

Markovian persuasion

EHUD LEHRER

Department of Economics, Durham University, UK

DIMITRY SHAIDERMAN

Department of Mathematics, Hebrew University of Jerusalem, Israel

In the classical Bayesian persuasion model, an informed player and an uninformed one engage in a static interaction. This work extends this classical model to a dynamic setting where the state of nature evolves according to a Markovian law, allowing for a more realistic representation of real-world situations where the state of nature evolves over time. In this repeated persuasion model, an optimal disclosure strategy of the sender must balance between obtaining a high-stage payoff and disclosing information that may have negative implications on future payoffs. We discuss optimal strategies under different discount factors and characterize when the asymptotic value achieves the maximal possible value.

KEYWORDS. Markovian persuasion, dynamic Bayesian persuasion, Markov chain, asymptotic value, absorbing set, homothety.

JEL CLASSIFICATION. D72, D82, D83, K40, M31.

1. INTRODUCTION

This paper focuses on a dynamic Bayesian persuasion model. This model is used to describe situations where a principal (sender) repeatedly persuades an agent (receiver) in a changing stochastic environment. Unlike the one-shot Bayesian persuasion problem (e.g., [Kamenica and Gentzkow \(2011\)](#)), the dynamic Bayesian persuasion framework incorporates a Markov chain that evolves stochastically in discrete time and is assumed to be ergodic.¹ At each period, the sender observes the Markov chain and sends a (stochastic) signal to the receiver. The sender commits to a signaling mechanism at the start of the interaction, and the receiver updates his daily posterior beliefs according to Bayes' law. The receiver, who acts myopically in his beliefs, determines both his and the sender's payoff. The sender's objective is to maximize his expected discounted daily payoffs.

In a dynamic setting, the sender's information provision policy has significant implications not only for the current payoff, but also for the future evolution of receiver's

Ehud Lehrer: ehud.m.lehrer@durham.ac.uk

Dimitry Shaiderman: dima.shaiderman@gmail.com

We thank the anonymous referees of this journal. The first author acknowledges the support of ISF 591/21 and DFG KA 5609/1-1.

¹That is, a Markov chain that is both irreducible and aperiodic.

beliefs. This is because the receiver's posterior belief, updated based on the sender's information, is shifted by the Markov transition matrix to establish a new prior belief on the following day. The sender must therefore strike an optimal balance between short-term and long-term objectives, balancing immediate payoff against future control of the receiver's beliefs, which may be in conflict.

The tension faced by the sender in the theoretical model is not just an abstract concept, it applies to real-life situations. For instance, [Ely \(2017\)](#) describes a scenario in which a CEO (sender) provides reports to a board of directors (receiver) to assess the firm's operational competence. The CEO wishes to send reports at the optimal time to minimize the number of costly audits ordered by the board. [Renault, Solan, and Vieille \(2017\)](#) consider the case of an investor (receiver) who relies on an informed advisor (sender) for market information. The advisor seeks to manipulate the investor's beliefs to his financial advantage, as he receives a fixed fee each time the investor opts in. Commercial banks (here, the receiver) receive signals from a central bank (sender), such as information about interest rates or bank stability. A forward-looking central bank must take into account the evolving stochastic nature of the markets when signaling to commercial banks to maximize economic stability. These examples illustrate how the sender's optimal balance between short-term goals and long-term control of the receiver's beliefs applies to a range of real-world scenarios.

Our contribution Our main result identifies a joint condition involving both the sender's payoff function and the Markov chain that enables a patient sender to ensure the concavification of his payoff function at the stationary distribution of the Markov chain. This condition characterizes situations in which a patient sender is not impacted by any additional effects of the stochastic environment compared to the case of the static environment. The finding is consistent with previous research by [Kamenica and Gentzkow \(2011\)](#) because the concavification just mentioned is the solution to the one-shot Bayesian persuasion problem at the stationary distribution.

Our joint condition can be explained as follows: let u be the payoff function, $(\text{Cav } u)$ be its concavification, and M be the Markov matrix, with π_M as its stationary distribution. Our condition asserts that the supporting hyperplane to $(\text{Cav } u)$ at π_M should contain a set of beliefs C that satisfy two key properties. First, for every $p \in C$, $u(p) = (\text{Cav } u)(p)$. Second, C should include all M shifts of its elements, meaning that the shift pM of a belief $p \in C$ remains within the convex hull of C .

This condition is sufficient due to two basic properties. First, $(\text{Cav } u)$ is affine on the convex hull of C . Second, starting from the prior π_M , the posteriors at any subsequent day have a mean of π_M . To guarantee an expected daily payoff of $(\text{Cav } u)(\pi_M)$, the sender first splits π_M to C . Since the joint condition ensures that the next day's prior (i.e., the shift of the first day realized posterior) is contained in C , the sender splits it back to C . By using these properties, the sender can repeat the same strategy on every subsequent day and receive an expected daily payoff of $(\text{Cav } u)(\pi_M)$.

If the sender starts from a different prior, the ergodicity of the Markov chain ensures that, with patience, the sender can wait until the prior approaches π_M and then utilize C in the same way.

While the necessity of this joint condition may not be immediately intuitive, its implications are significant. Specifically, if the joint condition on the payoff and the Markov matrix is not satisfied, then at best, a patient sender can obtain a strictly lower expected payoff than $(\text{Cav } u)(\pi_M)$. In this way, our main result characterizes precisely the cases in which the stochastic environment has an “additional effect” on a patient sender compared to the static environment. To prove the necessity of this condition, we draw on tools and techniques from the literature on repeated games with incomplete information (see Renault (2020)).

The contribution relative to Ely (2017) and Renault, Solan, and Vieille (2017) Both papers discovered that the dynamic Bayesian persuasion model can be simplified to a Markov decision problem over the belief state space. This insight allowed them to derive a Bellman equation for the optimal discounted values of the sender. However, it is often impossible to provide a more informative description of the value than the Bellman equation for a general payoff function, as Abreu, Pearce, and Stacchetti (1990) has shown.

Despite the difficulties associated with the Bellman equation, Ely (2017) was able to solve it for his “beeps” example. Renault, Solan, and Vieille (2017) examined a step function defined by a hyperplane and homothetic Markov chains. They showed that the optimal strategy for the sender corresponds to the greedy strategy from a certain (random) stage onward.² They also demonstrated that the greedy strategy is not always optimal for the sender.

Our result states that, when the joint condition is satisfied, there exists an optimal greedy strategy. However, when the condition is not satisfied, a greedy strategy can still be optimal, but it will guarantee a patient sender less than $(\text{Cav } u)(\pi_M)$.

Our innovation lies in being the first to provide a general result that goes beyond the Bellman equation. Our criterion reveals one of the crucial aspects of the relations between the payoff function u and the Markov matrix M , which is at the core of the dynamic Bayesian persuasion problem. While there may be other aspects in those relations, our criterion is an important step toward a deeper understanding of the problem.

The dynamic Bayesian persuasion model shares similarities with the model of *Markov chain games* described in Renault (2006). Both models explore long-term interactions between two players in a Markovian environment, where one player is informed of the realizations of the Markov chain. However, there are key differences between the two. Markov chain games involve a zero-sum game without commitment power, and both players are long-lived, meaning the uninformed player need not act myopically and has infinite memory. In Section 4, we provide a detailed comparison of the two models and clarify that our results are independent of those presented in Renault (2006) and Pęski and Toikka (2017), which specifically focus on Markov chain games.

Practical implications In dealing with a dynamic Bayesian persuasion problem, the sender (such as a CEO, central bank, etc.) should first verify whether the primitives of their dynamic problem satisfy the joint condition. This can be a challenging task, as

²The greedy strategy is the one in which the sender acts myopically and maximizes his stage payoffs.

even for just three states, there may be infinitely many candidates for such a set C (as discussed in Section 4). To assist senders in this process, we examine several cases in Section 5 that can simplify matters for the sender.

There are certain scenarios where verifying the joint condition can be made easier. For instance, when the Markov chain is a homothety, the joint condition holds for every payoff function (see Section 5.1). In two-state Markov chains, the joint condition is easy to verify and is phrased in terms of the dynamics induced by the M -shift around π_M (see Theorem 4 in Section 5.2).

Our paper is organized as follows. In Section 2, we present a formal exposition of the dynamic Bayesian persuasion model. Then, in Section 3, we provide a detailed glimpse into our main result by studying a two-state example. This section introduces essential notions and arguments required for the main theorem, which we state and discuss in Section 4. In Section 5, we give several applications of the main result. The proofs are presented in Section 6.

2. THE MODEL

Let $K = \{1, \dots, k\}$ 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 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 .³

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

The payoffs of the sender and the receiver at 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 . The sender is assumed to observe the actions of the receiver and, therefore, knows his and the receiver's stage payoffs. However, we assume that the receiver does not have access to his own or the sender's stage payoffs. Both the sender and the receiver discount their payoffs by a factor $\delta \in [0, 1]$. We denote this game by $\Gamma_\delta(p)$. As in the models of Renault, Solan, and Vieille (2017), Ely (2017), and Farhadi and Teneketzis (2022), the receiver obtains information only through the sender.

A *signaling strategy* σ of the sender in $\Gamma_\delta(p)$ is described by a sequence of stage strategies (σ_n) , where each σ_n is a mapping $\sigma_n : (K \times S)^{n-1} \times K \rightarrow \Delta(S)$. 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 $x_1, \dots, x_{n-1}$ and past signals $s_1, \dots, s_{n-1}$ together with the current state x_n . Let Σ be the space of all signaling strategies.⁴

³This assumption is in place to make sure the sender can disclose $(x_n)_n$.

⁴The model could also be interpreted as follows. As is conventionally assumed in the Bayesian persuasion literature, both players are initially *symmetrically* uninformed about the state. The sender makes a choice (signaling strategy) at the start of $\Gamma_\delta(p)$ when he has no information about the state yet.

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

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

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

$$p_n^\ell = \mathbb{P}_{p,\sigma}(X_n = \ell | s_1, \dots, s_n).$$

Set $p_n = (p_n^\ell)_{\ell \in K}$.

A second key assumption of our model is that the receiver's decision at any period n depends only on her posterior, p_n . Such an assumption is realistic in cases where the receiver maximizes her expected payoff based on her current belief. It is a natural assumption in several scenarios. One is when a sequence of transitory short-lived receivers (e.g., [Jackson and Kalai \(1997\)](#)) are involved in a recurring game, where a stage game is played sequentially by different players having social memory of the information provided (as is often the case in online markets). Another scenario is when at any stage a particular receiver is chosen from a large population of anonymous receivers, and the sender is a political party or a media outlet trying to persuade the general public.

Denote by $\theta : \Delta(K) \rightarrow B$ the decision policy of the receiver, that is, the mapping that 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 $u : \Delta(K) \rightarrow \mathbb{R}$ defined by $u(q) = \sum_{\ell \in K} q^\ell v(\ell, \theta(q))$ is upper semicontinuous. Such an assumption is important for two reasons. First, it is consistent with the assumption that whenever the receiver is indifferent between two actions he chooses that which benefits the sender. Second, the fact that upper semicontinuous functions defined on $\Delta(K)$ achieve their maxima will be useful throughout the paper.

To summarize, our assumptions imply that the signaling strategy σ of the sender determines his payoff at any period n . Moreover, the total expected payoff to the sender in $\Gamma_\delta(p)$ under the signaling strategy σ can now be written as

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

where $\mathbb{E}_{p,\sigma}$ is the expectation with respect to $\mathbb{P}_{p,\sigma}$. The value of the game $\Gamma_\delta(p)$ is $v_\delta(p) = \sup_{\sigma \in \Sigma} U_\delta(p, \sigma)$.

3. A DETAILED GLIMPSE AT THE MAIN RESULT

To facilitate the understanding of our main result, we provide a qualitative overview of the key ideas and concepts upon which the paper relies. We focus on the case of two states, denoted as $K = \{1, 2\}$, which allows us to simplify our presentation. Specifically, we can represent the set $\Delta(K)$ as the unit interval $[0, 1]$, where each point $p \in [0, 1]$ corresponds to the distribution $(p, 1 - p)$.

For the remainder of our discussion, we assume that the payoff function u is defined on $[0, 1]$. As an illustration, let us consider a specific example of such a payoff function:

$$u(p) = \begin{cases} 2p, & \text{for } p \in [0, 0.3], \\ 1.5 - 3p, & \text{for } p \in [0.3, 0.5], \\ 3p - 1.5, & \text{for } p \in [0.5, 0.75], \\ 3 - 3p, & \text{for } p \in [0.75, 1]. \end{cases}$$

Figure 1 provides a graphical representation of the function u , where we have set $q_1 = 0.3$, $q_2 = 0.75$, and $\pi_M = 0.5$. We assume that $(0.5, 0.5)$ is the stationary distribution of the Markov chain M . This notation will be useful in the subsequent discussion.

The one-shot Bayesian persuasion problem When considering the one-shot game (corresponding to $\lambda = 0$), the sender faces a classical Bayesian persuasion problem with payoff function u , as studied in [Kamenica and Gentzkow \(2011\)](#). The solution concept in this setting is based on $(\text{Cav } u)$, the so-called “concavification” of u . In the given example, $(\text{Cav } u)$ agrees with u on the intervals $[0, 0.3]$ and $[0.75, 1]$, where u is concave, while on the interval $[0.3, 0.75]$, $(\text{Cav } u)$ is given by the hyperplane $1/3p + 0.5$. (See the blue dashed line in Figure 1 for a graphical representation of $(\text{Cav } u)$.)




FIGURE 1. A two-state example.

The sender's optimal policy for any belief $p \in [0.3, 0.75]$ involves “splitting” p into two points where the hyperplane $1/3p + 0.5$ touches the function u , namely q_1 and q_2 in our example (as shown by the green points in Figure 1). For example, when the belief is π_M , this optimal policy yields an expected payoff of $(\text{Cav } u)(\pi_M) = 2/3$ in the one-shot game.

The asymptotic approach The law of the irreducible and aperiodic Markov chain $(X_n)_{n \geq 1}$ dictates that it converges with time to the stationary distribution π_M of M . Therefore, if the sender is sufficiently patient, regardless of the receiver's prior probability p over K , the sender can wait until the receiver's beliefs naturally approach π_M , and then act informatively. Building on this observation, Theorem 1 argues that $v_\lambda(p)$ converges uniformly (on $\Delta(K)$), as $\lambda \rightarrow 1$, to a number $v_\infty \in \mathbb{R}$, referred to as *the asymptotic value*. Thus, according to Theorem 1, the study of v_∞ can be accomplished solely through the mapping $\delta \mapsto v_\lambda(\pi_M)$.

The mean-consistency property at π_M The Markov dynamics dictate that the receiver's belief at the end of a stage, $q \in \Delta(K)$, is transformed by M into qM , which then becomes the receiver's belief at the beginning of the next stage. As a result, the elements of the sequence of posteriors $(p_n)_{n \geq 1}$, starting from the prior π_M , will have the same mean, namely π_M . This follows from the fact that splits are mean-preserving, coupled with the observation that if $\sum \alpha_i q_i = \pi_M$ for some convex weights (α_i) and beliefs (q_i) , then the linearity of the M -shift on $\Delta(K)$ implies that $\sum \alpha_i q_i M = \pi_M M = \pi_M$.

This mean-consistency property implies that the sender cannot achieve more than $(\text{Cav } u)(\pi_M)$ at any stage $n \geq 1$ when starting from the prior π_M . Hence, $v_\lambda(\pi_M) \leq (\text{Cav } u)(\pi_M)$ for any λ , and the asymptotic value v_∞ cannot exceed $(\text{Cav } u)(\pi_M)$ either.

The research question The goal of this paper is to investigate the conditions that guarantee that $v_\infty = (\text{Cav } u)(\pi_M)$. Specifically, we aim to determine when can a patient sender guarantee a payoff close to that in the one-shot Bayesian persuasion problem with prior π_M .

Keeping the posteriors supported on $\{q_1, q_2\}$ The mean consistency property suggests that if the sender could ensure that the belief is supported on $\{q_1, q_2\}$ for every n , then he would guarantee at every stage an expected payoff of $(\text{Cav } u)(\pi_M)$. Since he cannot guarantee more, this would show that $v_\lambda(\pi_M) = (\text{Cav } u)(\pi_M)$ for every $\lambda \in [0, 1]$, proving that the asymptotic value v_∞ also equals $(\text{Cav } u)(\pi_M)$.

Absorption under M -shifts of $\{q_1, q_2\}$ The sequence of posteriors (p_n) can only be supported on the set $\{q_1, q_2\}$ if q_1 and q_2 are shifted by M to the interval $[q_1, q_2]$ (e.g., see the red arrow demonstrated for q_1 in Figure 1). Indeed, only in that case, the sender could split the initial belief at any subsequent stage (being $q_1 M$ or $q_2 M$) back to q_1 and q_2 . In geometric terms, this property implies that the set $\{q_1, q_2\}$ is absorbed under M , in the sense that its shift under M lies within its convex hull, which is the interval $[q_1, q_2]$. We generalize this property naturally to an arbitrary state set K as follows. A set of beliefs $C \subseteq \Delta(K)$ is said to be *M -absorbing* if its M -shift lies within its convex boundaries (see Definition 1).

The revelation of information policy The information provision policy can be summarized as follows. In the first stage, the sender “splits” the prior belief π_M to $\{q_1, q_2\}$. At the beginning of each subsequent stage, the sender splits the realized shift, either q_1M or q_2M , back to $\{q_1, q_2\}$. This policy is illustrated in Figure 1 for q_1M using the blue arrows.

The main result—a characterization The previous arguments provided a sufficient condition for v_∞ to be equal to $(\text{Cav } u)(\pi_M)$. Specifically, the set $\{q_1, q_2\}$ needs to be M -absorbing. The main theorem of the paper, Theorem 2, states that this condition is also necessary. Moreover, the generality of Theorem 2 extends beyond the two-state case. In essence, the theorem asserts that a patient sender can only achieve $(\text{Cav } u)(\pi_M)$ if there exists an M -absorbing set of beliefs within the set of beliefs where the supporting hyperplane to $(\text{Cav } u)$ at π_M touches the graph of u . In the current example, this set is $\{q_1, q_2\}$.

4. THE MAIN THEOREM

4.1 The existence of asymptotic value

To present our first result, we must establish some notation. Let us begin by defining π_M to be the unique stationary distribution of M . Additionally, for any function $g : \Delta(K) \rightarrow \mathbb{R}$, define the function $(\text{Cav } g)$ by

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

Our first finding indicates that the effect of $p \in \Delta(K)$ on the value $v_\delta(\pi_M)$ of a sufficiently patient sender (i.e., with λ close to 1) is insignificant in comparison to the impacts of u and M . Moreover, as the patience level λ approaches 1, the sequence of functions $v_\lambda(\cdot)$ converges uniformly on $\Delta(K)$. To be precise, we state this outcome as follows.

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

$$|v_\delta(p) - v_\infty| < \varepsilon, \quad \forall \delta > \delta(\varepsilon), \forall p \in \Delta(K).$$

Interestingly, the upper bound on v_∞ presented in Theorem 1 is an exact bound. In the subsequent section, we will provide a geometric criterion for this upper bound to be achieved. This will require us to first introduce and examine the concept of M -absorbing sets.

4.2 M -Absorbing sets

DEFINITION 1. A nonempty set $C \subseteq \Delta(K)$ is said to be M -absorbing if $qM \in \text{conv}(C)$ for every $q \in C$.⁵

⁵ $\text{conv}(C)$ denotes the convex hull of the set C .

The reasoning behind our choice of terminology is as follows. Given that $q \mapsto qM$ is a linear operator, if a set C is M -absorbing, then so is its convex hull $\text{conv}(C)$. However, when it comes to $\text{conv}(C)$, M -absorption means that the image of $\text{conv}(C)$ under M is fully contained within (or absorbed by) $\text{conv}(C)$. Dually, if a closed convex set $C \subseteq \Delta(K)$ is M -absorbing, then by the Krein–Milman theorem we have that $\text{ext}(C)$ is also M -absorbing.⁶ As a consequence, since $\Delta(K)$ itself is M -absorbing, so is its set of extreme points (i.e., the set of all mass-point distributions). Finally, we observe that if both C_1 and C_2 are M -absorbing, then so is their union $C_1 \cup C_2$, because $\text{conv}(C_1) \cup \text{conv}(C_2) \subseteq \text{conv}(C_1 \cup C_2)$.

To provide further insight into the concept of M -absorbing sets, consider a scenario in which a football is passed in a straight line from any point $q \in \Delta(K)$ to the point qM . The trajectory of the football starting from q is given by 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 trajectory of the football starting from any point $q \in C$ remains inside $\text{conv}(C)$ and never leaves it. This can be thought of as a metaphorical “absorption” of the football’s trajectory by the set C .

Let us provide some basic examples of M -absorbing sets. The simplest ones are the singleton set $\{\pi_M\}$, consisting of the unique stationary distribution π_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 norm⁷ of M with respect to the ℓ_i -norm, $i \in \{1, 2, \infty\}$.

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

$$\|qM - \pi_M\|_i = \|qM - \pi_M M\|_i \leq \|M\|_i \|q - \pi_M\|_i. \quad (1)$$

It is well 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 [Horn and Johnson \(2013\)](#)), so $\|M\|_\infty = 1$. Additionally, since $\|M\|_2$ coincides with the maximal singular value of M (see, e.g., Example 5.6.6 on p. 346 in [Horn and Johnson \(2013\)](#)), $\|M\|_2 = \|M\|_\infty = 1$. Therefore, according to (1), any ball (either open or closed) centered at π_M with respect to the ℓ_2 or ℓ_∞ -norm is M -absorbing.

If M is doubly stochastic, it is known that $\|M\|_1 = 1$ since $\|M\|_1$ coincides with the maximal ℓ_1 -norm of a column of M (see, e.g., Example 5.6.5 on pp. 344–345 in [Horn and Johnson \(2013\)](#)). Hence, in this case, any ball (either open or closed) centered at π_M with respect to the ℓ_1 -norm is also M -absorbing. Examples of M -absorbing sets for these norms are illustrated in Figure 2.

In all of the examples above, the M -absorbing sets contain π_M , and this is not a coincidence. In fact, due to the linearity of the map M , for every M -absorbing set C , the image of $\text{conv}(C)$, is also contained in $\text{conv}(C)$. Therefore, by Brouwer’s fixed-point theorem, M possesses a fixed point in the closed convex hull $\text{cl } \text{conv}(C)$ of C . As the

⁶ $\text{ext}(C)$ denotes the set of extreme points of C , that is, points in C that cannot be expressed as a convex combination of two distinct points in C .

⁷ $\|M\|_i = \max_{\|x\|_i=1} \|xM\|_i$.




FIGURE 2. Absorbing sets.

only fixed point of M is π_M , we deduce that π_M lies in the closure of $\text{conv}(C)$ for every M -absorbing set C .

We conclude the discussion on M -absorbing sets with the following relevant proposition.

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

4.3 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 $g : \Delta(K) \rightarrow \mathbb{R}$ 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 $f_z : \mathbb{R}^k \rightarrow \mathbb{R}$ by $f_z(x) := (\text{Cav } u)(\pi_M) + \langle z, x - \pi_M \rangle$. Second, set

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


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 *supergradient* 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 $\text{Graph}[u]$ with the supporting hyperplane to $\text{Graph}[(\text{Cav } u)]$ at $(\pi_M, (\text{Cav } u)(\pi_M))$ parametrized by z . A visualization of A_z when $k = 3$ is given in Figure 3.

PROPOSITION 2. *We have the following:*

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

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

The strategy induced by an M -absorbing set M -absorbing sets contained in A_z are of particular importance due to their role in controlling the receiver's beliefs. 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 $\text{conv}(C)$. At this point, the sender may send messages that would induce posteriors within C , and in particular in A_z . As $(\text{Cav } u)$ is an affine function on $\text{conv}(A_z)$ (see Lemma 6 in Section 6), starting from π_M , the weighted average of the values of $(\text{Cav } u)$ evaluated at these posteriors, is equal to the value of $(\text{Cav } u)$ at π_M (see the mean-consistency property at π_M in Section 3).

The candidates for M -absorbing sets within A_z In Figure 3, we can observe that even in the 3-dimensional case (i.e., $|K| = 3$), the set A_z can be quite general and may include an infinite number of subsets. However, as Proposition 2 suggests, it is sufficient for only one of those subsets to be M -absorbing.

The main theorem summarizes the results of Theorem 1 and Proposition 2 and characterizes when a patient sender can achieve a value close to the maximum possible, as stated in Theorem 1. In other words, it describes the optimal policy of a patient sender when there exists $z \in \Lambda$ such that A_z contains an M -absorbing set: the sender should send messages that induce posteriors within this M -absorbing set of beliefs. The more technically challenging part of Theorem 2 shows that when condition (ii) is not satisfied, the static optimal value cannot be achieved. In such cases, the optimal strategy can be described using a Bellman-type condition, but typically not in terms of the static model.

THEOREM 2. *The following statements are equivalent:*

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

$$|v_\delta(p) - (\text{Cav } u)(\pi_M)| < \varepsilon, \quad \forall \delta > \delta(\varepsilon), \forall p \in \Delta(K).$$

(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.*

The grand question While the criterion given in Theorem 2 is more general than those previously addressed in the literature (such as Ely (2017) and Renault, Solan, and Vieille (2017)), it does not offer a solution for the optimal disclosure of information by a patient sender when the disclosure has future payoff implications. This question remains open and it seems challenging to tackle.

The effect of the discount factor The following proposition delineates the behavior exhibited by the mapping $\delta \mapsto v_\delta(\pi_M)$ when $v_\infty = (\text{Cav } u)(\pi_M)$.

PROPOSITION 3. *Assume that $v_\infty = (\text{Cav } u)(\pi_M)$. Then $v_\delta(\pi_M) = (\text{Cav } u)(\pi_M)$ for every $\delta \in [0, 1)$.*

The proposition states that whenever $v_\infty = (\text{Cav } u)(\pi_M)$, or equivalently, whenever A_z contains an M -absorbing set for some $z \in \Lambda$, the discount factor has no effect on the sender's value. Moreover, regardless of the discount factor, the sender should follow the strategy induced by the given M -absorbing set. In essence, the significance of the scenario where $v_\infty = (\text{Cav } u)(\pi_M)$ transcends mere asymptotic considerations. It has an important insight regarding the values $v_\delta(\pi_M)$ across the entire spectrum of discount factors δ .

EXAMPLE 1. Let us return to the example in Section 3. Since $\pi_M = 0.5$ and M is assumed to be irreducible and aperiodic, there must exist $0 < a < 1$ so that

$$M = \begin{pmatrix} a & 1-a \\ 1-a & a \end{pmatrix}.$$

In that case, the set $\{q_1, q_2\} = \{0.3, 0.75\}$ is M absorbing if and only if

$$0.3 \leq \langle (0.3, 0.7), (a, 1-a) \rangle \leq 0.75, \quad (2)$$

and

$$0.3 \leq \langle (0.75, 0.25), (a, 1-a) \rangle \leq 0.75. \quad (3)$$

A straightforward check shows that the inequalities (2) and (3) amount to $0.1 \leq a$. Hence, $v_\infty = (\text{Cav } u)(\pi_M)$ if and only if $0.1 \leq a$. $\diamond$

The informational interpretation The core of the model analyzed in this paper lies in the sender's optimal balance between disclosing information to achieve short-term goals (increasing today's payoff) and considering the potential impact of such information on long-term goals (future payoffs). Theorem 2 highlights a scenario where such a balance is not necessary. Specifically, it states that this occurs whenever there exists a $z \in \Lambda$ such that A_z contains an M -absorbing set. In other words, the sender can act myopically when releasing information in this case.

However, even in cases where A_z does not contain an M -absorbing set, it is possible that the optimal policy for the sender is still myopic. Although no such example has been identified in the literature, it raises a natural research question.

QUESTION 1. *What are the necessary and sufficient conditions for a greedy strategy to be optimal?*

The interrelationships between u and M It is worth noting that Theorem 2 provides a characterization of the sender's optimal policy based solely on the model primitives, namely u and M . The interplay between these primitives is crucial, as the sets A_z , $z \in \Lambda$, depend solely on u , while the M -absorbing sets are determined by M . This raises an intriguing question: how sensitive is the sender's optimal policy to the interplay between u and M ?

QUESTION 2. *Assume that A_z does not contain an M -absorbing set for some $z \in \Lambda$. Can one quantify the difference between $(\text{Cav } u)(\pi_M)$ and v_∞ in terms of u and M ?*

Comparison with Markov chain games The Markovian persuasion model analyzed in this paper shares similarities with the Markov chain game model introduced by Renault (2006). Both models can be represented as a stochastic game over the state space $\Delta(K)$, which in the case of the Markovian persuasion model reduces to an MDP (as discussed in Section 6). Renault (2006) proved that the limit of values of finite-duration Markov chain games with Cesàro valuation exists and is also the uniform value of the games. This result, coupled with Ziliotto's Tauberian theorem for stochastic games (see Section 3.3 of Ziliotto (2016)) establishes that the asymptotic value of Markov chain games exists when the discount factor approaches 1. This is similar to Theorem 1, which also shows that the Markovian persuasion model admits an asymptotic value.

Despite these similarities, there are significant differences between the two models. As mentioned in the Introduction, Markov chain games are zero-sum, and the sender does not need to commit to a strategy.⁸ Furthermore, the receiver need not act myopically.⁹ Along with these significant conceptual differences, there are also important technical distinctions between both models.

In the Markovian persuasion model, we consider a general upper semicontinuous payoff function u , unlike in Markov chain games, where payoffs are restricted to a specific family of functions. Secondly, the evolution of the state variable in Markov chain games is based on the posterior distribution on K at stage n , which is induced by the first $n - 1$ actions of the informed player. In contrast, in the Markovian persuasion game, the state variable p_n is the distribution on K induced by the first n (compared to $n - 1$) signals sent by the sender. This subtle difference has important technical implications and is central to the main result of this paper, Theorem 2.

Theorem 2 provides a unique insight into the interplay between the stage payoffs and the Markov transition rule. Specifically, we demonstrate how M -absorbing sets and stage payoffs jointly impact the asymptotic value and the optimal strategy of the sender. This result stands out from previous works such as Renault (2006) and Pęski and Toikka

⁸The sender in Markov chain games is called Player 1. Similar to the sender studied in this paper, he also observes the Markov chain and may reveal information to his opponent regarding the evolution of the chain.

⁹The receiver in Markov chain games is called Player 2.

(2017), which focus on the values of Markov chain games without considering the specific relation between the payoff function and the Markov transition rule. Accordingly, the techniques used in our paper differ from those employed in the existing literature.

5. APPLICATIONS OF THE MAIN RESULT

5.1 When the main theorem holds for every u : Homothety

The results obtained previously have provided insight into the relationship between M and u . In particular, the main theorem establishes the conditions under which $v_\infty = (\text{Cav } u)(\pi_M)$, based on M and u . This leads to a natural question: can we determine if, for a fixed value of M , this result holds for all possible u ? To address this question, we must first introduce the concept of homothety.

DEFINITION 2. A linear map $\psi : \Delta(K) \rightarrow \Delta(K)$ is said to be a *homothety* with respect to the pair $(v, \beta) \in \Delta(K) \times [0, 1]$ if ψ maps each point $x \in \Delta(K)$ into the point $\beta x + (1 - \beta)v$. The point v is called the *center* and β is called the *ratio*.

A homothety with respect to (v, β) fixes the point v and reduces the distance from any point x to v by a factor of β . Thus, an irreducible aperiodic stochastic matrix M is referred to as a homothety if the mapping $\psi : x \mapsto xM$ is a homothety with center π_M and ratio β for some $\beta \in [0, 1]$. When M is a homothety, the transition from one state to another follows the following rule: each state remains unchanged with probability $1 - \beta$ and moves to other states based on the distribution v with probability β .

We can describe an interesting class of M -absorbing sets of a homothety M by considering sets that are star-shaped around a point $p \in \Delta(K)$. A set $E \subseteq \Delta(K)$ is said to be *star-shaped around* p if $[p, q] \subseteq E$ for every $q \in E$. In simpler terms, if an observer is located at point p , then E is star-shaped around p if the line of sight, $[p, q]$, to any point $q \in E$ lies entirely in E .

Suppose now that E is star-shaped around π_M , and M is a homothety. Then for any $q \in E$, we have $qM \in [\pi_M, q] \subseteq E \subseteq \text{conv}(E)$. This means that every point in E transitions under the homothety M to a point within E . Therefore, when M is a homothety, every star-shaped set around any π_M is M -absorbing.

We present a result that characterizes when an irreducible and aperiodic matrix M is a homothety, in terms of v_∞ . To make this result clear, note that v_∞ is constant on $\Delta(K)$ by Theorem 1, and is a function of both u and M . In our characterization, we allow u to vary over the space of all functions defined on $\Delta(K)$, and thus v_∞ varies accordingly.

THEOREM 3. M is a homothety if and only if $v_\infty = (\text{Cav } u)(\pi_M)$ for every upper semi-continuous function $u : \Delta(K) \rightarrow \mathbb{R}$.

The practical implications Let $C \subseteq A_z$ be a set¹⁰ for which $\pi_M \in \text{conv}(C)$. Since $\text{conv}(C)$ is star-shaped around π_M , we can conclude that when M is a homothety, both $\text{conv}(C)$ and C are M -absorbing. Therefore, from a complexity standpoint, under the homothety assumption, the sender need not search for an M -absorbing set within A_z . Any set $C \subseteq A_z$ containing π_M in its convex hull is M -absorbing.

¹⁰Such a set is assured to exist by Carathéodory's theorem (e.g., Corollary 17.1.5 in Rockafellar (1970)).

5.2 A characterization of the two-state case

We return to the binary case $K = \{1, 2\}$. The irreducible and aperiodic stochastic matrix M is of the form

$$M = \begin{pmatrix} a & 1-a \\ b & 1-b \end{pmatrix},$$

where $a, b \in (0, 1)$. For simplicity, we identify a distribution $(p, 1-p) \in \Delta(K)$ with the mass p it assigns to state 1. With this identification in mind, one can easily verify that $\pi_M = b/(b-a+1)$. In the following theorem, we give an explicit characterization of the case when $v_\infty = (\text{Cav } u)(\pi_M)$.

THEOREM 4. $v_\infty = (\text{Cav } u)(\pi_M)$ if and only if one of the two following conditions holds:

[A] $a \geq b$;

[B] $b > a$ and for some $z \in \Lambda$ there exists $y \in A_z \cap [\pi_M, 1]$ such that

$$\left[\frac{b-y}{b-a}, b - (b-a)y \right] \cap A_z \neq \emptyset.$$

The next discussion is dedicated to the intuition behind Theorem 4.

The geometry of the M -shift in two states Consider the M -shift map $p \mapsto pM$. The geometric behavior of the M -shift can be of two types: (i) pM lies in the line segment $[p, \pi_M]$ or (ii) $pM \in [\pi_M, 1]$ when $p \leq \pi_M$ (or $[0, \pi_M]$) if $p \geq \pi_M$. In words, the M -shift either keeps the belief's orientation around π_M (case (i)) or flips the belief around π_M (case (ii)). An illustration of case (ii) is given in Figure 1.

The proof of Theorem 4 shows that in the two-state case, keeping the orientation around π_M is equivalent to $a \geq b$, which in turn is equivalent to M being a homothety with center π_M and ration $a-b$. Therefore, by Theorem 3, Case [A] in Theorem 4 is clearly sufficient to guarantee that $v_\infty = (\text{Cav } u)(\pi_M)$. The remaining task is to identify when $v_\infty = (\text{Cav } u)(\pi_M)$ in the case where M is flipping around π_M (i.e., $b > a$).

Suppose that M is flipping around π_M . Assume, for simplicity, that u has the property that there exists a unique pair q_1, q_2 such that $q_1 < \pi_M < q_2$ and $(\text{Cav } u)(\pi_M)$ is achieved by splitting π_M to $\{q_1, q_2\}$. In this case, Theorem 2 states that $v_\infty = (\text{Cav } u)(\pi_M)$ if and only if the set $\{q_1, q_2\}$ is M -absorbing. For this to happen, two inequalities must be satisfied: $q_1 M \leq q_2$ and $q_2 M \geq q_1$ (see Figure 1). However, these two inequalities are equivalent to $q_1 \in [(b-q_2)/(b-a), b - (b-a)q_2]$, as indicated in condition [B].

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 $q_1, \dots, q_{|S|} \in \Delta(K)$ 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

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

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

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

As in Renault, Solan, and Vieille (2017) and Ely (2017), the dynamic decision problem faced by the sender can be 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 , that is, $\mathcal{S}_q$. The payoff associated with the state q and the action $\{(q_i, \alpha_i)\}_{i=1}^{|S|} \in \mathcal{S}_q$ is $\sum_{i=1}^{|S|} \alpha_i u(q_i)$.

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 $\Gamma_\delta(p)$, 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_iM with probability α_i . We set $y_{n+1} = q_iM$ 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} : $y_{n+1} = q_iM$ with probability α_i , namely $y_{n+1} = p_nM$.

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 a 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:¹¹

$$\mathbb{E}_{p, \sigma}(p_{n+1} | p_n) = p_n M, \quad \forall n \geq 1. \quad (4)$$

Consequently, $\mathbb{E}_{p, \sigma} p_{n+1} = (\mathbb{E}_{p, \sigma} p_1) M^n = p M^n$ for every $n \geq 1$. In particular, by taking $p = \pi_M$ we obtain the *mean-consistency property at π_M* : $\mathbb{E}_{\pi_M, \sigma} p_n = \pi_M$ for every $n \geq 1$.

By reducing the problem to MDP and applying the dynamic program principle (e.g., Theorem 1.22 in Solan (2022)) we obtain the following recursive formula for $v_\delta(p)$:

$$v_\delta(p) = \max_{\{(q_i, \alpha_i)\}_{i=1}^{|S|} \in \mathcal{S}_p} \left\{ (1 - \delta) \sum_{i=1}^{|S|} \alpha_i u(q_i) + \delta \sum_{i=1}^{|S|} \alpha_i v_\delta(q_i M) \right\}$$

¹¹In 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 (1999)). 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 (1999) 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, $T(q) = qM$.

$$= \max_{\{(q_i, \alpha_i)\}_i \in \mathcal{S}_p} \sum_{i=1}^{|S|} \alpha_i \cdot \{(1 - \delta)u(q_i) + \delta v_\delta(q_i M)\} \quad (5)$$

where we note that the max in the right-hand side is attained by upper semicontinuity arguments. Moreover, Carathéodory's theorem (see, e.g., Corollary 17.1.5 in [Rockafellar \(1970\)](#)) implies that the expression on right-hand side of Equation (5) is the same for all signal sets $|S| \geq k$. Therefore, v_δ is the same for all signal sets S of cardinality at least k .

Moreover, v_δ may be described in terms of a functional recursive equation. Indeed, consider the operator $\phi : \Delta(K) \rightarrow \Delta(K)$ defined by $\phi(q) = qM$. Then Carathéodory's theorem (see, e.g., Corollary 17.1.5 in [Rockafellar \(1970\)](#)) implies that the expression on right-hand side of Equation (5) equals $(\text{Cav}\{(1 - \delta)u + \delta v_\delta \circ \phi\})(p)$. Thus the following key relation holds:

$$v_\delta(p) = (\text{Cav}\{(1 - \delta)u + \delta(v_\delta \circ \phi)\})(p). \quad (6)$$

In particular, this shows that the function $v_\delta : \Delta(K) \rightarrow \mathbb{R}$ is concave for every δ . As ϕ is linear, $v_\delta \circ \phi$ is also concave. Then, by the definition of Cav, we infer from Equation (6) the inequality

$$v_\delta(p) \leq (1 - \delta)(\text{Cav } u)(p) + \delta(v_\delta \circ \phi)(p). \quad (7)$$

Since the sender can always decide to reveal no information at p , that is, to choose the split $\{(q_i, \alpha_i)\}_i \in \mathcal{S}_p$, where $q_i = p$ for all $i = 1, \dots, |S|$, and thereafter play optimally in the game $\Gamma_\delta(pM)$, we also have that $v_\delta(p) \geq (1 - \delta)u(p) + \delta(v_\delta \circ \phi)(p)$. The latter combined with Equation (7) gives the following result.

LEMMA 5. *Assume that $p \in \Delta(K)$ satisfies $u(p) = (\text{Cav } u)(p)$. Then, for any $\delta \in [0, 1)$, the optimal signaling strategy σ_δ in $\Gamma_\delta(p)$ 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 (e.g., [Renault \(2020\)](#)). 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

$$U_N(p, \sigma) := \mathbb{E}_{p, \sigma} \left(\frac{1}{N} \sum_{n=1}^N u(p_n) \right), \quad \forall \sigma \in \Sigma.$$

The value of $\Gamma_N(p)$ will be denoted by $v_N(p)$. Standard continuity and compactness-based arguments (see, e.g., Theorem 1.15 in [Solán \(2022\)](#)) show that $v_N(p) = \max_{\sigma \in \Sigma} U_N(p, \sigma)$.

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

PROPOSITION 4. *We have the following:*

(i) *For every $N \geq 1$, the function $v_N : \Delta(K) \rightarrow \mathbb{R}$ is concave and satisfies*

$$v_N(p) = \left(\text{Cav} \left\{ \frac{1}{N} \cdot u + \frac{N-1}{N} \cdot (v_{N-1} \circ \phi) \right\} \right)(p), \quad \forall p \in \Delta(K), \quad (8)$$

where $v_0 : \Delta(K) \rightarrow \mathbb{R}$ is defined by $v_0(p) = 0$ for every $p \in \Delta(K)$.

(ii) *There exists $\varepsilon_0 > 0$ such that for every positive $\varepsilon < \varepsilon_0$ it holds that*

$$\|p - \pi_M\|_1 < \varepsilon \implies |v_N(p) - v_N(\pi_M)| \leq \frac{2\|u\|_\infty}{\varepsilon_0} \cdot \varepsilon, \quad \forall N \geq 1.$$

(iii) *The sequence $\{N v_N(\pi_M)\}_N$ is subadditive.*

(iv) *The sequence $\{v_N(\pi_M)\}_N$ converges.*

(v) *The sequence $\{v_{bN}(\pi_M)\}_N$ is nonincreasing for every $b \in \mathbb{N}$.*

(vi) *For every $N \geq 2$ and $p \in \Delta(K)$, it holds that*

$$|v_N(p) - v_{N-1}(pM)| \leq \frac{2\|u\|_\infty}{N-1}.$$

PROOF OF PROPOSITION 3. (i). Applying the dynamical programming principle for Markov decision problems (e.g., Theorem 1.17 in [Solán \(2022\)](#)), we obtain that for every $N \geq 1$ and any $p \in \Delta(K)$ it holds

$$\begin{aligned} v_N(p) &= \sup_{\{(q_i, \alpha_i)\}_{i \in \mathcal{S}_p}} \left\{ \frac{1}{N} \sum_{i=1}^{|S|} \alpha_i u(q_i) + \frac{N-1}{N} \sum_{i=1}^{|S|} \alpha_i v_{N-1}(q_i M) \right\} \\ &= \sup_{\{(q_i, \alpha_i)\}_{i \in \mathcal{S}_p}} \sum_{i=1}^{|S|} \alpha_i \cdot \left\{ \frac{1}{N} \cdot u(q_i) + \frac{N-1}{N} \cdot v_{N-1}(q_i M) \right\}. \end{aligned} \quad (9)$$

As $|S| \geq k$, Carathéodory's theorem (see, e.g., Corollary 17.1.5 in [Rockafellar \(1970\)](#)) implies that the expression on the right-hand side of Equation (9) equals

$$\left(\text{Cav} \left\{ \frac{1}{N} \cdot u + \frac{N-1}{N} \cdot (v_{N-1} \circ \phi) \right\} \right)(p),$$

which in turn with Equation (9) shows that the relation described in (8) holds. In particular, we obtain that v_N is concave for every $N \geq 1$, as by (8) it is the concavification of some function on $\Delta(K)$.

(ii). To show (ii), it will be sufficient to show that

$$\forall p, q \in B_{\ell_1}(\pi_M, \varepsilon) : |v_N(p) - v_N(q)| \leq \frac{2\|u\|_\infty}{\varepsilon_0} \cdot \varepsilon, \quad \forall N \geq 1, \quad (10)$$

where $B_{\ell_1}(\pi_M, \varepsilon) \subset \Delta(K)$ denotes the open ℓ_1 -ball of radius ε centered at π_M .

To show relation (10), set $\varepsilon_0 := (1/2) \cdot \min_{\ell \in K} \pi_M^\ell$, and note that $\varepsilon_0 > 0$ due to the ergodicity of M . Fix a positive $\varepsilon < \varepsilon_0$ and $p, q \in B_{\ell_1}(\pi_M, \varepsilon)$. Let $\partial\Delta(K)$ denote the topological boundary of $\Delta(K)$. Define $w \in \partial\Delta(K)$ to be the point in the intersection of $\partial\Delta(K)$ with the ray starting from q and going through p . By the concavity of v_N proved in (i), it holds that

$$v_N(p) \geq \frac{\|w - p\|_1}{\|w - q\|_1} \cdot v_N(q) + \frac{\|p - q\|_1}{\|w - q\|_1} \cdot v_N(w). \quad (11)$$

Since $\|\pi_M - z\|_1 \geq \min_{\ell \in K} \pi_M^\ell = 2\varepsilon_0$ for every $z \in \partial\Delta(K)$, and $q \in B_{\ell_1}(\pi_M, \varepsilon_0)$, we have that $\|w - q\|_1 \geq \varepsilon_0$. Indeed, otherwise, by the triangle inequality we would have $\|w - \pi_M\|_1 \leq \|w - q\|_1 + \|q - \pi_M\|_1 < 2\varepsilon_0$, in contradiction to the fact that $w \in \partial\Delta(K)$. Subtracting $v_N(q)$ from both sides of (11), and using $\|w - q\|_1 = \|w - p\|_1 + \|p - q\|_1$, we obtain

$$v_N(p) - v_N(q) \geq \frac{\|p - q\|_1}{\|w - q\|_1} \cdot [v_N(w) - v_N(q)] \geq -\frac{2\|u\|_\infty}{\varepsilon_0} \cdot \varepsilon.$$

As the reverse inequality follows by symmetry, relation (10) follows, and thus item (ii) as well.

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

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

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

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

Applying Jensen's inequality to the concave function v_L , and utilizing the mean-consistency property at π_M , we obtain that the above expression is at most

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

Therefore, we have shown that for every $\sigma \in \Sigma$,

$$U_{N+L}(\pi_M, \sigma) \leq (N v_N(\pi_M) + L v_L(\pi_M)) / (N + L).$$

Maximizing over $\sigma \in \Sigma$, and multiplying both sides by $(N + L)$ shows that $(N + L) v_{N+L}(\pi_M) \leq N v_N(\pi_M) + L v_L(\pi_M)$, proving that the sequence $\{N v_N(\pi_M) : N \geq 1\}$ is subadditive.

(iv). 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 subadditive sequence; then $\{a_n/n\}$ converges.

(v). By a repeated use of (iii), we see that if L divides N , then

$$Nv_N(\pi_M) \leq \left(\frac{N}{L} - 1\right)Lv_{\left(\frac{N}{L}-1\right)L}(\pi_M) + Lv_L(\pi_M) \leq \cdots \leq \frac{N}{L}(Lv_L(\pi_M)).$$

Therefore, $v_N(\pi_M) \leq v_L(\pi_M)$, which is sufficient to prove (v).

(vi). First, as the sender can reveal no information in the first stage of $\Gamma_N(p)$, and then follow the optimal strategy in $\Gamma_{N-1}(pM)$ from the second stage on, we have that

$$v_N(p) \geq \frac{u(p)}{N} + \frac{(N-1)v_{N-1}(pM)}{N} \geq v_{N-1}(pM) - \frac{2\|u\|_\infty}{N-1},$$

where we used the fact that $v_{N-1}(\cdot) \leq \|u\|_\infty$. The reverse inequality requires some additional preparations. Equation (9) together with Carathéodory's theorem implies that for any N , $v_N : \Delta(K) \rightarrow \mathbb{R}$ is the same for all signal sets S satisfying $|S| \geq k$. Let $\sigma_* \in \Sigma$ be an optimal strategy in $\Gamma_N(p)$. Assume without loss of generality that it uses k signals only. Let $\mu_2^* \in \Delta(\Delta(K))$ denote the distribution of p_2 under σ_* , starting from the prior p . We have that μ_2^* has mean pM , and is supported on at most k^2 points in $\Delta(K)$.

Consider now the game $\Gamma_{N-1}(pM)$, and define $\hat{\sigma} \in \Sigma$ as follows: at the first stage of $\Gamma_{N-1}(pM)$, split pM to μ_2^* (where we allow $\hat{\sigma}$ to use arbitrary many signals). Then, at any stage $n = 2, \dots, N-1$, given $p_{n-1} = \xi$, split the belief at the start of stage n , being ξM , to the conditional distribution of p_n under σ in $\Gamma_N(p)$, given the event $p_{n-1} = \xi$. The definition of $\hat{\sigma}$ implies that¹²

$$\mathbb{E}_{pM, \hat{\sigma}} u(p_n) = \mathbb{E}_{p, \sigma_*} u(p_{n+1}), \quad \forall n = 1, \dots, N-1.$$

Using the above property, we obtain that

$$\begin{aligned} v_{N-1}(pM) &\geq U_{N-1}(pM, \hat{\sigma}) \\ &= \frac{N}{N-1} \left[U_N(p, \sigma_*) - \frac{\mathbb{E}_{p, \sigma_*}(u(p_1))}{N} \right] \\ &= \frac{N}{N-1} \cdot v_N(p) - \frac{\mathbb{E}_{p, \sigma_*}(u(p_1))}{N-1} \\ &\geq v_N(p) - \frac{2\|u\|_\infty}{N-1}, \end{aligned}$$

where the second equality follows from the optimality of σ_* , and the last inequality uses the fact that $v_N(\cdot) \leq \|u\|_\infty$. $\square$

With the help of Proposition 4, we can proceed to prove Theorem 1.

PROOF OF THEOREM 1. By item (i) of Proposition 4 and the definition of the Cav operator, for any $N \geq 1$ and $p \in \Delta(K)$ we have that

$$\frac{1}{N} \min_{\Delta(K)} u + \frac{1}{N} (v_{N-1} \circ \phi)(p) \leq v_N(p) \leq \frac{1}{N} \max_{\Delta(K)} u + \frac{N-1}{N} (v_{N-1} \circ \phi)(p). \quad (12)$$

¹²Indeed, by construction, for each $n = 1, \dots, N-1$ the marginal distribution of p_n under $\hat{\sigma}$ in $\Gamma_{N-1}(pM)$ is the same as the marginal distribution of p_{n+1} under σ_* in $\Gamma_N(p)$.

Set $\underline{v}(p) := \liminf_{n \rightarrow \infty} v_N(p)$ and $\bar{v}(p) := \limsup_{n \rightarrow \infty} v_N(p)$. Equation (12) implies that $\underline{v}(p) = \underline{v}(pM)$ and $\bar{v}(p) = \bar{v}(pM)$ for every $p \in \Delta(K)$. In particular, we obtain that

$$\underline{v}(p) = \underline{v}(pM) = \dots = \underline{v}(pM^n), \quad \forall n \geq 1. \quad (13)$$

As by item (ii) of Proposition 4, $\underline{v}$ is continuous at π_M , and as $pM^n \rightarrow \pi_M$ by the convergence theorem for Markov chains (e.g., Theorem 4.9 in Levin and Peres (2017)), Equation (13) implies that $\underline{v}(p) = \underline{v}(pM^n) \rightarrow \underline{v}(\pi_M)$ as $n \rightarrow \infty$. Hence, $\underline{v}$ is a constant function. Similarly, as $\bar{v}$ is also continuous at π_M by item (ii) of Proposition 4, one may show using similar arguments that $\bar{v}(p) = \bar{v}(\pi_M)$ for every $p \in \Delta(K)$. Since by item (iv) of Proposition 4 we must have that $\underline{v}(\pi_M) = \bar{v}(\pi_M)$, we obtain that $\bar{v}(p) = \underline{v}(p) = \underline{v}(\pi_M)$ for every $p \in \Delta(K)$. We conclude that for every $p \in \Delta(K)$, $v_N(p) \rightarrow v_\infty$ as $N \rightarrow \infty$, where $v_\infty = \underline{v}(\pi_M)$.

Let us now show that v_N converges uniformly on $\Delta(K)$ to v_∞ . Fix $\varepsilon > 0$ and assume without loss of generality that $\varepsilon < \varepsilon_0$, where ε_0 was described in item (ii) of Proposition 4. By the convergence theorem for Markov chains (e.g., Theorem 4.9 in Levin and Peres (2017)), there exists $n_* \in \mathbb{N}$ such that

$$\|pM^{n_*} - \pi_M\|_1 < \varepsilon, \quad \forall p \in \Delta(K).$$

Next, let $N_* \in \mathbb{N}$ be sufficiently large such that it satisfies (a) $|v_N(\pi_M) - v_\infty| < \varepsilon$ for every $N \geq N_*$, and (b) $n_*/N_* < \varepsilon$. Define $N(\varepsilon) = N_* + n_*$. We have that for every $N > N(\varepsilon)$ and $p \in \Delta(K)$ it holds

$$\begin{aligned} |v_N(p) - v_\infty| &\leq |v_N(p) - v_{N-n_*}(pM^{n_*})| \\ &\quad + |v_{N-n_*}(pM^{n_*}) - v_{N-n_*}(\pi_M)| \\ &\quad + |v_{N-n_*}(\pi_M) - v_\infty| \\ &\leq \sum_{j=1}^{n_*} |v_{N-(j-1)}(pM^{j-1}) - v_{N-j}(pM^j)| \\ &\quad + \frac{2\|u\|_\infty}{\varepsilon_0} \cdot \varepsilon + \varepsilon \\ &\leq \sum_{j=1}^{n_*} \frac{2\|u\|_\infty}{N-j} + \frac{2\|u\|_\infty}{\varepsilon_0} \cdot \varepsilon + \varepsilon \\ &\leq \frac{2\|u\|_\infty}{N_*} \cdot n_* + \frac{2\|u\|_\infty}{\varepsilon_0} \cdot \varepsilon + \varepsilon \\ &\leq 2\|u\|_\infty \cdot \varepsilon + \frac{2\|u\|_\infty}{\varepsilon_0} \cdot \varepsilon + \varepsilon, \end{aligned}$$

proving that v_N converges uniformly to v_∞ , where the second inequality follows from the choice of $N(\varepsilon)$, N_* , and n_* , and from a use of item (ii) of Proposition 4, the third inequality uses item (vi) of Proposition 4, and the last inequality follows from the choice of N_* and the fact that $N - n_* > N_* - n_*$.

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 (1992)), we get that $v_\delta(\cdot)$ must also converge to v_∞ uniformly as $\delta \rightarrow 1^-$.

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

$$\mathbb{E}_{\pi_M, \sigma} u(p_n) \leq \mathbb{E}_{\pi_M, \sigma} (\text{Cav } u)(p_n) \leq (\text{Cav } u)(\mathbb{E}_{\pi_M, \sigma} p_n) = (\text{Cav } u)(\pi_M), \quad (14)$$

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

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

LEMMA 6. *For every $z \in \Lambda$, we have:*

- (i) $(\text{Cav } u)(q) = f_z(q)$ for every $q \in \text{conv}(A_z)$.
- (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 $\sum_i \alpha_i u(q_i) = (\text{Cav } u)(\pi_M)$, then $q_i \in A_z$ for every i .

PROOF OF LEMMA 6. Let $q \in \text{conv}(A_z)$. Take $(q_i) \in A_z$ and convex weights (α_i) such that $q = \sum_i \alpha_i q_i$. Since $(\text{Cav } u)$ is concave and $q_i \in A_z$, we have

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

where the last equality follows from the fact that f_z is affine. Since by the definition of Λ we have $(\text{Cav } u)(q) \leq f_z(q)$ 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

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

We reached a contradiction. $\square$

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*

$$v_\delta(\pi_M) = (\text{Cav } u)(\pi_M), \quad \forall \delta \in [0, 1). \quad (15)$$

PROOF OF PROPOSITION 5. We show that Equation (15) holds via a two-sided inequality. For the first side, by applying Jensen's inequality and using the mean-consistency

property at π_M , we get that, for every $\sigma \in \Sigma$ and every $\delta \in [0, 1)$,

$$\begin{aligned} U_\delta(\pi_M, \sigma) &\leq (1 - \delta) \sum_{n=1}^{\infty} \delta^{n-1} (\text{Cav } u)(\mathbb{E}_{\pi_M, \sigma} p_n) \\ &= (1 - \delta) \sum_{n=1}^{\infty} \delta^{n-1} (\text{Cav } u)(\pi_M) = (\text{Cav } u)(\pi_M). \end{aligned} \quad (16)$$

We therefore obtain that $v_\delta(\pi_M) \leq (\text{Cav } u)(\pi_M)$ for every $\delta \in [0, 1)$. Let us now show that the opposite inequality holds as well. Fix $\varepsilon > 0$. Since $(\text{Cav } u)$ and v_δ are continuous at π_M , there exists $\rho(\varepsilon) > 0$ such that

$$\begin{aligned} \|p - \pi_M\|_2 < \rho(\varepsilon) &\implies |(\text{Cav } u)(p) - (\text{Cav } u)(\pi_M)| < \varepsilon \quad \text{and} \\ |v_\delta(p) - v_\delta(\pi_M)| &< \varepsilon. \end{aligned} \quad (17)$$

Since C is M -absorbing, Brouwer's fixed-point theorem ensures that there exists $p_* \in \text{conv}(C)$ such that $\|p_* - \pi_M\|_2 < \rho(\varepsilon)$. Next, consider for each $q \in \text{conv}(C)$ the set $\mathcal{S}_q^C \subseteq \mathcal{S}_q$ defined by

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

Since $|S| \geq k$, Carathéodory's theorem shows that $\mathcal{S}_q^C \neq \emptyset$. Consider now the strategy σ^C defined as follows: at each $n \geq 1$, if $p_{n-1} = q \in \text{conv}(C)$, σ^C will choose an element in $\mathcal{S}_{qM}^C$; otherwise, if $p_{n-1} = q \in \Delta(K) \setminus \text{conv}(C)$, then σ^C will chose some element in $\mathcal{S}_{qM}$; As $p_* \in \text{conv}(C)$, and $\text{conv}(C)$ is M -absorbing, we have that under the strategy σ^C , $\text{supp}(p_n) \subseteq C$ for every $n \geq 1$. Indeed, we show this by induction on n . For $n = 1$, since $p_* \in \text{conv}(C)$, $\text{supp}(p_1) \subseteq C$ by the definition of σ^C . Assume now that $\text{supp}(p_n) \subseteq C$ for some $n \geq 2$. Since C is M -absorbing, $\text{supp}(p_n M) \subseteq \text{conv}(C)$, and thus by the definition of σ^C we see that $\text{supp}(p_{n+1}) \subseteq C$ as well. The latter, coupled with $C \subseteq A_z$ implies that the discounted payoff under σ^C can be computed as follows:

$$\begin{aligned} U_\delta(p_*, \sigma^C) &= (1 - \delta) \sum_{n=1}^{\infty} \delta^{n-1} \mathbb{E}_{p_*, \sigma^C} [u(p_n)] \\ &= (1 - \delta) \sum_{n=1}^{\infty} \delta^{n-1} \mathbb{E}_{p_*, \sigma^C} [f_z(p_n)] \\ &= (1 - \delta) \sum_{n=1}^{\infty} \delta^{n-1} f_z(\mathbb{E}_{p_*, \sigma^C} p_n) \\ &= (1 - \delta) \sum_{n=1}^{\infty} \delta^{n-1} f_z(p_* M^{n-1}) \\ &= (1 - \delta) \sum_{n=1}^{\infty} \delta^{n-1} (\text{Cav } u)(p_* M^{n-1}) \end{aligned} \quad (18)$$

where we note that the third equality holds because f_z is affine, and the last equality is a consequence of item (i) of Lemma 6. As

$$\|p_* M^{n-1} - \pi_M\|_2 = \|p_* M^{n-1} - \pi_M M^{n-1}\|_2 \leq \|M^{n-1}\|_2 \|p_* - \pi_M\|_2 = \|p_* - \pi_M\|_2,$$

we obtain from Equations (17) and (18) that $v_\delta(\pi_M) \geq U_\delta(p_*, \sigma^C) - \varepsilon \geq (\text{Cav } u)(\pi_M) - 2\varepsilon$. As $\varepsilon > 0$ is arbitrary, we deduce that $v_\delta(\pi_M) \geq (\text{Cav } u)(\pi_M)$, thus proving the opposite inequality. $\square$

We proceed with the proof of Proposition 2.

PROOF OF PROPOSITION 2. Assume that C is an M -absorbing subset of A_z where $z \in \Lambda$. Then, by Proposition 5, $v_\delta(\pi_M) \rightarrow (\text{Cav } u)(\pi_M)$ as $\delta \rightarrow 1^-$. Therefore, Theorem 1 ensures that $v_\infty = (\text{Cav } u)(\pi_M)$, as desired.

Let us continue with item (ii). Assume that $v_\infty = (\text{Cav } u)(\pi_M)$. We have (a) $\lim_{N \rightarrow \infty} v_N(\pi_M) = (\text{Cav } u)(\pi_M)$, (b) $v_N(\pi_M) \leq (\text{Cav } u)(\pi_M)$ for every N (by Equation (14)) and (c) $\{v_{b^N}(\pi_M)\}_N$ is nonincreasing for every $b \in \mathbb{N}$ (by item (v) of Proposition 4). A combination of (a), (b), and (c) shows that $v_N(\pi_M) = (\text{Cav } u)(\pi_M)$ 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^N and the prior probability π_M . By Jensen's inequality, $\mathbb{E}_{\pi_M, \sigma^N}[u(p_n^N)] \leq (\text{Cav } u)(\pi_M)$ for every n . Hence, as $U_N(\pi_M, \sigma^N) = (\text{Cav } u)(\pi_M)$, we obtain that $\mathbb{E}_{\pi_M, \sigma^N}[u(p_n^N)] = (\text{Cav } u)(\pi_M)$ for every $n = 1, \dots, N$. Fix $\delta \in (0, 1)$. We see that

$$U_\delta(\pi_M, \sigma^N) \geq (1 - \delta) \sum_{n=1}^N \delta^{n-1} (\text{Cav } u)(\pi_M) - \delta^N \|u\|_\infty$$

for every $N \geq 1$. Letting $N \rightarrow \infty$, we get that $v_\delta(\pi_M) \geq (\text{Cav } u)(\pi_M)$. Since, by Equation (14), the opposite inequality holds as well, and we deduce that $v_\delta(\pi_M) = (\text{Cav } u)(\pi_M)$. Therefore, there exists a strategy σ^δ such that $U_\delta(\pi_M, \sigma^\delta) = (\text{Cav } u)(\pi_M)$. Denote the sequence of posteriors induced by σ^δ and the prior probability π_M by $(p_n^\delta)_n$. By Jensen's inequality, $\mathbb{E}_{\pi_M, \sigma^\delta}[u(p_n^\delta)] \leq (\text{Cav } u)(\pi_M)$, and we therefore obtain that $\mathbb{E}_{\pi_M, \sigma^\delta}[u(p_n^\delta)] = (\text{Cav } u)(\pi_M)$ for every n . Moreover, as $\text{supp}(p_n^\delta)$ is finite and $\mathbb{E}_{\pi_M, \sigma^\delta} p_n^\delta = \pi_M$ for every n , item (ii) of Lemma 6 implies that $\text{supp}(p_n^\delta) \subseteq A_z$ for every $z \in \Lambda$ and every n . Set $C := \bigcup_{n \geq 1} \text{supp}(p_n^\delta)$. Then $C \subseteq A_z$ 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 $p_n^\delta = q$ with positive probability. Since $\mathbb{E}_{\pi_M, \sigma^\delta}(p_{n+1}^\delta | p_n^\delta = q) = qM$, we obtain that $qM \in \text{conv}(\text{supp}(p_{n+1}^\delta)) \subseteq \text{conv}(C)$. To summarize, C is a countable M -absorbing subset of A_z for every $z \in \Lambda$, as desired. $\square$

PROOF OF PROPOSITION 3. Proposition 3 follows immediately from Propositions 2 and 5. $\square$

The proof of Proposition 1 may enhance the understanding of absorbing sets.

PROOF OF PROPOSITION 1. By Carathéodory's theorem (see, e.g., Corollary 17.1.5 in Rockafellar (1970)), to each $q \in C$ we can assign k distributions $w_1(q), \dots, w_k(q) \in C$ such that $qM \in \text{conv}(\{w_1(q), \dots, w_k(q)\})$. Define a correspondence $\xi : C \rightarrow 2^C$ by $\xi(q) = \{w_1(q), \dots, w_k(q)\}$. In particular, $qM \in \text{conv}(\xi(q))$. 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 $wM \in \text{conv}(\xi(w)) \subseteq \text{conv}(\xi^n(q)) \subseteq \text{conv}(\mathcal{A}(q))$, as desired. $\square$

PROOF OF THEOREM 3. Suppose that M is a homothety and fix a function u . Carathéodory's Theorem (see, e.g., Corollary 17.1.5 in Rockafellar (1970)) implies that there exist points $q_1, \dots, q_m \in \Delta(K)$, $m \leq k$, and positive convex weights $(\alpha_i)_{i=1}^m$ such that $\pi_M = \sum_{i=1}^m \alpha_i q_i$ and $(\text{Cav } u)(\pi_M) = \sum_{i=1}^m \alpha_i u(q_i)$. Hence, by item (ii) of Lemma 6, $q_i \in A_z$ for every i and every $z \in \Lambda$. Therefore, $\pi_M \in \text{conv}(A_z)$ for every $z \in \Lambda$, and since $\text{conv}(A_z)$ is convex, we see that $\text{conv}(A_z)$ is star-shaped around π_M . Hence, since M is a homothety, we get that $\text{conv}(A_z)$ 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 $v_{\infty} = (\text{Cav } u)(\pi_M)$, proving the first direction of Theorem 3.

Suppose now that $v_{\infty} = (\text{Cav } u)(\pi_M)$ for every $u : \Delta(K) \rightarrow \mathbb{R}$. 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 $\pi_M \in \text{int}(\Delta(K))$, 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 $u_{i,n} : \Delta(K) \rightarrow \mathbb{R}$ by

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

By its definition, $u_{i,n}(q) = 1$ 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 have that $(\text{Cav } u)(\pi_M) = 1$ and that $0 \in \Lambda$ (for Λ corresponding to $u = u_{i,n}$), because the hyperplane $f_0(x) = (\text{Cav } u)(\pi_M) = 1$, $x \in \mathbb{R}^k$, supports $(\text{Cav } u)$ 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 see 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 $e_i^n \rightarrow \pi_M$ as $n \rightarrow \infty$, 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 $q \rightarrow Mq$ is a linear operator, to prove that M is a homothety it 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

$$\begin{aligned} \pi_M M &= (\alpha e_1 + (1 - \alpha) e_2) M \\ &= \alpha (e_1 M) + (1 - \alpha) (e_2 M) \\ &= \alpha (\beta_1 e_1 + (1 - \beta_1) \pi_M) + (1 - \alpha) (\beta_2 e_2 + (1 - \beta_2) \pi_M). \end{aligned} \tag{19}$$

By plugging $\pi_M = \alpha e_q + (1 - \alpha)e_2$ into the last expression in Equation (19) 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, we get 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 $|\text{supp}(v)| = 2$, whereas $|\text{supp}(\pi_M)| = k \geq 3$, 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 $\pi_M \in \text{int}(\Delta(K))$, 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 $n \geq N_0$ a function $u_n : \Delta(K) \rightarrow \mathbb{R}$ 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 $v_n \rightarrow \pi_M$ as $n \rightarrow \infty$, we see that $vM \in (v, \pi_M]$. Also, we have

$$\begin{aligned} vM &= \frac{1}{2}(e_i M + e_j M) \\ &= \frac{1}{2}(\beta_i e_i + (1 - \beta_i)\pi_M) + \frac{1}{2}(\beta_j e_j + (1 - \beta_j)\pi_M) \\ &= \left(1 - \frac{\beta_i}{2} - \frac{\beta_j}{2}\right)\pi_M + \beta_i v + \frac{1}{2}(\beta_j - \beta_i)e_j. \end{aligned} \quad (20)$$

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 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. $\square$

PROOF OF THEOREM 4. We begin by proving sufficiency. We argue first that condition [A] implies that M is homothety. Denote $\phi(p) := pa + (1 - p)b$. This is the probability mass assigned by pM to state 1. It turns out that

$$\phi(p) = b + (a - b)p = (1 - (a - b))\pi_M + (a - b)p,$$

proving that M is a homothety with center π_M and ratio $a - b$. Therefore, by Theorem 3, if condition [A] is satisfied, then $v_\infty = (\text{Cav } u)(\pi_M)$.

Consider now condition [B]. Let $y \in A_z \cap [\pi_M, 1]$ for which $[(b - y)/(b - a), b - (b - a)y] \cap A_z \neq \emptyset$ and take $x \in [(b - y)/(b - a), b - (b - a)y] \cap A_z$. Since $y \geq \pi_M$,

$$x \leq b - (b - a)y \leq b - (b - a)\pi_M = (1 + (b - a))\pi_M - (b - a)\pi_M = \pi_M \leq y.$$

Therefore, to show that $\{x, y\}$ is M -absorbing we need to verify that (i) $\phi(y) \geq x$ and (ii) $\phi(x) \leq y$. Elementary computation shows that (i) and (ii) are equivalent to $x \leq b - (b - a)y$ and $(b - y)/(b - a) \leq x$, respectively. Since the last two conditions are satisfied, we obtain that $\{x, y\}$ is an M -absorbing subset of A_z , and thus, by Theorem 2, $v_\infty = (\text{Cav } u)(\pi_M)$.

To show necessity, assume that neither condition [A] nor [B] hold. This implies that M is not a homothety. Moreover, by the sufficiency part, there are no $x \leq \pi_M \leq y \in A_z$ such that $\{x, y\}$ is M -absorbing. We claim that in this case, no subset of A_z is M -absorbing, which stands in contradiction with Theorem 2. Indeed, if some set $C \subset A_z$ is M -absorbing, then so is $\{x_C, y_C\}$, where $x_C = \inf_C p$ and $y_C = \sup_C p$. Then $x_C \leq \pi_M \leq y_C$, and since A_z is closed,¹³ we have $\{x_C, y_C\} \subset A_z$, a contradiction. $\square$

REFERENCES

- Abreu, Dilip, David Pearce, and Ennio Stacchetti (1990), “Toward a theory of discounted repeated games with imperfect monitoring.” *Econometrica*, 58 (5), 1041–1063. [0073]
- Aumann, Robert J. and Michael B. Maschler (1995), *Repeated Games With Incomplete Information*. MIT Press, Cambridge, MA. With the collaboration of Richard E. Stearns. [0085]
- Blackwell, David (1951), “Comparison of experiments.” In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, 1950, 93–102, University of California Press, Berkeley and Los Angeles, Calif. [0085]
- Ely, Jeffrey C. (2017), “Beeps.” *American Economic Review*, 107 (1), 31–53. [0072, 0073, 0074, 0082, 0086]
- Farhadi, Farzaneh and Demosthenis Teneketzis (2022), “Dynamic information design: A simple problem on optimal sequential information disclosure.” *Dynamic Games and Applications*, 12, 443–484. [0074]
- Horn, Roger A. and Charles R. Johnson, eds (2013), *Matrix Analysis*, second edition. Cambridge University Press, Cambridge, U.K. [0079]
- Jackson, Matthew O. and Ehud Kalai (1997), “Social learning in recurring games.” *Games and Economic Behavior*, 21 (1–2), 102–134. [0075]
- Kamenica, Emir and Matthew Gentzkow (2011), “Bayesian persuasion.” *American Economic Review*, 101 (6), 2590–2615. [0071, 0072, 0076]
- Lehrer, Ehud and Sylvain Sorin (1992), “A uniform Tauberian theorem in dynamic programming.” *Mathematics of Operations Research*, 17 (2), 303–307. [0092]
- Levin, David A. and Yuval Peres (2017), *Markov Chains and Mixing Times*, second edition. American Mathematical Society, Providence, RI. With contributions by Elizabeth L. Wilmer. With a chapter on “Coupling from the past” by James G. Propp and David B. Wilson. [0091]
- Neyman, Abraham and Elon Kohlberg (1999), “A strong law of large numbers for nonexpansive vector-valued stochastic processes.” *Israel J. Math.*, 111, 93–108. [0086]

¹³The fact that A_z is closed follows due to (i) u is upper semicontinuous, (ii) f_z is continuous, and (iii) $u \leq f_z$.

Pęski, Marcin and Juuso Toikka (2017), “Value of persistent information.” *Econometrica*, 85 (6), 1921–1948. [0073, 0083, 0084]

Renault, Jerome (2006), “The value of Markov chain games with lack of information on one side.” *Mathematics of Operations Research*, 31 (3), 433–648. [0073, 0083]

Renault, Jerome (2020), “Repeated games with incomplete information.” In *Complex Social and Behavioral Systems—Game Theory and Agent-Based Models*, Encycl. Complex. Syst. Sci., 157–184, Springer, New York. [0073, 0087]

Renault, Jerome, Eilon Solan, and Nicolas Vieille (2017), “Optimal dynamic information provision.” *Games and Economic Behavior*, 104, 329–349. [0072, 0073, 0074, 0082, 0086]

Rockafellar, Tyrrell R. (1970), *Convex Analysis*. Princeton University Press. [0084, 0087, 0088, 0095]

Solan, Eilon (2022), *A Course in Stochastic Game Theory*. Cambridge University Press. [0086, 0087, 0088]

Ziliotto, Bruno (2016), “A Tauberian theorem for nonexpansive operators and applications to zero-sum stochastic games.” *Mathematics of Operations Research*, 41 (4), 1522–1534. [0083]

Co-editor Simon Board handled this manuscript.

Manuscript received 17 July, 2022; final version accepted 14 September, 2023; available online 21 August, 2025.

All authors assume responsibility for all aspects of the paper.