Games and Economic Behavior 142 (2023) 292–314

Contents lists available at ScienceDirect

Games and Economic Behavior

journal homepage: www.elsevier.com/locate/geb

Markovian persuasion with two states ☆

Galit Ashkenazi-Golan [a,*], Penélope Hernández [b], Zvika Neeman [c], Eilon Solan [d]

[a] London School of Economics and Political Science, United Kingdom of Great Britain and Northern Ireland
[b] ERI-CES, Department of Economics, University of Valencia, Spain
[c] Berglas School of Economics, Tel Aviv University, Israel
[d] The School of Mathematical Sciences, Tel Aviv University, Israel

A R T I C L E I N F O

Article history:
Received 29 September 2022
Available online 7 August 2023

JEL classification:
D82
D83

Keywords:
Bayesian persuasion
Information design
Markov games
Repeated games with incomplete information

A B S T R A C T

This paper addresses the question of how to best communicate information over time in order to influence an agent's belief and induced actions in a model with a binary state of the world that evolves according to a Markov process, and with a finite number of actions. We characterize the sender's optimal message strategy in the limit, as the length of each period decreases to zero. We show that the limit optimal strategy is myopic for beliefs smaller than the invariant distribution of the underlying Markov process. For beliefs larger than the invariant distribution, the optimal policy is more elaborate and involves both silence and splitting of the receiver's beliefs; it is not myopic.

1. Introduction

This paper addresses the question of how to best communicate information over time in order to influence an agent's beliefs and induced actions. We consider a model in which a binary state of the world evolves according to a Markov process. In every period, a sender (she) observes the state and sends a message to a myopic receiver (he). The message that is sent by the sender induces the myopic receiver's belief and so also her action in that period, and in addition it also affects the future beliefs of the receiver, and so the way in which the receiver would respond to future messages. The question is how should the sender balance the current and the future implications of her messages.

Ely (2017) and Renault et al. (2017) have studied such models (we provide a detailed discussion of their work below), and have characterized the sender's optimal strategy when the receiver has only two actions. They showed that in this case the sender's optimal strategy is myopic. That is, the sender's optimal policy ignores the effect of the sender's messages on the receiver's future beliefs. In contrast, we allow for any finite number of actions, and find that the larger set of actions calls for a non-myopic sender's optimal strategy.

☆ Ashkenazi-Golan acknowledges the support of the Israel Science Foundation, Grants 217/17 and 722/18, and NSFC-ISF Grant 2510/17. Hernández acknowledges the support of the Consellería d'Innovació, Universitats, Ciència i Societat Digital, Generalitat Valenciana, grant number AICO/2021/257, and the Spanish Ministry of Economics and Innovation, grant number PID2021-128228NB-I00. Neeman acknowledges the support of the Israel Science Foundation, Grant 1465/18. Solan acknowledges the support of the Israel Science Foundation, Grants 217/17 and 211/22.

* Corresponding author.
E-mail addresses: galit.ashkenazi@gmail.com (G. Ashkenazi-Golan), penelope.hernandez@uv.es (P. Hernández), zvika@tauex.tau.ac.il (Z. Neeman), eilons@tauex.tau.ac.il (E. Solan).

https://doi.org/10.1016/j.geb.2023.08.001

For simplicity, we assume that the receiver's action is an increasing function of his belief, and the sender's payoff is an increasing function of the receiver's action. This assumption implies that the receiver's optimal strategy is piecewise-constant in his beliefs. That is, the space of beliefs, which is represented by the unit interval, can be divided into finitely many subintervals, and in each such subinterval the optimal strategy of the receiver is a constant. This in turn implies that the sender's stage payoff, which is a function of the receiver's belief and the receiver's action, is discontinuous in the receiver's belief. We focus on the case in which the sender's indirect payoff, as a function of the receiver's beliefs about the state, satisfies a condition that resembles concavity. Such concavity arises naturally when the sender's marginal benefit from the receiver's action is decreasing.

For example, consider a seller of an experience good such as wine, whose quality changes stochastically over time depending on local climate. The seller prefers that buyers buy as much wine as possible, but obtains a decreasing marginal benefit from each case sold. The seller may disclose information about the wine quality to buyers who decide what quantity of wine to purchase in each period.

The main conceptual contribution of the paper is the understanding of the driving forces behind the sender's optimal strategy. Standard results imply that, in every period, the sender can induce any distribution of posterior beliefs whose mean is equal to the belief in the previous period plus the one-period drift in the Markov process that governs the evolution of the state of the world. Such distributions are said to be "Bayes plausible" (Kamenica and Gentzkow, 2011). We show that the optimal strategy for the sender involves only two types of distributions of induced beliefs. The first distribution arises as a consequence of the sender's silence. In this case, the receiver's belief "slides" toward the invariant distribution of the Markov process. The second distribution requires simple communication and consists of a binary split of the receiver's posterior belief.[1]

Suppose that $p < p'$ are two beliefs that lie below the invariant distribution of the Markov process. Suppose that the current receiver's belief is p. The dynamics pushes the current belief toward p'. The observation mentioned in the previous paragraph suggests two strategies that the sender can use to facilitate this change in the receiver's beliefs: (a) the sender reveals no information until the belief becomes p' ("silence"), and (b) the sender repeatedly splits the receiver's belief between p and p', until the belief finally coincides with p'. It turns out that when comparing the expected discounted time it takes the belief to reach p' under these two strategies, the latter strategy is quicker. A similar result holds when $p > p'$ and the two beliefs lie above the invariant distribution of the Markov process.

Since the sender's payoff is monotone in the receiver's belief, the sender would like the belief to be as high as possible. This "speed-based" argument suggests that when the current receiver's belief is below the invariant distribution, repeated splitting would be better for the sender because it is quicker in generating receiver's beliefs that are more favorable for the sender; and when the current receiver's belief is above the invariant distribution, silence would be better for the sender because it is slower in generating receiver's beliefs that are less favorable for the sender.

However, these two strategies also generate different instantaneous payoffs for the sender: if the sender repeatedly splits the receiver's belief between p and p', then her instantaneous payoff is a weighted average of her instantaneous payoffs at p and p'; and if the sender reveals no information, then her instantaneous payoff is the instantaneous payoff as the beliefs slide from p to p'. The sender's instantaneous payoff is increasing discontinuously in the receiver's beliefs. Payoffs to the left of a discontinuity point are significantly smaller than the payoff at the discontinuity point. This "payoff-based" argument suggests that repeated splitting yields higher instantaneous payoffs than no revelation of information both below and above the invariant distribution of the Markov process.

For receiver's beliefs that lie below the invariant distribution, the speed-based and payoff-based forces are in agreement, and so for such receiver's beliefs, the sender's optimal strategy involves repeated splitting of the receiver's belief between the discontinuity points of the sender's payoff function, and it is myopic. But for receiver's beliefs that lie above the invariant distribution, the speed-based and payoff-based forces work in opposite directions. We show that for such beliefs the sender's optimal strategy is not myopic: at beliefs that are slightly above a discontinuity point, the difference between the instantaneous payoffs of the two strategies is small, the payoff-based argument dominates, and the sender reveals no information[2]; while at beliefs that are slightly below a discontinuity point, the difference between the instantaneous payoffs of the two strategies is large, the speed-based argument dominates, and the sender repeatedly splits the receiver's belief between the belief at the discontinuity point and a well-chosen belief below it.

In the context of the example of the wine seller, when buyers believe that wine quality is lower than average, then the seller's optimal strategy is myopic. In this case, the seller need not worry about the buyers' future beliefs. But when buyers believe that wine quality is higher than average, the seller's optimal strategy is not myopic. In this case the seller's optimal policy is more elaborate and it involves both silence and splitting of the buyers' beliefs.

[1] Ely et al. (2015) state that "a period generates more suspense if the variance of the next period's belief is greater. A period generates more surprise if the current belief is further from last period's belief." Accordingly, periods of "silence," in which the receiver's belief slides predictably towards a belief in which they will be split involve suspense, because a discontinuous change of beliefs is anticipated; and periods in which the belief is split involve surprise, because splitting generates a discrete jump in the receiver's belief. As explained below, the optimal sender's message strategy switches between suspense and surprise, depending on whether the current receiver's belief is above or below the invariant distribution.

[2] Thus, in Ely et al.'s terminology, this silence generates suspense.

Related literature

Our work relates to two distinct research directions that should probably be more closely linked to each other: the first is about information design and Bayesian persuasion, and the second is about repeated games under incomplete information. For recent surveys of these two directions, see Kamenica (2019) and Mertens et al. (2016), respectively.

The model studied in this paper is a dynamic extension of the static Bayesian persuasion model of Kamenica and Gentzkow (2011). In the first part of his paper, Ely (2017) studied a special case of our model. In the second part, he considered a general model without any assumptions on the payoff functions, and characterized the value function as a fixed point of a certain functional equation, which expresses the dynamic programming principle. He showed that iterating this functional equation converges, so that, in principle, it is possible to solve for the value function by starting from some arbitrary function, and iterating it by using the functional equation. However, computing the value function in this way is analytically intractable and it does not allow us to identify the qualitative properties of the solution. In the case of binary actions and states, where one of the two states is absorbing, Ely (2017) characterized the optimal strategy, and showed that it reveals the (absorbing) state with delay. A delay of time T implies that, starting at time zero, the receiver's beliefs that the state has switched slides upwards with time. Denote the belief that the state is absorbing at time t by p^t. Under the sender's optimal strategy, at time T, the receiver still hasn't learned anything about the state, and so his belief p^T reflects the knowledge that no switch has occurred until at least T moments ago. For $t > T$, the fact that the sender reveals that the state has switched with delay T implies that the receiver's beliefs are either given by $p^t = 1$ or $p^t = p^T$, because a receiver who is told that the state has switched updates his belief to $p^t = 1$, and a receiver who hasn't heard anything yet knows only that the state was not absorbing T moments ago, and so has beliefs $p^t = p^T$. Thus, the optimal strategy identified by Ely (2017) also combines sliding and splitting between the two beliefs $p = p^T$ and $p = 1$.

Silence in our model can also be interpreted as delay in Ely's (2017) model. But unlike in Ely's model, the silences that are part of the optimal strategy identified here vary in length, and are punctuated by messages that induce beliefs that reflect different levels of certainty about the current state.

Another related paper is that of Renault et al. (2017), who consider a similar model, again with just two actions. They show that with two states, a greedy or myopic strategy is optimal for the sender (the optimal strategy in Ely's (2017) model is myopic as well). As explained above, this is not the case here. Renault et al. (2017) also show that with more than two states, the optimal strategy for the sender need not be myopic, and provide a sufficient condition on the Markov dynamic process that ensures it is myopic.

Recent papers by Ball (2019), Ely and Szydlowski (2020), and Smolin (2021) obtain results about the timing of the optimal revelation of information in specific settings. Their focus is more on the optimal time to reveal information, rather than what information to reveal as in this paper. There is also a large literature in economics on the design of information feedback in dynamic principal-agent problems and games (see, e.g., the literature review in Ely (2017)). However, as noted by Ely, with a few exceptions, these papers generally consider exogenous information structures or compare a few policies such as full, public, and no disclosure.

The two key methods that are used in the Bayesian persuasion literature are Bayes plausibility and the geometric characterization of the optimum through a concavification of the sender's indirect payoff function. Both of these ideas were adapted from the work of Aumann et al. (1995), who studied repeated games with one-sided incomplete information. In Aumann et al. (1995), one of two players learns which of several two-player zero-sum normal form stage games is to be played, and then this game is played repeatedly. Their analysis has been extended by Renault (2006) to cover Markov games. In the setting with incomplete information studied by Aumann et al. (1995), the sender reveals information only once, in the first period of the game, and then continues to play in a way that is uninformative for the rest of the game. In contrast, in the Markov game studied by Renault (2006), the state of the world changes over time, and so it is optimal for the sender to continue to reveal information about the state as it evolves. Cardaliaguet et al. (2016) and Gensbittel (2019) show that the value of a continuous-time Markov game is given by the solution of a differential equation (but stop short of obtaining explicit solutions). Ashkenazi-Golan et al. (2020) present an algorithm that converges to a solution of this differential equation, but they only apply it to a few examples of two-player zero-sum two-state Markov games with one-sided information. They provide the important insight that it is possible to characterize the optimal strategy by using the derivative of the putative value function given a specific suggested split of the uninformed player's beliefs. The optimal strategy in our set up is similar to their optimal strategy.

As mentioned above, our assumptions about the structure of the family of the sender's indirect payoff function allow us to obtain an explicit characterization of the optimal information strategy for the sender in a class of dynamic Bayesian persuasion problems. In addition, we also compute the expected discounted time it takes to switch from one induced posterior belief to another, which is relevant also to two-player Markov games. We rely on similar ideas to those in Ashkenazi-Golan et al. (2020) to obtain an explicit description of the solution for a class of dynamic persuasion games. However, while the entire literature on repeated games under incomplete information has restricted its attention to the special case where the informed player's indirect payoff function is *continuous* in the uninformed player's beliefs, we study *discontinuous* indirect payoff functions. And the finiteness of the receiver's set of actions assumed here necessarily implies that the sender's indirect payoff function is indeed discontinuous.[3]

The rest of the paper proceeds as follows. In Section 2, we present the model. In Section 3, we present our main results and discuss possible extensions. All proofs are relegated to the Appendix.

2. Model and main result

We consider a discrete-time game with two players: a sender (she) and a receiver (he). In every period $n \in \{1, 2, \ldots\}$, the sender observes the state of the world $\omega^n \in \{0, 1\}$ and sends a message $m^n \in M$ to the receiver, who takes an action $a^n \in A$. The set A is assumed to be finite and the set M contains at least two messages.

Markov transitions.

The probability that the initial state is 0 is given, and denoted by p^0.

In each period there is a constant probability, which may depend on the current state but is independent of the history of the play prior to the current period, that the state switches to the other state in the next period. To eliminate trivial cases, we assume that the transition probabilities are positive and less than one. The switches between the states give rise to a Markov chain, which has a stationary distribution. We denote the probability that the state is 1 according to the stationary distribution by p^*.

Posterior beliefs.

We assume that the sender is committed to her message strategy, and the receiver is aware of this commitment. As a result, at the beginning of each period n, the receiver updates the probability p^n that the state is 1 given the message he received in stage $n-1$ while taking into account the Markov transition. Standard results in the theory of Markov chains imply that starting with any initial probability p^0, when the sender reveals no information about the state, the probability p^n converges to p^* as n increases.

Stage payoffs.

In any period n, the (period-n) payoffs of both the sender and receiver are functions of the state ω^n and the receiver's action a^n in period n. The receiver is assumed to be myopic: in every period n, he chooses the action a^n that maximizes his payoff given his belief p^n over the states.[4] It follows that the sender's (indirect) payoff in period n, denoted $u : [0, 1] \to \mathbb{R}$, may be viewed as a function of the receiver's belief in that period, p^n.

We assume that the receiver's action is increasing in his belief, and the sender's payoff is increasing in the receiver's action. Monotonicity together with the fact that the set of actions A is finite imply that u is increasing and piecewise constant. We assume, in addition, that the function u has a concave linear interpolation. That is, the piecewise linear function that connects all the discontinuity points of u is concave (see Fig. 1).

Assumption 1. There exist[5] two numbers $m \geq 0$ and $m' \geq 1$, $m + m' + 1$ thresholds $0 = p_{-m} < p_{-m+1} < \cdots < p_0 < p_1 < \cdots < p_{m'} = 1$, and $m + m'$ payoffs $h_{-m} < h_{-m+1} < \cdots < h_{m'-1}$ such that $u(p) = h_i$ for $p \in [p_i, p_{i+1})$ if $i \in \{-m, \ldots, m'-1\}$, and $u(1) = h_{m'-1}$. Moreover, the line segment that connects the points (p_{i-1}, h_{i-1}) and (p_{i+1}, h_{i+1}) lies below the point (p_i, h_i) for $i \in \{-m+1, \ldots, m'-2\}$.

We refer to the intervals $[p_i, p_{i+1})$ mentioned in Assumption 1 above as "continuity intervals" of u. Without loss of generality, we assume that $p^* \in [p_0, p_1)$.

Objective of the game.

The length of a period is denoted by $\Delta > 0$. The sender's objective is to maximize her discounted payoff, calculated with respect to her discount rate $r > 0$. The value at the initial belief p^0 is

$$v_\Delta(p^0) := \max_\sigma \mathbb{E}_{p^0,\sigma}\left[(1 - e^{-r\Delta}) \sum_{n=1}^{\infty} e^{-r\Delta n} u(p^n)\right], \tag{1}$$

where $\mathbb{E}_{p^0,\sigma}$ is the probability distribution over plays induced by p^0 and σ, and the maximum is over all sender's message strategies σ. We denote the game described above by $G_\Delta(u)$.

[3] Moreover, even if the receiver has an infinite number of actions available, strong continuity requirements need to be imposed on the payoffs to ensure that the sender's indirect payoff function would be continuous.

[4] We assume that when indifferent, the receiver chooses the action that is better for the sender. This assumption does not affect the existence of the value, and ensures the existence of a sender's optimal message strategy (rather than an ε-optimal strategy). If, when indifferent, the receiver chooses the action that is less favorable to the sender with a positive probability, then the sender would ensure that the receiver is never indifferent between these two actions.

[5] The case $m = 0$ is understood as $0 = p_0 < p_1 < \ldots$.

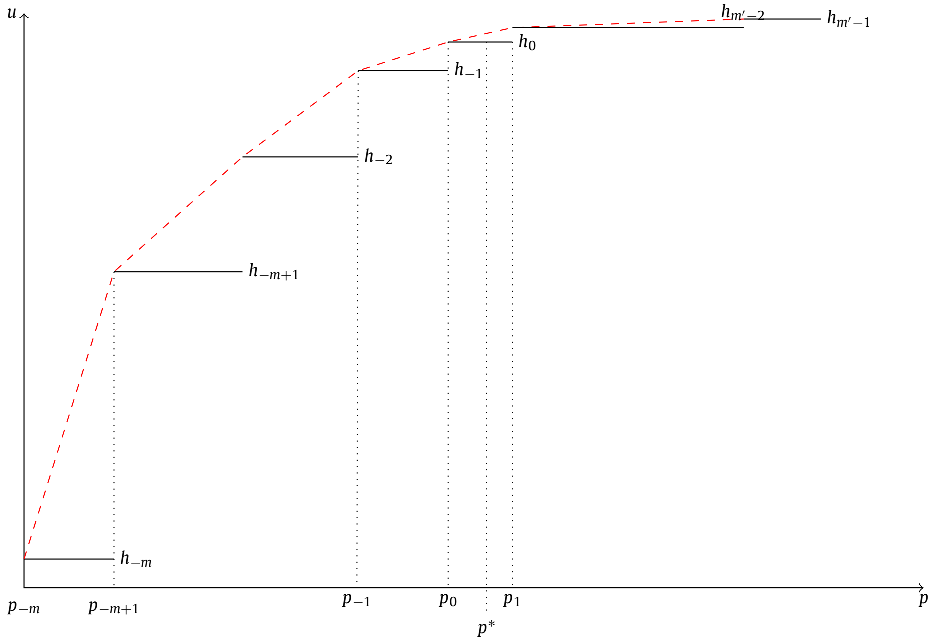

Fig. 1. The function u (solid line) and its linear interpolation (dashed line).

Continuous-time game. We are interested in characterizing the value and the sender's optimal message strategy when the length of a period Δ is small. To this end, we study the corresponding continuous-time game denoted $G_{cont}(u)$.

To properly relate the games in discrete time to the game in continuous time, we assume that in $G_\Delta(u)$, the per-period probability of switching from state 1 to 0 is $1 - e^{-\lambda_1 \Delta}$, and the per-period probability of switching from state 0 to 1 is $1 - e^{-\lambda_0 \Delta}$, where both $\lambda_0, \lambda_1 > 0$. In the game $G_{cont}(u)$, the generator of the Markov chain is

$$R = \begin{pmatrix} -\lambda_0 & \lambda_0 \\ \lambda_1 & -\lambda_1 \end{pmatrix},$$

and the stationary probability of state 1 is[6]

$$p^* = \frac{\lambda_0}{\lambda_0 + \lambda_1}.$$

In the game $G_{cont}(u)$, the state variable is the receiver's belief p^t, and this belief determines the receiver's and sender's instantaneous payoffs. When no information is revealed, the belief changes as a result of the Markov transition as follows:

$$\frac{\partial p^t}{\partial t} = -\lambda_1 p^t + \lambda_0(1 - p^t) = (\lambda_0 + \lambda_1)(p^* - p^t). \tag{2}$$

This implies that for every strategy of the sender, the process $(p^t)_{t \geq 0}$ satisfies

$$\mathbb{E}[p^{t+h} \mid p^t] = p^* + (p^t - p^*)e^{(-(\lambda_0+\lambda_1)h)}, \quad \forall t, h \geq 0. \tag{3}$$

The set of sender's message strategies can therefore be identified with the set of càdlàg processes $(p^t)_{t \geq 0}$ that satisfy Eq. (3), where p^0 is the initial belief.

Denote by v_{cont} the value function of $G_{cont}(u)$:

$$v_{cont}(p) := \sup_\sigma \mathbb{E}\left[\int_0^\infty re^{-rt} u(p^t) \mathrm{d}t\right],$$

where σ ranges over all message strategies of the sender. We will characterize the value and the sender's optimal message strategy σ^* in $G_{cont}(u)$. We will then prove that $v_{cont} = \lim_{\Delta \to 0} v_\Delta$, and that σ^* is approximately optimal in $G_\Delta(u)$, provided Δ is sufficiently small.

[6] For this, as well as all the other results on Markov chains used in this paper, see, e.g., Norris (1998).

Our contribution. Our main result is Theorem 1 below, which provides a characterization of the optimal sender's message strategy in the continuous-time game $G_{cont}(u)$, from which we can calculate the value function v_{cont}. Theorem 1 also implies that the value function in the continuous-time problem approximates the value function in the discrete-time problem, and that a discrete-time approximation of the optimal sender's message strategy in the continuous-time game is approximately optimal in the discrete-time game $G_\Delta(u)$.

We show that the optimal sender's message strategy is Markovian: the play at each time instance t depends only on the receiver's belief at that time. In addition, the optimal strategy involves only two types of behaviors: either the sender reveals no information, or the sender sends one of two messages, which split the receiver's belief into two possible beliefs. This property is a consequence of the fact that there are two states.

The theorem shows that the optimal sender's message strategy is different for receiver's beliefs that lie below p^* (where the sender's payoff is smaller than her payoff at p^*) and for beliefs that lie above p^* (where the sender's payoff is larger than her payoff at p^*). For any receiver's belief that belongs to a continuity interval $[p_{-j}, p_{-j+1})$ that lies to the left of the invariant distribution p^* as well as for beliefs that belong to the continuity interval $[p_0, p_1)$ that contains the invariant distribution (when $p^* \in (p_0, p_1)$), the sender splits the receiver's belief between the endpoints of the interval that contains it. For receiver's beliefs that belong to continuity intervals $[p_j, p_{j+1})$ that lie above the interval $[p_0, p_1)$, the sender's optimal behavior is different: for each such continuity interval, there is a cutoff $q_j \in (p_j, p_{j+1})$ such that at beliefs in $[p_j, q_j]$ the sender reveals no information, while at receiver's beliefs in $[q_j, p_{j+1})$ the senders split the belief between q_j and p_{j+1}. Moreover, the optimal strategy in the continuous-time problem is almost optimal in the discrete-time problem, provide Δ is sufficiently small.

Theorem 1. *Suppose that the indirect payoff function u satisfies Assumption 1. The game in continuous time $G_{cont}(u)$ admits a value function v_{cont}, and the following Markovian message strategy σ^* of the sender is optimal in the continuous-time game:*

- *If $p^* = p_0$, then at p^* the sender reveals no information.*
- *If $p^* \in (p_0, p_1)$, then at every $p \in [p_0, p_1]$ the sender splits the belief into p_0 and p_1.*
- *For every $j \in \{0, \dots, m-1\}$ and every $p \in [p_{-(j+1)}, p_{-j})$, at p the sender splits the belief into $p_{-(j+1)}$ and p_{-j}.*
- *For every $j \in \left\{1, \dots, m'-2\right\}$, there is $q_j \in (p_j, p_{j+1})$ such that*
  - *at every receiver's belief $p \in [p_j, q_j]$, the sender reveals no information, and*
  - *for every receiver's belief $p \in (q_j, p_{j+1}]$, at p the sender splits the receiver's belief p into q_j and p_{j+1}.*

Moreover, for every $\varepsilon > 0$ there is $\Delta_0 > 0$ such that σ^* is ε-optimal when the gap between stages is Δ, for every $\Delta \in (0, \Delta_0)$.

The intuition for this result is as follows. Suppose that the receiver's current belief is $p < p^*$, and the sender would like to have the belief p reach some belief $p' \in (p, p^*]$. There are two simple ways in which the sender can achieve this goal: (i) she can reveal no information, and let the belief slide towards p' because of the Markov transition, or (ii) she can let the belief move slightly towards p^* because of the Markov transition and immediately reveal information to the receiver in such a way that the receiver's belief is split between p and p'. It turns out that when discounting is taken into account, repeated splitting of the receiver's belief achieves a faster convergence to the belief p' than sliding. When $p < p^*$, the monotonicity of the sender's payoff implies that repeated splitting is superior to sliding, and the concavity of the payoffs implies that it is optimal for the sender to split the belief between p and the discontinuity point of u to its right. When $p > p' \geq p^*$, repeatedly splitting the belief between p and p' still converges to p' faster than sliding, but now the monotonicity of u does not imply that sliding is better than splitting. On the one hand, to reach from p_{j+1} to p_j, sliding yields the sender the payoff h_j for as long as possible. On the other hand, repeated splitting allows the sender to obtain the higher payoff $h_{j+1} = u(p_{j+1})$. For beliefs p on the continuity interval $[p_j, p_{j+1}]$ that are close to p_j, delaying the arrival to p_j by revealing no information turns out to be optimal. The situation is reversed for beliefs in this continuity interval that are close to p_{j+1}. For such beliefs, splitting between p_j and p_{j+1} generates the belief p_{j+1} with high probability, which implies that splitting is better than sliding.

Comparison with existing literature. Ely (2017) and Renault et al. (2017) studied the model with a single discontinuity point (where concavity has no bite). Their model corresponds to our model: if p^* is below the discontinuity point, then $m = 0$ and $m' = 2$; if p^* is at least the discontinuity point and less than 1, then $m = m' = 1$; and if $p^* = 1$, then $m = 2$ and $m' = 0$. These authors proved that the myopic strategy is optimal. In our setup, the myopic message strategy of the sender uses binary splits below p_0 and can be either sliding or splitting above p_0. Theorem 1 shows that when $m' > 1$, the sender's optimal message strategy is myopic at p^*, and at all continuity intervals to the left of p^*, but is not myopic at the continuity intervals that lie to the right of p^*. Ashkenazi-Golan et al. (2020) studied the model when u is continuous (rather than piecewise constant and concave), and provided an algorithm for calculating the value function and the sender's optimal message strategy in the continuous-time game. Cardaliaguet et al. (2016) studied the model with finitely many states and continuous payoff function, characterized the value function of the continuous-time game as the viscosity solution of a certain equation, and proved that the value of the discrete-time game converges to the value of the continuous-time game as the inter-stage duration goes to 0. Theorem 1 extends the results of Ashkenazi-Golan et al. (2020) and the approximation result of Cardaliaguet et al. (2016) to discontinuous u.

2.1. Sketch of the proof

In this subsection we highlight the main ideas behind the proof of Theorem 1. The detailed proof appears in Section 3, and technical aspects are relegated to the appendix.

The value function v_{cont} was studied and characterized by Cardaliaguet et al. (2016), for the case in which the payoff function u is continuous. When the payoff function is not continuous, as is the case here, their characterization of v_{cont} is not valid, and the existence of a sender optimal strategy is not guaranteed.

We use the result of Cardaliaguet et al. (2016) to prove the existence of the value for indirect payoff functions u that satisfy Assumption 1, and to characterize the sender's optimal message strategy. This is done by bounding the discontinuous function u by continuous functions $\overline{u}_\delta$, which approach u from above as δ decreases to zero. We show that as δ decreases to zero the respective values, $\overline{v}_\delta$ converge to the value that is obtained by using the strategy σ^* when the payoff function is the discontinuous function u. The monotonicity of the value implies that this is also the value for the discontinuous payoff function u.

The proof, including the construction of σ^*, is presented in Section 3. In Section 3.1, we introduce continuous payoff functions $\overline{u}_\delta$, that are higher than or equal to u and approximate u as δ decreases to 0. In Section 3.2, we characterize the optimal strategies for games with payoff functions $\overline{u}_\delta$. Section 3.3 returns to the value function u and specifies the optimal strategy in continuous time. In Section 3.4, we show that the sender's optimal message strategy for the continuous-time game has a close strategy which is approximately optimal for the discrete-time game.

3. Analysis

In this section we present the detailed proof of Theorem 1. As mentioned above, the results of Cardaliaguet et al. (2016) hold when the indirect payoff function u is continuous. When the payoff function is upper semi-continuous, as in the case here, the existence of an optimal strategy follows from the arguments of Cardaliaguet et al. (2016), yet it is not known whether their characterization of v_{cont} is valid. Since our proof uses the characterization of v_{cont} by Cardaliaguet et al. (2016), we will approximate u from above by continuous functions.

3.1. The approximating continuous functions

For every $\delta > 0$ such that $\delta < p_{j+1} - p_j$ for every $j = -m, -m+1, \ldots, m' - 1$, define a payoff functions $\overline{u}_\delta$ as follows (see Fig. 2):

$$\overline{u}_\delta(p) := \begin{cases} h_j, & p \in [p_j, p_{j+1} - \delta],\ j \in \{-m, \ldots, m'-2\}, \\ \left(\frac{h_{j+1} - h_j}{\delta}\right) \cdot (p - p_{j+1}) + h_{j+1}, & p \in [p_{j+1} - \delta, p_{j+1}],\ j \in \{-m, \ldots, m'-2\}, \\ h_{m'-1}, & p \in [p_{m'-1}, p_{m'}]. \end{cases}$$

Since u has a concave linear interpolation, so does the function $\overline{u}_\delta$. The sequence $(\overline{u}_\delta)_{\delta>0}$ is nonincreasing (as δ goes to 0) and converges pointwise to u. Denote by $\overline{v}_\delta$ the value function of the game $G_{cont}(\overline{u}_\delta)$. Since the sequence $(\overline{u}_\delta)_{\delta>0}$ is nonincreasing, the sequence $(\overline{v}_\delta)_{\delta>0}$ of value functions is nonincreasing as well (as δ goes to 0). Denote the limit value function by

$$\overline{v}_0(p) := \lim_{\delta \to 0} \overline{v}_\delta(p), \quad \forall p \in [0, 1].$$

Since $\overline{u}_\delta \geq u$, we have $\overline{v}_0 \geq v_{cont}$, that is,

$$\overline{v}_0(p) \geq v_{cont}(p), \quad \forall p \in [0, 1]. \tag{4}$$

We will prove that in fact Eq. (4) holds with equality: the limit of the value functions of the approximating games is the value function of the original problem in continuous time.

3.2. Characterizing the optimal strategy in $G_{cont}(\overline{u}_\delta)$

The heart of the proof is the characterization of the sender's optimal message strategy $\overline{\sigma}_\delta$ in the game $G_{cont}(\overline{u}_\delta)$, for $\delta > 0$ sufficiently small, which is displayed in Fig. 3. The strategy $\overline{\sigma}_\delta$ is Markovian; for each $j \in \{1, \ldots, m\}$ there is a real number $\overline{q}_{-j}(\delta) \in (p_{-j} - \delta, p_{-j})$ such that in the interval $[p_{-j-1}, \overline{q}_{-j}(\delta)]$ the sender splits the receiver's belief between the endpoints of the interval, and in the interval $[\overline{q}_{-j}(\delta), p_{-j}]$ the sender reveals no information. Similarly, for each $j = 1, \ldots, m' - 1$ there is a real number $\overline{q}_j(\delta) \in (p_j, p_{j+1})$ such that in the interval $[p_j, \overline{q}_j(\delta)]$ the sender reveals no information, and in the interval $[\overline{q}_j(\delta), p_{j+1}]$ the sender splits the receiver's belief between the endpoints of the interval.

The formal statement follows.

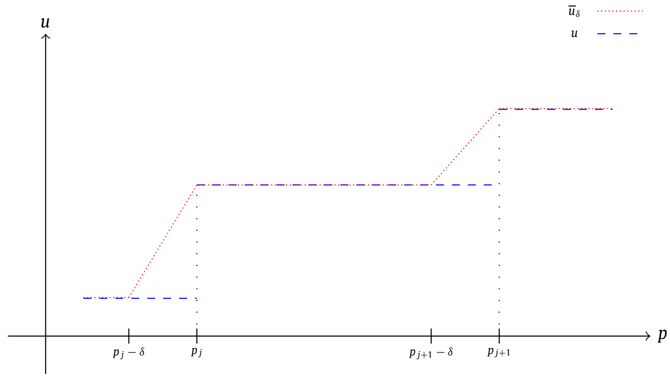

Fig. 2. The continuous payoff functions $\overline{u}_\delta$, approximating u from above.

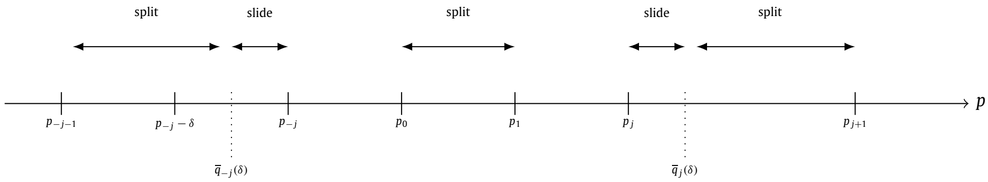

Fig. 3. The characterization of the optimal strategy in $G_{cont}(\overline{u}_\delta)$.

Lemma 1. *Let $\delta > 0$ be sufficiently small. For every $j \in \{0, 1, \dots, m-1\}$ there exists $\overline{q}_{-j}(\delta) \in (p_{-j} - \delta, p_{-j})$, and for every $j \in \{1, 2, \dots, m'-1\}$ there exists $\overline{q}_j(\delta) \in (p_j, p_{j+1})$, such that the sender's optimal message strategy in $G_{cont}(\overline{u}_\delta)$, denoted $\overline{\sigma}_\delta$, is as follows:*

- *For $p \in [p_0, p_1]$, split the belief between p_0 and p_1.*
- *For $p \in [p_{-j-1}, \overline{q}_{-j}(\delta)]$, split the belief between p_{-j-1} and $\overline{q}_{-j}(\delta)$, for every $j \in \{1, \dots, m-1\}$ and $j = 0$ if $p^* > p_0$.*
- *For $p \in [\overline{q}_{-j}(\delta), p_{-j}]$, reveal no information, for every $j \in \{1, \dots, m-1\}$ and $j = 0$ if $p^* > p_0$.*
- *If $p^* = p_0$, then for $p \in [p_{-1}, p_0)$ split the belief between p_{-1} and p_0, for $p \in (p_0, p_1]$ split the belief between p_0 and p_1, and for $p = p^*$ reveal no information.*
- *For $p \in [p_j, \overline{q}_j(\delta)]$ reveal no information, for every $j \in \{1, \dots, m'-1\}$.*
- *For $p \in [\overline{q}_j(\delta), p_{j+1}]$, split the belief between $\overline{q}_j(\delta)$ and p_{j+1}, for every $j \in \{1, \dots, m'-1\}$.*

The proof of Lemma 1 requires a careful analysis of the characterization of the value function due to Cardaliaguet et al. (2016) and Gensbittel (2019), and relies on the special structure of the payoff function that is implied by Assumption 1. The proof is relegated to the Appendix.

3.3. The optimal strategy in $G_{cont}(u)$

As discussed in Lemma 1, the strategy $\overline{\sigma}_\delta$ is determined by the cut-offs $(\overline{q}_{-j}(\delta))_{j=0}^{m-1}$ and $(\overline{q}_j(\delta))_{j=1}^{m'-1}$. By taking a subsequence, we can assume w.l.o.g. that the limits

$$q_j := \lim_{\delta \to 0} \overline{q}_j(\delta), \quad j \in \{-m+1, \dots, m'-1\},$$

exist. Moreover, for $j \in \{-m+1, \cdots, 0\}$ we have $p_j - \delta < \overline{q}_j(\delta) < p_j$, and hence $\lim_{\delta \to 0} \overline{q}_j(\delta) = p_j$. Let σ^* be the strategy that is defined by these limits, see Fig. 4.

- For $p \in [p_0, p_1]$ split the belief between p_0 and p_1.
- For $p \in [p_{-j-1}, p_{-j}]$, split the belief between p_{-j-1} and p_{-j}, for every $j \in \{0, \dots, m-1\}$.
- For $p \in [p_j, q_j]$ reveal no information, for every $j \in \{1, \dots, m'-1\}$.
- For $p \in [q_j, p_{j+1}]$, split between q_j and p_{j+1}, for every $j \in \{1, \dots, m'-1\}$.

As the following result states, the payoff under σ^* in $G_{cont}(u)$ is $\overline{v}_0$. This holds because the behavior of the sender under $\overline{\sigma}_\delta$ converges to her behavior under σ^*.

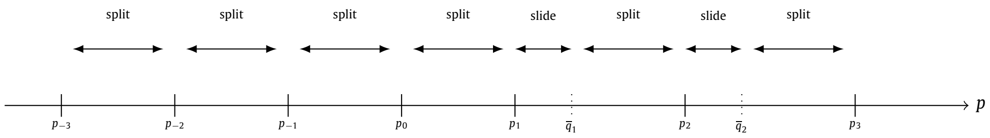

Fig. 4. The strategy σ^* in $G_{cont}(u)$.

For every message strategy σ of the sender and every $p \in [0,1]$, denote the putative value under σ at p in $G_{cont}(u)$ by

$$w(p,\sigma) := \mathbb{E}_{p^0,\sigma}\left[\int_{t=0}^{\infty} e^{-rt}u(p^t)dt\right], \tag{5}$$

Lemma 2. *For every $p \in [0,1]$ we have $\overline{v}_0(p) = w(p,\sigma^*)$.*

Since there is a sender's message strategy that guarantees the payoff $\overline{v}_0$, we have

$$v_{cont}(p) \geq \overline{v}_0(p), \quad \forall p \in [0,1].$$

Together with Eq. (4) this implies that

$$v_{cont} = \overline{v}_0, \tag{6}$$

and that σ^* is an optimal strategy in $G_{cont}(u)$.

The proof is inductive: We show that $\overline{v}_0 = w(\cdot,\sigma^*)$ on $[p_0, p_1]$, and continue to show the same on continuity intervals $[p_{-j-1}, p_{-j}]$ and $[p_j, p_{j+1}]$ with larger j's. The proof appears in Section 6.9.1.

3.4. The value function in continuous time as an approximation of the value function in discrete time

So far we characterized the value function and the sender's optimal message strategy in the continuous-time game. Here we complete the proof of Theorem 1, by showing that σ^*, when interpreted as a strategy in the discrete-time game $G_\Delta(u)$, is approximately optimal, provided Δ is sufficiently small.

Cardaliaguet et al. (2016) proved that if the instantaneous payoff function u is Lipschitz for the L^1-norm, then the sender's optimal message strategy in the continuous-time game is approximately optimal in $G_\Delta(u)$, provided Δ is sufficiently small. In our model u is not continuous, hence we cannot apply the results of Cardaliaguet et al. (2016). Our proof is divided into two steps. Lemma 3 states that the payoff under the strategy σ^* in $G_\Delta(u)$ approaches v_{cont}, the payoff under σ^* in $G_{cont}(u)$ as Δ goes to 0. Lemma 4 then implies that v_{cont} is close to the value of the discrete-time game.

The strategy σ^* is a strategy in the continuous-time game $G_{cont}(u)$: For each receiver's belief p it indicates whether the sender reveals no information, or whether she splits the receiver's belief between two beliefs. However, it can be viewed also as a strategy in the discrete-time game $G_\Delta(u)$: At every stage n, as a function of the current receiver's belief p^n, it reveals no information if σ^* reveals no information at p^n, and otherwise it splits the belief as σ^* does. To avoid cumbersome notations, we denote the strategy induced in the discrete-time game by σ^* as well.

Denote the putative value obtained by the sender under the message strategy σ^* in the discrete-time game $G_\Delta(u)$ by $w_\Delta(p,\sigma^*)$. The next lemma, which states that the putative value under σ^* in the discrete-time game converges, as Δ goes to 0, to the putative value σ in the continuous-time game, follows by a standard limiting argument.

Lemma 3. *For every $p \in [0,1]$, $\lim_{\Delta\to 0} w_\Delta(p,\sigma^*) = w(p,\sigma^*) = v_{cont}(p)$.*

The proof of Lemma 3 appears in Section 6.9.2.

Lemma 3 implies that $\lim_{\Delta\to 0} v_\Delta \geq v_{cont}$. The next lemma states that $\lim_{\Delta\to 0} v_\Delta \leq v_{cont}$, thereby completing the proof of Theorem 1.

Lemma 4. *For every $p \in [0,1]$ and every sender's message strategy σ, $\lim_{\Delta\to 0} v_\Delta(p) \leq v_{cont}(p)$.*

The proof of Lemma 4 appears in Section 6.9.3. The intuition for the lemma is that, since $u \leq \overline{u}_\delta$ for every δ, the function v_Δ is at most the value of the discrete-time game with indirect payoff function $\overline{u}_\delta$. By Cardaliaguet et al. (2016), the value of the discrete-time game with indirect payoff function $\overline{u}_\delta$ converges, as Δ goes to 0, to the value of the continuous-time game with indirect payoff $\overline{u}_\delta$. By Eq. (6) the latter limit coincides with v_{cont}, and the proof is complete.

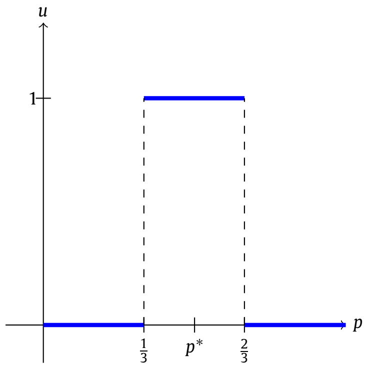

Fig. 5. A nonmonotone function u.

4. Discussion

Assumption 1 requires that the payoff function is monotone, piecewise constant, and has concave linear interpolation. How does the characterization of the sender's optimal message strategy change when this assumption is weakened?

The case that u is continuous (rather than piecewise constant) and concave falls under the model studied by Cardaliaguet et al. (2016), who showed that in this case the sender's optimal message strategy is to never reveal information. This can be viewed as a limit case of our model, when the set of discontinuity points of u becomes dense in $[0, 1]$, in which case the splits of beliefs become narrow (that is, the sender splits the receiver's belief into nearby beliefs), and at the limit no splitting is done.

When u is monotone and piecewise constant but does not have a concave linear interpolation, it is no longer true that for every $p \in [p_{-j-1}, p_{-j})$, the optimal split is to p_{-j-1} and p_{-j}. Indeed, if in the interval $[p_{-j-1}, p_{-j})$ the indirect payoff function u lies strictly below the concave linear interpolation of u, the interval will be "skipped" and the receiver's belief will never lie in this interval (after the initial split).

We discuss next the case when u is piecewise constant with a concave linear interpolation that is not monotone. Consider for example the case where $p^* = \frac{1}{2}$, and u is given by (see Fig. 5):

$$u(p) = \begin{cases} 0, & p \in [0, \frac{1}{3}), \\ 1, & p \in [\frac{1}{3}, \frac{2}{3}), \\ 0, & p \in [\frac{2}{3}, 1]. \end{cases}$$

Over the belief interval $[0, \frac{1}{2}]$, the function fits the model analyzed in the paper, and therefore the optimal strategy for $p \in [0, \frac{1}{3}]$ is to split the belief between 0 and $\frac{1}{3}$. For symmetric reason, the optimal strategy for $p \in [\frac{2}{3}, 1]$ is to split the belief between $\frac{2}{3}$ and 1. This is unlike the case when u is increasing, where to the right of p^* the optimal strategy involves sliding as well as splitting.

In this example, the maximum of u is attained at p^*. When the maximum of u is attained, say, at p_j for $j \geq 1$, the sender's optimal message strategy on $[0, p_j]$ coincides with σ^*, yet the sender's optimal message strategy on $[p_j, 1]$ may be more intricate than σ^*.

For a general piecewise constant u, some continuity intervals will be skipped altogether, in some the receiver's belief will be split between the two endpoints of the interval (as happens in our model for continuity intervals below p^*), and some will be divided into two (as happens in our model for continuity intervals above p^*): in one part no information will be revealed, and in the other the belief will be split between the interval's cutoff point and some discontinuity point of u, which may or may not be the endpoint of the continuity interval.

The sender's optimal message strategy is such that the interval $[0, 1]$ is divided into a finite number of subintervals; in some of these subintervals, the sender reveals no information, and in the others the sender splits the belief into the endpoints of the interval. This structure of optimal strategy appeared already in Ashkenazi-Golan et al. (2020). We conjecture that this is also the structure of the optimal strategy whenever the indirect payoff function u is semi-algebraic. This is because semi-algebraic functions have finitely many discontinuity points, and hence can be approximated by continuous functions similar to the approximation done in this paper. We conjecture that in the limit, as the continuous indirect utility functions approach the semi-algebraic one, the optimal strategy would be constructed of the same building blocks as those employed by optimal sender's strategy described above.

A natural question is whether the sender has a *uniformly ε-optimal* message strategy; that is, a strategy that is ε-optimal for every discount rate r sufficiently close to 0. In this case, only the far future matters. By Theorem 1, if $p^* \in (p_0, p_1)$, under σ^* the receiver's belief after the sender sends her message converges to $\{p_0, p_1\}$ with probability 1, and any strategy under which the belief after the sender sends her message converges to $\{p_0, p_1\}$ with probability 1 is uniformly ε-optimal.

If $p^* = p_0$, the same holds for any message strategy under which the receiver's belief after the sender sends her message converges to $\{p_0\}$.

Another interesting question concerns a variation of the model, where the receiver obtains information about the state at random times, independently of the sender's choices. We conjecture that the sender's optimal message strategy will have the same structure as σ^*, yet the cutoffs $(q_j)_{j=1}^{m'-1}$ will be higher than the ones we identified, to compensate for the lower significance of the instantaneous payoff.

5. Conclusion

Our result is part of the growing literature devoted to the study of optimal strategies in persuasion games. There are two aspects that single out our work. First, as in Ely (2017) and Renault et al. (2017), the payoff function in our model is not continuous, but piecewise constant. The first part of Ely (2017) and Renault et al. (2017) focused their attention on situations where the belief space is divided into two convex regions and the payoff in each region is constant, and asked whether a specific strategy, namely, the myopic strategy, is optimal. In contrast, we allow for more than two continuity regions. It turns out that in the interval $[0, p_1]$ (or $[0, p_0]$, if $p^* = p_0$) the sender's optimal message strategy is myopic, while on the interval $[p_1, 1]$ it is not. Our work highlights the interplay between the Markov transition and the monotonicity of payoffs: When the Markov transition leads to beliefs with higher (resp. lower) payoff, the myopic strategy is optimal (resp. not optimal). Second, our study combines tools provided by the literature on continuous-time games, like the approach taken by Gensbittel and Rainer (2021) with geometric intuitions.

We studied the model with two states of nature. A natural question is whether similar analysis can be carried out in the presence of more than two states of nature. Unfortunately, the answer is negative. Renault et al. (2017) presented a discrete-time example with three states of nature where the function u is piecewise constant and attains two values, and the optimal strategy is quite involved. A similar phenomenon occurs in continuous-time games, as studied by Gensbittel and Rainer (2021).

6. Appendix

The most challenging part of the proof of Theorem 1 is Lemma 1. The proof of the lemma is organized as follows: In Section 6.1 we present two useful Markovian strategies: one reveals no information in a certain range $[p', p'']$ of beliefs, and the other splits the receiver's beliefs between p' and p'' whenever the current belief is in $[p', p'']$. We study the payoff under these strategies in the continuous-time game. In Section 6.2 we present results about the monotonicity of the value function. Section 6.3 presents useful results from the literature and provides the sender's optimal message strategy for beliefs in $[p_0, p_1]$. Section 6.4 introduces a certain function, denoted $\overline{g}_\delta$, and presents its relation with the derivative of the value function. Section 6.5 characterizes the derivative of the value function using $\overline{g}_\delta$. Section 6.6 connects the strategies presented in Section 6.1 to the derivative found in Section 6.5. Section 6.7 concludes the proof. Section 6.8 contains the proof of auxiliary results used in earlier sections, and Section 6.9 presents the proofs of Lemmas 2, 3, and 4.

6.1. Two useful strategies

In this section we present two strategies of the sender that induce different ways for the belief of the receiver to get from one value p' to another value p'', where either $p' < p'' \le p^*$ or $p' > p'' \ge p^*$. We then compute the expected discounted time it takes for the belief to get from p' to p'' under each of the two strategies. The purpose of this computation is threefold. First, this will allow us to study properties of the sender's optimal message strategy. Second, the analysis supports an intuitive explanation for the sender's optimal message strategy provided in the introduction. Third, this analysis is general and does not depend on the function $u(p)$, so it might be of independent interest.

For every two distinct beliefs $p', p'' \in [0, 1]$, let $\sigma^{split}_{p',p''}$ be a Markovian sender's message strategy that splits the receiver's belief between p' and p'' whenever the receiver's belief is in $[p', p'']$. Let $\sigma^{slide}_{p',p''}$ be a Markovian sender's message strategy that reveals no information whenever the belief is in the interval $[p', p'')$.

We now compare the time it takes for each of the strategies $\sigma^{split}_{p',p''}$ and $\sigma^{slide}_{p',p''}$ to make the receiver's belief move from p' to p''. Denote by $\tau_{p''} := \min\{t \ge 0\colon p^t = p''\}$ the first time when the receiver's belief is p'', by $Y^{split}_{p',p''} := \mathbb{E}_{p^0,\sigma^{split}_{p',p''}}\left[1 - e^{-r\tau_{p''}}\right]$ the expected discounted time to reach belief p'' from belief p' under $\sigma^{split}_{p',p''}$, and by $Y^{slide}_{p',p''}$ the corresponding quantity under $\sigma^{slide}_{p',p''}$.

Suppose that $p' > p'', p^*$ or $p' < p'', p^*$. Under the strategy $\sigma^{split}_{p',p''}$, when $p^0 = p'$, the stopping time $\tau_{p''}$ has exponential distribution with parameter $\Lambda := \frac{\lambda_0 - p'(\lambda_0+\lambda_1)}{p''-p'}$. Using the definition of p^* and defining $\mu := \frac{r}{\lambda_0+\lambda_1}$, simple algebraic manipulations yield that

$$Y^{split}_{p',p''} = \frac{\mu \cdot (p'' - p')}{p^* - p' + \mu \cdot (p'' - p')}. \tag{7}$$

We turn to calculate the analog of Eq. (7) for $\sigma^{slide}_{p',p''}$. When $p' < p^* < p''$ or $p'' < p^* < p'$, under the strategy $\sigma^{slide}_{p',p''}$, when $p^0 = p'$, the belief will converge to p^* and never reach p''. For the following computation we therefore assume that $p' < p'' < p^*$ or $p^* < p'' < p'$. Under $\sigma^{slide}_{p',p''}$, when $p^0 = p'$, we have

$$\tau_{p''} = -\frac{1}{\lambda_0 + \lambda_1} \ln\left(\frac{p^* - p''}{p^* - p'}\right), \tag{8}$$

hence

$$Y^{slide}_{p',p''} = 1 - \left(\frac{p^* - p''}{p^* - p'}\right)^{\frac{r}{\lambda_0+\lambda_1}} = 1 - \left(\frac{p^* - p''}{p^* - p'}\right)^{\mu}. \tag{9}$$

Observe that the validity of Eqs. (7) and (9) does not depend on the behavior of σ outside the interval $[p', p'']$ (for $p' < p''$) or $[p'', p']$ (for $p'' < p'$).

Lemma 5. *$Y^{slide}_{p',p''} \geq Y^{split}_{p',p''}$ for every $p' < p'' \leq p^*$ or $p' > p'' \geq p^*$.*

Proof. By Eqs. (7) and (9), we need to show that

$$1 - \left(\frac{p^* - p''}{p^* - p'}\right)^{\mu} \geq \frac{\mu\left(1 - \frac{p^*-p''}{p^*-p'}\right)}{1 + \mu\left(1 - \frac{p^*-p''}{p^*-p'}\right)}.$$

Denoting $k := \frac{p^*-p''}{p^*-p'} \in (0, 1)$, we need to show that

$$1 - k^{\mu} \geq \frac{\mu(1-k)}{1 + \mu(1-k)}.$$

Simple algebraic manipulations show that the above inequality is equivalent to

$$1 + \mu k^{\mu+1} - (\mu + 1)k^{\mu} \geq 0.$$

For $k = 0$ the inequality is strict and for $k = 1$ it is weak. Finally, the derivative of the left-hand side with respect to k is negative. This completes the proof. □

Conclusion 1. *Let σ be a sender's message strategy, where the belief is split between p' and p'' for every $p \in [p', p'']$, where $p' \leq p^* \leq p''$. Then the resulting putative value function is*

$$w(p,\sigma) = u(p') \cdot \frac{p'' \cdot (\mu+1) - p^*}{(p''-p')(\mu+1)} + u(p'') \cdot \frac{p^* - p' \cdot (\mu+1)}{(p''-p')(\mu+1)} + p\mu \cdot \frac{u(p'') - u(p')}{(p''-p')(\mu+1)}. \tag{10}$$

To see why Conclusion 1 holds, note that from Eq. (7) and using the fact that σ is Markovian for beliefs in $[p', p'']$, we obtain that

$$w(p', \sigma) = Y^{split}_{p',p''} \cdot u(p') + (1 - Y^{split}_{p',p''}) \cdot w(p'', \sigma). \tag{11}$$

Eq. (11) is employed twice, once as it appears, and once with the roles of p' and p'' exchanged. By definition, $w(\cdot, \sigma)$ is linear on $[p', p'']$, and Eq. (10) follows.

6.2. Monotonicity of the value function

It is well known that the value functions in $G_{cont}(u)$ and $G_{\Delta}(u)$ are concave, whether or not u is continuous. In this section we explore monotonicity properties of the value function under the assumption that u is continuous. We argue that when u is nondecreasing, v_{cont} is nondecreasing as well, and if $u(1) > u(p^*)$, then the value function is strictly increasing.

Lemma 6. *Suppose that the indirect payoff function u is continuous and nondecreasing. Then v_{cont} is nondecreasing. If $u(1) > u(p^*)$, then v_{cont} is increasing.*[7]

[7] Lemma 6 is valid even without the assumption that u is continuous, yet we will use it only for the approximating functions $(\overline{u}_{\delta})_{\delta>0}$, which are continuous.

Proof. If $p^* = 1$, then $v_{cont}(1) = u(1)$. Since u is nondecreasing, we moreover have $v_{cont}(p) \le u(1)$ for every $p \in [0, 1)$. This implies that the maximum of v_{cont} is attained at 1, and the concavity of v_{cont} implies that v_{cont} is nondecreasing.

Assume next that $p^* < 1$. Since v_{cont} is concave, to prove that it is nondecreasing it is sufficient to verify that it is nondecreasing on $[p^*, 1]$. Let $p^* \le p' < p''$. We will prove that $v_{cont}(p') \le v_{cont}(p'')$. We distinguish between two cases: $v_{cont}(p') \le u(p'')$ and $v_{cont}(p') > u(p'')$.

Case 1: $v_{cont}(p') \le u(p'')$**.** Suppose that the initial belief is p'', and consider a message strategy σ that splits the belief of the receiver between p' and p'' for all beliefs in $(p', p'']$, and plays optimally once the belief is p'. As long as $p^t = p''$, the instantaneous payoff is $u(p'') \ge v_{cont}(p')$, and once the belief is $p^t = p'$, the continuation payoff is $v_{cont}(p')$. Therefore,

$$v_{cont}(p'') \ge w(p'', \sigma) \ge v_{cont}(p').$$

Case 2: $v_{cont}(p') > u(p'')$**.**

Let σ be an optimal message strategy in $G_{cont}(u)$, which exists since u is continuous, and let τ be the first time t when $p^t \ge p''$. For every stopping time τ denote by $W(\tau, \sigma)$ the expected payoff under σ from time τ and on. This quantity is a random variable, measurable with respect to the information at time τ. Since σ is optimal, $W(\tau, \sigma) = v_{cont}(p_\tau)$ with probability 1 under the initial belief and σ. Hence,

$$u(p'') < v_{cont}(p') = w(p', \sigma) \tag{12}$$

$$= \mathbb{E}_{p',\sigma}\left[\int_0^\tau re^{-rt}u(p_t)\mathrm{d}t + (1 - e^{-r\tau})W(\tau, \sigma)\right] \tag{13}$$

$$= \mathbb{E}_{p',\sigma}\left[\int_0^\tau re^{-rt}u(p_t)\mathrm{d}t + (1 - e^{-r\tau})v_{cont}(p_\tau)\right]. \tag{14}$$

Since u is monotone, the integral within the expectation is at most $u(p_2) < v_{cont}(p_1)$. Therefore, there is a belief $p''' \ge p''$ such that $v_{cont}(p''') > v_{cont}(p')$. The concavity of v_{cont} implies that $v_{cont}(p'') > v_{cont}(p')$.

We turn to prove the second claim. Assume that $u(p^*) < u(1)$, and suppose, by way of contradiction, that v_{cont} is not increasing. Since v_{cont} is concave and nondecreasing, this implies that there is $p' \in (p^*, 1)$ such that v_{cont} is increasing on $[0, p]$ and constant on $[p', 1]$. Since $u(p^*) < u(1)$, since u is continuous, and since $\lim_{t\to\infty} \mathbb{E}_{p',\sigma}[p^t] = p^*$ uniformly over the sender's strategies, we have $v_{cont}(p') < u(1)$. Let σ be a sender's strategy that splits the belief of the receiver between 1 and p' for all beliefs in $(p', 1)$, and plays optimally once the belief reaches p'. Then

$$u(1) \ge v_{cont}(1) \ge w(1, \sigma) = Y^{split}_{1,p'} \cdot u(1) + (1 - Y^{split}_{1,p'}) \cdot v_{cont}(p').$$

Since $Y^{split}_{1,p'}$ is less than 1, this implies that $v_{cont}(1) \ge v_{cont}(p')$, a contradiction. □

6.3. *Strategies in continuous time — previous results*

Cardaliaguet et al. (2016) studied our game when the indirect payoff function u is continuous, characterized the value function, proved that the sender has an optimal message strategy, and characterized such a strategy. Gensbittel (2019) further studied the game when u is continuous. In this section we present two results from these papers.

Recall that the *hypograph* of a function $f : [0, 1] \to \mathbb{R}$ is the set of all points that lie on or below the graph of the function. When f is concave, its hypograph is a convex set, and its set of extreme points coincides with the set of points on the graph of f where f is not affine, plus the corner points $(0, f(0))$ and $(1, f(1))$.

For simplicity of presentation, define

$$\mu := \frac{r}{\lambda_0 + \lambda_1}.$$

This is the ratio between the discount rate and the rate at which the state changes.

Theorem 6.1 *(Theorem 2.12 in Gensbittel (2019), and Theorem 2.3 in Ashkenazi-Golan et al. (2020)). Provided the indirect payoff function u is continuous, the value function v_{cont} in $G_{cont}(u)$ is the unique continuous, concave function $v : [0, 1] \to \mathbb{R}$ that is differentiable on* $[0, 1]$*, except, possibly, at* p^**, and satisfies the following conditions:*

G.1 $v_{cont}(p^*) \ge u(p^*)$*, with equality if* $(p^*, v_{cont}(p^*))$ *is an extreme point of the hypograph of* v_{cont}*.*
G.2 *For every* $p \in [0, 1] \setminus \{p^*\}$ *we have* $v'(p)(p - p^*) + \mu \cdot (v_{cont}(p) - u(p)) \ge 0$*.*

G.3 *For every extreme point $(p, v_{cont}(p))$ of the hypograph of v_{cont} such that $p \neq p^*$ we have*

$$v'(p)(p - p^*) + \mu \cdot (v_{cont}(p) - u(p)) = 0, \tag{15}$$

where for $p = 0$ (resp. $p = 1$), $v'(p)$ stands for the right (resp. left) derivative of v_{cont} at p.

Observe that points $(p, v_{cont}(p))$ that are not extreme points of the hypograph of v_{cont} lie on a line segment connecting two extreme points of the hypograph; that is, they are convex combinations of these two extreme points, denoted $(p', v_{cont}(p'))$ and $(p'', v_{cont}(p''))$. This implies that the value at such belief p can be obtained by a split of the belief between p' and p''.

We will use Theorem 6.1 to obtain the optimal message strategy for beliefs outside the continuity interval $[p_0, p_1]$. For beliefs in the continuity interval $[p_0, p_1]$ we use the following characterization of the optimal message strategy. This result follows from Cardaliaguet et al. (2016) and applies to both $G_{cont}(u)$ and $G_{cont}(\overline{u}_\delta)$.

Lemma 7. *If p^* is a discontinuity point of u (so that $p^* = p_0$), then the sender's optimal message strategy at receiver's belief p^* for both u and $\overline{u}_\delta$, is to reveal no information.*

If $p^* \in (p_0, p_1)$, then for both u and $\overline{u}_\delta$, for every $p \in [p_0, p_1]$, the optimal message strategy at receiver's belief p is to split the belief between p_0 and p_1.

Proof. The result follows from Lemma 3 in Cardaliaguet et al. (2016), which states that if $(p^*, (\text{cav } u)(p^*))$ lies on the line segment that connects $(p', u(p'))$ and $(p'', u(p''))$, for some $p', p'' \in [0, 1]$ that satisfy $p' \le p^* \le p''$, then the value function is linear on $[p', p'']$, and the optimal message strategy at each belief $p \in [p', p'']$ is to split the belief between p' and p'' (and to reveal no information if $p' = p'' = p^*$).

Since u has a concave linear interpolation, if $p^* = p_0$, then $u(p) = (\text{cav } u)(p)$, and then the result follows by setting $p' = p'' := p_0$. If $p^* \in (p_0, p_1)$, then the result follows by setting $p' := p_0$ and $p'' := p_1$. The same reasoning holds for $\overline{u}_\delta$.

While Cardaliaguet et al. (2016) analyze a model where u is continuous, their Lemma 3 does not depend on the continuity of u. □

The intuition behind Lemma 7 is as follows. When the initial belief p^0 is p^*, for every message strategy the unconditional expectation $\mathbb{E}[p^t]$ is equal to p^*. The expected instantaneous payoff is $\mathbb{E}[u(p^t)]$, which, by Jensen's inequality, is smaller than $(\text{cav } u)(\mathbb{E}[p^t]) = (\text{cav } u)(p^*)$.

Consider now the message strategy σ^* described in Theorem 1. If $p^* = p_0$, then $p^t = p_0$ for every $t \ge 0$ and $(\text{cav } u)(p^*) = u(p^*)$. It follows that the sender's payoff under σ^* is $(\text{cav } u)(p^*)$, which is the best possible payoff. If $p^* \in (p_0, p_1)$, then $p' = p_0$ and $p'' = p_1$, and the posterior belief p^t is either p' or p'': when, say, $p^t = p''$, the Markov transition makes the belief slide toward p^*, and then the sender splits the belief again between p' and p''. Since the unconditional expectation of p^t is p^*, the unconditional probability α that the belief at period n is p' satisfies $\alpha p_0 + (1 - \alpha)p_1 = p^*$. As a result, this message strategy guarantees to the sender the payoff $\alpha u(p') + (1 - \alpha)u(p'')$, which is equal to $(\text{cav } u)(p^*)$. Hence, in this case as well, σ^* guarantees to the sender the highest possible payoff.

The above discussion provides the optimal message strategy for the continuity interval $[p_0, p_1]$. In the next subsections we handle the other continuity intervals.

6.4. The functions $(\overline{g}_\delta)_{\delta>0}$

Inspired by Theorem 6.1, for every $\delta > 0$ sufficiently small define a function $\overline{g}_\delta : [0, 1] \setminus \{p^*\} \to \mathbb{R}$ by

$$\overline{g}_\delta(p) := \mu \cdot \left(\frac{\overline{u}_\delta(p) - \overline{v}_\delta(p)}{p - p^*} \right), \quad \forall p \in [0, 1] \setminus \{p^*\}. \tag{16}$$

In this section we will study the function $\overline{g}_\delta$. Since $\overline{u}_\delta$ is continuous and $\overline{v}_\delta$ is Lipshitz, $\overline{g}_\delta$ is continuous.

By Theorem 6.1, the function $\overline{g}_\delta$ is related to the derivative of $\overline{v}_\delta$. Indeed, by (G.2), $\overline{v}'_\delta(p) \le \overline{g}_\delta(p)$ for every $p < p^*$, and $\overline{v}'_\delta(p) \ge \overline{g}_\delta(p)$ for every $p > p^*$. By (G.3), $\overline{g}_\delta(p) = \overline{v}'_\delta(p)$ for $p \neq p^*$ such that $(p, \overline{v}_\delta(p))$ is an extreme point of the hypograph of v_{cont}. Furthermore, for $p \neq p^*$ the function $\overline{v}_\delta(p)$ is linear over segments $[p', p'']$ where all $p \in (p', p'')$ are not extreme points of the hypograph of $\overline{v}_\delta(p)$. Hence, the derivative $\overline{v}'_\delta$ is constant over such interval (p', p'') and satisfies $\overline{g}_\delta(p') = \overline{g}_\delta(p'')$. We deduce the following.

Lemma 8. *If $(p', \overline{v}_\delta(p'))$ and $(p'', \overline{v}_\delta(p''))$ are extreme points of the hypograph of $\overline{v}_\delta$, and none of the points $(p, \overline{v}_\delta(p))$, for $p \in (p', p'')$, is an extreme point of the hypograph of $\overline{v}_\delta$, then $\overline{g}_\delta(p') = \overline{g}_\delta(p'')$. Moreover, if $p'' < p' \le p^*$ or $p'' > p' \ge p^*$, then*

$$\overline{v}_\delta(p'') = \frac{\mu \cdot (p' - p'')}{p^* - p'' + \mu \cdot (p' - p'')} \cdot \overline{u}_\delta(p'') + \frac{p^* - p''}{p^* - p'' + \mu \cdot (p' - p'')} \cdot \overline{v}_\delta(p').$$

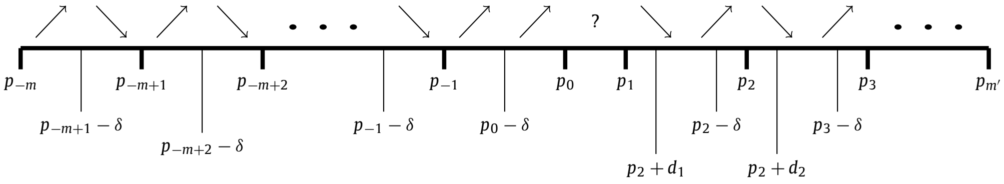

Fig. 6. The regions where $\overline{g}_\delta$ increases and decreases.

Proof. The first claim holds since $\overline{v}_\delta$ is linear on $[p', p'']$. From this and by Eq. (16) we conclude that

$$\frac{\overline{v}_\delta(p'')-\overline{v}_\delta(p')}{p''-p'}=\frac{\mu\cdot\left(\overline{u}_\delta(p'')-\overline{v}_\delta(p'')\right)}{p''-p^*}. \tag{17}$$

The second claim follows from Eq. (17) and simple algebraic manipulations. □

The next result describes the graph of $\overline{g}_\delta(p)$ on the segments $[0, p_0]$ and $[p_1, 1]$. Its proof is not inspiring and is relegated to Section 6.8.1. Note that we do not[8] describe $\overline{g}_\delta$ on $[p_0, p_1]$ if $p^* > p_0$.

Lemma 9. *For every $\delta > 0$ sufficiently small, the function $\overline{g}_\delta$ satisfies the following properties, see Fig. 6:*

(a) *$\overline{g}_\delta$ increases on $(p_{-j-1}, p_{-j}-\delta)$, for $j \in \{0, \dots, m-1\}$.*
(b) *$\overline{g}_\delta$ decreases on $(p_{-j}-\delta, p_{-j})$, for $j \in \{1, \dots, m-1\}$ (if $p^* = p_0$) or $j \in \{0, 1, \dots, m-1\}$ (if $p^* > p_0$).*
(c) *If $p^* = p_0$, then:*
  (i) *$\overline{g}_\delta$ increases on $(p_0-\delta, p_0)$, and*
  (ii) *$\overline{g}_\delta$ is smaller than or equal[9] to $\overline{g}_\delta(p_1)$ on $(p_0, p_1-\delta)$.*
(d) *For each $j \in \{1, \dots m'-1\}$ there is $d_j \in (p_j, p_{j+1}-\delta)$ such that $\overline{g}_\delta$ is positive and decreasing on $[p_j, p_j+d_j)$, and, if it is zero at p_j+d_j, then it remains nonpositive on $[p_j+d_j, p_{j+1}-\delta)$.*
(e) *$\overline{g}_\delta$ increases on $(p_j-\delta, p_j)$, for $j \in \{1, \dots m'\}$.*

To complete the description of the function $\overline{g}_\delta$ we compare the values that $\overline{g}_\delta$ attains at the discontinuity points of u.

Lemma 10. *For every $\delta > 0$ sufficiently small, the function $\overline{g}_\delta$ satisfies the following properties:*

(a) *$\overline{g}_\delta(p_{-j}) > \overline{g}_\delta(p_{-j+1})$ for $j \in \{2, \dots, m\}$, and if $p^* > p_0$, then $\overline{g}_\delta(p_{-1}) > \overline{g}_\delta(p_0)$.*
(b) *$\overline{g}_\delta(p_j) > \overline{g}_\delta(p_{j+1})$ for $j \in \{1, \dots, m'-1\}$.*
(c) *If $p^* = p_0$, then*
  (i) *$\overline{g}_\delta(p_{-1}) < \lim_{\eta\to 0}\overline{g}_\delta(p_0-\eta)$, and*
  (ii) *$\overline{g}_\delta(p_1) \ge \lim_{\eta\to 0}\overline{g}_\delta(p_0+\eta)$.*

The proof of Lemma 10 is relegated to Section 6.8.2. Figs. 7 and 8 summarize Lemmas 9 and 10. In these figures, the graph of the function $\overline{g}_\delta$ is the dashed line. The continuity of $\overline{g}_\delta$ on $[0,1]\setminus\{p^*\}$ ensures that for every $j \in \{1, \dots, m-1\}$ there exists $\overline{q}_{-j}(\delta) \in (p_{-j}-\delta, p_{-j})$ such that $\overline{g}_\delta(\overline{q}_{-j}(\delta)) = \overline{g}_\delta(p_{-j})$, and if $p^* > p_0$, then this conclusion holds for $j = 0$ as well, see Fig. 7. Similarly, for every $j \in \{1, \dots, m'-1\}$ there exists $\overline{q}_j(\delta) \in (p_j, p_{j+1}-\delta)$ such that $\overline{g}_\delta(\overline{q}_j(\delta)) = \overline{g}_\delta(p_{j+1})$, see Fig. 8. Note that the function $\overline{g}_\delta$ is *not* piecewise constant.

6.5. The derivative of the value function $\overline{v}_\delta$

Theorem 6.1 and the results so far allow us to describe the structure of $\overline{v}_\delta$, and specifically, its derivative. The value function $\overline{v}_\delta$ is concave, and by Lemma 6 it is nondecreasing. Hence, $\overline{v}'_\delta$ is nonnegative and nonincreasing. By Theorem 6.1(G.2), $\overline{v}'_\delta \le \overline{g}_\delta$ on $[0, p^*)$, and $\overline{v}'_\delta \ge \overline{g}_\delta$ on $(p^*, 1]$. In intervals where $\overline{v}'_\delta$ is constant, $\overline{v}_\delta$ is linear, and the two endpoints of such intervals are extreme points of the hypograph of $\overline{v}_\delta$. If there exists no $\varepsilon > 0$ such that $\overline{v}'_\delta$ is constant on $(p-\varepsilon, p+\varepsilon)$, then $(p, \overline{v}_\delta(p))$ is an extreme point of the hypograph of $\overline{v}_\delta$, and hence $\overline{v}'_\delta(p) = \overline{g}_\delta(p)$ (Theorem 6.1(G.3)).

The unique function that satisfies these properties is the function that is displayed in red in Figs. 7 and 8:

[8] When $p^* \in (p_0, p_1)$, simple computations yield that $\overline{g}_\delta(p) = \frac{h_1-h_0}{(p_1-p_0)(\mu+1)}\left(\frac{(p_0-p^*)(\mu+1)}{p-p^*}-\mu\right)$ for $p \in (p_0, p_1-\delta)$. Note that in this case $\overline{g}_\delta$ is increasing on both segments (p_0, p^*) and (p^*, p_1).

[9] We cannot determine whether $\overline{g}_\delta$ increases or decreases on this interval using simple observations like is done in this lemma. Later on we will be able to conclude that it is actually constant on this interval.

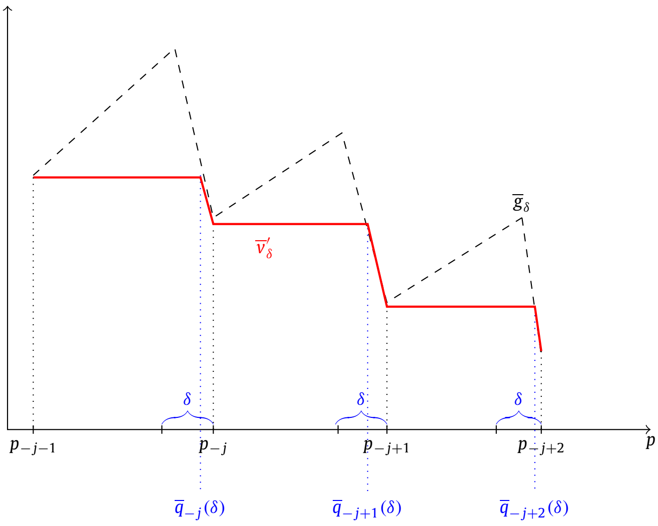

Fig. 7. The functions $\overline{g}_\delta$ (dashed) and $\overline{v}'_\delta$ (red) for $p < p_0$. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

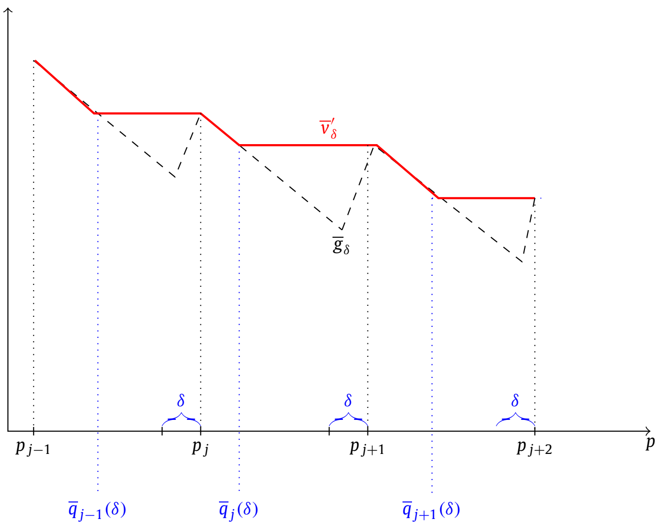

Fig. 8. The functions $\overline{g}_\delta$ (dashed) and $\overline{v}'_\delta$ (solid, red) for $p > p_1$.

- Since $(0, \overline{v}_\delta(0))$ is an extreme point of the hypograph of $\overline{v}_\delta$, we have $\overline{v}'_\delta(0) = \overline{g}_\delta(0)$.
- On the interval $[0, \overline{q}_{-m+1}(\delta)]$ the function $\overline{g}_\delta$ is at least $\overline{v}'_\delta(0)$, hence $\overline{v}'_\delta$ must be constant on this interval.
- On the interval $[\overline{q}_{-m+1}(\delta), p_{-m+1}]$, the only function that is (a) at most $\overline{g}_\delta$ and (b) coincides with it when it is not constant, is $\overline{g}_\delta$. Hence, $\overline{v}'_\delta = \overline{g}_\delta$ on this interval, and so on.

Thus, for every $j \in \{2, \dots, m-1\}$, and for $j = 1$ in case $p^* > p_0$,

$$\overline{v}'_\delta(p) = \begin{cases} \overline{g}_\delta(p_{-j}), & p \in [p_{-j}, \overline{q}_{-j+1}(\delta)), \\ \overline{g}_\delta(p), & p \in [\overline{q}_{-j+1}(\delta), p_{-j+1}]. \end{cases} \tag{18}$$

Similarly,

- Since $(p_1, \overline{v}_\delta(p_1))$ is an extreme point of the hypograph of $\overline{v}_\delta$, we have $\overline{v}'_\delta(p_1) = \overline{g}_\delta(p_1)$.
- On the interval $[p_1, \overline{q}_1(\delta)]$, the only function that is at least $\overline{g}_\delta$ and coincides with it when it is not constant is $\overline{g}_\delta$. Therefore, $\overline{v}'_\delta(p) = \overline{g}_\delta(p)$ on this interval.
- On the interval $[\overline{q}_1(\delta), p_2]$, the function $\overline{g}'_\delta$ is at most $\overline{v}'_\delta(\overline{q}_1(\delta))$ (which is equal to $\overline{g}_\delta(p_0)$), hence $\overline{v}_\delta(p)$ is constant on this interval, and so on.

Thus, for every $j \in \{1, \ldots, m' - 1\}$,

$$\overline{v}'_\delta(p) = \begin{cases} \overline{g}_\delta(p), & p \in [p_j, \overline{q}_j(\delta)), \\ \overline{g}_\delta(p_{j+1}), & p \in [\overline{q}_j(\delta), p_{j+1}]. \end{cases} \tag{19}$$

For $j = 1$ and $p^* = p_0$, we have $\overline{g}_\delta(p_{-1}) = \overline{v}'_\delta(p)$ for all $p \in [p_{-1}, p_0]$, as discussed in Remark 1.

Remark 1. For the case where $p_0 = p^*$, we need a further observation to describe the value function. The function $\overline{g}_\delta(p)$ is not defined at $p_0 = p^*$. Theorem 6.1 nonetheless holds. By Lemma 9(a) and Lemma 9(c.i), for $p \in (p_{-1}, p_0)$ we have $\overline{g}_\delta(p) > \overline{g}_\delta(p_{-1})$. Since $\overline{v}_\delta(p) \le \overline{g}_\delta(p)$ for $p \in [p_{-1}, p_0)$, and since the derivative of v_{cont} is nonincreasing, there is no $p \in (p_{-1}, p_0)$ such that $\overline{v}'_\delta(p) = \overline{g}_\delta(p)$. We conclude that $(p_{-1}, \overline{v}'_\delta(p_{-1}))$ and $(p_0, \overline{v}'_\delta(p_0))$ are extreme points of the hypograph of $\overline{v}_\delta$. Similar arguments using Lemma 9(c.ii) and Lemma 9(e) lead to the conclusion that $(p_1, v_{cont}(p_1))$ is an extreme point of that hypograph as well.

6.6. Strategies and the derivative of the putative value they generate

Once the derivative of the value function is characterized by Eqs. (18) and (19), we study the derivative of the putative value generated by the strategies $\sigma^{split}_{p',p''}$ and $\sigma^{slide}_{p',p''}$.

Let $\overline{w}_\delta(p, \sigma)$ denote the putative value obtained under sender's message strategy σ at belief p in the game $G_{cont}(\overline{u}_\delta)$.

The next result characterizes the derivative of the putative value of $\sigma^{slide}_{p',p''}$ and $\sigma^{split}_{p',p''}$.

Lemma 11. *Let $p', p'' \in [0, 1]$ be such that either $p' < p'' < p^*$ or $p^* < p'' < p'$. For every p that lies strictly between p' and p'',*

$$\overline{w}_\delta{}'(p, \sigma^{slide}_{p',p''}) = \mu \cdot \frac{\overline{u}_\delta(p) - \overline{w}_\delta(p, \sigma^{slide}_{p',p''})}{p - p^*} \tag{20}$$

and

$$\overline{w}_\delta{}'(p, \sigma^{split}_{p',p''}) = \mu \cdot \frac{\overline{u}_\delta(p') - \overline{w}_\delta(p', \sigma^{split}_{p',p''})}{p' - p^*}. \tag{21}$$

If $p = p'$ or $p = p''$, then the directional derivative at p (the left-derivative if $p = \max\{p', p''\}$ or the right-derivative if $p = \min\{p', p''\}$) is equal to the quantity given above.

The proof of Eq. (20) involves differentiation of the putative value function, and the proof of Eq. (21) uses Eqs. (7) and (11). Both calculations are uninspiring and omitted.

Lemma 11 implies that when σ is an optimal strategy, if at sender's belief p the sender reveals no information, then $\overline{w}_\delta{}'_-(p, \sigma) = \overline{g}_\delta(p)$.

6.7. Proof of Lemma 1

In this section we prove Lemma 1 by collecting the results we described so far. Let $(\overline{q}_{-j}(\delta))_{j=1}^{m-1}$ and $(\overline{q}_j(\delta))_{j=1}^{m'-1}$ be the constants that are defined at the end of Section 6.4. Let $\overline{\sigma}^*_\delta$ be the sender's message strategy defined in the statement of Lemma 1 with these constants.

By Eqs. (18) and (19), the function $\overline{v}_\delta$ is a solution of the following piecewise linear differential equation:

$$f'(p) = \begin{cases} \mu \cdot \left(\frac{\overline{u}_\delta(p_{-j}) - f(p_{-j})}{p_{-j} - p^*} \right), & p \in [p_{-j}, \overline{q}_{-j+1}(\delta)), 2 \le j \le m-1, p^* \in (p_0, p_1), \\ \mu \cdot \left(\frac{\overline{u}_\delta(p_{-j}) - f(p_{-j})}{p_{-j} - p^*} \right), & p \in [p_{-j}, \overline{q}_{-j+1}(\delta)), 1 \le j \le m-1, p^* = p_0, \\ \mu \cdot \left(\frac{\overline{u}_\delta(p) - f(p)}{p - p^*} \right), & p \in [\overline{q}_{-j+1}(\delta), p_{-j+1}], 2 \le j \le m-1, \\ \mu \cdot \left(\frac{\overline{u}_\delta(p) - f(p)}{p - p^*} \right), & p \in [p_j, \overline{q}_j(\delta)], 1 \le j \le m'-1, \\ \mu \cdot \left(\frac{\overline{u}_\delta(p_{j+1}) - f(p_{j+1})}{p_{j+1} - p^*} \right), & p \in [\overline{q}_j(\delta), p_{j+1}], 1 \le j \le m'-1. \end{cases}$$

By Lemma 11, $\overline{w}_\delta(\cdot, \overline{\sigma}^*_\delta)$ is also a solution of this differential equation. By Lemma 7, $\overline{\sigma}^*_\delta$ is optimal on $[p_0, p_1]$ (if $p^* \in (p_0, p_1)$) or at p_0 (if $p^* = p_0$), and therefore $\overline{w}_\delta(\cdot, \overline{\sigma}^*_\delta) = \overline{v}_\delta$ on $[p_0, p_1]$ (if $p^* \in (p_0, p_1)$) or on $[p_{-1}, p_1]$ (if $p^* = p_0$). By the existence and uniqueness theorem for ordinary differential equations, $\overline{w}_\delta(\cdot, \overline{\sigma}^*_\delta) = \overline{v}_\delta$ on $[0, 1]$.

6.8. Proofs of auxiliary results for Lemma 1

6.8.1. Proof of Lemma 9

To prove Lemma 9 we need an auxiliary result, which states that at the discontinuity points to the right of p^*, the value is strictly below that indirect payoff. This result follows by (a) the monotonicity of $\overline{u}_\delta$, (b) the assumption that the linear interpolation of $\overline{u}_\delta$ is concave, and (c) since we consider discontinuity points to the right of the invariant distribution.

Lemma 12. *For every $j \in \{1, ..., m'\}$ we have $\overline{v}_\delta(p_j) < \overline{u}_\delta(p_j)$.*

Proof. Fix $j \in \{1, ..., m'\}$. Denote $\widetilde{u} := \text{cav}\ (\overline{u}_\delta)$. Since in particular $\widetilde{u} \geq \overline{u}_\delta$, it follows that the value function $\widetilde{v}$ of $G_{cont}(\widetilde{u})$ satisfies $\widetilde{v} \geq \overline{u}_\delta$. Assumption 1 implies that $\widetilde{u}(p_j) = \overline{u}_\delta(p_j)$. The function $\widetilde{u}$ is continuous and concave, hence by Corollary 4 in Cardaliaguet et al. (2016), the optimal sender's message strategy in $G_{cont}(\widetilde{u})$ is to never reveal information. Hence,

$$\overline{v}_\delta(p_j) \leq \widetilde{v}(p_j) = \int_0^\infty re^{-rt}\widetilde{u}(p^t)\mathrm{d}t < \widetilde{u}(p_j) = \overline{u}_\delta(p_j),$$

where the process $(p^t)_{t\geq 0}$ under the integral term is given that the initial belief is p_j and that the sender reveals no information. The strict inequality holds because $\widetilde{u}$ is strictly increasing, and p^t is decreasing in t. The claim follows. □

Proof of Lemma 9. Recall that

$$\overline{g}_\delta(p) = \mu \cdot \frac{\overline{v}_\delta(p) - \overline{u}_\delta(p)}{p^* - p}, \quad \forall p \in [0,1] \setminus \{p^*\}. \tag{22}$$

Proof of (a): Fix $p \in (p_{-j-1}, p_{-j} - \delta)$ for $j \in \{0, \ldots, m-1\}$. Then

$$\overline{u}_\delta(p) = h_{-j-1} < \overline{w}_\delta(p, \sigma^{slide}_{p,p^*}) \leq \overline{v}_\delta(p).$$

Thus, on the interval $(p_{-j-1}, p_{-j} - \delta)$, the numerator in Eq. (22) is positive and by Lemma 6 it is nondecreasing. The denominator in Eq. (22) is positive on this interval and decreasing, and therefore $\overline{g}_\delta$ is increasing.

Proof of (b): The derivative of $\overline{g}_\delta$ is

$$\overline{g}'_\delta(p) = \mu \cdot \frac{(\overline{v}'_\delta(p) - \overline{u}'_\delta(p))(p^* - p) + (\overline{v}_\delta(p) - \overline{u}_\delta(p))}{(p^* - p)^2}, \quad \forall p \in (0,1) \setminus \{p^*\}. \tag{23}$$

Let $p \in (p_{-j} - \delta, p_{-j})$, where $j \in \{1, \ldots, m-1\}$ (if $p^* = p_0$) or $j \in \{0, 1, \ldots, m-1\}$ (if $p^* \in (p_0, p_1)$). We have $\overline{u}_\delta(p) = \left(\frac{h_{-j} - h_{-j-1}}{\delta}\right)(p - p_{-j}) + h_{-j}$. Hence, on this interval $\overline{u}'_\delta(p) = \frac{h_{-j} - h_{-j-1}}{\delta} > 0$, which is large for a small δ. Since $\overline{v}_\delta$ is concave and, by (G.3) for $p = 0$ the derivative $\overline{v}'_\delta$ is bounded, positive, and at most $\overline{g}_\delta$. The functions $\overline{u}_\delta$ and $\overline{v}_\delta$ are bounded as well. Therefore, provided δ is sufficiently small, $\overline{g}'_\delta(p) < 0$ for every $p \in (p_{-j} - \delta, p_{-j})$.

Proof of (c), item (i): Suppose that $p^* = p_0$. On the interval $(p_0 - \delta, p_0)$ we have $\overline{u}_\delta(p) = \frac{h_0 - h_{-1}}{\delta} \cdot (p - p_0) + h_0$. Hence, on this interval,

$$\overline{g}_\delta(p) = \mu \cdot \frac{\frac{h_0 - h_{-1}}{\delta} \cdot (p - p_0) + h_0 - \overline{v}_\delta(p)}{p - p_0} = \mu \cdot \left(\frac{h_0 - h_{-1}}{\delta} + \frac{h_0 - \overline{v}_\delta(p)}{p - p_0}\right).$$

By Lemma 7, $\overline{v}_\delta(p_0) = h_0$. Therefore, $\frac{h_0 - \overline{v}_\delta(p)}{p - p_0} = -\frac{\overline{v}_\delta(p_0) - \overline{v}_\delta(p)}{p_0 - p}$. The concavity of $\overline{v}_\delta$ implies that $\overline{g}_\delta$ is increasing on this interval.

Proof of (c), item (ii): Suppose again that $p^* = p_0$. We need to show that for $\eta \in [0, p_1 - p_0 - \delta]$,

$$\frac{\overline{u}_\delta(p_0 + \eta) - \overline{v}_\delta(p_0 + \eta)}{\eta} < \frac{\overline{u}_\delta(p_1) - \overline{v}_\delta(p_1)}{p_1 - p_0}.$$

By definition, $\overline{u}_\delta(p_0 + \eta) = h_0$. Let $\widehat{\sigma}$ be a strategy that splits the receiver's belief between p_0 and p_1 for all beliefs in $[p_0, p_1]$. Since $p^* \in [p_0, p_1]$, the payoff under $\widehat{\sigma}$ for initial beliefs in $[p_0, p_1]$ depends only on the initial belief, and is independent of the definition of $\widehat{\sigma}$ outside this interval. Plainly, $\overline{v}_\delta(p_0 + \eta) \geq \overline{w}_\delta(p_0 + \eta, \widehat{\sigma})$. By Conclusion 1, for $p = p_0 + \eta$, $p' = p_0$, and $p'' = p_1$ we have

$$\begin{aligned}
&\overline{w}_\delta(p_0+\eta,\widehat{\sigma})\\
&=h_0\cdot\frac{p_1(\mu+1)-p_0}{(p_1-p_0)(\mu+1)}+h_1\cdot\frac{p_0-p_0(\mu+1)}{(p_1-p_0)(\mu+1)}+(p_0+\eta)\cdot\mu\cdot\frac{h_1-h_0}{(p_1-p_0)(\mu+1)}\\
&=h_0\cdot\frac{p_1(\mu+1)-p_0-p_0\mu-\eta\mu}{(p_1-p_0)(\mu+1)}+h_1\cdot\frac{-h_0\mu+p_0\mu+\eta\mu}{(p_1-p_0)(\mu+1)}\\
&=h_0\cdot\frac{(p_1-p_0)(\mu+1)-\eta\mu}{(p_1-p_0)(\mu+1)}+h_1\cdot\frac{\eta\mu}{(p_1-p_0)(\mu+1)}.
\end{aligned}$$

Hence,

$$\frac{h_0-\overline{v}_\delta(p_0+\eta)}{\eta}\le\frac{h_0-\overline{w}_\delta(p_0+\eta,\widehat{\sigma})}{\eta}=\frac{\eta\mu(h_0-h_1)}{\eta(\mu+1)(p_1-p_0)}=\frac{\mu(h_0-h_1)}{(\mu+1)(p_1-p_0)}.$$

It is therefore sufficient to show that

$$\frac{\mu(h_0-h_1)}{(\mu+1)(p_1-p_0)}<\frac{h_1-\overline{v}_\delta(p_1)}{p_1-p_0}.$$

Canceling out the term (p_1-p_0) and rearranging the remaining terms, we see that it is sufficient to show that

$$h_1+\frac{\mu(h_1-h_0)}{(\mu+1)}>\overline{v}_\delta(p_1),$$

which holds by Lemma 12.

Proof of (d): Let $j\in\{1,\dots,m'\}$. By Lemma 12, $\overline{g}_\delta(p_j)>0$. Moreover, $\overline{u}_\delta$ and $\overline{v}_\delta$ are continuous, hence so is $\overline{g}_\delta$ on $[p_j,p_{j+1}-\delta]$. Therefore, there exists $d_j>0$ such that $\overline{g}_\delta$ is positive on $[p_j,p_j+d_j)$. In particular, $\overline{u}_\delta>\overline{v}_\delta$ on $[p_j,p_j+d_j)$.

To see that $\overline{g}_\delta$ is decreasing on $[p_j,p_j+d_j)$, consider its derivative, given in Eq. (23). On $[p_j,p_{j+1}-\delta]$ we have $\overline{u}'_\delta=0$. By Lemma 6, on this interval $\overline{v}_\delta$ is increasing, and hence $\overline{v}'_\delta>0$. Since $p^*<p$, it follows that $(\overline{v}'_\delta(p)-\overline{u}'_\delta(p))(p^*-p)<0$ on $[p_j,p_j+d_j)$ and $\overline{u}_\delta>\overline{v}_\delta$, and therefore $\overline{g}'_\delta$ is negative on $[p_j,p_j+d_j)$.

Suppose that there is $d_j\in(0,p_{j+1}-p_j-\delta)$ such that $\overline{g}_\delta(p_j+d_j)=0$. As above, $\overline{g}'_\delta(p_j+d_j)<0$, and therefore $\overline{g}_\delta$ keeps decreasing. Since $\overline{u}_\delta$ is constant on $[p_j,p_{j+1})$ and $\overline{v}_\delta$ is increasing on this interval, $\overline{g}_\delta$ remains negative on $[p_j+d_j,p_{j+1}-\delta]$.

Proof of (e): The proof is similar to the proof of item (b). □

6.8.2. Proof of Lemma 10

Proof of (a): Fix $j\in\{2,\dots,m\}$ (if $p^*=p_0$), or $j\in\{1,\dots,m\}$ (if $p^*\in(p_0,p_1)$). We need to show that

$$\frac{\overline{v}_\delta(p_{-j})-\overline{u}_\delta(p_{-j})}{p^*-p_{-j}}>\frac{\overline{v}_\delta(p_{-j+1})-\overline{u}_\delta(p_{-j+1})}{p^*-p_{-j+1}},$$

or, equivalently,

$$\overline{v}_\delta(p_{-j})>\overline{u}_\delta(p_{-j})+\frac{p^*-p_{-j}}{p^*-p_{-j+1}}\cdot\left(\overline{v}_\delta(p_{-j+1})-\overline{u}_\delta(p_{-j+1})\right). \tag{24}$$

Let σ be a strategy that splits the receiver's belief between p_{-j} and p_{-j+1} for all beliefs in (p_{-j},p_{-j+1}), and, once the belief becomes p_{-j+1}, continues optimally (that is, $\overline{w}_\delta(p_{-j+1},\sigma)=\overline{v}_\delta(p_{-j+1})$). By Eqs. (7) and (11) we have

$$\begin{aligned}
\overline{v}_\delta(p_{-j})&\ge\overline{w}_\delta(p_{-j},\sigma)\\
&\ge\frac{\mu\cdot(p_{-j+1}-p_{-j})}{p^*-p_{-j}+\mu\cdot(p_{-j+1}-p_{-j})}\cdot\overline{u}_\delta(p_{-j})\\
&\quad+\left(1-\frac{\mu\cdot(p_{-j+1}-p_{-j})}{p^*-p_{-j}+\mu\cdot(p_{-j+1}-p_{-j})}\right)\cdot\overline{v}_\delta(p_{-j+1}).
\end{aligned}\tag{25}$$

Eqs. (24) and (25) imply that it is sufficient to show that

$$\begin{aligned}
&\frac{\mu\cdot(p_{-j+1}-p_{-j})}{p^*-p_{-j}+\mu\cdot(p_{-j+1}-p_{-j})}\cdot\overline{u}_\delta(p_{-j})+\left(1-\frac{\mu\cdot(p_{-j+1}-p_{-j})}{p^*-p_{-j}+\mu\cdot(p_{-j+1}-p_{-j})}\right)\cdot\overline{v}_\delta(p_{-j+1})\\
&>u(p_{-j})+\frac{p^*-p_{-j}}{p^*-p_{-j+1}}\left(\overline{v}_\delta(p_{-j+1})-\overline{u}_\delta(p_{-j+1})\right).
\end{aligned}$$

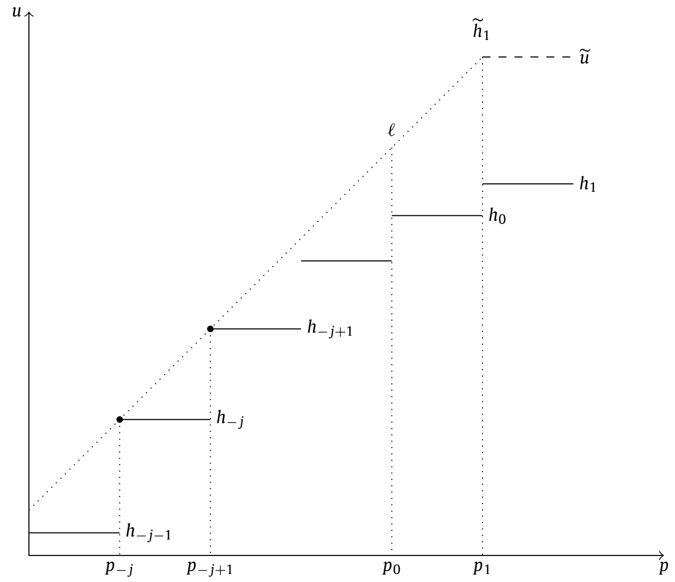

Fig. 9. The line ℓ and the function $\widetilde{u}$ in Case (a).

Simple algebraic manipulations show that this inequality is equivalent to

$$\overline{u}_\delta(p_{-j+1}) \cdot \frac{p^* - p_{-j} + \mu \cdot (p_{-j+1} - p_{-j})}{(\mu+1)(p_{-j+1} - p_{-j})} - \overline{u}_\delta(p_{-j}) \cdot \frac{p^* - p_{-j+1}}{(\mu+1)(p_{-j+1} - p_{-j})} > \overline{v}_\delta(p_{-j+1}). \tag{26}$$

Recall that $\overline{u}_\delta(p_{-j}) = h_{-j}$ and $\overline{u}_\delta(p_{-j+1}) = h_{-j+1}$. To prove that Eq. (26) holds, we will use a geometric argument rather than a long list of mathematical derivations. Consider the line ℓ that passes through the points (p_{-j}, h_{-j}) and (p_{-j+1}, h_{-j+1}), see Fig. 9. Since u has a concave linear interpolation, the graph of u lies below ℓ, except at p_{-j} and p_{-j+1}. Denote by $\widetilde{h}_1$ the unique real number such that $(p_1, \widetilde{h}_1)$ lie on ℓ. Then $\widetilde{h}_1 > h_1$.

Let $\widetilde{u}_\delta : [0,1] \to \mathbb{R}$ be the function that coincides with $\overline{u}_\delta$ except on $[p_1, p_2 - \delta)$ where it is equal to $\widetilde{h}_1$, and for $p \in [p_2 - \delta, p_2]$, where it is equal to $\frac{h_2 - \widetilde{h}_1}{\delta} \cdot (p - p_2) + h_2$. Restrict attention to beliefs in $[0, p_1]$. Since the line ℓ lies above the graph of u, the concavification of $\widetilde{u}_\delta$ at p^* is on ℓ. Lemma 7 implies that on the interval $[p_{-j}, p_1]$ the optimal message strategy $\widetilde{\sigma}^*$ in $G_{cont}(\widetilde{u}_\delta)$ is to split the receiver's belief between p_{-j} and p_1, and the value function $\widetilde{v}_\delta$ of $G_{cont}(\widetilde{u}_\delta)$ exists and is linear on this interval. We argue that on the interval $[p_{-j}, p_1]$ we have

$$\begin{aligned}\widetilde{v}_\delta(p) = h_{-j} \cdot \frac{p_1 \cdot (\mu+1) - p^*}{(p_1 - p_{-j})(\mu+1)} + \widetilde{h}_1 \cdot \frac{p^* - p_{-j} \cdot (\mu+1)}{(p_1 - p_{-j})(\mu+1)} \\ + p\mu \cdot \frac{\widetilde{h}_1 - h_{-j}}{(p_1 - p_{-j})(\mu+1)}.\end{aligned}$$

By Conclusion 1,

$$\begin{aligned}\widetilde{v}_\delta(p_{-j+1}) = \overline{u}_\delta(p_{-j+1}) \cdot \frac{p^* - p_{-j} + \mu \cdot (p_{-j+1} - p_{-j})}{(\mu+1)(p_{-j+1} - p_{-j})} \\ - \overline{u}_\delta(p_{-j}) \cdot \frac{p^* - p_{-j+1}}{(\mu+1)(p_{-j+1} - p_{-j})}.\end{aligned} \tag{27}$$

Since $\widetilde{v}_\delta \geq \overline{v}_\delta$, Eq. (26) holds with weak inequality. When the initial belief p is in the interval $[p_{-j+1}, p_1]$, the only optimal strategy in $G_{cont}(\widetilde{u}_\delta)$ is the strategy that splits the receiver's belief between p_{-j+1} and p_1. Since $\widetilde{h}_1 > h_1$, this strategy yields in $G_{cont}(\overline{u}_\delta)$ a payoff lower than $\widetilde{v}_\delta(p)$. It follows that $\widetilde{v}_\delta(p) > \overline{v}_\delta(p)$ for every $p < p_1$, and the claim follows.

Proof of (b): Fix $j \in \{1, \ldots, m' - 1\}$. We want to show that

$$\frac{\overline{u}_\delta(p_j) - \overline{v}_\delta(p_j)}{p_j - p^*} > \frac{u(p_{j+1}) - \overline{v}_\delta(p_{j+1})}{p_{j+1} - p^*}.$$

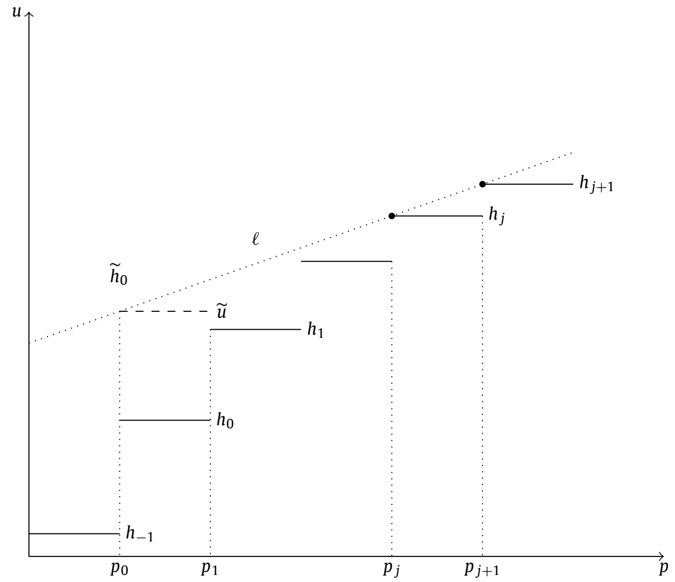

Fig. 10. The line ℓ and function $\widetilde{u}$ in Case (b).

As for item (a), it is sufficient to show that (compare this equation with Eq. (26))

$$\overline{v}_\delta(p_j) < -\frac{p_j - p^*}{(\mu+1)(p_{j+1}-p_j)} \cdot \overline{u}_\delta(p_{j+1}) + \frac{p_{j+1} - p^* + \mu \cdot (p_{j+1}-p_j)}{(\mu+1)(p_{j+1}-p_j)} \cdot \overline{u}_\delta(p_j). \tag{28}$$

As in item (a), we consider an auxiliary problem. Let ℓ be the line that passes through (p_j, h_j) and (p_{j+1}, h_{j+1}), and let $\widetilde{h}_0$ be the unique real number such that $(p_0, \widetilde{h}_0)$ lies on ℓ, so that $\widetilde{h}_0 = \left(\frac{h_{j+1}-h_j}{p_{j+1}-p_j}\right)(p_0 - p_j) + h_j$, see Fig. 10. Since $\overline{u}_\delta$ has a concave linear interpolation, $\widetilde{h}_0 \geq h_0$. Let $\widetilde{u}_\delta : [0,1] \to \mathbb{R}$ be the function that coincides with $\overline{u}_\delta$, except on $[p_0, p_1 - \delta)$, where it is equal to $\widetilde{h}_0$.

Assumption 1 implies that when the initial belief is in $[p_0, 1]$, the value function of $G_{cont}(\widetilde{u}_\delta)$, denoted by $\widetilde{v}_\delta$, is the line ℓ. By Lemma 8,

$$\widetilde{v}(p_j) = \widetilde{h}_0 \cdot \frac{p_{j+1}(\mu+1) - p^* - \mu p_j}{(p_{j+1}-p_0)(\mu+1)} + h_{j+1} \cdot \frac{p^* - p_0(\mu+1) + \mu p_j}{(p_{j+1}-p_0)(\mu+1)}. \tag{29}$$

Since $\widetilde{u}_\delta \geq \overline{u}_\delta$ on $[p_0, 1]$, we have $\widetilde{v}_\delta \geq \overline{v}_\delta$. In particular, $\widetilde{v}_\delta(p_j) \geq \overline{v}_\delta(p_j)$. Therefore, Eq. (28) will hold as soon as we show that

$$\widetilde{v}(p_j) \leq -\frac{p_j - p^*}{(\mu+1)(p_{j+1}-p_j)} \cdot h_{j+1} + \frac{p_{j+1} - p^* + \mu \cdot (p_{j+1}-p_j)}{(\mu+1)(p_{j+1}-p_j)} \cdot h_j. \tag{30}$$

Plugging the expression in Eq. (29) in Eq. (30), and using the definition of $\widetilde{h}_0$, it is sufficient to show that

$$\begin{aligned}\left(\left(\frac{h_{j+1}-h_j}{p_{j+1}-p_j}\right)(p_0 - p_j) + h_j\right) &\frac{p_{j+1}\cdot(\mu+1) - p^* - \mu p_j}{(p_{j+1}-p_0)(\mu+1)} \\ &+ h_{j+1} \cdot \frac{p^* - p_0\cdot(\mu+1) + \mu p_j}{(p_{j+1}-p_0)(\mu+1)} \\ &\leq -\frac{p_j - p^*}{(\mu+1)(p_{j+1}-p_j)} \cdot h_{j+1} + \frac{p_{j+1} - p^* + \mu \cdot (p_{j+1}-p_j)}{(\mu+1)(p_{j+1}-p_j)} \cdot h_j. \end{aligned} \tag{31}$$

Canceling out the term $\mu + 1$ and multiplying both sides of Eq. (31) by $(p_{j+1} - p_j)$, we see that we need to verify that

$$\begin{aligned}&\big((h_{j+1}-h_j)(p_0-p_j) + h_j \cdot (p_{j+1}-p_j)\big) \cdot \frac{p_{j+1}\cdot(\mu+1) - p^* - \mu p_j}{(p_{j+1}-p_0)} \\ &+ h_{j+1} \cdot (p_{j+1}-p_j) \cdot \frac{p^* - p_0 \cdot (\mu+1) + \mu p_j}{(p_{j+1}-p_0)} \\ &\leq -(p_j - p^*)h_{j+1} + \big(p_{j+1} - p^* + \mu \cdot (p_{j+1}-p_j)\big)h_j.\end{aligned} \tag{32}$$

The coefficients of h_{j+1} in Eq. (32) cancel out, as do the coefficients of h_j. Therefore Eq. (32) holds as an equality, which implies that Eq. (28) holds with weak inequality. The proof that Eq. (28) holds with strict inequality uses the same arguments as for Part (a).

Proof of (c) items (i) and (ii): These items are direct consequences of items (c)(i) and (c)(ii) of Lemma 9, respectively.

6.9. Proofs of Lemmas 2, 3, and 4

6.9.1. Proof of Lemma 2

The proof is by induction over the continuity intervals of u.

Step 1: The interval $[p_0, p_1]$ when $p^* \in (p_0, p_1)$.

Suppose that $p^* \in (p_0, p_1)$. On the interval $[p_0, p_1]$ the strategies $(\overline{\sigma}^*_\delta)_{\delta>0}$ and σ^* coincide: they both split the receiver's belief between p_0 and p_1. It follows that on this interval $\overline{v}_\delta$ is independent of δ, and hence on $[p_0, p_1]$,

$$w(\cdot, \sigma^*) = w(\cdot, \overline{\sigma}_\delta) = \overline{v}_\delta = \overline{v}_0.$$

Step 2: The interval $[p_{-1}, p_1]$ when $p^* = p_0$.

If $p^* = p_0$, then the strategies σ^* and $(\overline{\sigma}_\delta)_{\delta>0}$ coincide and instruct splitting the receiver's belief between p_{-1} and $p^* = p_0$ (on the interval $[p_{-1}, p_0]$ and between p_0 and p_1 (on the interval $[p_0, p_1]$). The argument proceeds as in Step 1.

Step 3: The intervals to the left of p^*.

Suppose by induction that $\overline{v}_0(p) = w(p, \overline{\sigma}^*)$ for $p \in [p_{-j}, p_0]$, where $j \in \{1, 2, \dots, m-1\}$ (if $p^* \in (p_0, p_1)$) or $j \in \{2, \dots, m-1\}$ (if $p^* = p_0$). Consider the interval $[p_{-j-1}, p_{-j}]$. On this interval, the strategy σ^* splits the receiver's belief between p_{-j-1} and p_{-j}; and for each $\delta > 0$ sufficiently small, the strategy $\overline{\sigma}_\delta$ splits the receiver's belief between p_{-j-1} and $\overline{q}_{-j}(\delta)$, and reveals no information between $\overline{q}_{-j}(\delta)$ and p_{-j}. Since $\lim_{\delta\to 0} \overline{q}_{-j}(\delta) = p_{-j}$,

$$\lim_{\delta\to 0} \overline{w}_\delta(\overline{q}_{-j}(\delta), \overline{\sigma}_\delta) = \lim_{\delta\to 0} \overline{w}_\delta(p_{-j}, \overline{\sigma}^*_\delta) = \overline{v}_0(p_{-j}) = w(p_{-j}, \overline{\sigma}^*).$$

As a result, $\overline{v}_\delta$ converges to $w(\cdot, \overline{\sigma}^*)$ on $[p_{-j}, p_{-j+1})$.

Step 4: The intervals to the right of p^*.

Suppose by induction that $\overline{v}_0(p) = w(p, \overline{\sigma}^*)$ for $p \in [p_1, p_j]$, where $j \in \{1, 2, \dots, m'-1\}$, and consider the interval $[p_j, p_{j+1}]$. On this interval, the strategy σ^* reveals no information between p_j and q_j, and splits the receiver's belief between q_j and p_{j+1}. For each $\delta > 0$ sufficiently small, the strategy $\overline{\sigma}_\delta$ reveals no information between p_j and $\overline{q}_j(\delta)$, and splits the receiver's belief between $\overline{q}_j(\delta)$ and p_{j+1}. Since $\lim_{\delta\to 0} \overline{q}_j(\delta) = q_j$, the functions $\overline{w}_\delta(\cdot, \overline{\sigma}_\delta)$ converge to $w(\cdot, \overline{\sigma}^*)$ on $[p_j, q_j)$, and by monotonicity, the same holds at q_j. As in Step 3, $\overline{w}_\delta(\cdot, \overline{\sigma}_\delta)$ converge to $w(\cdot, \overline{\sigma}^*)$ on $[q_j, p_{j+1}]$.

6.9.2. Proof of Lemma 3

Recall that $Y^{split}_{p',p''}$ and $Y^{slide}_{p',p''}$ are the expected discounted time to reach belief p'' when the initial belief is p' under $\sigma^{split}_{p',p''}$ and $\sigma^{slide}_{p',p''}$, respectively, in the continuous-time game. Denote by $Y^{\Delta,split}_{p',p''}$ and $Y^{\Delta,slide}_{p',p''}$ the corresponding quantities in the discrete-time game with length of period Δ. The reader can verify that for every distinct $p', p'' \in [0, 1]$,

$$\lim_{\Delta\to 0} Y^{\Delta,split}_{p',p''} = Y^{split}_{p',p''},$$

and for every $p' < p'' < p^*$ and every $p^* < p'' < p'$,

$$\lim_{\Delta\to 0} Y^{\Delta,slide}_{p',p''} = Y^{slide}_{p',p''}.$$

The proof now follows similar arguments to those used in the proof of Lemma 2.

6.9.3. Proof of Lemma 4

In this proof only we compare the value function for games with different indirect payoff functions: u and $\overline{u}_\delta$. We therefore write $v_\Delta(u)$, $v_\Delta(\overline{u}_\delta)$, $v_{cont}(u)$, and $v_{cont}(\overline{u}_\delta)$ for the various value functions.

Since $u \le \overline{u}_\delta$, for every $\Delta > 0$ and every $\delta > 0$ sufficiently small we have $v_\Delta(u) \le v_\Delta(\overline{u}_\delta)$. Taking the limit as Δ goes to 0, we have $\lim_{\Delta\to 0} v_\Delta(u) \le \lim_{\Delta\to 0} v_\Delta(\overline{u}_\delta)$. By Theorem 1 in Cardaliaguet et al. (2016), for every $\delta > 0$ we have $\lim_{\Delta\to 0} v_\Delta(\overline{u}_\delta) = v_{cont}(\overline{u}_\delta)$. We conclude that for every $\delta > 0$ sufficiently small, $\lim_{\Delta\to 0} v_\Delta(u) \le v_{cont}(\overline{u}_\delta)$. Since this inequality holds for every sufficiently small $\delta > 0$, taking the limit as δ goes to 0 yields $\lim_{\Delta\to 0} v_\Delta(u) \le \lim_{\delta\to 0} v_{cont}(\overline{u}_\delta) = v_{cont}(u)$, where the last equality follows from Eq. (6).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

References

Ashkenazi-Golan, Galit, Rainer, Catherine, Solan, Eilon, 2020. Solving two-state Markov games with incomplete information on one side. Games Econ. Behav. 122, 83–104.
Aumann, Robert J., Maschler, Michael, Stearns, Richard E., 1995. Repeated Games with Incomplete Information. MIT Press.
Ball, Ian, 2019. Dynamic information provision: rewarding the past and guiding the future. Available at SSRN 3103127.
Cardaliaguet, Pierre, Rainer, Catherine, Rosenberg, Dinah, Vieille, Nicolas, 2016. Markov games with frequent actions and incomplete information—the limit case. Math. Oper. Res. 41 (1), 49–71.
Ely, Jeffrey C., 2017. Beeps. Am. Econ. Rev. 107 (1), 31–53.
Ely, Jeffrey C., Frankel, Alexander, Kamenica, Emir, 2015. Suspense and surprise. J. Polit. Econ. 123 (1), 215–260.
Ely, Jeffrey C., Szydlowski, Martin, 2020. Moving the goalposts. J. Polit. Econ. 128 (2), 468–506.
Gensbittel, Fabien, 2019. Continuous-time Markov games with asymmetric information. Dyn. Games Appl. 9 (3), 671–699.
Gensbittel, Fabien, Rainer, Catherine, 2021. Optimal dynamic provision on continuous time. In preparation.
Kamenica, Emir, 2019. Bayesian persuasion and information design. Annu. Rev. Econ. 11 (1), 249–272.
Kamenica, Emir, Gentzkow, Matthew, 2011. Bayesian persuasion. Am. Econ. Rev. 101 (6), 2590–2615.
Mertens, Jean-François, Sorin, Sylvain, Zamir, Shmuel, 2016. Repeated Games. Cambridge University Press.
Norris, James R., 1998. Markov Chains. Cambridge University Press.
Renault, Jérôme, 2006. The value of Markov chain games with lack of information on one side. Math. Oper. Res. 31 (3), 490–512.
Renault, Jérôme, Solan, Eilon, Vieille, Nicolas, 2017. Optimal dynamic information provision. Games Econ. Behav. 104, 329–349.
Smolin, Alex, 2021. Dynamic evaluation design. Am. Econ. J. Microecon. 13 (4), 300–331.