increases linearly with the probability Bob assigns to the state* H *and attains its maximum when Bob is convinced that the quality of the good is high* $(p = 1)$*. Notice that* $\min_p v(H,p) > \max_p v(L,p)$*, meaning, in particular, that Anna always prefers her good to be of a high quality.*

We obtain that $u((p,1-p)) = pv(H,p) + (1-p)v(L,p) = p(2-3|p-\frac{1}{2}|) + (1-p)p/10$*. The graph of* u*, which for the sake of convenience is plotted as a function of* p*, is exhibited on the left panel of Figure 1. Note that* u *is convex on the interval* $[0, 0.5]$ *and concave on the interval* $[0.5, 1]$*. This graph illustrates the short-term strategic incentives of Anna. Indeed, the minimal possible stage payoff,* 0*, occurs when the Anna reveals that the quality is low (corresponding to* $p = 0$*). A revelation of the state* H *(corresponding to* $p = 1$*) would result in a payoff of* 0.5*. The maximal payoff for Anna,* 1.045*, is attained when* $p = 0.581$*. As explained below (see Example 3), for that* p *the optimal signaling strategy would instruct Anna to not reveal any information regarding the realized state.*

In a dynamic model the amount of information revealed by the sender is a result of an interplay between the one-shot payoff, u, and the transition law governing the evolution of future states. The tension between these two factors is discussed in the rest of this paper.

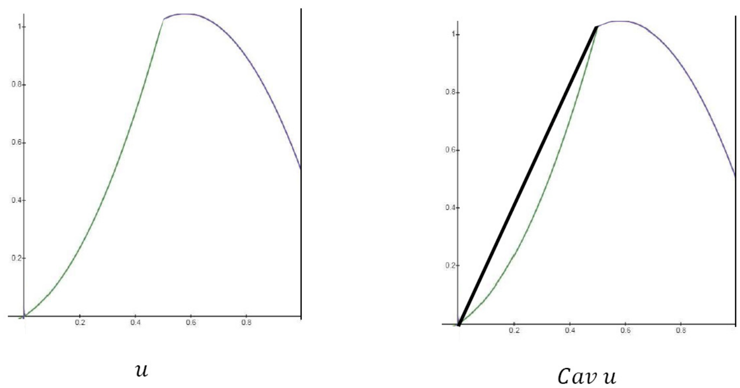

FIGURE 1. The graphs of u and $(\mathrm{Cav}\, u)$

3. THE MAIN THEOREM

3.1. **The Existence of Asymptotic Value.** To state our first result we need to introduce some notations. First, let π_M be the unique stationary distribution of M. Second, for any function $g : \Delta(K) \to \mathbb{R}$, define the function $(\mathrm{Cav}\, g)$ by

$$(\mathrm{Cav}\, g)(q) := \inf\{h(q) \,:\, h : \Delta(K) \to \mathbb{R} \text{ concave, } h \geq g\}, \quad \forall q \in \Delta(K).$$

To showcase the Cav operator in action, we show the graph of $(\mathrm{Cav}\, u)$, where u is that given in Example 1, on the right-hand panel of Figure 1.

Our first result reveals that the influence of $p \in \Delta(K)$ on the value $v_\lambda(p)$ of a sufficiently patient sender, i.e., with λ close to 1, is negligible compared to the influences of u and M. Moreover, as the patience level λ gets closer to 1 the sequence of functions $v_\lambda(\cdot)$ converges uniformly on $\Delta(K)$. Formally, this result is stated as follows.

Theorem 1. *There exists a scalar $v_\infty \in \mathbb{R}$, $v_\infty \leq (\mathrm{Cav}\, u)(\pi_M)$, such that for every $\varepsilon > 0$ there exists $0 < \delta < 1$ such that*

$$|v_\lambda(p) - v_\infty| < \varepsilon, \qquad \forall \lambda > \delta, \quad \forall p \in \Delta(K). \tag{4}$$

As it turns out, the upper bound on v_∞ described in Theorem 1 is tight. In our next section we shall give a geometric criterion for this upper bound to be attained. To do so we begin by introducing and studying the notion of M-absorbing sets.

3.2. M-absorbing sets.

Definition 1. *A non-empty set $C \subseteq \Delta(K)$ is said to be M-*absorbing *if*[7] $qM \in \mathrm{conv}(C)$ *for every* $q \in C$.

The intuition behind this choice of terminology is the following. Since $q \mapsto qM$ is a linear operator, if C is M-absorbing, so is $\mathrm{conv}(C)$. However, for $\mathrm{conv}(C)$ M-absorption exactly describes the situation in which the image of $\mathrm{conv}(C)$ under M lies inside (is absorbed in) $\mathrm{conv}(C)$. In a dual fashion, if $C \subseteq \Delta(K)$ is a closed convex M-absorbing set, then since by Krein-Milman Theorem[8] $\mathrm{conv}(\mathrm{ext}(C)) = C$, we get that $\mathrm{ext}(C)$ is also M-absorbing. This implies, in particular, that since $\Delta(K)$ is M-absorbing, so is the set of its extreme points (i.e., the set of all mass-point distributions). Lastly, note that since $\mathrm{conv}(C_1) \cup \mathrm{conv}(C_2) \subseteq \mathrm{conv}(C_1 \cup C_2)$, if C_1 and C_2 are both M-absorbing, then so is $C_1 \cup C_2$.

Example 2. [Example 1, continued] *Suppose that the quality of the divisible good produced by Anna evolves on a day to day basis. This variation may be caused by random effects, such as the quality of raw material* (*which may depend on exogenous factors such as weather, market dynamics, etc.*)*. We model the random evolution by a Markov transition rule, given by the stochastic matrix*

$$M = \begin{pmatrix} 0.1 & 0.9 \\ 0.6 & 0.4 \end{pmatrix}. \tag{5}$$

Here, for instance, a day in which the quality of the good was high, will be followed by a day with a low quality with probability 0.9*. Also, it is more likely than not* (*with probability* 0.6) *that a day with low quality will be followed by one with high quality.*

Note that the stationary distribution π_M of M equals $(0.4, 0.6)$ (*which corresponds to $p = 0.4$ in Figure* 1)*. Also, we learn from Figure 1 that the pair* $((0.4, 0.6), (\mathrm{Cav}\, u)((0.4, 0.6)))$ *lies on the straight segment of the graph of* $(\mathrm{Cav}\, u)$ (*see the right panel of Figure* 1).

[7] $\mathrm{conv}(C)$ is the convex hull of the set C.

[8] $\mathrm{ext}(C)$ is the set of the extreme points of C, i.e., points in C that cannot be expressed as a convex combination of two distinct points in C.

Moreover, the set consisting of the points $(0,1)$ *and* $(0.5,0.5)$ (*presented in Figure 1 as the points* 0 *and* 0.5) *is not* M*-absorbing. The reason is that* $(0,1)$ *is mapped by* M *to* $(0,1)M = (0.6,0.4)$*, which does not lie in the convex hull of* $(0,1)$ *and* $(0.5,0.5)$.

In order to enhance the intuition about M-absorbing sets, assume that at any point $q \in \Delta(K)$ a football is passed in a straight line to the point qM. The orbit generated by the football at q is the union of line segments $\bigcup_{n\geq 1}[qM^{n-1}, qM^n]$, where $[x,y] = \{\alpha x + (1-\alpha)y : 0 \leq \alpha \leq 1\}$. When the set C is M-absorbing, the orbit generated by the football with starting point at $q \in C$ never exits $\mathrm{conv}(C)$.

We proceed with some basic examples of M-absorbing sets. The simplest are the singleton $\{\pi_M\}$ and the entire set, $\Delta(K)$. To describe additional examples, consider the ℓ_1-,ℓ_2-, and ℓ_∞-norms on $\Delta(K)$, denoted by $\|q\|_1 := \sum_{\ell\in K}|q^\ell|$, $\|q\|_2 := \sqrt{\sum_{\ell\in K}(q^\ell)^2}$ and $\|q\|_\infty := \max_{\ell\in K}|q^\ell|$, respectively, for $q \in \Delta(K)$. Denote by $\|M\|_i$ the operator norm[9] of M w.r.t. the ℓ_i-norm, $i \in \{1,2,\infty\}$.

For every $i \in \{1,2,\infty\}$ we have,

$$\|qM - \pi_M\|_i = \|qM - \pi_M M\|_i \leq \|M\|_i \|q - \pi_M\|_i. \tag{6}$$

It is known that $\|M\|_\infty$ coincides with the largest ℓ_1-norm of a row of M (see, e.g., Example 5.6.5 on p. 345 in [19]), and therefore $\|M\|_\infty = 1$. Also,[10] $\|M\|_2 = \|M\|_\infty = 1$. Thus, in view of (6), any ball (either open or closed) w.r.t. to the ℓ_2 or ℓ_∞-norm centered at π_M is M-absorbing. Moreover, if M is doubly stochastic it is known that[11] $\|M\|_1 = 1$ and therefore in that case any ball (either open or closed) w.r.t. the ℓ_1-norm, centered at π_M, is also M-absorbing. See Figure 2.

In all the examples above the M-absorbing sets contain π_M. This is not a coincidence. Indeed, for every M-absorbing set C, the image of $\mathrm{conv}(C)$ under the linear map M is also contained in $\mathrm{conv}(C)$. Therefore, by Brouwer's fixed-point theorem, M possesses a fixed point in[12] $\mathrm{cl}\,\mathrm{conv}(C)$. As the only fixed point of M is π_M, we deduce that $\pi_M \in \mathrm{cl}\,\mathrm{conv}(C)$ for every M-absorbing set C.

[9] $\|M\|_i = \max_{\|x\|_i=1}\|xM\|_i$.

[10] This follows from the fact that $\|M\|_2$ coincides with the maximal singular value of M (see, e.g., Example 5.6.6 on p. 346 in [19]).

[11] This follows from the fact that $\|M\|_1$ with the maximal ℓ_1-norm of a column of M (see, e.g., Example 5.6.5 on p. 344–345 in [19])).

[12] $\mathrm{cl}\,\mathrm{conv}(C)$ is the closure of $\mathrm{conv}(C)$.

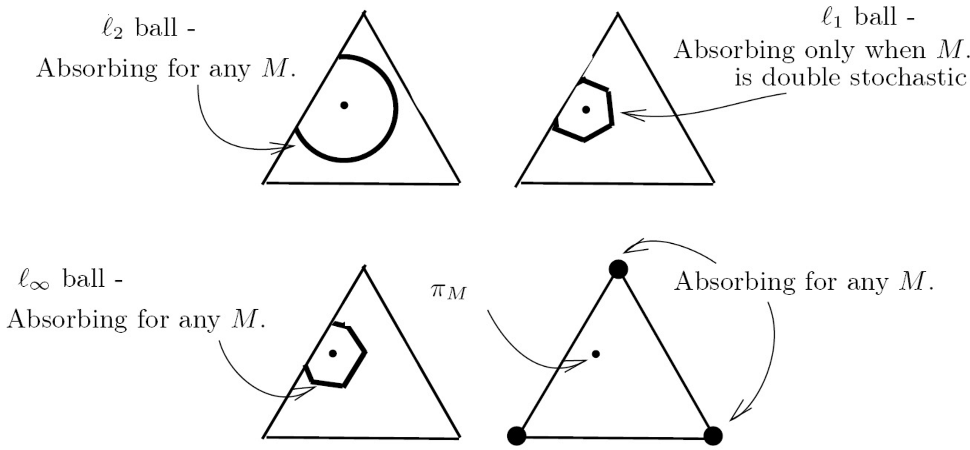

FIGURE 2. Absorbing sets.

We end the discussion on M-absorbing sets with the following proposition whose content and proof (provided in Section 6) may enhance the intuition about absorbing sets.

Proposition 1. *Let C be an M-absorbing set. Then, C contains a countable M-absorbing set.*

3.3. **The Main Theorem.** To state our main result we begin with a review of some basic concepts from the theory of concave functions. First, for each $g : \Delta(K) \to \mathbb{R}$ let $\text{Graph}[g] := \{(q, g(q)) : q \in \Delta(K)\}$. Since $(\text{Cav}\, u)$ is a concave function, $\text{Graph}[(\text{Cav}\, u)]$ can be supported at $(\pi_M, (\text{Cav}\, u)(\pi_M))$ by a hyperplane. We may parametrize each such supporting hyperplane by a point in $\mathbb{R}^k$ as follows; first, for every $z \in \mathbb{R}^k$ define $f_z : \mathbb{R}^k \to \mathbb{R}$ by $f_z(x) := (\text{Cav}\, u)(\pi_M) + \langle z, x - \pi_M \rangle$. Second, set

$$\Lambda := \{z \in \mathbb{R}^k : (\text{Cav}\, u)(q) \le f_z(q), \forall q \in \Delta(K)\}.$$

As $f_z(\pi_M) = (\text{Cav}\, u)(\pi_M)$ for every z, the set Λ corresponds to all supporting hyperplanes of $\text{Graph}[(\text{Cav}\, u)]$ at $(\pi_M, (\text{Cav}\, u)(\pi_M))$. In convex theory terminology, the set Λ is termed the *supper gradient of* $(\text{Cav}\, u)$ *at* π_M. For every $z \in \Lambda$ let

$$A_z := \{q \in \Delta(K) : u(q) = f_z(q)\}.$$

The set A_z can thus be interpreted as the projection to the first k coordinates of the intersection of Graph$[u]$ with the supporting hyperplane to Graph$[(\mathrm{Cav}\, u)]$ at $(\pi_M, (\mathrm{Cav}\, u)(\pi_M))$ parametrized by z. A visualization of A_z when $k = 3$ is given in Figure 3.

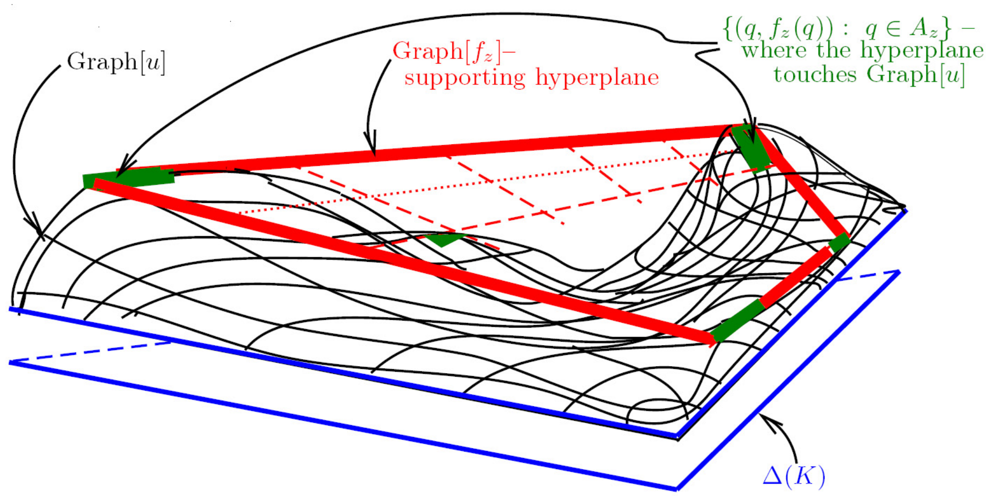

FIGURE 3. A visualization of A_z for $z \in \Lambda$.

Proposition 2. *We have the following:*

(i) *If A_z contains an M-absorbing set for some $z \in \Lambda$, then $v_\infty = (\mathrm{Cav}\, u)(\pi_M)$.*
(ii) *If $v_\infty = (\mathrm{Cav}\, u)(\pi_M)$, then for every $z \in \Lambda$, A_z contains a countable M-absorbing set.*

Why M-absorbing sets contained in the A_z's are of importance? This has to do with the control the sender has on the receiver's beliefs. Indeed, once a belief is in the convex hull of an M-absorbing subset $C \subseteq A_z$, $z \in \Lambda$, its shift under M, which describes the evolution of the posterior in one time period, also lies in $\mathrm{conv}(C)$. At this point in time the sender may send messages that would induce posteriors within C, and in particular in A_z. As $(\mathrm{Cav}\, u)$ is an affine function on $\mathrm{conv}(A_z)$ (see Lemma 4 in Section 6), the weighted average of the values of $(\mathrm{Cav}\, u)$ evaluated at these posteriors, is equal to the value of $(\mathrm{Cav}\, u)$ at π_M.

In the main theorem we summarize the results of Theorem 1 and Proposition 2. This theorem characterizes when patient senders can obtain a value close to the maximum possible, the upper bound stated in Theorem 1.

Theorem 2. *The following statements are equivalent:*

(i) *For every $\varepsilon > 0$ there exists $0 < \delta < 1$ such that*

$$|v_\lambda(p) - v_\infty| < \varepsilon, \qquad \forall \lambda > \delta, \quad \forall p \in \Delta(K). \tag{4}$$

(ii) *There exists $z \in \Lambda$ such that A_z contains an M-absorbing set.*

(iii) *For every $z \in \Lambda$, A_z contains a countable M-absorbing set.*

Note that this characterization is stated in terms of the primitives of the model: u and M. Moreover, it unravels the sensitivity of a patient sender to the interrelationships between u and M, as the sets A_z, $z \in \Lambda$, are fully determined by the former, whereas M-absorbing sets are clearly determined by the latter. An interesting question that arises naturally in this context is how sensitive is a patient sender to such interrelationships. Assume, for instance, that A_z does not contain an M-absorbing set for some $z \in \Lambda$. Can one quantify the difference between $(\operatorname{Cav} u)(\pi_M)$ and v_∞ in terms of u and M?

Example 3. [Example 2, continued] *Consider again the matrix M in* Eq. (5) *and recall that π_M equals $(0.4, 0.6)$. For the function u given in Example 1 the set A_z consists of the two points $(0, 1)$ and $(0.5, 0.5)$ in $\Delta(K)$ (corresponding to the points 0 and 0.5 in Figure 1) for every supporting hyperplane $z \in \Lambda$. Since this set is not M-absorbing, we conclude by Theorem 2 that for a sufficiently patient sender the value is strictly less than $(\operatorname{Cav} u)(\pi_M)$. However, if*

$$M = \begin{pmatrix} 1/2 & 1/2 \\ 1/6 & 5/6 \end{pmatrix}, \tag{8}$$

then $\pi_M = (0.25, 0.75)$. Therefore, as $A_z = \{(0, 1), (0.5, 0.5)\}$ coincides with the set of extreme points of the ball of radius 0.25 w.r.t. the ℓ_∞-norm around π_M, we deduce that A_z is M-absorbing for every $z \in \Lambda$. Thus, Theorem 2 ensures that the value $v_\lambda(p)$ of a sufficiently patient sender (i.e., when λ is close to 1) is close to $(\operatorname{Cav} u)(\pi_M) = 0.512$.

Under both transition matrices, the maximal payoff possible is obtained at $p = 0.581$. This point is located in the region where u is equal to $(\operatorname{Cav} u)$. The proof of Lemma 1 in Section 6 shows that at this point the sender has no incentive to alter the prior belief of the receiver. He therefore reveals no information to the latter. Such a result holds for any continuous u and any point p on which u agrees with $(\operatorname{Cav} u)$.

4. Additional Results: A Non-Asymptotic Approach and a Strong Law

4.1. **The value v_λ for every λ.** As it turns out, the case where $v_\infty = (\operatorname{Cav} u)(\pi_M)$ encompasses information about the behavior of v_λ across all $\lambda \in [0,1)$. To showcase this, we begin by recalling that as any union of M-absorbing sets is also M-absorbing, by Proposition 2, if $v_\infty = (\operatorname{Cav} u)(\pi_M)$, one may associate with each $z \in \Lambda$ a maximal (w.r.t. inclusion) M-absorbing set $B_z \subseteq A_z$. Set $D := \bigcup_{z\in\Lambda} \operatorname{conv}(B_z)$ (see Figure 4).

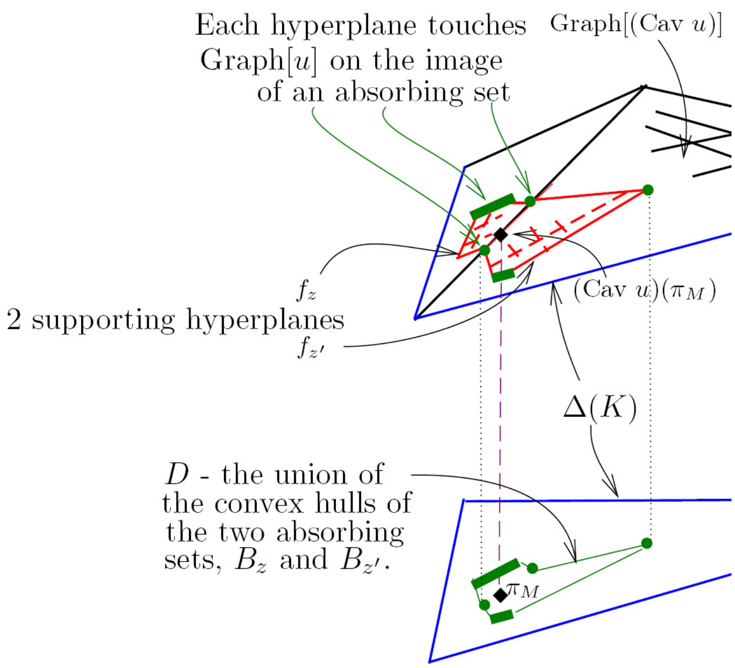

Figure 4. A visualization of D.

In our next theorem we give an exact formula for $v_\lambda(p)$ for all $p \in D$ and all $\lambda \in [0,1)$ in the case where $v_\infty = (\operatorname{Cav} u)(\pi_M)$. Therefore, the case $v_\infty = (\operatorname{Cav} u)(\pi_M)$ is not of a mere asymptotic nature, but rather reveals the full information regarding v_λ for any discount factor λ on the domain $D \subseteq \Delta(K)$. However, outside of D, the exact behavior of v_λ, even in the case $v_\infty = (\operatorname{Cav} u)(\pi_M)$, remains an open problem.

Theorem 3. *Assume that $v_\infty = (\operatorname{Cav} u)(\pi_M)$. Then, if $p \in \operatorname{cl} D$ we have*

$$v_\lambda(p) = (\operatorname{Cav} u)\left((1-\lambda)p\,(\mathrm{Id}_k - \lambda M)^{-1}\right), \quad \forall \lambda \in [0,1), \tag{9}$$

where Id_k *is the* $k \times k$ *identity matrix.*

This theorem implies that if $v_\infty = (\mathrm{Cav}\, u)(\pi_M)$ and the prior probability is π_M, then for every discount factor the value equals $(\mathrm{Cav}\, u)(\pi_M)$. Formally,

Corollary 1. *Assume that* $v_\infty = (\mathrm{Cav}\, u)(\pi_M)$. *Then,* $v_\lambda(\pi_M) = (\mathrm{Cav}\, u)(\pi_M)$ *for every* $\lambda \in [0, 1)$.

Let us now give some intuition regarding the formula given in Theorem 3. First, recall that the sum of an infinite geometric series with common ratio $r \in [0, 1)$ and scale factor 1 equals $1/(1 - r)$. As M is stochastic, we will argue in the proof of Theorem 3 that we apply a matrix version of the formula for the sum of an infinite geometric series so that

$$(1-\lambda)p\left(\mathrm{Id}_k - \lambda M\right)^{-1} = (1-\lambda)\sum_{n=1}^{\infty} \lambda^{n-1} p M^{n-1}.$$

Moreover, as the proof section shows (e.g., Eq. (15)), for any signaling strategy σ and prior belief $p \in \Delta(K)$ we have that $\mathbb{E}_{p,\sigma} p_n = pM^{n-1}$. Roughly speaking, the proof of Theorem 3 is divided into two parts. In the first, we come up with a signaling strategy σ that exploits the M-absorbing property of the $\mathrm{conv}(B_z)$'s. On the sets $\mathrm{conv}(B_z)$ the stage payoffs achieve their maximal possible value, shown to be equal to $(\mathrm{Cav}\, u)(pM^{n-1})$ for $n \geq 1$. The second part connects $(1-\lambda)\sum_{n=1}^{\infty} \lambda^{n-1} (\mathrm{Cav}\, u)(pM^{n-1})$ with the formula given in Eq. (9). Such an argument is valid on D, because $(\mathrm{Cav}\, u)$ is an affine function on each $\mathrm{conv}(B_z)$, $z \in \Lambda$. The extension of the result to $\mathrm{cl}\, D$ follows by continuity arguments.

It turns out that whenever $v_\infty = (\mathrm{Cav}\, u)(\pi_M)$, the geometric structure of the set D around the point π_M may be utilized to derive two-sided bounds on $v_\lambda(p)$ for every discount factor $\lambda \in \Delta(K)$. Let

$$n_D(p) := \min\{n \geq 0 \, : pM^n \in \mathrm{cl}\, D\};$$

$n_D(p)$ is the first time the orbit $\{pM^n\}_{n \geq 0}$ visits the set $\mathrm{cl}\, D$. Assume now that π_M is in the interior[13] of $\mathrm{cl}\, D$. By the Convergence Theorem for Markov chains (e.g., Theorem 4.9 in [28]) $pM^n \to \pi_M$, implying that $n_D(p)$ is finite, for every $p \in \Delta(K)$.[14]

[13] $\mathrm{int}(\mathrm{cl}\, D)$ is the interior of $\mathrm{cl}\, D$.

[14] $n_D(p)$ can be bounded uniformly (over all $p \in \Delta(K)$) from above by $c\lceil \log_2 r_D^{-1} \rceil$, where $c > 0$ is a constant (which may depend on M) and $r_D := \sup\{r > 0 : B_{\ell_1}(\pi_M, r) \subseteq \mathrm{cl}\, D\}$, where $B_{\ell_1}(\pi_M, r)$ is the ball of radius r (w.r.t. the ℓ_2-norm) centered at π_M (see, e.g., Eq. (4.34) in Section 4.5 in [28]).

This geometric property has an important behavioral implication. The sender may keep mute (i.e., disclose no private information) for $n_D(p)$ periods (until the first time the receiver's belief is in $\operatorname{cl} D$). Once the receiver's belief is in $\operatorname{cl} D$, the sender plays optimally from that time period on. Since, by Theorem 3, we can quantify the payoff for the optimal strategy at any point in $\operatorname{cl} D$ for any discount factor λ, the described strategy leads to the following theorem, providing two-sided bounds on $v_\lambda(p)$ for every discount factor λ.

Theorem 4. *Assume that $v_\infty = (\operatorname{Cav} u)(\pi_M)$ and that $\pi_M \in \operatorname{int}(\operatorname{cl} D)$. Then, for every $p \in \Delta(K)$ and every $\lambda \in [0, 1)$ we have the following bounds:*

$$(10) \quad (1-\lambda) \sum_{n=1}^{n_D(p)} \lambda^{n-1} u(pM^{n-1}) \leq v_\lambda(p) - \mathcal{I}_\lambda(p) \leq (1-\lambda) \sum_{n=1}^{n_D(p)} \lambda^{n-1} (\operatorname{Cav} u)(pM^{n-1}),$$

where

$$\mathcal{I}_\lambda(p) := \lambda^{n_D(p)} (\operatorname{Cav} u) \left((1-\lambda) p M^{n_D(p)} \left(\operatorname{Id}_k - \lambda M\right)^{-1} \right).$$

Example 4. [Example 3, continued] *Consider again the matrix M in* Eq. (8) *and recall that $\pi_M = (0.25, 0.75)$. For the function u given in Example 1 we have that $B_z = A_z$ for all $z \in \Lambda$. In words, the set of points in $\Delta(K)$ that correspond to the points where the hyperplane touches the graph of u, namely $(0, 1)$ and $(0.5, 0.5)$, is M-absorbing. Thus, $\operatorname{cl} D = D = \operatorname{conv}(A_z) = [(0, 1), (0.5, 0.5)]$. Since $(\operatorname{Cav} u)$ on D equals the linear function determined by the points $(0, 1, 0)$ and $(0.5, 0.5, 1.05)$ (in $\mathbb{R}^3$), we have that $(\operatorname{Cav} u)((p, 1-p)) = 2.05p$ for every $p \in [0, 0.5]$. A simple operation on matrices implies that*

$$(\operatorname{Id}_2 - \lambda M)^{-1} = \frac{3}{(\lambda-3)(\lambda-1)} \begin{pmatrix} 1 - 5\lambda/6 & \lambda/2 \\ \lambda/6 & 1 - \lambda/2 \end{pmatrix}.$$

This formula, coupled with Theorem 3, yields the following formula for v_λ over the range D:

$$(11) \qquad v_\lambda((p, 1-p)) = \frac{2.05\,(3(\lambda-1)p - \lambda/2)}{\lambda - 3}, \quad \forall p \in [0, 0.5].$$

As for points $(p, 1-p) \notin D$, $(p, 1-p)M \in D$, implying that $n_D((p, 1-p)) = 1$ for every $p \in (0.5, 1]$. Since $u = (\operatorname{Cav} u)$ on this range, when the initial prior is $p \in (0.5, 1]$, at the first period the sender discloses no information and the payoff is $u((p, 1-p))$. At the second period, *the belief shifts to $(p, 1-p)M$ which is in D. Now one may apply the formula (11) to obtain,*

$$v_\lambda((p,1-p)) = (1-\lambda)u((p,1-p)) + \lambda v_\lambda((p,1-p)M) = (1-\lambda)u((p,1-p)) + \lambda \frac{2.05\,((\lambda-1)p-1/2)}{\lambda-3}, \quad \forall p \in (0.5,1]. \tag{12}$$

Figure 5 illustrates the behavior of $v_\lambda((p,1-p))$ for different values of $p \in (0,1)$. Two phenomena are demonstrated here. First, as stated in Theorem 2, for every prior p the values v_λ converge to $(\mathrm{Cav}\, u)(\pi_M)$ as λ tends to 1. Second, as stated in Corollary 1, $v_\lambda(\pi_M)$ is constant across λ.

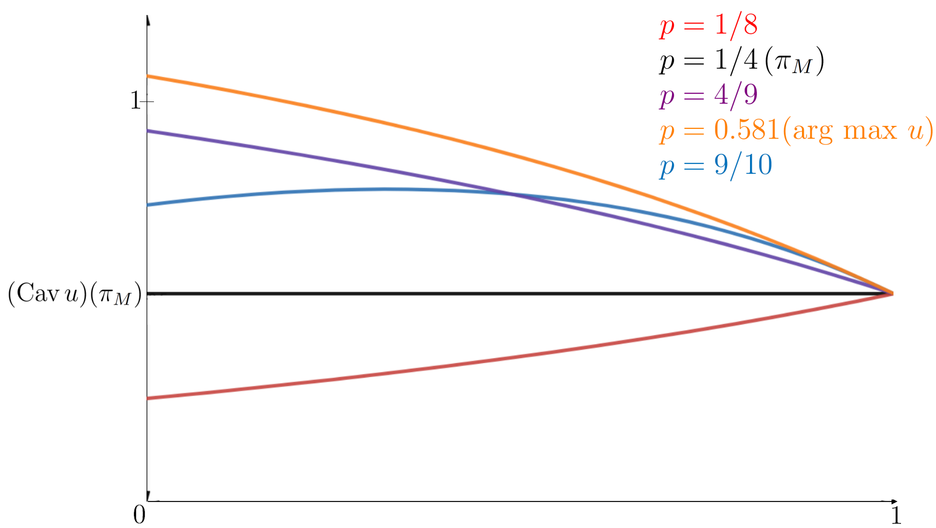

FIGURE 5. The graphs of $v_\lambda((p,1-p))$ for different values of p.

4.2. **A Strong Law.** The next result is concerned with convergence in the strong sense, namely with the $\mathbb{P}_{p,\sigma}$-almost-surely behavior of the random variable $\lim_{\lambda\to 1^-}(1-\lambda)\sum_{n=1}^{\infty} \lambda^{n-1} u(p_n)$ for a specific prior p and a strategy σ. The weak-type of convergence employed so far is concerned with the limit of the expectations. In contrast, the next theorem deals with the almost-surely convergence of the payoff, corresponding to a sender interested in the payoffs he actually obtains. These are the payoffs he truly gets along realized paths, not just their expected values.

As it turns out, finite M-absorbing sets can be used to deduce a strong law for the distribution of $\lim_{\lambda \to 1^-}(1-\lambda)\sum_{n=1}^{\infty}\lambda^{n-1}u(p_n)$ when $v_\infty = (\operatorname{Cav} u)(\pi_M)$.

Theorem 5. *Assume that A_z contains a finite M-absorbing set C for some $z \in \Lambda$. Then, there exist an M-absorbing subset $Q \subseteq C$ and a strategy $\sigma \in \Sigma$ such that for every $p \in \operatorname{conv}(Q)$ it holds that*

$$\lim_{\lambda \to 1^-}(1-\lambda)\sum_{n=1}^{\infty}\lambda^{n-1}u(p_n) = (\operatorname{Cav} u)(\pi_M), \quad \mathbb{P}_{p,\sigma}\text{-a.s.} \tag{13}$$

Moreover, if $\pi_M \in \operatorname{int}(Q)$, then Eq. (13) *holds for every $p \in \Delta(K)$.*

In words, under the assumptions of Theorem 5, for almost every infinite sequences of realizations $(x_1, x_2, ...)$ of the Markov chain $(X_n)_{n\geq 1}$, if the sender is patient enough, then by following σ he can guarantee himself a payoff close to $(\operatorname{Cav} u)(\pi_M)$.

5. When the Main Theorem Holds for every u: Homothety.

The previous results shed light on the connection between M and u. Specifically, the main theorem characterizes when $v_\infty = (\operatorname{Cav} u)(\pi_M)$ in terms of M and u. A natural question arises as to when this result holds for a fixed M and for every u. To answer this question, we need to introduce the notion of a homothety.

Definition 2. *A linear map $\psi : \mathbb{R}^k \to \mathbb{R}^k$ is said to be a* homothety *with respect to the pair $(v, \beta) \in \mathbb{R}^k \times [0, 1)$ if ψ maps each point $x \in \mathbb{R}^k$ into the point $\beta x + (1-\beta)v$. The point v is called the* center *and β is called the* ratio.

It is clear that when ψ is a homothety with respect to (v, β), the point v is a fixed point of ψ. Moreover, ψ reduces the distance from any point x to v by a factor of β. When $v \in \operatorname{int}(\Delta(K))$, or equivalently[15] $|\operatorname{supp}(v)| = k$, and $\beta \in [0, 1)$, the matrix M^ψ defined by the homothety ψ is an irreducible aperiodic stochastic matrix. In particular, the stationary distribution of M is v. Therefore, we shall say that an irreducible aperiodic stochastic matrix M is a homothety if the mapping $\phi : x \mapsto xM$ is a homothety of $\mathbb{R}^k$ with center π_M and ratio β, for some $\beta \in [0, 1)$.

When a stochastic matrix is a homothety, the transition from one state to another follows this law: each state stays unchanged with

[15] For any measure μ defined on a finite probability space, we denote by $\operatorname{supp}(\mu)$ its support, i.e., the set of elements to which μ assigns a positive probability.

probability $1-\beta$ and moves according to the distribution v to other states with probability β.

We proceed by describing an interesting class of M-absorbing sets of a homothety M. A set $E \subseteq \Delta(K)$ is said to be *star shaped around* $p \in \Delta(K)$ if $[p,q] \subseteq E$ for every $q \in E$. In words, assume that an observer is located at the point $p \in \Delta(K)$. Then E is star shaped around p if the line of sight, $[p,q]$, to any point $q \in E$ lies entirely in E. Assume now that E is star shaped around π_M. If M is a homothety, then $qM \in [\pi_M, q) \subseteq E \subseteq \text{conv}(E)$ for every $q \in E$. Hence, when M is a homothety, every star shaped set around any π_M is M-absorbing.

Let M be irreducible and aperiodic. Our next result gives a characterization of when M is a homothety in terms of v_∞. To make the result transparent, note that, by Theorem 1, v_∞ is constant on $\Delta(K)$, and as such it is simply a function of u and M. In our new characterization we let u vary over the space of all continuous functions defined on $\Delta(K)$, and so v_∞ also varies accordingly.

Theorem 6. *M is a homothety if and only if $v_\infty = (\text{Cav}\, u)(\pi_M)$ for every continuous function u.*

6. Proofs

We start this section by reviewing the notion of a *split*, a cornerstone in the field of Bayesian persuasion. This can be described informally as follows: Given a lottery X over K with law $p \in \Delta(K)$, to which extent can the sender manipulate (split) p using his signals. The answer, given by Blackwell (1951) and Aumann and Maschler (1995), is that for every choice of distributions $q_1, ..., q_{|S|} \in \Delta(K)$ and convex weights $(\alpha_i)_{i=1}^{|S|} \in \Delta(S)$ such that $\sum_{i=1}^{|S|} \alpha_i q_i = p$, the sender can correlate his lottery over signals, Y, with the lottery X, so that on the event that $s^i \in S$ is chosen (having marginal probability α_i) the posterior belief over states becomes q_i. This lottery Y will obey the rule

$$\mathbb{P}(Y = s^i \,|\, X = \ell) = \frac{\alpha_i q_i^\ell}{p^\ell}, \quad \forall i = 1, ..., |S|, \forall\, \ell \in K. \tag{14}$$

Let us denote by $\mathcal{S}_p$ the set of splits at p. Formally,

$$\mathcal{S}_p = \left\{ \{(q_i, \alpha_i)\}_{i=1}^{|S|} \,:\, q_i \in \Delta(K)\ \forall i,\ (\alpha_i)_{i=1}^{|S|} \in \Delta(S),\ \text{s.t.} \sum_{i=1}^{|S|} \alpha_i q_i = p \right\}.$$

As in [32], the dynamic decision problem faced by the sender can reformulated as a Markov decision problem (MDP). For the sake of completeness we briefly describe the structure of this MDP. The state space is $\Delta(K)$, and the action set at a state $q \in \Delta(K)$ consists of all possible splits at q, i.e., $\mathcal{S}_q$. The payoff associated with the state q and the action $\{(q_i, \alpha_i)\}_{i=1}^{|S|} \in \mathcal{S}_q$ is $\sum_{i=1}^{|S|} \alpha_i u(q_i)$.

In order to describe the transition rule, denote by y_n the state at time n, while the initial state is $y_1 = p$. Recall that the posterior belief after observing $n-1$ messages, namely, at the end of the $(n-1)$'th stage of the game $\Gamma_\lambda(p)$, is p_{n-1}. Due to the underlying Markovian dynamics, the receiver's belief at the start of stage n (before obtaining the n'th signal from the sender) is $p_{n-1}M$. We set $y_n = p_{n-1}M$. Now assume that at this stage (i.e., n) of the game the sender uses the split $\{(q_i, \alpha_i)\}_{i=1}^{|S|} \in \mathcal{S}_{y_n}$: $y_n = \sum \alpha_i q_i$. This implies that the posterior p_n (after observing also the n-th message) is equal to the result of this split: q_i with probability α_i. The belief at the start of stage $(n+1)$ is therefore $q_i M$ with probability α_i. We set $y_{n+1} = q_i M$ with probability α_i. Stated differently, the state y_n and the action $\{(q_i, \alpha_i)\}_{i=1}^{|S|} \in \mathcal{S}_{y_n}$ determine the stochastic transition to y_{n+1}: $y_{n+1} = q_i M$ with probability α_i, namely, $y_{n+1} = p_n M$.

The transition rule, together with the fact that any split $\{(q_i, \alpha_i)\}_{i=1}^{|S|}$ of a given $q \in \Delta(K)$ has mean q (i.e., $\sum_{i=1}^{|S|} \alpha_i q_i = q$) implies that the sequence of posteriors (p_n) of the receiver satisfies the following important distributional law[16]:

$$\mathbb{E}_{p,\sigma}\left(p_{n+1} \,|\, p_n\right) = p_n M, \qquad \forall n \geq 1. \tag{15}$$

Consequently, $\mathbb{E}_{p,\sigma}\, p_{n+1} = (\mathbb{E}_{p,\sigma} p_1) M^n = pM^n$ for every $n \geq 1$. In particular, if $p = \pi_M$, then $\mathbb{E}_{\pi_M,\sigma}\, p_n = \pi_M$ for every $n \geq 1$.

By reducing the problem to MDP and applying the dynamic program principle (e.g., Theorem 2.20 in [36]) we obtain the following recursive formula for $v_\lambda(p)$:

$$v_\lambda(p) = \sup_{\{(q_i,\alpha_i)\}_i \in \mathcal{S}_p} \left\{ (1-\lambda) \sum_{i=1}^{|S|} \alpha_i u(q_i) + \lambda \sum_{i=1}^{|S|} \alpha_i v_\lambda(q_i M) \right\}. \tag{16}$$

Consider the operator $\phi : \Delta(K) \to \Delta(K)$ defined by $\phi(q) = qM$. Since $|S| \geq k$, Carathéodory's Theorem (see, e.g., Corollary 17.1.5 in

[16] In the literature, when T is a mapping from a space to itself this law is referred to as a T-martingale (see, e.g., Neyman and Kohlberg [26]). An integrable sequence of random variables $(\xi_n)_{n\geq 1}$ is called a *T-martingale* if $\mathbb{E}(\xi_{n+1} \,|\, \xi_n) = T(\xi_n)$. Neyman and Kohlberg [26] provide sufficient conditions for different forms of convergence of the sequence (ξ_n/n), whenever (ξ_n) is a T-martingale. In the current context $T(q) = qM$.

[34]) implies that the expression on right-hand side of Eq. (16) equals $(\mathrm{Cav}\,\{(1-\lambda)u + \lambda v_\lambda \circ \phi\})(p)$. Thus the following key relation holds:

$$v_\lambda(p) = (\mathrm{Cav}\,\{(1-\lambda)u + \lambda\,(v_\lambda \circ \phi)\})(p). \tag{17}$$

In particular, this shows that the function $v_\lambda : \Delta(K) \to \mathbb{R}$ is concave for every λ. As ϕ is linear, $v_\lambda \circ \phi$ is also concave. Then, by the definition of Cav, we infer from Eq. (17) the inequality

$$v_\lambda(p) \le (1-\lambda)(\mathrm{Cav}\,u)(p) + \lambda(v_\lambda \circ \phi)(p). \tag{18}$$

Since the sender can always decide to not reveal any information at p, i.e., to choose the split $\{(q_i, \alpha_i)\}_i \in \mathcal{S}_p$, where $q_i = p$ for all $i = 1, ..., |S|$, and thereafter play optimally in the game $\Gamma_\lambda(pM)$, we also have that $v_\lambda(p) \ge (1-\lambda)u(p) + \lambda(v_\lambda \circ \phi)(p)$. The latter combined with Eq. (18) gives the following result:

Lemma 1. *Assume that $p \in \Delta(K)$ satisfies $u(p) = (\mathrm{Cav}\,u)(p)$. Then, for any $\lambda \in [0,1)$, the optimal signaling strategy σ_λ in $\Gamma_\lambda(p)$ would instruct the sender to reveal no information at p.*

We move on with the goal of proving Theorem 1. As it turns out, this requires classical tools and techniques from the field of repeated games with incomplete information [33]. We begin by introducing, for every $N \in \mathbb{N}$ and $p \in \Delta(K)$, the N-stage game $\Gamma_N(p)$ over the strategy space Σ with payoff given by the formula

$$\gamma_N(p,\sigma) = \mathbb{E}_{p,\sigma}\left(\frac{1}{N}\sum_{n=1}^{N} u(p_n)\right). \tag{19}$$

The value of $\Gamma_N(p)$ will be denoted by $v_N(p)$. Standard continuity and compactness based arguments (see, e.g., Theorem 2.14 on p. 15 in [36]) show that $v_N(p) = \max_{\sigma\in\Sigma} \gamma_N(p,\sigma)$.

The following proposition establishes a number of fundamental properties of $v_N(p)$ which will play an important role in our future proofs.

Proposition 4. *We have the following:*

(i) *$v_N : \Delta(K) \to \mathbb{R}$ is concave for every $N \in \mathbb{N}$.*
(ii) *For every $N \in \mathbb{N}$, the function $v_N : \Delta(K) \to \mathbb{R}$ is Lipschitz (w.r.t. the ℓ_1-norm) with constant $\|u\|_\infty$.*
(iii) *The sequence $\{N v_N(\pi_M)\}_N$ is sub-additive.*
(iv) *The sequence $\{v_N(\pi_M)\}_N$ converges.*
(v) *The sequence $\{v_{b^N}(\pi_M)\}_N$ is non-increasing for every $b \in \mathbb{N}$.*
(vi) *For every $N \in \mathbb{N}$ and every $p \in \Delta(K)$,*

$$v_{N+1}(q) = \left(\mathrm{Cav}\left\{\frac{1}{N+1}u + \frac{N}{N+1}(v_N \circ \phi)\right\}\right)(p). \tag{20}$$

Proof of Proposition 4. The proof of (i) uses the following neat classical argument. Let $q_1, q_2 \in \Delta(K)$ and $\alpha \in (0,1)$ such that $p = \alpha q_1 + (1-\alpha)q_2$. Prior to the start of $\Gamma_N(p)$, the sender writes a computer program Z which draws either the digit 1 or 2 according to a random lottery. This lottery, denoted also by Z, is dependent of X_1, and independent of $(X_n)_{n\geq 2}$. The conditional laws are given as follows:

$$\mathbb{P}(Z = 1 \,|\, X_1 = \ell) = 1 - \mathbb{P}(Z = 2 \,|\, X_1 = \ell) = \frac{\alpha q_1^\ell}{p^\ell}, \quad \forall \ell \in K. \tag{21}$$

Then, at the first stage the sender sends a messenger[17] to collect on his behalf the realized value of X_1, type it into his program Z, and then tell him the output of the program. If the outcome of Z is 1 (2), the sender plays his optimal strategy in $\Gamma(q_1)$ ($\Gamma(q_2)$). We now argue that this strategy would guarantee him $\alpha v_N(q_1) + (1-\alpha)v_N(q_2)$ in $\Gamma_N(p)$. Indeed, the latter follows by Bayes' law, which implies that the outcome of the program is 1 (resp., 2) with probability α (resp., $1-\alpha$) and that the posterior distribution of X_1 is q_1 (resp., q_2). The last argument required to depict the situation in which the sender faces $\Gamma_N(q_1)$ (resp., $\Gamma_N(q_2)$) based on the message of the messenger is that after hearing the message (1 or 2), the sender must ask the messenger to inform him of the value of X_1 that was realized.

For the proof of (ii), let us observe that by conditioning on the outcome of X_1 one has for every $\sigma \in \Sigma$ that

$$\gamma_N(p,\sigma) = \sum_{\ell \in K} p^\ell \, \mathbb{E}_{p,\sigma}\left(\frac{1}{N} \sum_{n=1}^N u(p_n) \,\middle|\, X_1 = \ell \right). \tag{22}$$

Since the distribution of the sequence (p_n) given the outcome of X_1, depends only on σ and M, and not on p, we have that the function $p \mapsto \mathbb{E}_{p,\sigma}\left(\frac{1}{N}\sum_{n=1}^N u(p_n) \,|\, X_1 = \ell\right)$ is constant for any $\ell \in K$. Thus, by the triangle inequality, for every $p, q \in \Delta(K)$ and $\sigma \in \Sigma$ it holds that

$$\begin{aligned} |\gamma_N(p,\sigma) - \gamma_N(q,\sigma)| &\leq \sum_{\ell \in K} |p^\ell - q^\ell| \left| \mathbb{E}_{p,\sigma}\left(\frac{1}{N} \sum_{n=1}^N u(p_n) \,\middle|\, X_1 = \ell \right) \right| \\ &\leq \|u\|_\infty \sum_{\ell \in K} \|p - q\|_1. \end{aligned} \tag{23}$$

As $\Gamma_N(p)$ can be viewed as a finite game in extensive form, it admits a normal-form game description, and thus (ii) is a consequence from the

[17] The messenger cannot be the receiver.

following basic inequality: for every two zero-sum matrix games A and B of equal dimensions, $|\mathrm{val}(A) - \mathrm{val}(B)| \leq \|A - B\|_\infty$.

We move on to (iii). For every $\sigma \in \Sigma$ and every $N, L \in \mathbb{N}$, the $(N + L)$'th stage game payoff $\gamma_{N+L}(\pi_M, \sigma)$ equals

$$\frac{N}{N+L}\,\mathbb{E}_{\pi_M,\sigma}\left[\frac{1}{N}\sum_{n=1}^{N} u(p_n)\right] + \frac{L}{N+L}\,\mathbb{E}_{\pi_M,\sigma}\left[\frac{1}{L}\sum_{n=N+1}^{N+L} u(p_n)\right]. \tag{24}$$

As the belief of the receiver at the start of the $(N + 1)$'st time period, prior to obtaining the signal s_{N+1}, equals $p_N M$, we may bound the latter from above by

$$\frac{N}{N+L} v_N(\pi_M) + \frac{L}{N+L}\mathbb{E}_{\pi_M,\sigma}\left[v_L(p_N M)\right]. \tag{25}$$

By Jensen's inequality applied for the functions $\{v_N\}$, which are known to be concave by (i), in conjunction with the fact that at the initial belief π_M, all the steps of the sequence (p_n) have expectation π_M, we obtain that the latter is at most

$$\frac{N}{N+L} v_N(\pi_M) + \frac{L}{N+L} v_L(\mathbb{E}_{\pi_M,\sigma}(p_N M)) = N v_N(\pi_M) + L v_L(\pi_M). \tag{26}$$

Since this holds for every σ, we conclude that $(N + L)v_{N+L}(\pi_M) \leq N v_N(\pi_M) + L v_L(\pi_M)$, which completes the proof of (iii).

The proof of (iv) can be deduced directly from (iii) based on a basic result from analysis which states that if $\{a_n\}$ is a sub-additive sequence, then $\{a_n/n\}$ converges. Next, by a repeated use of (iii) we see that if L divides N, then

$$N v_N(\pi_M) \leq (\frac{N}{L} - 1) v_{\frac{N}{L}-1}(\pi_M) + L v_L(\pi_M) \leq \cdots \leq \frac{N}{L}(L v_L(\pi_M)). \tag{27}$$

Therefore, $v_N(\pi_M) \leq v_L(\pi_M)$, which is sufficient to prove (v). Lastly, by the dynamical programming principle for Markov decision problems (e.g., Theorem 2.17 in [36]) we have for any $N \in \mathbb{N}$ and any $p \in \Delta(K)$:

$$v_{N+1}(p) = \sup_{\{(q_i,\alpha_i)\}_{i \in \mathcal{S}_p}} \left\{ \frac{1}{N+1}\sum_{i=1}^{|S|} \alpha_i u(q_i) + \frac{N}{N+1}\sum_{i=1}^{|S|} \alpha_i v_N(q_i M) \right\}. \tag{28}$$

As $|S| \geq k$, Carathéodory's Theorem (see, e.g., Corollary 17.1.5 in [34]) shows that the expression on the right-hand side of Eq. (16) equals $\left(\mathrm{Cav}\left\{\frac{1}{N+1}u + \frac{N}{N+1}v_N \circ \phi\right\}\right)(p)$, thus proving item (vi). Note that item (vi) gives an alternative (partial) proof of item (i), as it shows that $v_N(\cdot)$ is concave for any $N \geq 2$. $\square$

We are now ready to prove Theorem 1.

Proof of Theorem 1. By items (i) and (vi) of Proposition 4 and the definition of the Cav operator, for any $N \geq 1$ and $p \in \Delta(K)$ we have that

$$\frac{1}{N+1} \min_{\Delta(K)} u + \frac{1}{N+1} (v_N \circ \phi)(p) \leq v_{N+1}(p) \leq \frac{1}{N+1} \max_{\Delta(K)} u + \frac{N}{N+1} (v_N \circ \phi)(p). \tag{29}$$

Set $\underline{v}(p) := \liminf_{n\to\infty} v_n(p)$ and $\bar{v}(p) := \limsup_{n\to\infty} v_n(p)$. Eq. (29) implies that $\underline{v}(p) = \underline{v}(pM)$ and $\bar{v}(p) = \bar{v}(pM)$ for every $p \in \Delta(K)$. Moreover, by item (ii) of Proposition 4, both $\underline{v}$ and $\bar{v}$ must be Lipschitz (w.r.t. the ℓ_1-norm) with the constant $\|u\|_\infty$. In particular these functions are continuous at π_M. Hence, by Convergence Theorem for Markov Chains (e.g., Theorem 4.9 in [28]), $pM^n \to \pi_M$. We obtain that

$$\underline{v}(p) = \underline{v}(pM) = \cdots = \underline{v}(pM^n) \to \underline{v}(\pi_M) \quad \text{as} \quad n \to \infty, \tag{30}$$

implying that $\underline{v}$ is constant on $\Delta(K)$. Similarly, it follows that $\bar{v}$ is constant on $\Delta(K)$. Since by item (iv) of Proposition 4 $\underline{v}(\pi_M) = \bar{v}(\pi_M)$, we obtain that $\bar{v} = \underline{v} = \underline{v}(\pi_M)$. We conclude that for every $p \in \Delta(K)$, $v_N(p) \to v_\infty$ as $N \to \infty$, where $v_\infty = \underline{v}(\pi_M)$.

We claim that $\{v_N(\cdot)\}$ converges uniformly to v_∞. Indeed, take a finite ε-net $\{q_1, ..., q_J\}$ in $\Delta(K)$ w.r.t. the ℓ_1-norm. In other words, for every $p \in \Delta(K)$ there exists a $j(p) \in \{1, ..., J\}$ such that $\|p - q_{j(p)}\|_1 \leq \varepsilon$. Since $\{v_N(\cdot)\}$ converge to v_∞ and $\mathbb{R}$ is complete, there exists an $R \in \mathbb{N}$ such that $|v_L(q_j) - v_N(q_j)| \leq \varepsilon$ for every $j = 1, ..., J$, and every $N, L \geq R$. By employing item (ii) of Proposition 4 we obtain that for every $N, L \geq R$ and every $p \in \Delta(K)$ it holds

$$\begin{aligned} |v_L(p) - v_N(p)| &\leq |v_L(p) - v_L(q_{j(p)})| \\ &\quad + |v_L(q_{j(p)}) - v_N(q_{j(p)})| + |v_N(q_{j(p)}) - v_N(p)| \\ &\leq \|u\|_\infty \varepsilon + \varepsilon + \|u\|_\infty \varepsilon. \end{aligned} \tag{31}$$

Hence, letting $L \to \infty$ we obtain that $|v_N(p) - v_\infty| \leq (1 + 4\|u\|_\infty)\varepsilon$ for each $p \in \Delta(K)$ and every $N \geq R$, which proves our claim. By a uniform Tauberian theorem for Markov decision problems over Borel state spaces (e.g., Theorem 1 and discussion in Section 6 in Lehrer and Sorin [23]), we get that $v_\lambda(\cdot)$ converges to v_∞ uniformly as $\lambda \to 1^-$.

Finally, we show that $v_\infty \leq (\operatorname{Cav} u)(\pi_M)$. Indeed, for any $\sigma \in \Sigma$, Jensen's inequality shows that the expected payoff at any time period

n satisfies

$$\mathbb{E}_{\pi_M,\sigma} u(p_n) \leq \mathbb{E}_{\pi_M,\sigma} (\operatorname{Cav} u)(p_n) \leq (\operatorname{Cav} u)(\mathbb{E}_{\pi_M,\sigma}\, p_n) = (\operatorname{Cav} u)(\pi_M), \tag{32}$$

implying that $v_\lambda(\pi_M) \leq (\operatorname{Cav} u)(\pi_M)$ for every $\lambda \in [0,1)$. Hence, $v_\infty \leq (\operatorname{Cav} u)(\pi_M)$, as desired. □

As the first step towards the proof of Proposition 2 and Theorems 3 and 4, we show the following basic lemma.

Lemma 4. *For every $z \in \Lambda$ we have*

(i) $\operatorname{Cav} u(q) = f_z(q)$ *for every* $q \in \operatorname{conv}(A_z)$.
(ii) *Let* $(\alpha_i)_{i=1}^m$, $\alpha_i > 0$ *for every* i, $\sum_i \alpha_i = 1$ *and* $(q_i)_{i=1}^m \in \Delta(K)$ *such that* $\sum_i \alpha_i q_i = \pi_M$. *If* $\sum_i \alpha_i u(q_i) = (\operatorname{Cav} u)(\pi_M)$, *then* $q_i \in A_z$ *for every* i.

Proof of Lemma 4. Let $q \in \operatorname{conv}(A_z)$. Take $(q_i) \in A_z$ and convex weights (α_i) such that $q = \sum_i \alpha_i q_i$. Since $\operatorname{Cav} u$ is concave and $q_i \in A_z$, we have

$$(\operatorname{Cav} u)(q) \geq \sum_i \alpha_i (\operatorname{Cav} u)(q_i) \geq \sum_i \alpha_i u(q_i) = \sum_i \alpha_i f_z(q_i) = f_z(q), \tag{33}$$

where the last equality follows from the fact that f_z is affine. Since by the definition of Λ we have $(\operatorname{Cav} u)(q) \leq f_z(q)$ for every $q \in \Delta(K)$, we have shown (i). For (ii) assume that there exists $q_{i_0} \notin A_z$. Then $u(q_{i_0}) < f_z(q_{i_0})$, and since $\alpha_{i_0} > 0$ and $z \in \Lambda$, we have

$$(\operatorname{Cav} u)(\pi_M) = \sum_i \alpha_i u(q_i) < \sum_i \alpha_i f_z(q_i) = f_z(\pi_M) = (\operatorname{Cav} u)(\pi_M). \tag{34}$$

We reached a contradiction. □

Next we prove the following proposition, which demonstrates the special advantages of M-absorbing subsets of A_z for $z \in \Lambda$.

Proposition 5. *Assume that for some $z \in \Lambda$, the set A_z admits an M-absorbing subset C. Then,*

$$v_\lambda(p) = (\operatorname{Cav} u)\left(\lambda p \left(\mathrm{Id}_k - \lambda M\right)^{-1}\right), \quad \forall \lambda \in [0,1), \tag{35}$$

for every $p \in \operatorname{conv}(C)$, where Id_k is the k-dimensional identity matrix.

Proof of Proposition 5. We show that Eq. (35) holds via a two-sided inequality. First, by applying the Jensen's inequality in the same fashion as in Eq. (32), we get that, for every $\sigma \in \Sigma$ and every $\lambda \in [0,1)$,

$$\begin{aligned}\gamma_\lambda(p,\sigma) &\leq (1-\lambda)\sum_{n=1}^{\infty}\lambda^{n-1}(\operatorname{Cav} u)(\mathbb{E}_{p,\sigma}\, p_n)\\ &= (1-\lambda)\sum_{n=1}^{\infty}\lambda^{n-1}(\operatorname{Cav} u)(pM^{n-1})\\ &\leq (\operatorname{Cav} u)\left((1-\lambda)\sum_{n=1}^{\infty}\lambda^{n-1}pM^{n-1}\right),\end{aligned} \tag{36}$$

where the last inequality is due to the concavity of $(\operatorname{Cav} u)$. Next, since M is a stochastic matrix, we have the following well-know formula (e.g., Theorem 2.29 in [36]):

$$(1-\lambda)\sum_{n=1}^{\infty}\lambda^{n-1}pM^{n-1} = \lambda p\left(\mathrm{Id}_k - (1-\lambda)M\right)^{-1}, \quad \forall \lambda \in [0,1). \tag{37}$$

A combination of Eqs. (36) and (37) yields that $v_\lambda(p)$ is at most $(\operatorname{Cav} u)\left((1-\lambda)p\left(\mathrm{Id}_k - \lambda M\right)^{-1}\right)$ for every $p \in \operatorname{conv}(C)$ and every $\lambda \in [0,1)$.

Let us now show that the opposite inequality holds as well. We start by defining for each $q \in \Delta(K)$ the set $\mathcal{S}_q^C \subseteq \mathcal{S}_q$ by

$$\mathcal{S}_q^C := \{\{(q_i,\alpha_i)\}_{i=1}^{|S|} \,:\, q_i \in C \;\; \forall i = 1,...,|S|\}. \tag{38}$$

Since $|S| \geq k$, Carathéodory's Theorem shows that $\mathcal{S}_q^C \neq \emptyset$ whenever $q \in \operatorname{conv}(C)$. Consider now the strategy σ^C defined as follows: at each $n \geq 1$, if $p_{n-1} = q \in \operatorname{conv}(C)$, σ^C will chose an element in $\mathcal{S}_{qM}^C$; otherwise, if $p_{n-1} = q \in \Delta(K) \setminus \operatorname{conv}(C)$, then σ^C will chose some element in $\mathcal{S}_q$. As $p \in \operatorname{conv}(C)$, and $\operatorname{conv}(C)$ is M-absorbing, we have that under the strategy σ^C, $\operatorname{supp}(p_n) \subseteq C$ for every $n \geq 1$. Indeed, we show this by induction on n. For $n = 1$, since $p \in \operatorname{conv}(C)$, $\operatorname{supp}(p_1) \subseteq C$ by the definition of σ^C. Assume now that $\operatorname{supp}(p_n) \subseteq C$ for some $n \geq 1$. Since C is M-absorbing, $\operatorname{supp}(p_n M) \subseteq \operatorname{conv}(C)$, and thus by the definition of σ^C we see that $\operatorname{supp}(p_{n+1}) \subseteq C$ as well. The latter, coupled with $C \subseteq A_z$ implies that the discounted payoff under

σ^C can be computed as follows:

$$
\begin{aligned}
\gamma_\lambda(p,\sigma^C) &= (1-\lambda)\sum_{n=1}^{\infty}\lambda^{n-1}\mathbb{E}_{p,\sigma^C}[u(p_n)] \\
&= (1-\lambda)\sum_{n=1}^{\infty}\lambda^{n-1}\mathbb{E}_{p,\sigma^C}\left[f_z(p_n)\right] \\
&= (1-\lambda)\sum_{n=1}^{\infty}\lambda^{n-1}f_z\left(\mathbb{E}_{p,\sigma^C}\, p_n\right) \\
&= (1-\lambda)\sum_{n=1}^{\infty}\lambda^{n-1}f_z\left(pM^{n-1}\right) \\
&= f_z\left((1-\lambda)\sum_{n=1}^{\infty}\lambda^{n-1}pM^{n-1}\right) \\
&= (\operatorname{Cav} u)\left((1-\lambda)\sum_{n=1}^{\infty}\lambda^{n-1}pM^{n-1}\right).
\end{aligned}
\tag{39}
$$

where we note that the third and fifth equalities hold because f_z is affine, and the last equality is a consequence of item (i) of Lemma 4. Hence, in view of (37), we have shown that $(\operatorname{Cav} u)\left((1-\lambda)p\left(\mathrm{Id}_k-\lambda M\right)^{-1}\right)$ is not greater than $v_\lambda(p)$, as needed. □

We proceed with the proof of Proposition 2.

Proof of Proposition 2. Assume that C is an M-absorbing subset of A_z where $z\in\Lambda$. Let $q\in\operatorname{conv}(C)$. First, by Proposition 5,

$$v_\lambda(q) = (\operatorname{Cav} u)\left((1-\lambda)q\left(\mathrm{Id}_k-\lambda M\right)^{-1}\right), \quad \forall\lambda\in[0,1).$$

Next, since $qM^n\to\pi_M$ as $n\to\infty$, we have that $(1/n)\sum_{\ell=1}^{n} qM^{\ell-1}\to\pi_M$ as $n\to\infty$. Thus, by a Tauberian Theorem (see, e.g., Theorem 3.1. in [36]) we obtain that $(1-\lambda)\sum_{n=1}^{\infty}\lambda^{n-1}qM^{n-1}\to\pi_M$ as $\lambda\to1^-$. Hence, from the identity (37) we obtain that $(1-\lambda)q\left(\mathrm{Id}_k-\lambda M\right)^{-1}\to\pi_M$ as $\lambda\to1^-$. Moreover, since $(\operatorname{Cav} u)$ is continuous on $\Delta(K)$, it must be continuous at π_M. A combination of the above arguments with Theorem 1 yields

$$
\begin{aligned}
v_\infty = \lim_{\lambda\to1^-} v_\lambda(q) = \lim_{\lambda\to1^-}(\operatorname{Cav} u)\left((1-\lambda)q\left(\mathrm{Id}_k-\lambda M\right)^{-1}\right) \\
= (\operatorname{Cav} u)(\pi_M),
\end{aligned}
\tag{40}
$$

thus proving item (i) of Proposition 2.

Let us continue with item (ii). Assume that $v_\infty=(\operatorname{Cav} u)(\pi_M)$. We have *(i)* $\lim_{N\to\infty} v_N(\pi_M)=(\operatorname{Cav} u)(\pi_M)$, *(ii)* $v_N(\pi_M)\le(\operatorname{Cav} u)(\pi_M)$

for every N (by Eq. (32)) and *(iii)* $\{v_{b^N}(\pi_M)\}_N$ is non-increasing for every $b \in \mathbb{N}$ (by item (v) of Proposition 4). A combination of *(i)*, *(ii)*, and *(iii)* shows that $v_N(\pi_M) = (\operatorname{Cav} u)(\pi_M)$ for every N. Let σ^N be an optimal strategy in $\Gamma_N(\pi_M)$. Denote by $(p_n^N)_n$ the sequence of posteriors induced by σ^λ and the prior probability π_M. By Jensen's inequality, $\mathbb{E}_{\pi_M,\sigma^N}[u(p_n^N)] \leq (\operatorname{Cav} u)(\pi_M)$ for every n. Hence, as $\gamma_N(\pi_M, \sigma^N) = (\operatorname{Cav} u)(\pi_M)$, we obtain that $\mathbb{E}_{\pi_M,\sigma^N}[u(p_n^N)] = (\operatorname{Cav} u)(\pi_M)$ for every $n = 1, ..., N$. Fix $\lambda \in (0,1)$. We see that

$$\gamma_\lambda(\pi_M, \sigma^N) \geq (1-\lambda)\sum_{n=1}^{N} \lambda^{n-1}(\operatorname{Cav} u)(\pi_M) - \lambda^N \|u\|_\infty \tag{41}$$

for every $N \geq 1$. Letting $N \to \infty$ we get $v_\lambda(\pi_M) \geq (\operatorname{Cav} u)(\pi_M)$. Since by Eq. (32) the opposite inequality holds as well, we deduce that $v_\lambda(\pi_M) = (\operatorname{Cav} u)(\pi_M)$. Therefore, there exists a strategy σ^λ such that $\gamma_\lambda(\pi_M, \sigma^\lambda) = (\operatorname{Cav} u)(\pi_M)$. Denote the sequence of posteriors induced by σ^λ and the prior probability π_M by $(p_n^\lambda)_n$. By Jensen's inequality, $\mathbb{E}_{\pi_M,\sigma^\lambda}[u(p_n^\lambda)] \leq (\operatorname{Cav} u)(\pi_M)$, and we therefore obtain that $\mathbb{E}_{\pi_M,\sigma^\lambda}[u(p_n^\lambda)] = (\operatorname{Cav} u)(\pi_M)$ for every n. Moreover, as $\operatorname{supp}(p_n^\lambda)$ is finite and $\mathbb{E}_{\pi_M,\sigma^\lambda} p_n^\lambda = \pi_M$ for every n, item (ii) of Lemma 4 implies that $\operatorname{supp}(p_n^\lambda) \subseteq A_z$ for every $z \in \Lambda$ and every n. Set $C := \bigcup_{n\geq 1} \operatorname{supp}(p_n^\lambda)$. We have $C \subseteq A_z$ for every $z \in \Lambda$. Moreover, as the set of signals S of the receiver is finite, C is a countable union of finite sets, and thus is countable.

We claim that C is M-absorbing. Indeed, if $q \in C$, then there exists an n such that $p_n^\lambda = q$ with positive probability. Since $\mathbb{E}_{\pi_M,\sigma^\lambda}(p_{n+1} \mid p_n^\lambda = q) = qM$, we obtain that $qM \in \operatorname{conv}(\operatorname{supp}(p_{n+1}^\lambda)) \subseteq \operatorname{conv}(C)$. To summarize, C is a countable M-absorbing subset of A_z for every $z \in \Lambda$, as desired. □

We can now complete the proofs of Theorems 3 and 4.

Proofs of Theorems 3 and 4. By Proposition 2, the sets B_z, $z \in \Lambda$, are well defined (i.e., non-empty). As they are M-absorbing, we can apply Proposition 5 to any point $p \in \operatorname{conv}(B_z)$ to get the result of Theorem 3 for any $p \in D$. The result extends to any $p \in \operatorname{cl} D$ by the continuity of the functions[18] $v_\lambda(\cdot)$, $(\operatorname{Cav} u)$, and $q \mapsto (1-\lambda)q\,(\mathrm{Id}_k - \lambda M)^{-1}$.

As for the proof of Theorem 4, if the sender does not use his private information along the first $n_D(p)$ time periods, and then plays optimally

[18] The functions $v_\lambda(\cdot)$ are continuous for every λ as one might prove in similar fashion to the proof of item (ii) of Proposition 4 that $v_\lambda(\cdot)$ is Lipschitz (w.r.t. the ℓ_1-norm) with constant $\|u\|_\infty$.

from time $n_D(p)+1$ and on, he guarantees

$$(1-\lambda)\sum_{n=1}^{n_D(p)} \lambda^{n-1} u(pM^{n-1}) + \lambda^{n_D(p)} v_\lambda(pM^{n_D(p)}).$$

As $pM^{n_D(p)} \in \operatorname{cl} D$, Theorem 3 ensures that $\lambda^{n_D(p)} v_\lambda(pM^{n_D(p)}) = \mathcal{I}_\lambda(p)$, thus proving the left hand side of Eq. (10). On the other hand, by using Jensen's inequality for both $\operatorname{Cav} u$ and v_λ, we see that for every $\sigma \in \Sigma$ and $p \in \Delta(K)$

$$\gamma_\lambda(p,\sigma) \leq (1-\lambda)\sum_{n=1}^{n_D(p)} \lambda^{n-1} \operatorname{Cav} u(pM^{n-1}) + \lambda^{n_D(p)} v_\lambda(pM^{n_D(p)}),$$

and thus, since $\lambda^{n_D(p)} v_\lambda(pM^{n_D(p)}) = \mathcal{I}_\lambda(p)$ by Theorem 3, we obtain the right-hand side of Eq. (10). □

Proof of Corollary 1. By Proposition 2, the sets B_z, $z \in \Lambda$, are well defined (i.e., non-empty). As they are M-absorbing, $\pi_M \in \operatorname{cl}\operatorname{conv}(B_z) \subseteq \operatorname{cl} D$ for any $z \in \Lambda$. Thus, the result follows from Theorem 3 and the fact that, by Eq. (37),

$$\lambda \pi_M (\mathrm{Id}_k - (1-\lambda)M)^{-1} = (1-\lambda)\sum_{n=1}^{\infty} \lambda^{n-1} \pi_M M^{n-1} = \pi_M.$$

□

We move on to Theorem 5.

Proof of Theorem 5. Let $C = \{q_1, ..., q_r\}$. Since C is M-absorbing, we can assign for each $i = 1, ..., r$ a distribution $\alpha^i \in \Delta(C)$ so that $q_i M = \sum_{j=1}^{r} \alpha^i_j q_j$. Moreover, by Carathéodory's Theorem, we can choose α^i so that $|\operatorname{supp}(\alpha^i)| \leq k$ for every i. Define the $r \times r$ matrix W by $W_{i,j} := \alpha^i_j$; W is a stochastic matrix. Also, define the $r \times k$ matrix P by $P_{i,\ell} := q^\ell_i$, where $i \in \{1, ..., r\}$ and $\ell \in K$. We have the following algebraic relation:

$$PM = WP. \tag{42}$$

Let $\mathcal{R}$ be a communicating class of W whose states are recurrent (see, e.g., Lemma 1.26 in [28]). Denote by $W_\mathcal{R}$ the restriction of the matrix W to the set of states $\ell \in \mathcal{R}$. The $|\mathcal{R}| \times |\mathcal{R}|$ matrix $W_\mathcal{R}$ is clearly stochastic. Moreover, since $\mathcal{R}$ is a communication class, $W_\mathcal{R}$ is also irreducible. Next, let us denote by $P_\mathcal{R}$ the $k \times |\mathcal{R}|$ matrix with entries $(P_\mathcal{R})_{i,\ell} := q^\ell_i$, where $i \in \mathcal{R}$ and $\ell \in K$. It follows from Eq. (42) that

$$P_\mathcal{R} M = W_\mathcal{R} P_\mathcal{R}. \tag{43}$$

As $W_{\mathcal{R}}$ is stochastic, Eq. (43) thus implies that the set $Q = \{q_i\}_{i\in\mathcal{R}}$ is M-absorbing.

Consider the following strategy $\sigma \in \Sigma$. If $p \in \text{conv}(Q)$, split p into k elements of Q according to some prior $\mu_p \in \Delta(Q)$. Otherwise, ignore private information and send the fixed signal $s_0 \in S$. Assume that $p_n = q$ for some $n \geq 1$. If $q \notin \text{conv}(Q)$, the sender ignores his private information and sends the fixed signal $s_0 \in S$. Next, if $q \in \text{conv}(Q) \setminus Q$, σ instructs the sender to choose a split from $\mathcal{S}^Q_{qM}$ (see Eq. (38) for the definition of $\mathcal{S}^Q_{qM}$).

Finally, if $q = q^j \in Q$, σ instructs the sender to split qM into $\{(q_i, w^j_i)\}_{i\in\mathcal{R}}$, where $(w^j_1, ..., w^j_{|\mathcal{R}|})$ is the j'th row of $W_{\mathcal{R}}$. Note that such a split is available to the sender because each row in W contains at most k non-zero elements, and $k \leq |S|$.

It follows from the definition of σ that for each $p \in \text{conv}(Q)$, the sequence of posteriors (p_n) follows a Markov chain over the state space Q with initial probability μ_p and transition rule given by the stochastic matrix $W_{\mathcal{R}}$. Since the latter is irreducible, we may employ the Ergodic Theorem for Markov chains (e.g., Theorem C.1 in [28]) to obtain that

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} u(p_n) = \sum_{i=1}^{|\mathcal{R}|} \nu_i u(q_i), \quad \mathbb{P}_{p,\sigma}\text{-a.s.} \tag{44}$$

where $\nu = (\nu_i)_{i=1}^{|\mathcal{R}|}$ is the unique stationary distribution of $W_{\mathcal{R}}$. Furthermore, since $Q \subseteq A_z$, we have that

$$\sum_{i=1}^{|\mathcal{R}|} \nu_i u(q_i) = \sum_{i=1}^{|\mathcal{R}|} \nu_i f_z(q_i) = f_z\left(\sum_{i=1}^{|\mathcal{R}|} \nu_i q_i\right) = (\text{Cav}\, u)(\sum_{i=1}^{|\mathcal{R}|} \nu_i q_i), \tag{49}$$

where the last equality follows from item (i) of Lemma 4. Next, by multiplying by ν from the left both sides of Eq. (43) we obtain

$$\nu P_{\mathcal{R}} M = \nu W_{\mathcal{R}} P_{\mathcal{R}} = \nu P_{\mathcal{R}}, \tag{45}$$

which together with the uniqueness of π_M implies that $\nu P_{\mathcal{R}} = \pi_M$. However, as $\pi_M = \nu P_{\mathcal{R}} = \sum_{i=1}^{|\mathcal{R}|} \nu_i q_i$, we deduce that

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} u(p_n) = (\text{Cav}\, u)(\pi_M), \quad \mathbb{P}_{p,\sigma}\text{-a.s.} \tag{46}$$

Therefore, by a Tauberian theorem (see, e.g., Theorem 3.1. in [36]),

$$\lim_{\lambda\to 1^-} (1-\lambda) \sum_{n=1}^{\infty} \lambda^{n-1} u(p_n) = (\text{Cav}\, u)(\pi_M), \quad \mathbb{P}_{p,\sigma}\text{-a.s.} \tag{47}$$

Finally, assume that $\pi_M \in \operatorname{int}(\operatorname{conv}(Q))$. Then, since M is irreducible and aperiodic, $pM^N \in \operatorname{conv}(Q)$ for some finite time period N for every $p \in \Delta(K)$. Then, by the definition of σ, we see that Eq. (46) holds for every $p \in \Delta(K)$ and thus so does Eq. (47). □

The proof of Proposition 1 may enhance the intuition about absorbing sets.

Proof of Proposition 1. By Carathéodory's Theorem (see, e.g., Corollary 17.1.5 in [34]), to each $q \in C$ we can assign k distributions $w_1(q), ..., w_k(q) \in C$ such that $qM \in \operatorname{conv}(\{w_1(q), ..., w_k(q)\}$. Define a correspondence $\xi : C \to 2^C$ by $\xi(q) = \{w_1(q), ..., w_n(q)\}$. In particular, $qM \in \xi(q)$. The countable set $\mathcal{A}(q) := \bigcup_{n=1}^{\infty} \xi^{n-1}(q)$ is M-absorbing for every $q \in C$, where ξ^{n-1} is the $(n-1)$-fold composition of ξ with itself. Indeed, let $w \in \mathcal{A}(q)$, and let $n \geq 1$ be such that $w \in \xi^{n-1}(q)$. By the definition of ξ we have that $wM \in \operatorname{conv}(\xi(w)) \subseteq \operatorname{conv}(\xi^n(q)) \subseteq \operatorname{conv}(\mathcal{A}(q))$, as desired. □

We end the proof section with the proof of Theorem 6.

Proof of Theorem 6. Suppose that M is a homothety and let us fix[19] $u \in C(\Delta(K))$. Since u is continuous, Carathéodory's Theorem (see, e.g., Corollary 17.1.5 in [34]) implies that there exist points $q_1, ..., q_m \in \Delta(K)$, $m \leq k$, and positive convex weights $(\alpha_i)_{i=1}^m$ such that $\pi_M = \sum_{i=1}^m \alpha_i q_i$ and $(\operatorname{Cav} u)(\pi_M) = \sum_{i=1}^m \alpha_i u(q_i)$. Hence, by item (ii) of Lemma 4, $q_i \in A_z$ for every i and every $z \in \Lambda$. Therefore, $\pi_M \in \operatorname{conv}(A_z)$ for every $z \in \Lambda$, and since $\operatorname{conv}(A_z)$ is convex, we see that $\operatorname{conv}(A_z)$ is star shaped around π_M. Hence, since M is a homothety we get that $\operatorname{conv}(A_z)$ is M-absorbing for any $z \in \Lambda$. By the definition of an M-absorbing set, A_z must also be M-absorbing for any $z \in \Lambda$. By Proposition 2, we deduce that $v_\infty = (\operatorname{Cav} u)(\pi_M)$, proving the first direction of Theorem 6.

Suppose now that $v_\infty = (\operatorname{Cav} u)(\pi_M)$ for every $u \in C(\Delta(K))$. For each $i \in K$, let $e_i \in \Delta(K)$ be the Dirac measure concentrated on the i'th coordinate of $\mathbb{R}^k$. Fix $i \in K$ and consider for each $n \geq 1$ the vector $e_i^n = \pi_M + (\pi_M - e_i)/n$. Clearly, $\pi_M \in [e_i, e_i^n]$ for all n. Next, as $\pi_M \in \operatorname{int}(\Delta(K))$, there exists N_i such that $e_i^n \in \Delta(K)$ for every $n \geq N_i$. For each $n \geq N_i$ we define $u_{i,n} : \Delta(K) \to \mathbb{R}$ by

$$u_{i,n}(q) := 1 - \max\{\|q - e_i\|_2, \|q - e_i^n\|_2\}. \tag{48}$$

Clearly, $u_{i,n} \in C(\Delta(K))$. By its definition we have that $u_{i,n}(q) = 1$ for $q \in \{e_i, e_i^n\}$ and $u_{i,n}(q) < 1$ for $q \in \Delta(K) \setminus \{e_i, e_i^n\}$. Next, we

[19] We denote by $C(\Delta(K))$ the space of continuous real-valued functions defined on $\Delta(K)$.

have that $(\mathrm{Cav}\, u)(\pi_M) = 1$ and that $0 \in \Lambda$ (for Λ corresponding to $u = u_{i,n}$), because the hyperplane $f_0(x) = (\mathrm{Cav}\, u)(\pi_M) = 1$ for all $x \in \mathbb{R}^k$, supports $(\mathrm{Cav}\, u)$ at π_M. As $A_0 = \{e_i, e_i^n\}$, Proposition 2 shows that $\{e_i, e_i^n\}$ contains an M-absorbing subset. However, as M has a unique stationary distribution $\pi_M \notin \{e_i, e_i^n\}$, we get that neither $\{e_i\}$ nor $\{e_i^n\}$ is M-absorbing. Therefore, $\{e_i, e_i^n\}$ must be M-absorbing for every $n \geq N_i$. In particular, $e_i M \in (e_i, e_i^n]$ for every $n \geq N_i$. Since $e_i^n \to \pi_M$ as $n \to \infty$, we obtain that $e_i M \in (e_i, \pi_M]$. Thus, as i was arbitrary, we have shown that for each $i \in K$ there exists $\beta_i \in [0,1)$ such that $e_i M = \beta_i e_i + (1-\beta_i)\pi_M$. Since $q \to Mq$ is a linear operator, to prove that M is a homothety is suffices to show that $\beta_i = \beta_j$ for all $i \neq j \in K$.

The proof now bifurcates according to the dimension of $\Delta(K)$. First, let us assume that $k = 2$. Since M is irreducible, there exists a unique $\alpha \in (0,1)$ such that $\pi_M = \alpha e_1 + (1-\alpha)e_2$. We have

$$\begin{aligned} \pi_M M &= (\alpha e_1 + (1-\alpha)e_2)M \\ &= \alpha(e_1 M) + (1-\alpha)(e_2 M) \\ &= \alpha(\beta_1 e_1 + (1-\beta_1)\pi_M) + (1-\alpha)(\beta_2 e_2 + (1-\beta_2)\pi_M). \end{aligned} \tag{49}$$

By plugging $\pi_M = \alpha e_q + (1-\alpha)e_2$ into the last expression in Eq. (49) and using simple algebraic manipulations, we get that the convex weight ρ of e_1 in the convex decomposition of $\pi_M M$ with respect to e_1 and e_2 equals

$$\alpha\beta_1 + \alpha^2(1-\beta_1) + (1-\alpha)(1-\beta_2)\alpha.$$

However, as $\pi_M = \pi_M M$, we must have that $\rho = \alpha$. After some further simple algebraic manipulations, one gets that the equality $\rho = \alpha$ is equivalent to $\beta_1 - \beta_2 = \alpha(\beta_1 - \beta_2)$. As $\alpha \in (0,1)$, we obtain that $\beta_1 = \beta_2$, thus proving that M is a homothety whenever $k = 2$.

Next, let $k \geq 3$. Assume that $\beta_i < \beta_j$ for some $i \neq j \in K$. Define $v = (e_i + e_j)/2$. Since $|\mathrm{supp}(v)| = 2$, whereas $|\mathrm{supp}(\pi_M)| = k \geq 3$, we have that $\pi_M \neq v$. Consider for each $n \in \mathbb{N}$ the vector $v_n = \pi_M + (\pi_M - v)/n$. Then $\pi_M \in [v, v_n]$ for every n. Moreover, since $\pi_M \in \mathrm{int}(\Delta(K))$, there exists N_0 such that $v_n \in \Delta(K)$ for every $n \geq N_0$. As at the beginning of the proof, we take for each $n \geq N_0$ an element $u_n \in C(\Delta(K))$ satisfying $u_n(q) = 1$ for $q \in \{v, v_n\}$ and $u_n(q) < 1$ for $q \in \Delta(K) \setminus \{v, v_n\}$. Hence, by arguing as before for e_i, e_i^n and u_i^n, only this time for v, v_n and u_n, we obtain that $vM \in (v, v_n]$ for every $n \geq N_0$. As $v_n \to \pi_M$ as $n \to \infty$, we see that $vM \in (v, \pi_M]$. On the other hand,

$$\begin{aligned} vM &= \frac{1}{2}\left(e_i M + e_j M\right) \\ &= \frac{1}{2}\left(\beta_i e_i + (1-\beta_i)\pi_M\right) + \frac{1}{2}\left(\beta_j e_j + (1-\beta_j)\pi_M\right) \\ &= \left(1 - \frac{\beta_i}{2} - \frac{\beta_j}{2}\right)\pi_M + \beta_i v + \frac{1}{2}(\beta_j - \beta_i)e_j. \end{aligned} \tag{50}$$

As $0 \leq \beta_i < \beta_j < 1$, this implies that vM lies in the relative interior of the triangle with the vertices π_M, v, and e_j. This of course contradicts the fact that $vM \in (v, \pi_M]$. Hence, $\beta_i = \beta_j$ for every $i \neq j \in K$, thus proving that M is a homothety. □

References

[1] Abreu, D., Pearce, D., and Stacchetti, E. (1990), Toward a Theory of Discounted Repeated Games with Imperfect Monitoring. *Econometrica*, 58(5), 1041–1063.

[2] Arieli, I. and Babichenko, Y. (2019), Private Bayesian Persuasion. *Journal Economic Theory*, 182, 185–217.

[3] Arieli, I., Babichenko, Y., Smorodinsky, R. and Yamashita, T. (2020), Optimal Persuasion via Bi-Pooling. *In:The Twenty-First ACM Conference on Economics and Computation.*

[4] Athey, S. and Bagwell, K. (2008), Collusion with Persistent Cost Shocks. *Econometrica*, 76(3), 493-540.

[5] Au, P.H. (2015), Dynamic Information Disclosure. *The RAND Journal of Economics*, 46(4): 791–823.

[6] Augenblick, N. and Bodoh-Creed, A. (2018), To Reveal or not to Reveal: Privacy Preferences and Economic Frictions. *Games and Economic Behavior*, 110, 318–329.

[7] Aumann, R.J. and Maschler, M. (1995), *Repeated Games with Incomplete Information.* With the collaboration of R. Stearns, Cambridge, MA: MIT Press.

[8] Blackwell, D. (1951), Comparison of Experiments. *In: Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, 1950*, pp. 93–102. University of California Press, Berkeley and Los Angeles, Calif.

[9] Bressaud, X. and Quas, A. (2017), Dynamical Analysis of a Repeated Game with Incomplete Information. *Math. Oper. Res.* 42(4), 1085–1105.

[10] Dziuda, W. and Gradwohl, R. (2015), Achieving Cooperation Under Privacy Concerns. *Am. Econ. J. Microecon.*, 7, 142–173

[11] Ely, J.C. (2017), Beeps. *American Economic Review*, 107(1), 31–53.

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

[13] Farhadi, F. and Teneketzis, D. (2021), Dynamic Information Design: A Simple Problem on Optimal Sequential Information Disclosure. *Dyn Games Appl.*

[14] Forges, F. (2020). Games with Incomplete Information: From Repetition to Cheap Talk and Persuasion. *Annals of Economics and Statistics*, 137, 3–30.

[15] Fudenberg, D. and Maskin, E. (1986), The Folk Theorem in Repeated Games with Discounting or with Incomplete Information. *Econometrica*, 54(3), 533–554.

[16] Ganglmair, B. and Tarantino, E. (2014), Conversation With Secrets. *Rand J. Econ.*, 45, 273–302.

[17] Guo,Y. and Shmaya, E. (2018), The Value of Multistage Persuasion. *Working paper.*

[18] Honryo, T. (2018), Dynamic Persuasion. *J. Econom. Theory*, 178, 36–58.

[19] Horn, R.A. and Johnson, C.R. (2013), *Matrix analysis. Second edition.* Cambridge University Press, Cambridge.

[20] Hörner, J., Rosenberg, D., Solan, E., and Vieille, N. (2010), On a Markov Game with One-Sided Information. *Oper. Res.* 58 (4), part 2, 1107–1115.

[21] Kamenica, E. and Gentzkow, M. (2011), Bayesian Persuasion. *American Economic Review*, 101 (6), 2590–2615.

[22] Kolotilin, A., Mylovanov, T., Zapechelnyuk, A., and Li, M. (2017), Persuasion of a Privately Informed Receiver. *Econometrica*, 85(6), 1949–1964.

[23] Lehrer, E. and Sorin, S. (1992), A Uniform Tauberian Theorem in Dynamic Programming. *Math. Oper. Res.* 17, no. 2, 303–307.

[24] Lingenbrink, D. and Iyer, K. (2019) Optimal Signaling Mechanisms in Unobservable Queues with Strategic Customers. *Oper. Res.* 67, no. 5, 1397-1416.

[25] Mailath, G. and Samuelson, L. (2001), Who Wants a Good Reputation? *Rev. Econom. Stud.*, 68, no. 2, 415–441.

[26] Neyman, A. and Kohlberg, E. (1999), A strong Law of Large Numbers for Nonexpansive Vector-Valued Stochastic Processes. *Israel J. Math.*, 111, 93–108.

[27] Orlov, D., Skrzypacz, A., and Zryumov, P. (2020), Persuading the Principal to Wait. *Journal of Political Economy*, 128(7), 2542–2578.

[28] Levin, D. A. and Peres, Y. (2017), *Markov Chains and Mixing Times, second edition.* With contributions by Elizabeth L. Wilmer. With a chapter on "Coupling from the past" by James G. Propp and David B. Wilson. American Mathematical Society, Providence, RI.

[29] Peşki, M. and Toikka J. (2017), Value of Persistent Information. *Econometrica*, 85(6), 1921–1948.

[30] Phelan, C. (2006), Public Trust and Government Betrayal. *J. Econom. Theory*, 127(1), 27–43.

[31] Renault, J. (2006), The value of Markov Chain Games with Lack of Information on One Side. *Mathematics of Operations Research*, 31, 3: 433–648.

[32] Renault, J., Solan, E., and Vieille, N. (2017), Optimal Dynamic Information Provision. *Games Econ. Behav.*, 104, 329–349.

[33] Renault, J. (2018), *Repeated Games with Incomplete Information.* In: Meyers R. (eds) Encyclopedia of Complexity and Systems Science. Springer, Berlin, Heidelberg.

[34] Rockafellar, R. T. (1970), *Convex Analysis.* Princeton University Press.

[35] Guo, Y. and Shmaya, E. (2019), The Interval Structure of Optimal Disclosure. *Econometrica* 87 (2), 653–675.

[36] Solan, E. (2021), *A Course in Stochastic Game Theory.* Cambridge University Press. To appear.

[37] Wiseman, T. (2008), Reputation and Impermanent Types. *Games Econ. Behav.*, 62, 190–210.

Tel Aviv University, Tel Aviv 69978, Israel.
https://sites.google.com/site/profelehrer/home
E-mail address: lehrer@post.tau.ac.il

Tel Aviv University, Tel Aviv 69978, Israel.
E-mail address: dima.shaiderman@gmail.com