LUISS

IRIS - Luiss Guido Carli Research Institutional Repository | Catalogo istituzionale della ricerca Luiss Guido Carli

Regularity of dynamic opinion games

Pre-print version of the following publication: | Versione pre-print della seguente pubblicazione:

Original Citation/Citazione:
Venel, Xavier Mathieu Raymond. (2021). Regularity of dynamic opinion games. GAMES AND ECONOMIC BEHAVIOR, (ISSN: 0899-8256), 126: 305-334. Doi: 10.1016/j.geb.2020.12.007.

Availability/Disponibilità:
This version is available at: 11385/207655 since: 2026-10-05T11:10:47Z - Questa versione è disponibile alla pagina: 11385/207655 dal: 2026-10-05T11:10:47Z

Publisher/Casa editrice:

Published version/Pubblicato:
DOI: https://dx.doi.org/10.1016/j.geb.2020.12.007

License/Licenza:
Tutti i diritti riservati

Availability/Termini d'uso:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. Works made available under a Creative Commons license can be used according to the terms and conditions of said license. For all terms of use and more information see the publisher's website. | I termini e le condizioni relativi al riutilizzo della presente versione della pubblicazione sono disciplinati dalla politica editoriale. Le opere messe a disposizione con licenze Creative Commons possono essere utilizzate conformemente ai termini e alle condizioni previste da tali licenze. Per l'insieme delle condizioni di utilizzo e per ulteriori informazioni si rinvia al sito web dell'editore.

This item was downloaded from IRIS Luiss (https://iris.luiss.it/). When citing, please refer to the published version. | Questo documento è stato scaricato da IRIS Luiss (https://iris.luiss.it/). Per la citazione, fare riferimento alla versione pubblicata sul sito dell'editore.

(Article begins on next page | Il contributo inizia nella pagina successiva)

18 November 2026

Regularity of dynamic opinion games

Xavier Venel *

Abstract

We analyze a class of stochastic games where two lobbies compete by influencing the opinions in a society. We assume that the opinions of the agents in the society evolve according to De Groot opinion formation (1974) and that the decisions of the lobbies change the structure of the network representing the society. We show that the regularity of discounted Nash equilibrium payoffs when players become patient is highly model-dependent. We provide two extreme cases. First, we present an example where the sequence of discounted Nash equilibrium payoffs does not converge. Hence, the solution is highly sensitive to the discount factor and a modeler needs to know the discount factor precisely in order to compute equilibria. Second, we focus on a subclass of problems where lobbies are restricted and show the existence of a uniform equilibrium in this subclass. Hence, in this restricted framework, a modeler can predict some approximate equilbria with an imprecise knowledge of the discount factor.

Keywords: Game theory, Social Network, Dynamic games, Targeting, Stochastic games

Acknowledgment: Venel acknowledges the support of the Agence National de la Recherche [ANR CIGNE ANR-15-CE38-0007-01].

1 Introduction

Social networks have now become a major factor in information diffusion. In politics, Barack Obama's campaign for the 2008 U.S. presidential election illustrated the tremendous potential influence embedded in social networks, as shown by Cogburn and Espinoza [7]. As it plays an important part in the way agents form opinions, lobbies and advertisers have an incentive to use these social networks to propagate their opinions or to advertise their products. How should such lobbies behave? Hartline et al. [12], when analyzing the marketing strategy of a firm through a social network, emphasize the complementarity between pricing, timing, and targeting: "In general, it is advantageous to get influential buyers to buy the item early in the sequence; it even makes sense to offer such buyers smaller prices to get them to buy the item."

This question is closely related to the older question of identifying central agents in a network. Once these central agents have been identified, they may be used as a good proxy to influence the rest of society. The notion of a key agent is adapted to the study of one lobby trying to impact the network. Initiated by Bonacich [5] and Freeman[9], much work has been done in this direction, such as Ballester et al. [1] or Banerjee et al. [3]. As pointed by Borgatti [6] and Valente [19], the notion of central agent depends on the objective: speed of diffusion, cohesion of the network, etc. Hence, the notion of centrality and the optimal choice depends on the aim of the lobbies. For example, Ballester et al. [2] study delinquent networks and try to identify who should be removed to diminish criminality. In these articles, one principal influences the network once. There are several natural generalizations of this question: What happens if there are several lobbies? What happens if the lobbies can intervene several times on the network? The present article aims to show that answering both questions at the same time may introduce some irregularity in

*Paris School of Economics, Centre d'économie de la Sorbonne, Université Paris 1 Panthéon-Sorbonne, 106-110 Boulevard de de l'hôpital, 75647 Paris Cedex 13, France E.mail: xavier.venel@univ-paris1.fr

the behavior of the solution with respect to the parameters. Therefore, researchers and modelers need to be careful when using such models.

The first question has been approached in several articles. In these models, the lobbies take some decisions at the beginning of the game, yielding a homogeneous dynamic of opinions. The dynamic then induces the payoffs of the lobbies. Hence, the decisions and the dynamic are separated in time. Depending on the framework, the set of decisions and the definition of the dynamic change. For example, Bimpikis et al. [4] study two firms that choose how to allocate their marketing budget at the beginning of the game. Then, the agents in the society exchange some piece of information in each period. Goyal and al. [10] choose a contagion framework, where the choice of opinions of each agent is binary and each lobby chooses a seed from which their opinion propagates in the network. The closest models to our own are the ones studied by Lever [13] and by Grabisch et al. [11] based on De Groot's influence model [8]. De Groot chooses to model information transmission between the players inside a society in a non-Bayesian way: the opinion of an agent is represented by a number between 0 and 1 and the opinion of each agent is computed at every stage as a weighted average of the opinions of its neighbors at the preceding stage.

This is a particularly easy representation of word-of-mouth communication about a new product or a candidate for an election. Once the network is given, each agent listens to his neighbors and the propagation of opinion is not strategic. Lever [13] and Grabisch et al. [11] introduced lobbies that want to influence the opinions of agents in such a way as to drive the global opinion of the society in the direction of their own interests, for example in an election. In order to do so, there is a network representing society and the lobbies exert some influence to perturb this network at the outset of the game, informally transforming the network; and then in a second step, the opinions propagate according to the new perturbed network.

The assumption that the lobbies cannot change their target during the game is controversial. This is reasonable if there is a high cost attached to the intervention of the lobby, like creating a new network between the agents. It may not be reasonable if the action is simply to convince some agents. During an electoral campaign, the candidates travel between states to attend several meetings and hence they change their intervention on the network as time goes by. In a preceding paper, Mandel and Venel [14] proposed a dynamic version of the model of Grabisch et al. [11] where the players are no longer restricted to keep targeting the same agent. At every stage, a player can decide to switch at no cost and to target a new player. They focus on the case where the underlying network between the agents is primitive, i.e., there exists a minimal length l such that between each pair of agents there exists a path of any length greater than l[1]. They prove the existence of a strong solution concept called the uniform value and characterize the optimal strategy in some special classes of networks.

In this article, we will first be interested in a general model where the lobbies can heavily influence the network, relaxing the assumption that lobbies only perturb an existing network. We assume no restriction on the influence that the lobbies can have on the network: they can visit your home, organize a meeting, create a new association where people can debate, prevent two persons from communicating. We even assume that the actions of the lobbies may interact with each other before inducing the new network. Formally, at every stage, each player chooses some action. The pair of actions then induces three different effects:

- a direct influence of Player 1 on some agents,
- a direct influence of Player 2 on some agents,
- an indirect influence through the interaction between the agents.

In particular, in contrast to the model of Grabisch et al. [11], there may be no underlying network that is perturbed. Also, at every stage, we assume that each player obtains a payoff depending on the actions played and on the vector of opinions in the society. This will lead to a general class of games that we will

[1] This requirement is stronger than the connectivity of the graph. The connectivity of the graph only requires that every pair of agents is connected through some path.

call *dynamic opinion games*. This is a special class of stochastic games.

To study the long-run behavior of such a stochastic game, several solution concepts are used in the literature. The two main approaches are the asymptotic approach and the uniform approach. In the former, we first fix the length of the game (resp. the discount factor) and consider the Nash equilibrium payoffs in the game where the payoff is the n-stage average payoff (resp. the discounted Abel average). We then study the convergence of these sets when the players become patient. In the previous definition, the players know the length or the discount factor when choosing their strategies. The regularity at the limit is an important property for a modeler interested in the long-run equilibrium properties of the model. Indeed, if this is not the case then the predictions become highly sensitive to the discount factor of a patient player. Hence, the modeler needs to know precisely the discount factor to predict the equilibria of the game. In the second approach, called the uniform approach, we now consider what a player can guarantee in the long run without knowing the length of the game in advance. We will focus on the uniform approach.

It was shown by Mertens and Neyman [15] that every two-player zero-sum stochastic game such that the set of states and the sets of actions are finite admits a unique uniform Nash equilibrium. Moreover, it is the limit of the n-stage equilibria and the discounted equilibria. The existence of the uniform value ensures that an equilibrium computed by the modeler for a fixed small discount factor (patient players) is, therefore, robust to small perturbations of the discount factor. Vieille [22, 23] extended the existence result to two-player non-zero-sum games. The link between discounted equilibria and uniform equilibria is more complicated in non-zero-sum games and will be discussed later. We will need to introduce the notion of an approximate equilibrium.

In both articles, the set of states is finite, in contrast to our set of states, which will only be compact. This is a key assumption, as shown in Ziliotto [24]. By relaxing this assumption, the author constructed an example of a two-player zero-sum game with a compact state space and finite action spaces, where the sequence of discounted Nash-equilibrium payoffs does not converge. Since the game is zero-sum, it also implies that there exist no uniform equilibrium payoffs and even approximate Nash equilibria are irregular with respect to the discount factor. This article aims to raise awareness that opinion models based on a De Groot approach may have the same irregularity. They may be highly sensitive to the discount factor if one does not restrict the possible influence of the lobbies on the networks and yet well behaved under reasonable assumptions. We will not address the question of the characterization of optimal strategies. As highlighted in Mandel and Venel [14], this question is already very hard in simpler cases and it should be the focus of further research.

The paper has five parts. In the first part, we introduce the class of dynamic opinion games. In the second part, we present five results. We establish the existence of a uniform equilibrium in two simple frameworks inspired by the literature on stochastic games and we state the three other results. The last three section are dedicated to the proofs of the more advanced results. In the third section, we exhibit an example of a dynamic opinion game where the sequence of discounted Nash equilibrium payoffs does not converge. Without any restriction on the actions of the lobbies, the solution may be oscillating for patient lobbies. In the fourth section, we focus on a restricted case where the lobbies have limited power on the rest of the network: they can directly impact the opinions of agents, but the structure of interactions between the agents in the network is always the same. It includes in particular the framework of Mandel and Venel [14] where we relax the condition that the network is primitive. There may now be several connected components and some periodicity in the network (for example, cycles). We prove the existence of a uniform equilibrium. Hence, under this constraint, there exists a Nash equilibrium in a very strong sense: the lobbies can choose an almost optimal strategy without knowing the strategy of the opponent nor the precise discount factor. The proof relies on the reduction of the problem to finite stochastic games and the utilization of Vieille's existence result [22, 23]. Finally, in the sixth section, we provide two examples to show that a uniform equilibrium only gives information on ε-Nash discounted equilibria and not on Nash discounted equilibria.

2 Model

2.1 Definition

Given a positive natural number n, we denote the set $\{1, ..., n\}$ by $[n]$. Given a finite set K, we denote by $|K|$ the cardinality of K. Given a metric set X, we denote by $\Delta(X)$ the set of probability distributions over X. For $x \in X$, we denote by δ_x the Dirac measure at x, i.e., the probability distribution that is a single atom at x. Given $y \in X$, $\delta_x(\{y\})$ will be simply denoted by $\delta_x(y)$.

A dynamic opinion game $\Gamma = (K, I, J, q, g^1, g^2, x_1)$ is defined by

- a finite set of non-strategic agents denoted by K that induces a set of states $X = [0,1]^{|K|}$ called the set of vectors of opinions,
- a finite set of actions for Player 1 denoted by I,
- a finite set of actions for Player 2 denoted by J,
- a transition $q : X \times I \times J \to X$ defined such that for every $x \in X$, $i \in I$, and $j \in J$,

$$\begin{pmatrix} 1 \\ 0 \\ q(x,i,j) \end{pmatrix} = M(i,j) \begin{pmatrix} 1 \\ 0 \\ x \end{pmatrix} := \begin{pmatrix} 1 & 0 & \mathbf{0} \\ 0 & 1 & \mathbf{0} \\ A_1(i,j) & A_0(i,j) & B(i,j) \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ x \end{pmatrix},$$

  where $(A_1(i,j), A_0(i,j), B(i,j))$ is a row-stochastic matrix,
- two continuous payoff functions denoted respectively by $g^1 : X \times I \times J \to [0,1]$ and $g^2 : X \times I \times J \to [0,1]$,
- an initial vector of opinions $x_1 \in X$.

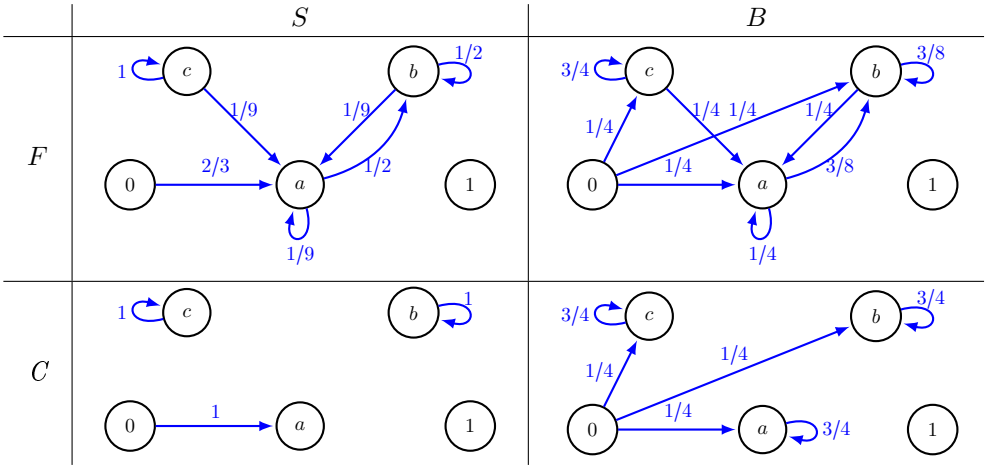
The matrix $M(i,j)$ indicates how the opinions propagate in the network if the players chose respectively i and j. The matrix $A_1(i,j)$ stands for the direct influence from Player 1 on the agents, whereas $A_0(i,j)$ stands for the direct influence from Player 2. The higher the weights $A_1(i,j)(k)$ (resp. $A_0(i,j)(k)$), the more agent k follows the opinion of Player 1 (resp. Player 2). The matrix $B(i,j)$ stands for the interaction between the agents. It is still a function of the actions i and j and, therefore, represents an indirect influence. For example, if $A_1(i,j) = 0$ and $A_0(i,j) = 0$, players can organize the society to manipulate the final opinion in the society, but they cannot directly express their opinions to the agents.

Example 2.1 *Consider the following example where there are three agents in the society $K = \{a, b, c\}$ and the action spaces of the two players are $I = \{F, C\}$ and $J = \{B, S\}$, respectively. The transition matrices are given by*

$$M(F,S) = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 2/3 & 1/9 & 1/9 & 1/9 \\ 0 & 0 & 1/2 & 1/2 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \qquad M(F,B) = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 1/4 & 1/4 & 1/4 & 1/4 \\ 0 & 1/4 & 3/8 & 3/8 & 0 \\ 0 & 1/4 & 0 & 0 & 3/4 \end{pmatrix}$$

$$M(C,S) = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \qquad M(C,B) = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 1/4 & 3/4 & 0 & 0 \\ 0 & 1/4 & 0 & 3/4 & 0 \\ 0 & 1/4 & 0 & 0 & 3/4 \end{pmatrix}$$

They can be represented by the following graphs. We only draw arrows involving agents in the society.



A possible interpretation of this game is the following. We assume that the agents in K are organized in a network M. The first lobby (with opinion 1) is a government agency that can take two actions $F(ree)$ or $C(ut)$. If it decides to choose C, then communication between the agents in the society is blocked, whereas if it chooses F, agents can communicate through M. The second lobby is a foreign state (with opinion 0) that has the choice between two actions $B(road)$ or $S(ingle)$. If it chooses B, it directly influences every agent, but with a small weight, whereas if it chooses S, it strongly influences agent a and no other agent.

Let us now comment on the matrix $M(F,B)$. By definition, we have

$$A_0(F,B) = \begin{pmatrix} 1/4 \\ 1/4 \\ 1/4 \end{pmatrix},\ A_1(F,B) = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} \text{ and } B(F,B) = \begin{pmatrix} 1/4 & 1/4 & 1/4 \\ 3/8 & 3/8 & 0 \\ 0 & 0 & 3/4 \end{pmatrix}.$$

This means that the opinion of an agent $k \in \{a,b,c\}$ at stage $t+1$ is obtained as the convex combination of the opinion of Player 2 with weights $1/4$, the opinion of Player 1 with weight 0, and the opinion of agent l with weight $B(F,B)_{kl}$ at the previous stage. For example,

$$x_{t+1}(b) = \frac{1}{4} * 0 + 0 * 1 + \frac{3}{8} x_t(b) + \frac{3}{8} x_t(a).$$

Remark 2.2 *Contrary to the usual representation of a Markov Chain, the directions of the arrows show how the opinions propagate in the networks and not how the Markov Chain evolves as a stochastic process. This choice was made to be consistent with the focus on diffusion of opinions.*

Let us now introduce the dynamic and strategies. For every $n \geq 1$, we denote by $H_n = (I \times J)^{n-1}$ the set of n-stage histories and by $H = \cup_{n=1}^{+\infty} H_n$ the set of histories. Let $H_\infty = (I \times J)^{+\infty}$ be the set of plays. Note that an initial vector of opinions and a play induce a unique sequence of vectors of opinions since the transition q is deterministic. A strategy for Player 1 is a function $\sigma : H \to \Delta(I)$ that prescribes a probability distribution over the action of Player 1 for every history. Let Σ be the set of strategies for Player 1. If for every history, Player 1 chooses a Dirac measure, the strategy is said to be *pure*. Informally, Player 1 is not randomizing. A strategy for Player 2 is a function $\tau : H \to \Delta(J)$. We denote the set of strategies for Player 2 by $\mathcal{T}$.

The dynamic of the game Γ is as follows. At stage 1, the opinions are initially fixed equal to the vector x_1. Then, Player 1 chooses a probability distribution a_1 over his set of actions and an action i_1 is chosen along a_1. Simultaneously, Player 2 chooses a probability distribution b_1 over his set of actions and an action j_1 is chosen along b_1. Then, each player observes the action of the other, a new vector of opinions is chosen such that $x_2 = q(x_1, i_1, j_1)$, and a stage payoff $g^1(x_1, i_1, j_1)$ (resp. $g^2(x_1, i_1, j_1)$) is obtained by Player 1 (resp. Player 2). More generally, at a given stage t, the state is x_t. Player 1 chooses a probability distribution a_t over his set of actions and an action i_t is randomly chosen along a_t. Simultaneously, Player 2 chooses a probability distribution b_t over his set of actions and an action j_t is chosen along b_t. Then, each player observes the other's action, a new state is chosen such that $x_{t+1} = q(x_t, i_t, j_t)$, and a stage payoff $g^1(x_t, i_t, j_t)$ (resp. $g^2(x_t, i_t, j_t)$) is obtained by Player 1 (resp. Player 2).

For the moment, we have described the dynamic of opinions. An initial state x_1 and a play induce an infinite sequence of vectors of opinions, so they generate two infinite streams of stage payoffs. There are several ways to transform these infinite streams of stage payoffs into a global payoff. For every $n \geq 1$, we define the *n-stage average* payoff for Player $m \in \{1, 2\}$ induced by the strategy profile (σ, τ) given the initial state x_1 by

$$\gamma_n^m(x_1, \sigma, \tau) = \frac{1}{n} \sum_{t=1}^{n} g^m(x_t, i_t, j_t),$$

and the payoff vector by $\gamma_n(x_1, \sigma, \tau) = (\gamma_n^1(x_1, \sigma, \tau), \gamma_n^2(x_1, \sigma, \tau)) \in \mathbb{R}^2$. For every $\lambda \in (0, 1]$, we define the *λ-discounted average* payoff for Player m induced by the strategy profile (σ, τ) given the initial state x_1 by

$$\gamma_\lambda^m(x_1, \sigma, \tau) = \lambda \sum_{t=1}^{+\infty} (1 - \lambda)^{t-1} g^m(x_t, i_t, j_t),$$

and the payoff vector by $\gamma_\lambda(x_1, \sigma, \tau) = (\gamma_n^1(x_1, \sigma, \tau), \gamma_n^2(x_1, \sigma, \tau)) \in \mathbb{R}^2$. To every evaluation, we associate a game and the set of Nash equilibria.

Definition 2.3 *Let $n \geq 1$. Let x_1 be an initial state and $\varepsilon \geq 0$. A profile of strategies (σ^*, τ^*) is an* ε-Nash equilibrium profile in the n-stage average game *if*

$$\forall \tau \in \mathcal{T},\ \gamma_n^2(x_1, \sigma^*, \tau) \leq \gamma_n^2(x_1, \sigma^*, \tau^*) + \varepsilon,$$
$$\forall \sigma \in \Sigma,\ \gamma_n^1(x_1, \sigma, \tau^*) \leq \gamma_n^1(x_1, \sigma^*, \tau^*) + \varepsilon.$$

We say that $\gamma_n(x_1, \sigma^*, \tau^*)$ is the ε-Nash equilibrium payoff *associated with (σ^*, τ^*).*

Definition 2.4 *Let $\lambda \in (0, 1]$. Let x_1 be an initial state and $\varepsilon \geq 0$. A profile of strategies (σ^*, τ^*) is an* ε-Nash equilibrium profile in the λ-discounted game *if*

$$\forall \tau \in \mathcal{T},\ \gamma_\lambda^2(x_1, \sigma^*, \tau) \leq \gamma_\lambda^2(x_1, \sigma^*, \tau^*) + \varepsilon,$$
$$\forall \sigma \in \Sigma,\ \gamma_\lambda^1(x_1, \sigma, \tau^*) \leq \gamma_\lambda^1(x_1, \sigma^*, \tau^*) + \varepsilon.$$

We say that $\gamma_\lambda(x_1, \sigma^*, \tau^*)$ is the ε-Nash equilibrium payoff *associated with (σ^*, τ^*).*

When $\varepsilon = 0$, ε-Nash equilibria are simply called Nash equilibria. We denote by $E_n(x_1)$ (resp. $E_\lambda(x_1)$) the set of Nash equilibrium payoffs in the n-stage average game (resp. λ-discounted game). We now introduce the main solution concept that we will use in this article.

Definition 2.5 *Let x_1 be an initial state. A vector $\gamma(x_1) = (\gamma^1(x_1), \gamma^2(x_1)) \in \mathbb{R}^2$ is* a uniform Nash equilibrium payoff *if, for every $\varepsilon > 0$, there exist (σ^*, τ^*) and $N \geq 1$ such that, for every $n \geq N$,*

$$\forall \tau \in \mathcal{T},\ \gamma_n^2(x_1, \sigma^*, \tau) \leq \gamma^2(x_1) + \varepsilon,$$
$$\forall \sigma \in \Sigma,\ \gamma_n^1(x_1, \sigma, \tau^*) \leq \gamma^1(x_1) + \varepsilon,$$

and

$$\|\gamma_n(x_1, \sigma^*, \tau^*) - \gamma(x_1)\|_\infty \leq \varepsilon.$$

We say that (σ^*, τ^*) is a uniform ε-Nash equilibrium profile *associated with $\gamma(x_1)$.*

We denote the set of uniform Nash equilibrium payoffs by $E_\infty(x_1)$.

In the special case where $g^1 = -g^2$, the game is said to be zero-sum. In this case, the sets $E_n(x_1)$, $E_\lambda(x_1)$, and $E_\infty(x_1)$ are either empty or reduced to a singleton. If not empty, their element can be represented by the payoff of Player 1, called the value and denoted by $v_n(x_1)$, $v_\lambda(x_1)$, and $v_\infty(x_1)$, respectively. The definitions can then be reformulated as follows:

$$v_n(x_1) = \sup_{\sigma\in\Sigma}\inf_{\tau\in\mathcal{T}} \gamma_n^1(x_1,\sigma,\tau) = \inf_{\tau\in\mathcal{T}}\sup_{\sigma\in\Sigma} \gamma_n^1(x_1,\sigma,\tau),$$

and $v_\infty(x_1)$, called the uniform value, is the unique number (if it exists) such that, for every $\varepsilon > 0$, there exist (σ^*,τ^*) and $N \geq 1$ such that, for every $n \geq N$,

$$\forall \tau \in \mathcal{T},\ \gamma_n^1(x_1,\sigma^*,\tau) \geq v_\infty(x_1) - 2\varepsilon,$$
$$\forall \sigma \in \Sigma,\ \gamma_n^1(x_1,\sigma,\tau^*) \leq v_\infty(x_1) + 2\varepsilon.$$

In zero-sum games, the existence of the uniform value implies that the sequence of n-stage values $(v_n(x_1))_{n\geq 1}$ (resp. $(v_\lambda)_{\lambda\in(0,1]}$) converges to $v_\infty(x_1)$ when n goes to ∞ (resp. λ goes to 0). Hence, the uniform value is a good reference point for a modeler who knows that the agents are patient but who does not know exactly the discount factor of the players. On the contrary, if the sequence of discounted value does not converge, then the value and the optimal strategy become sensitive to the discount factor. Then, the modeler needs to know precisely the discount factor to predict the payoffs at equilibrium and the behavior of the players.

In the non-zero-sum framework, the situation is more complicated. The existence of a uniform equilibrium does not imply the convergence of the set of discounted equilibria (Renault and Ziliotto [16]) and, even in the case of convergence, it may converge to a set disjoint from the uniform equilibrium set (Sorin [18]). Nevertheless, the conditions in Definition 2.5 imply that, for every $n \geq N$, (σ^*,τ^*) is a 2ε-Nash equilibrium profile in the n-stage game. It also implies that there exists λ_0 such that, for every $\lambda < \lambda_0$, it is a 2ε-Nash equilibrium profile in the λ-discounted game. Hence, the existence of a uniform Nash-equilibrium implies the existence of a ε-Nash equilibrium in the discounted game, close to the uniform equilibrium, for ε arbitrarily small and patient players. From a uniform Nash equilibrium, one can deduce approximate Nash equilibria in discounted games.

Remark 2.6 *It is interesting to point out that stochastic games with a compact state space and finite sets of actions appear naturally in the study of partial observation Markov decision processes and in the study of finite repeated games with symmetric information. In these games, instead of perfectly observing the state variable, the players observe past actions and a public signal. Applying Bayes rules, they may deduce a belief in the state variable. As such, if the initial set of states is K, then the belief is a probability over K and can be seen as an element in $\mathbb{R}^{|K|}$. Moreover, one can define an auxiliary stochastic game with finite sets of actions where the new state is the set of beliefs. In the case where the players only observe past actions but do not receive any additional signal, this auxiliary stochastic game is also deterministic and linear. The properties of the transitions in the dynamic opinion game and in this auxiliary stochastic game are not the same, but the two models share some similarities. In particular, both transitions are 1-Lipschitz for two different norms. As such, the methods of the proofs in this article are related to methods used in stochastic games with partial information and Partial Observation Markov Decision Processes, in particular in Rosenberg et al. [17] and in Venel [20].*

3 Results

We start with two easy applications of the theory of stochastic games. Then, we state our three main results.

Theorem 3.1 *Let Γ be an opinion game such that there is only one active player (J is reduced to a singleton). For every $x_1 \in X$, the game $\Gamma(x_1)$ has a uniform value in pure strategies.*

The proof of this result is omitted. The key argument is that the transition is 1-Lipschitz for the L_∞-norm:

$$\forall x, x' \in X, \forall i, j \in I \times J, \|q(x,i,j) - q(x',i,j)\|_\infty \leq \|x - x'\|_\infty,$$

where, for every $p \in X$, $\|p\|_\infty = \max_{k \in K} |p_k|$. It is an immediate consequence of Venel and Ziliotto [21] after reformulating the model with suitable notation. Similarly, one can obtain the following result by applying Venel [20].

Theorem 3.2 *Let Γ be an opinion game such that for every (i,j) in $I \times J$ and for every (i',j') in $I \times J$,*

$$\forall x \in X,\ q(q(x,i,j),i',j') = q(q(x,i',j'),i,j).$$

For every $x_1 \in X$, $\Gamma(x_1)$ has a uniform Nash equilibrium payoff.

Indeed, the transition is deterministic, commutative, and 1-Lipschitz for $\|.\|_\infty$, which is a polyhedral norm (i.e., such that the unit ball has a finite number of extreme points). Therefore, Theorem 4 in Venel [20] applies and the game has a uniform Nash equilibrium payoff.

In the general case, there may not exist a uniform Nash equilibrium payoff. Adapting one of the counterexamples from Ziliotto [24], we can prove the following theorem.

Theorem 3.3 *There exists a two-person zero-sum opinion game Γ and an initial state x_1 such that $v_\lambda(x_1)$ does not converge when λ goes to 0. This implies that there is no uniform value.*

The proof of this result is given in Section 4.

We now focus on a framework inspired by Mandel and Venel [14] where players are restricted. The idea is to define each matrix $M(i,j)$ as a perturbation of a common matrix M. Formally, an opinion game with limited lobbies is given by $(K, I, J, M, \rho, \mu, g^1, g^2)$ where

- K is a finite set of non-strategic agents,
- I is an action set for Player 1,
- J is an action set for Player 2,
- M is a row stochastic matrix of size $K \times K$ describing the network among the agents without external influence,
- $\rho : I \times J \times K \to [0,1]$ and $\mu : I \times J \times K \to [0,1]$ are two non-negative influence functions such that $\rho + \mu \leq 1$, representing the lobbying powers of Player 1 and Player 2, respectively,
- g^1 and g^2 are two continuous payoff functions from $X \times I \times J$ to $[0,1]$.

The induced opinion game is then given by the following transition matrix

$$B(i,j)_{kl} = (1 - \rho(i,j,k) - \mu(i,j,k))M_{kl},$$

and $A_1(i,j)$ and $A_0(i,j)$ the two vectors such that

$$A_1(i,j)_k = \rho(i,j,k)$$

and

$$A_0(i,j)_k = \mu(i,j,k).$$

We see that the influence of the players is much more constrained than in the general formulation. Indeed, each lobby can directly influence several agents, but the structure of interactions between the non-strategic agents is always induced by the structure of the network M. We obtain the following result.

Theorem 3.4 *Let Γ be an opinion game with limited lobbies. Then, for every $x_1 \in X$ the opinion game $\Gamma(x_1)$ has a uniform Nash equilibrium payoff.*

Mandel and Venel [14] studied the zero-sum case where the network induced by M is primitive, the weights functions for the lobbies are restricted, each player can only influence one agent at each period, and the payoffs for the players are the average of the opinions in the society:

$$g^1(x,i,j) = -g^2(x,i,j) = \frac{1}{|K|}\sum_{k\in K} x_k.$$

They show the existence of the uniform value by noticing that the value does not depend on the initial vector of opinions in the society. When the matrix is not primitive, this is no longer true, as shown by Example 3.5, and therefore the proof used in Mandel and Venel [14] is not possible. The new proof relies mainly on the reduction to a stochastic game with a finite set of states and finite action spaces. We can then deduce the existence of the uniform value by using the result of Mertens and Neyman [15], or the existence of a uniform equilibrium by using Vieille [22, 23]. These strong results allow us to relax many additional assumptions besides the primitivity of M.

Example 3.5 *We consider the following example $\Gamma = (K, I, J, M, \rho, \mu, g^1, g^2)$. The set of agents is $K = \{a, b, c\}$ and the set of actions are $I = J = K$. The matrix M is the identity matrix represented in Figure 1. Moreover, the influence functions are given, for all $i \in I$ and $j \in J$, by*

$$\rho(i,j,k) = \delta_i(k) \text{ and } \mu(i,j,k) = \delta_j(k) \text{ if } i \neq j,$$
$$\rho(i,j,k) = \frac{1}{2}\delta_i(k) \text{ and } \mu(i,j,k) = \frac{1}{2}\delta_j(k) \text{ if } i = j.$$

Informally, each strategic player targets one non-strategic agent and has an infinite convincing power over his target. If the two players choose the same target, then each of them has an influence with weight 1/2. Finally, we assume that $g^1(x) = -g^2(x) = \frac{1}{3}(x_1 + x_2 + x_3)$ is the average opinion in the society.

Let us prove that, if the initial vector of opinions is (α, β, γ) with $\alpha \leq \beta \leq \gamma$, then we have

$$\lim_{n\to+\infty} v_n(\alpha,\beta,\gamma) = (\beta+1)/3.$$

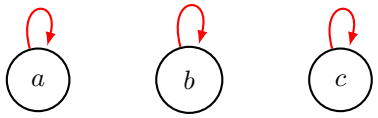


Figure 1: Example where the asymptotic value depends on the original vector of opinions

Proof. First, let us consider that both players target, at every stage, the agent whose opinion is furthest away from their own opinion: for example if the vector of opinions is (α, β, γ), then Player 1 targets the agent with opinion α, whereas Player 2 targets the agent with opinion γ. If there are several such agents, then they uniformly chose one of them. Since they have infinite power of persuasion, this yields the new configuration of opinion $(0, \beta, 1)$ that is invariant (up to relabeling of the agents) and we obtain an average opinion of $(\beta + 1)/3$.

Let us show that no player can do better by deviating. In order to prove this, we will use the following two facts. First, the limit value is invariant to permutation of the original vector of opinions. Second, if a vector of opinions x is greater than x' coordinate by coordinate and h is a fixed sequence of actions, then

the vector of opinions at every stage n starting at x along h is greater coordinate by coordinate than the vector of opinions starting at x' along h. This implies in particular that the value from x is greater than the value from x'.

We focus on the vector of opinions $(0, \beta, 1)$. The one-stage opinion is $(\beta+1)/3$. Moreover, by targeting the agent with opinion 0, Player 1 is sure to reach one of the following three vectors of opinions:

- $(1/2, \beta, 1)$ if Player 2 targets 0,
- $(1, 0, 1)$ if Player 2 targets β,
- $(1, \beta, 0)$ if Player 2 targets 1.

In the three cases, we obtain a new vector that is greater (up to permutation) than $(0, \beta, 1)$, coordinate by coordinate. By the previous facts, this implies that, from $(0, \beta, 1)$, Player 1 is sure to obtain at every stage a vector of opinions greater (up to permutation) coordinate by coordinate than $(0, \beta, 1)$. Hence, he ensures a stage payoff at every stage greater than $(\beta + 1)/3$.

We will now show that Player 1 can guarantee $(\beta + 1)/3$ from (α, β, γ). First, by targeting the agent with the lowest opinion, Player 1 ensures that the median value increases. Second, if Player 1 and Player 2 target different agents, then the vector of opinions is of the form $(0, \beta, 1)$. Therefore along any play with probability one, one reaches a vector of opinions $(x, \beta', 1)$ with $\beta' \geq \beta$, thus yielding the result stated above. By symmetry, Player 2 ensures that Player 1 gets at most $(\beta + 1)/3$, hence the result. ■

We previously pointed out that the existence of a uniform Nash equilibrium in non-zero-sum games only provides some information about approximate discounted equilibria. The sequence of sets of discounted equilibria may still be irregular. To complete the previous result, we formulate variants of the counterexample in Sorin [18] and the counterexample in Renault and Ziliotto [16] that we can then express as dynamic opinion games with limited lobbies. In Renault and Ziliotto [16], the authors show in addition that the sequence of discounted equilibria payoffs obtained by stationary strategies always converges. The analogous result in our framework is an open question.

Proposition 3.6 *We have the following counterexamples:*

- *There exists a dynamic opinion game Γ with limited lobbies and an initial vector of opinions x_1 such that $(E_\lambda(x_1))_{\lambda \in (0,1]}$ converges when λ goes to 0 to a set disjoint from $E_\infty(x_1)$.*
- *There exists a dynamic opinion game Γ' with limited lobbies and an initial vector of opinions x_1 such that $(E_\lambda(x_1))_{\lambda \in (0,1]}$ does not converge when λ goes to 0.*

4 Proof of Theorem 3.3

To prove the result, we formulate a variant of the counterexample of Ziliotto [24]. Let us recall some facts about the structure of this example. At every stage, only one of the players has an active role; he will be called the active player. When Player 1 is the active player, he obtains a stage payoff of 0 (the worst payoff in the game). Moreover, Player 1 then has two possible actions. He can either decide to "Continue" or to "Quit". When he decides to "Quit", Player 1 takes a risk and a lottery is drawn between two outcomes, a bad one and a neutral one. The bad outcome is the absorption of the game in a state with payoff 0 and it happens with a certain probability p depending on the history. The neutral outcome is for Player 2 to become the active player and it happens with the remaining probability $1 - p$. When Player 1 chooses to "Continue", he remains the active player and plays again. Notice that since the history has changed, the probability p changes. The situation is similar for Player 2. When Player 2 is the active player, the stage payoff is 1 (the worst payoff for him) and he has to decide to "Continue" or to "Quit". Hence, for a fixed discount factor, the players need to balance the current stage payoff and the improvement on the risk for the future. By construction, there is an asymmetry between the two players

such that they cannot take the same risks. One of them has more choice than the other. This implies non-convergence when the discount factor becomes close to 0. For some discount factor, the optimal risk is possible for both players, and hence the game is essentially fair, leading to a value of 1/2. For some other discount factors, the optimal risk can only be taken by one of them and the value is bounded away from 1/2.

We first present a repeated game with symmetric information that is a variation of the example in Ziliotto [24]. We argue that the modifications do not change the discounted values of the game. Hence the sequence of discounted values in this variant game does not converge. We then introduce an equivalent stochastic game with an infinite number of states. Finally, we express this stochastic game as a dynamic opinion game.

4.1 Model with signal

We consider the following variation of the example in Ziliotto [24]. The players play a repeated game with symmetric information Γ, with state space $K = \{1^*, 1^{++}, 1_T^{++}, 1^+, 1_T^+, 0^*, 0^{++}, 0^+\}$, action sets $I = J = \{C, Q\}$, and signal set $A = \{D, D'\}$. As explained in the introduction, in a repeated game with symmetric information, at every stage of the game, the players choose their actions and a pair consisting of a new state and a public signal is chosen according to the transition function. The players then observe the public signal and past actions.

The transitions are depicted in Figure 2. Let us comment on this figure. First, an arrow between a state k and a state k' denoted by $(C, 1/2, D')$ is interpreted as follows: if, at state k, the active player takes action C, then with probability 1/2 the new state is k' and the players observe the signal D'. When the transition is deterministic, the number 1 is omitted. Note that "Quit" is always associated with the signal D and leads to either absorption or a switch from the state $\{0^*, 0^{++}\}$ to $\{1^+, 1_T^+, 1^{++}, 1_T^{++}\}$ (or the opposite). In contrast, the action "Continue" may be associated with signal D or D' with a fixed probability $(1/2, 1/2)$.

The transitions are the same as in Ziliotto [24], except for two changes. First, the players do not observe the result of the coin toss when they quit. Hence, they do not know if the game is absorbed or if the other player became active. Formally, we have

$$q(0^{++}, Q) = \delta_{0^*, D},$$
$$q(1^{++}, Q) = \delta_{1^*, D}.$$

Second, we rename the state 1_T as 1_T^{++}. Then, we duplicate the state 1^+ of the original game when Player 2 is the active player into two copies 1^+ and 1_T^+ to keep track of whether the preceding state was 1^{++} or 1_T^{++}. Formally, we have

$$q(1^{++}, C) = \frac{1}{2}\delta_{1_T^{++}, D} + \frac{1}{2}\delta_{1^{++}, D'},$$
$$q(1_T^{++}, C) = \frac{3}{8}\delta_{1^+, D} + \frac{1}{8}\delta_{1^{++}, D} + \frac{1}{2}\delta_{1^{++}, D'},$$
$$q(1^+, C) = \frac{1}{2}\delta_{1_T^+, D} + \frac{1}{2}\delta_{1^{++}, D'},$$
$$q(1_T^+, C) = \frac{1}{2}\delta_{1^+, D} + \frac{1}{2}\delta_{1^{++}, D'}.$$

This last change implies that the following sequence of transitions in the original model

$$(1^+, C, D), (1^+, C, D), (1^+, C, D), (1^+, C, D), ...$$

becomes

$$(1^+, C, D), (1_T^+, C, D), (1^+, C, D), (1_T^+, C, D), ...$$

Moreover, we assume that quitting in this copy yields the absorbing payoff 1 that is the worst for Player 2.

$$q(1_T^+, C) = \delta_{1^*,D},$$

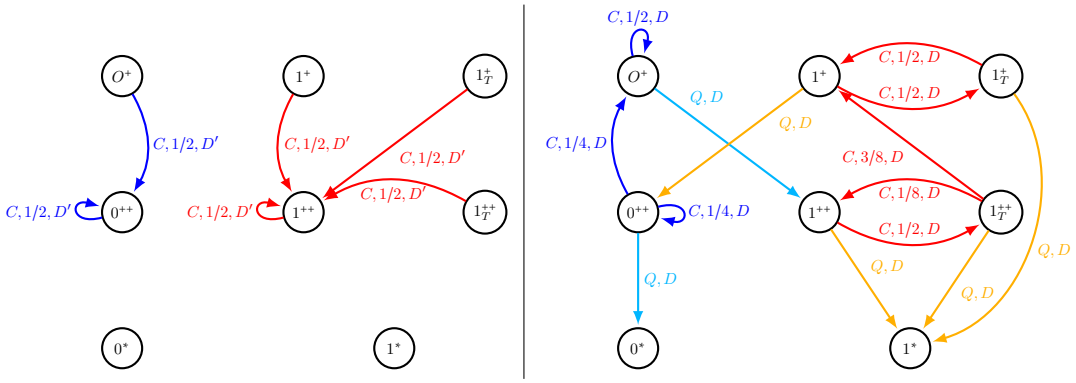


Figure 2: Transition of the repeated game (signal D' on the left and signal D on the right).

These two transformations do not change the value of the game. The first change implies that the players do not observe when absorption occurs. Nevertheless, since they know the actions played and since the payoffs in state 0^* and in state 1^* do not depend on the actions, the players can always play conditionally on the event that the process has not absorbed.

The second change does not change the value. Indeed, in the original game, Player 2 never decides to quit when the state can be 1_T^{++} since it would be dominant to quit at the previous stage. In our version, it is even worse to quit since, if the state is 1_T^{++} with positive probability, then with the remaining probability, the state is 1_T^+, where quitting yields the worst payoff possible. Hence, Player 2 never quits in the T-version. Since we changed the transition in an action that was not used and that is still not used, the value does not change.

4.2 Equivalent game with a countable set of states

To any repeated game with signals such that the players observe the actions and a public signal, one can associate a stochastic game defined on an auxiliary state space, the belief space, where the players observe everything. Let us describe the game associated with the previous repeated game.

Instead of using the set of probability distributions over K as the new state variable, we will encode the state as follows. The set of states will be $X = \{(\alpha, \beta, 1_n),\ \alpha \in [0,1],\ \beta \in [0,1] \text{ and } n \in \mathbb{N}\} \cup \{(\alpha, \beta, 0_n),\ \alpha \in [0,1],\ \beta \in [0,1] \text{ and } n \in \mathbb{N}\}$. The state $(\alpha, \beta, 1_n)$ (resp. $(\alpha, \beta, 0_n)$) is interpreted as follows:

- with probability α, the state is 1^*,
- with probability $1-\beta$, the state is 0^*,
- with the remaining probability, the process has not been absorbed yet and the active player is Player 2 (resp. Player 1),
- since the last change of the active player, the active player played "Continue" n times and each time the signal D was observed.

Note in particular that, when $\alpha = \beta$, the process has been absorbed with probability 1.

Formally, we define

$$(\alpha, \beta, 1_{2n}) = \alpha \cdot 1^* + (1-\beta) \cdot 0^* + (\beta - \alpha) 2^{-2n} \cdot 1^{++} + (\beta - \alpha)(1 - 2^{-2n}) \cdot 1^+, \tag{4.1}$$

$$(\alpha, \beta, 1_{2n+1}) = \alpha \cdot 1^* + (1-\beta) \cdot 0^* + (\beta - \alpha) 2^{-2n} \cdot 1_T^{++} + (\beta - \alpha)(1 - 2^{-2n}) \cdot 1_T^+, \tag{4.2}$$

$$(\alpha, \beta, 0_n) = \alpha \cdot 1^* + (1-\beta) \cdot 0^* + (\beta - \alpha) 2^{-n} \cdot 0^{++} + (\beta - \alpha)(1 - 2^{-n}) \cdot 0^+. \tag{4.3}$$

Moreover, given $(\alpha, \beta) \in [0,1] \times [0,1]$ and $n \in \mathbb{N}$, we define

$$\begin{aligned}
\alpha_n &= \alpha + (\beta - \alpha)\frac{1}{2^n} = \frac{1}{2^n}\beta + \left(1 - \frac{1}{2^n}\right)\alpha, \\
\beta_n &= \beta, \\
\alpha'_n &= \alpha, \\
\beta'_n &= \beta - (\beta - \alpha)\frac{1}{2^n} = \beta\left(1 - \frac{1}{2^n}\right) + \alpha\frac{1}{2^n}.
\end{aligned}$$

It is easier to understand the last equation by noting that $1-\beta$ is the probability that the state has been absorbed in state 0^* and $(\beta - \alpha)$ that the process has not been absorbed yet. Hence, the last equation is equivalent to

$$1 - \beta'_n = (1-\beta) + (\beta - \alpha)\frac{1}{2^n}.$$

We obtain the transitions pictured in Figure 3.

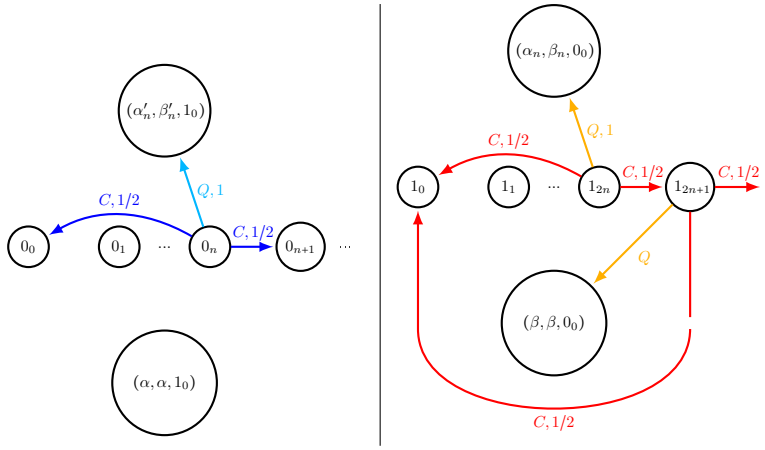


Figure 3: Transition corresponding to Player 1 being active (left) and to Player 2 being active (right) from $(\alpha, \beta, ...)$

.

Let us first explain the left-hand picture that describes the situation when Player 1 is active. Given that Player 1 is the active player at $(\alpha, \beta, 0_n)$, he can take two actions:

- If he plays C, then no additional absorption occurs and the new state is of the form $(\alpha, \beta, .)$. With probability $1/2$, players observe D' and the current state (if not absorbed) resets to 0^{++}, so the distribution is 0_0. With probability $1/2$, players observe D and the probability to be in 0^{++} is divided by 2, so the distribution changes to 0_{n+1}.
- If he plays Q, the players do not observe anything but they know how to update the probabilities:

– the probability of being absorbed in 0^* conditionally on not being absorbed before is $\frac{1}{2^n}$, so the total absorption becomes

$$(\beta - \alpha)\frac{1}{2^n} + (1 - \beta) = 1 - \beta'_n,$$

– the probability of being absorbed in 1^* does not change and is equal to $\alpha = \alpha'_n$,

– with the remaining probability the state is 1^{++}, hence 1_0.

A similar story yields the graph for Player 2. The main difference is the presence of two versions for some states: $\{1^+, 1^{++}\}$ and $\{1^+_T, 1^{++}_T\}$. Consider the last switch/reset as stage 0 and that the process has not been absorbed. Then, the distribution is supported by $\{1^+, 1^{++}\}$ on even stages, whereas the distribution is supported by $\{1^+_T, 1^{++}_T\}$ on odd stages. Given that Player 2 is the active player at $(\alpha, \beta, 1_{2n})$, he can take two actions:

- If he plays C, then no additional absorption occurs and the new state is of the form $(\alpha, \beta, .)$. With probability 1/2, the players observe D' and the current state (if not absorbed) resets to 1^{++}, so the distribution is 1_0. With probability 1/2, the players observe D and the state changes in a deterministic fashion from 1^{++} to 1^{++}_T and from 1^+ to 1^+_T. Hence, the new distribution is $(\alpha, \beta, 1_{2n+1})$.
- If he plays Q, then players do not observe anything but they know how to update the probabilities:
  – the probability of being absorbed in 0^* does not change and is equal to $\beta = \beta_n$,
  – the probability of being absorbed in 1^* conditionally on not being absorbed before is $\frac{1}{2^{2n}}$, so the total absorption becomes
  $$\alpha + (\beta - \alpha)\frac{1}{2^{2n}} = \alpha_{2n},$$
  – with the remaining probability the state is 0^{++}, hence 0_0.

Given that Player 2 is the active player at $(\alpha, \beta, 1_{2n+1})$, he can take two actions:

- If he plays C, then no additional absorption occurs and the new state is of the form $(\alpha, \beta, .)$. With probability 1/2, players observe D' and the current state (if not absorbed) resets to 1^{++}, so the distribution is 1_0. With probability 1/2, players observe D and the state transfers from $\{1^{++}_T, 1^+_T\}$ to $\{1^{++}, 1^+\}$. Moreover, the probability of being in 1^{++} is equal to the probability of being in 1^{++}_T divided by 4, so is equal to $\frac{1}{2^{2(n+1)}}$. It follows that the new distribution is $(\alpha, \beta, 1_{2n+2})$.
- If he plays Q, then players do not observe anything but they know how to update the probabilities:
  – the probability of being absorbed in 0^* does not change and is equal to $\beta = \beta_n$,
  – the probability of being absorbed in 1^* is the complement, so the new distribution is $(\beta, \beta, 0_0)$.

4.3 Opinion game

We will now describe an opinion game that has the same transitions as the game in Figure 3. We first describe the game informally. It is composed of $14+N$ non-strategic agents. We will focus on the changing opinions of the 14 agents. The N additional agents will all be similar and only play a role in controlling the payoff.

The set of non-strategic agents is $K = \{A_T, A_B, O_L, O_R, O_B, O_{BB}, O_T, O_{TT}, U, U_B, U_T, R_T, C, R_B\}$. Each player will have two actions, C and Q. We will design the transition such that each player can have two roles, "active" or "inactive," with the following properties. First, in every reachable configuration, there is only one "active" player. Second, the "inactive" player has a dominated action that is to continue, and hence is temporarily irrelevant. When an active player decides to play action Q, he becomes the "inactive" player, while the other player becomes the active player. Finally, the opinions of the N agents are all equal to the opinion of Agent C.

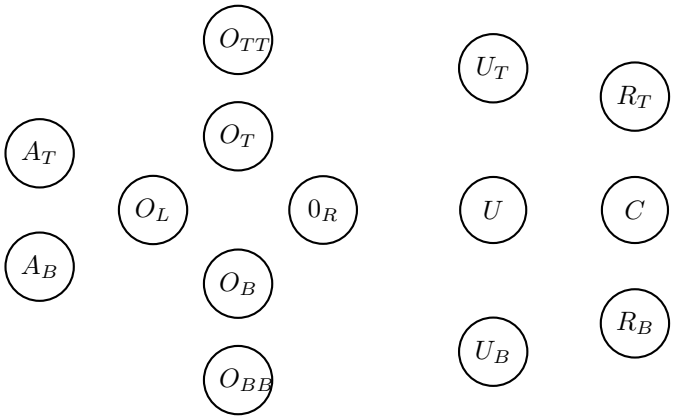


Figure 4: Set of non-strategic agents.

The transitions are given by Figures 5, 6, 7, and 8: the agent at the head of the arrow obtains the weighted average of the opinions of the agents at the origins of the arrows. An agent without any incident arrow does not change his opinion. We will only consider some special configurations of opinions depicted in Figure 9.

Remark 4.1 In the definition of an opinion game, we assumed that the transitions were deterministic. In our example, we assumed that if each player decides to play C then the transition is stochastic: with probability 1/2 the transition is given by $M_1(C,C)$ and with probability 1/2 the transition is given by $M_2(C,C)$. This can be replicated with a deterministic transition by using a jointly controlled lottery: each action C is duplicated into two actions C_1 and C_2 and the transitions are given as follows:

- if (C_1,C_1) or (C_2,C_2) is played, then $M_1(C,C)$ is applied,
- if (C_2,C_1) or (C_1,C_2) is played, then $M_2(C,C)$ is applied.

Now assume that each time a player is supposed to play C, he plays C_1 with probability 1/2 and C_2 with probability 1/2. Then the distribution of transitions is independent of the other player. No player controls the distribution and we obtain the right dynamic.

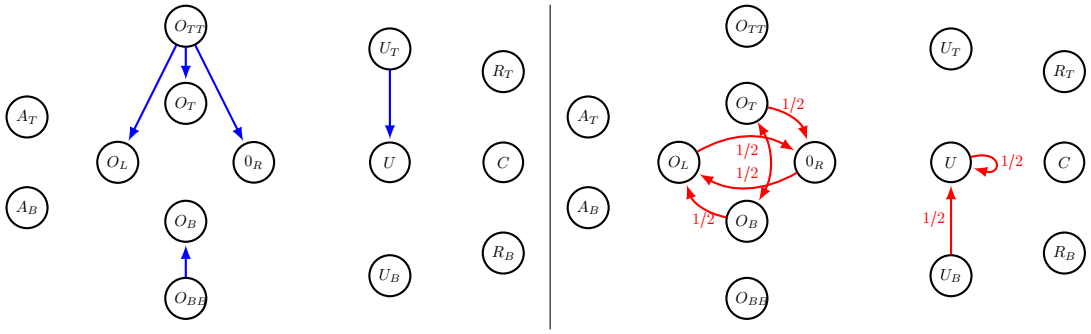


Figure 5: If players play (C, C): with probability 1/2 the transition is given by the left network and with probability 1/2 by the right matrix.

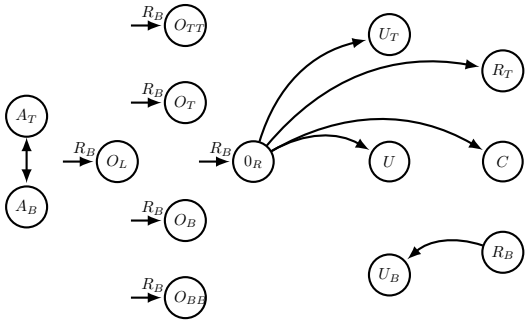


Figure 6: If players play (C, Q)

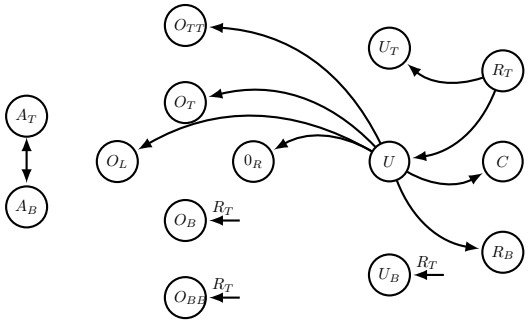


Figure 7: If players play (Q, C)

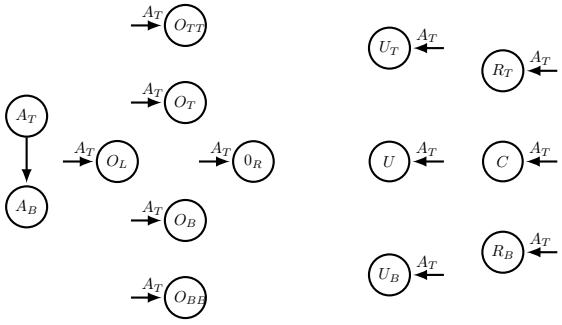


Figure 8: If players play (Q, Q).

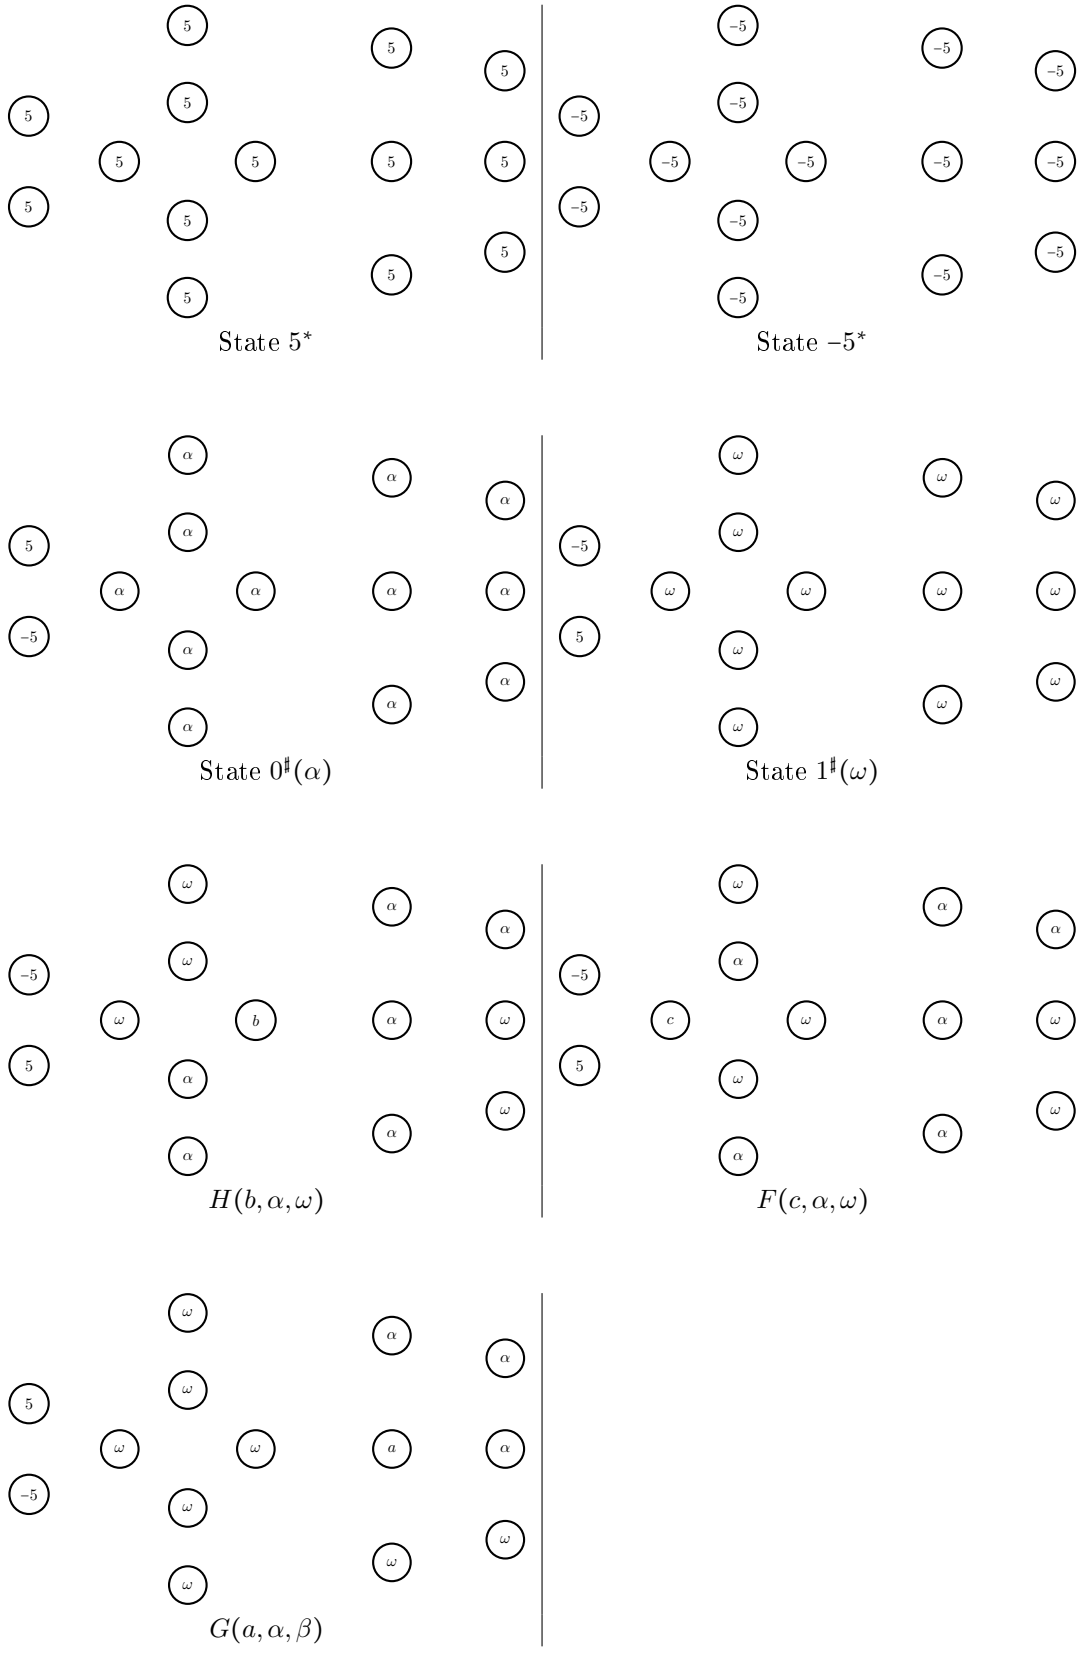


Figure 9: The subset of opinion configurations

The transitions and the configurations are designed such that for every $a, b, c, \alpha, \omega \in [0,1]$, we have:

- from 5^*, any transition yields again 5^*,
- from -5^*, any transition yields again -5^*,
- from $0^\sharp(\alpha)$:
  - $(C,C) \to 0^\sharp(\alpha)$,
  - $(Q,C) \to 1^\sharp(\alpha)$,
  - $(C,Q) \to 1^\sharp(\alpha)$,
  - $(Q,Q) \to 5^*$,
- from $1^\sharp(\omega)$:
  - $(C,C) \to 1^\sharp(\omega)$,
  - $(Q,C) \to 0^\sharp(\omega)$,
  - $(C,Q) \to 0^\sharp(\omega)$,
  - $(Q,Q) \to 5^*$,
- from $G(a,\alpha,\omega)$:
  - $(C,C) \to 1/2\delta_{G(\frac{a+\omega}{2},\alpha,\omega)} + 1/2\delta_{G(\alpha,\alpha,\omega)}$,
  - $(C,Q) \to 1^\sharp(\omega)$,
  - $(Q,C) \to H(a,\alpha,a)$,
  - $(Q,Q) \to 5^*$,
- from $H(b,\alpha,\omega)$:
  - $(C,C) \to 1/2\delta_{F(\frac{b+\alpha}{2},\alpha,\omega)} + 1/2\delta_{H(\omega,\alpha,\omega)}$,
  - $(C,Q) \to G(b,b,\omega)$,
  - $(Q,C) \to 0^\sharp(\alpha)$,
  - $(Q,Q) \to -5^*$,
- from $F(b,\alpha,\omega)$:
  - $(C,C) \to 1/2\delta_{H(\frac{b+\alpha}{2},\alpha,\omega)} + 1/2\delta_{H(\omega,\alpha,\omega)}$,
  - $(C,Q) \to 0^\sharp(\omega)$,
  - $(Q,C) \to 0^\sharp(\alpha)$,
  - $(Q,Q) \to -5^*$.

We now make the following identifications between the new states and the states of the preceding game:

$$(\alpha,\beta,1_{2n}) = H(\alpha_{2n},\alpha,\beta), \tag{4.4}$$
$$(\alpha,\beta,1_{2n+1}) = F(\alpha_{2n+1},\alpha,\beta), \tag{4.5}$$
$$(\alpha,\beta,0_n) = G(\beta'_n,\alpha,\beta), \tag{4.6}$$

with $\alpha'_0 = \alpha$ and $\beta_0 = \beta$.

Let us make several comments on the discounted values. State 5^* (resp. -5^*) is absorbing and each agent in the society has an opinion equal to 5 (resp. -5), hence the discounted values are equal to 5 (resp. -5). For $x \in [0,1]$, both players can ensure that the state stays in $\{1^\sharp(x), 0^\sharp(x)\}$ by playing "Continue," so the discounted values from these states are equal to x. Finally, if $x \in [\alpha,\omega]$, then from $G(x,\alpha,\omega)$, $H(x,\alpha,\omega)$, and $F(b,\alpha,\omega)$, each stage payoff in the remaining game lies between α and ω.

We can now deduce that playing Q for Player 2 in a state of the form $(\alpha,\beta,0_n)$ is dominated. Indeed, by playing Q, he obtains either β or 5, which are the worst payoffs possible, so Q is dominated. Hence, Player 2 plays C and can be assumed to be "inactive." Similarly, Player 1 does not play Q in a state of the form $(\alpha,\beta,1_n)$.

We now prove that the transitions are equal to the transitions of the previous section. It relies on the following simple observations:

$$\alpha_{n+1} = \alpha + (\beta-\alpha)\frac{1}{2^{n+1}} = \frac{1}{2}\alpha_n + \frac{1}{2}\alpha, \tag{4.7}$$
$$\beta'_{n+1} = \beta - (\beta-\alpha)\frac{1}{2^{n+1}} = \frac{1}{2}\beta'_n + \frac{1}{2}\beta. \tag{4.8}$$

For example, if Player 1 plays "Continue" from $(\alpha,\beta,0_n)$, the state is $G(\alpha,\alpha,\beta) = G(\beta'_0,\alpha,\beta) = (\alpha,\beta,0_0)$ with probability 1/2 and the state is $(\alpha,\beta,0_{n+1})$ with probability 1/2.

This concludes the reduction. By choosing N very large, the mean-average opinion in the society is almost equal to the opinion of C. Moreover, the opinion of C at a given state in the opinion game is equal to the payoff at the corresponding state in the stochastic game of Section 4.2. Hence, the sequence of discounted values in the opinion game does not converge.

5 Proof of Theorem 3.4

The proof goes as follows. For every $\varepsilon > 0$, we introduce two auxiliary games, Γ'_ε and $\widetilde{\Gamma}_\varepsilon$. The first game, Γ'_ε, is a reformulation of Γ where actions are grouped into blocks. The originality of this construction is that the size of these blocks will vary, depending on the past history. The second game, $\widetilde{\Gamma}_\varepsilon$, is an approximation of Γ'_ε with a finite set of states and finite sets of actions. Applying the result due to Vieille [22, 23], we obtain the existence of a uniform Nash equilibrium in $\widetilde{\Gamma}_\varepsilon$. Then, we can deduce the existence of a uniform Nash equilibrium in the original game.

The section is decomposed into five subsections. In the first, we recall some definitions of Markov chain theory and we explain how they are related to the diffusion of opinions. In the second and third, we define Γ'_ε and $\widetilde{\Gamma}_\varepsilon$, respectively. In the fourth, we prove several results comparing the payoffs in the three games. Finally, in the fifth, we use these results to conclude.

5.1 Reminder of Markov chain theory

When a player chooses an action that impacts agent k at stage t, he influences the opinion of agent k at stage $t+1$. The key element of the proof is to understand how this influence propagates into the network. In order to do so, we will rely on the standard vocabulary for Markov chains.

Let M be a Markov chain. Given $k \in K$, we define *the period of node* k denoted by $p(k)$ as the greatest common divisor of $\{n \geq 1, (M^n)_{kk} > 0\}$. For example, if $p(k) = 2$, then the node is periodic of period 2 and the opinions at odd stages and even stages are independent of each other. We define *the period* U of the chain as the lowest common denominator of $\{p(k), k \in K\}$. If $U = 1$, then the chain is said to be *aperiodic*.

Let us now focus on communicating classes. Given $n \geq 0$, k in K, and l in K, we say that l is *accessible in* n *stages* from k if

$$(M^n)_{kl} > 0,$$

with the convention that $M^0 = Id$. If an agent l is accessible in n stages from an agent k, then the opinion of l at stage 1 has some impact on the opinion of agent k at stage $n+1$. Hence, agent l influences agent k in n stages. Note that the opinions diffuse in the networks in reverse flow compared to the Markov chain. We say that l is *accessible* from k if there exists some natural number n such that l is accessible in n stages from k. The nodes k and l are *communicating* if l is accessible from k and k is accessible from l. It is clear that communication defines an equivalence relation whose equivalence classes are the maximal communicating sets. A maximal communicating set will be called *a communicating class*. A communicating class C is then said to be *absorbing* if the probability of remaining in the class is equal to 1:

$$\forall k \in C,\ \sum_{l \in C} M_{kl} = 1.$$

In terms of opinions, opinions in C at the previous stage are sufficient to compute the opinions of agents in C at the current stage. If the class is not absorbing, then it is said to be *transient*.

Consider a communicating class such that every element in C has a period equal to 1, the class is said to be *aperiodic* and there exists some natural number $S(C)$ such that

$$\forall k, l \in C,\ \forall n \geq S(C),\ (M^n)_{kl} > 0.$$

This is a minimal time for opinions inside class C to be properly mixed. If C is non-absorbing, we also have

$$\forall k \in C, \exists l \notin C, \ \forall n \geq S(C)+1, \ (M^n)_{kl} > 0. \tag{5.1}$$

We denote by S the sum of $S(C)+1$ over all communicating classes C. It ensures that after S stages the opinions inside a communicating class are properly mixed, and if the communicating class is transient, the opinions after S stages are influenced by the initial opinions on an absorbing class. Formally, we obtain the following proposition whose proof is omitted.

Proposition 5.1 *Let M be an aperiodic chain and C be a communicating class. Then,*

$$\forall k, l \in C, \ \forall n \geq S, \ (M^n)_{kl} > 0.$$

If the class is not absorbing, then there also exists an absorbing class C' such that

$$\forall k \in C, \exists l \in C', \ \forall n \geq S, \ (M^n)_{kl} > 0.$$

Let us now explain the link between the structure of the Markov chain and the propagation of the influence of a player. For the purposes of explanation, we assume that there is only one player, say Player 1, that impacts agent l only once, at stage t. Then, there is no longer any outside influence on the network and the opinions propagate in the network free from outside influence. We want to study how this one-stage influence propagates in the network, in relation to the previous vocabulary. Let k be another agent in the network. We assume that both l and k are in absorbing classes and we distinguish several cases:

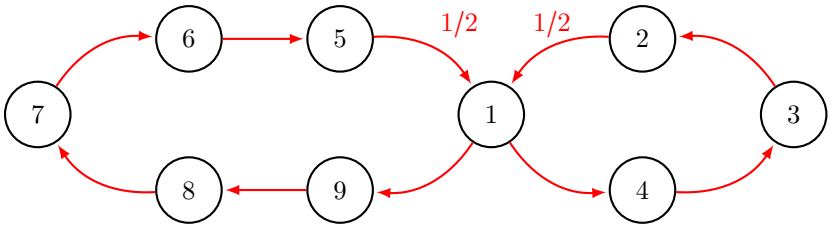
- If k and l are not communicating, then the action of Player 1 will never have any influence on k.
- If k and l are communicating and their class is aperiodic, then the influence of Player 1 will eventually impact the opinion of k at every stage. Nevertheless, it may take some time for the influence to reach agent k.
- If k and l are communicating but their communicating class is periodic, the influence will eventually influence the opinion of k but only at some stages. These stages will depend on the length of the communicating path between l and k and the period of the class. As in the previous case, the propagation of influence may also take some time.

One key element of the previous remark is that the actions of targeting agent l in a periodic communicating class at stage t and at stage $t+1$ are essentially very different: they will influence the opinion of an agent k at different stages. Equivalently, these two actions will influence different sets of states at a given stage $t' \geq t$. We now consider an example to illustrate the third case.

Example 5.2 Consider the Markov chain defined on $M = \{1,2,3,4,5,6,7,8\}$ with transition matrix

$$M = \begin{pmatrix} 0 & 1/2 & 0 & 0 & 1/2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

This Markov chain is made by concatenating two cycles of respective lengths 4 and 6 with 1 as a joined state.



First, the Markov chain has a unique communicating class, since there is a path from any state to any state. Second, the Markov chain is periodic of period 2. Indeed,

$$\{n \in \mathbb{N}, P(X_n = 1 | X_0 = 1)\} = \{6 * l + 4 * m, (l, m) \in \mathbb{N}\}.$$

Hence, the greatest common divisor is 2 and the period of 1 is 2. It is a classic result that all members in the same communicating class have the same period. So the Markov chain is 2-periodic and irreducible.

Let us consider the matrix M^2. It is equal to

$$M^2 = \begin{pmatrix} 0 & 0 & 1/2 & 0 & 0 & 1/2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1/2 & 0 & 0 & 1/2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1/2 & 0 & 0 & 1/2 & 0 & 0 & 0 & 0 \end{pmatrix}$$

We see that there is a separation into two communicating classes as pictured in the following figure. Informally, if we consider even stages, then the opinions of non-strategic agents in $\{1, 3, 6, 8\}$ only depend on the original opinions in $\{1, 3, 6, 8\}$. We have a similar result concerning $\{2, 4, 5, 7, 9\}$. Moreover, both classes are now aperiodic, so the opinions are completely mixed inside each class, provided we wait long enough. For example, if we consider the following matrix, which is the restriction of M^2 to the subset $\{1, 3, 6, 8\}$,

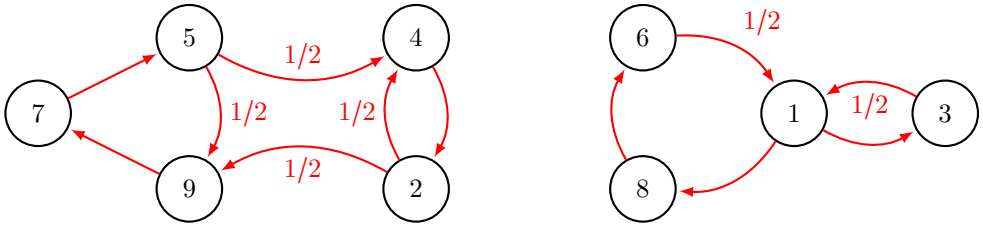
$$B = \begin{pmatrix} 0 & 1/2 & 1/2 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix}$$

then we need to wait until stage 7 (stage 13 in the original timing) for the initial opinion of every agent to have an influence on every other agent forever. Indeed, there is no path of length 5 between agent 8 and agent 6:

$$\{n \in \mathbb{N}, B^n_{68} > 0\} = \{2 + 3k + 2l, (k, l) \in \mathbb{N}^2\} = \{1, 4, 6, 7, 8, 9, ...\}.$$

One can check that every coefficient of B^6 is strictly positive, which implies that it is also true for every $n \geq 6$. Recall that the opinion at stage 7 is given by the image of the initial vector of opinions by B^6:

$$B^6 = \begin{pmatrix} 3/8 & 1/4 & 1/4 & 1/8 \\ 1/2 & 1/8 & 1/8 & 1/4 \\ 1/4 & 1/4 & 1/4 & 1/4 \end{pmatrix}$$



Let us come back to the original diffusion process. If a player targets agent 1 at stage t, we know that it has some impact on the opinion of agent 1 at stage $t+1$. By the previous remark, it follows that it has some impact on the opinions of each agent in $\{1,3,6,8\}$ at stages $t+14, t+16, ...$ and on the opinions of each agent in $\{2,4,5,7,9\}$ at stages $t+15, t+17, ...$

Hence, we see in the previous example that we face two challenges in the proof related to time. The first is to decompose the networks into subgraphs that are not communicating. A naive approach based on the communicating classes of M does not work because we need to take into account the period of the class. The second is that the influence of a player may take some time to propagate into the network. A natural solution used in Mandel and Venel [14] is simply to group actions into blocks of actions, but we will see that, contrary to traditional approaches in the stochastic game literature, one needs to use blocks with varying sizes.

5.2 Definition of Γ'_ε

Given the Markov chain M with period U, define $M' = M^U$. M' is aperiodic, and in particular one can apply the result on communicating aperiodic classes to M'. Define S such that Proposition 5.1 is true, i.e., opinions are mixed in S stages. Let $M'' = M^{SU}$, where SU stages have been concatenated. We obtain:

- M'' is aperiodic,
- for every communicating absorbing class C of M''

$$\forall k, l \in C,\ \forall n \geq 1,\ ((M'')^n)_{kl} > 0.$$

Let $\{K^1, \dots, K^D\}$ be the absorbing communicating classes associated with M'' and K^0 the union of all transient communicating classes of M''. Finally, we know that the sequence $(M'')^n$ converges to some matrix B_0 such that $M''B_0 = B_0M'' = B_0$. More precisely, for $\varepsilon \in (0, 1/2]$, there exists $R \geq 1$ such that for every $n \geq R$,

$$\|(M'')^n - B_0\|_\infty \leq \varepsilon,$$

where matrices of dimension $|K| \times |K|$ are seen as elements of $\mathbb{R}^{|K|\times|K|}$.

Let $T = |K|$ and N a natural number such that $N \geq \max(R, T/\varepsilon)$. This implies in particular that $T/N \leq \varepsilon$. We will construct large blocks of actions which can be decomposed into small blocks of size SU. By introducing U, we take into account the periodicity. By introducing S, we take into account the mixing of the opinions inside a communicating class. The size of the large blocks will be a multiple of SU varying between NSU and $(N+T)SU$. The choice of N ensures that the relative difference in size between the large blocks will be small compared to the total length of a block (only a fraction ε). This remark will only be used in the last step of the proof when comparing the payoff functions in Γ_ε and Γ'_ε.

Let us define the auxiliary game $\Gamma'_\varepsilon = (X, A, B, q', g')$. X is independent of ε, but A, B, q', and g' depend on ε. For clarity of notation, we will omit this dependence in the notation.

5.2.1 Definition of the sets of actions

Recall that the informal idea is to group actions in blocks. In the original game, the players may base their decisions on the history up to the current stage. Once the stages are grouped into a block, this information becomes separated into two parts: the history up to the beginning of the block and the history between the start of the block and the current stage. In order, to obtain a relevant auxiliary game, the new action space needs to incorporate the dependence on the actions previously played in the current block.

Let $\Omega = \cup_{l=1}^{(N+T)SU} (I \times J)^{l-1}$ be the set of histories of length strictly less than $(N+T)SU$, i.e., the set of all partial histories possible inside a block of $(N+T)SU$ stages. The sets of actions are defined by $A = I^{\Omega}$ and $B = J^{\Omega}$.

We denote by $H'_n = (A \times B)^{n-1}$ the set of n-stage histories in the auxiliary game, by $H' = \cup_{n \geq 1} H'_n$ the set of histories and by $H'_\infty = (A \times B)^{+\infty}$ the set of plays. We denote the set of strategies of Player 1 by Σ' and the set of strategies of Player 2 by $\mathcal{T}'$.

5.2.2 Definition of the payoff function

A state x and a pair of actions (a, b) in the auxiliary game induce a history in the original game of length $(N+T)SU$. The payoff function $g'(x, a, b)$ is defined as the mean average of the payoffs along this virtual history on $t(a,b)SU$ stages with $N \leq t(a,b) \leq N+T$. The novelty here is that the number of stages that is considered in the original game is variable: every large block consists of $t(a,b)$ small blocks of SU stages, for some natural number $t(a,b)$ depending on a and b. In particular, each large block has a length between NSU and $(N+T)SU$.

The intuition for having blocks of varying sizes is the following. On each block, we want to partition communicating classes into two groups. On one hand, we have the communicating classes where no agent is influenced by the players. In these classes, the opinions evolve according to M. On the other hand, we have the communicating classes where every agent is influenced by the players. In these classes, the transition is contracting since every agent is impacted by the external influence with a strictly positive weight. To obtain such a decomposition, it may be necessary to wait at least SU stages for the external influences to propagate inside a class that has been targeted by a player. During these waiting stages, the player may influence a new communicating class, so this waiting time may need to be renewed. Let us insist that it is impossible to obtain such a decomposition if we fix the size of the blocks independently of the actions that are played. Indeed, it is always possible for a player to target a new class in the last stage and the influence will not have time to propagate.

We first define $t(h)$ for a finite sequence h of LSU pairs of actions with $L \geq N$. $t(h)$ will be equal to an integer or to $+\infty$. Fix a sequence h of LSU pairs of actions denoted by $(i_1, j_1, ..., i_{LSU}, j_{LSU})$ such that $L \geq N$. For every $i \in I$, $j \in J$ and $t \leq LSU$, we define three sets of agents $K^1(i,j,t), \overline{K}^1(i,j,t)$, and $\overline{\overline{K}}^1(i,j,t)$ as follows. First $K^1(i,j,t) = \{k,\ \rho(k,i,j) > 0\}$, then

$$\overline{K}^1(i,j,t) = \cup_{l \in K^1(i,j,t)} \left\{ k \in K,\ \left(M^{\lceil \frac{t+1}{U} \rceil U - (t+1)} \right)_{kl} > 0 \right\},$$

and

$$\overline{\overline{K}}^1(i,j,t) = \cup_{l \in K^1(i,j,t)} \left\{ k \in K,\ \left(M^{SU + \lceil \frac{t+1}{U} \rceil U - (t+1)} \right)_{kl} > 0 \right\}.$$

$K^1(i,j,t)$ is the set of agents whose opinions at stage $t+1$ is influenced by the action of Player 1 at stage t, $\overline{K}^1(i,j,t)$ is the set of agents whose opinions are influenced by this action at the next stage that is a multiple of U, and $\overline{\overline{K}}^1(i,j,t)$ is the communicating class corresponding to $\overline{K}^1(i,j,t)$, so it is the set of all agents whose opinion will eventually be influenced by this action at stages that are multiples of U. We define $K^2(i,j,t)$, $\overline{K}^2(i,j,t)$, and $\overline{\overline{K}}^2(i,j,t)$ similarly for Player 2. We have $K^2(i,j,t) = \{k,\ \mu(k,i,j) > 0\}$,

$$\overline{K}^2(i,j,t) = \cup_{l \in K^2(i,j,t)} \left\{ k \in K,\ \left(M^{\lceil \frac{t+1}{U} \rceil U - (t+1)} \right)_{kl} > 0 \right\},$$

and

$$\overline{\overline{K}}^2(i,j,t) = \cup_{l \in K^2(i,j,t)} \left\{ k \in K,\ \left(M^{SU + \lceil \frac{t+1}{U} \rceil U - (t+1)} \right)_{kl} > 0 \right\}.$$

For every $n \leq L$, we define $\overline{\overline{K}}_n(h)$ by

$$\overline{\overline{K}}_n(h) = \cup_{t=1}^{nSU} \left(\overline{\overline{K}}^1(i_t, j_t, t) \cup \overline{\overline{K}}^2(i_t, j_t, t) \right).$$

$\overline{\overline{K}}_n(h)$ is the set of agents that eventually, at stages that are multiples of U, come under the influence of the actions of the players before stage nSU (included). The sequence $\left(\overline{\overline{K}}_m(h)\right)_{N-1\leq m\leq L}$ is a non-decreasing sequence of $L-N+2$ sets. Let $\mathcal{Z}$ be the set of integers $t\in\{N,...,L\}$ such that

$$\overline{\overline{K}}_{t-1}(h)=\overline{\overline{K}}_t(h). \tag{5.2}$$

Informally, the actions induced during the tth blocks do not open up a new communicating class compared to the classes already under influence. If this set is not empty, we define $t(h)$ to be the smallest natural number t satisfying Equation 5.2. If $\mathcal{Z}$ is empty, then $t(h)$ is set to $+\infty$. If $L\geq N+|K|$, $\mathcal{Z}$ is non-empty and $t(h)$ is a natural number. Indeed, the sequence $\left(\overline{\overline{K}}_t(h)\right)_{N-1\leq m\leq L}$ has $L-N+2$ elements which is strictly larger than $|K|$ and it is bounded from above by K (for inclusion order), so it may only strictly increase at most K times. This leads to a contradiction so $\mathcal{Z}$ is non-empty and $t(h)$ must be a natural number.

We now apply this function to the history induced by the actions played. For every $a\in A$ and $b\in B$, we define the finite sequence of actions $H(a,b)=(i_1,j_1,...,i_{(N+T)SU},j_{(N+T)SU})$ such that $i_1=a(\varnothing)$, $j_1=b(\varnothing)$ and for every $t\geq 1$, $i_t=a(i_1,j_1,...,i_{t-1},j_{t-1})$ and $j_t=b(i_1,j_1,...,i_{t-1},j_{t-1})$. Since $L=N+T=N+|K|$, we can define the strictly positive integer $t(a,b):=t(H(a,b))$ and also

$$(g')^m(x,a,b):=\frac{1}{t(a,b)SU}\sum_{t=1}^{t(a,b)SU}g^m(x_t,i_t,j_t),\text{ for every }m\in\{1,2\}.$$

By construction $t(a,b)SU\leq(N+T)SU$ so the payoff is well defined. We also define $K(a,b):=\overline{\overline{K}}_{t(a,b)}(H(a,b))$.

5.2.3 Definition of the transition q'

We now define the auxiliary transition by considering the original transition over $t(a,b)SU$ stages.

First, we extend q to finite histories as follows: for every $n\geq 1$, we define inductively

$$\forall x\in X,\ \forall(h_n,i,j)\in H_{n+1},\ q(x,(h_n,i,j))=q(q(x,h_n),i,j).$$

Informally, given $x\in X$ and $h\in H_{m+1}$, $q(x,h)$ is the vector of opinions obtained after playing the sequence of actions h from the initial vector of opinions x.

We now define the new transition as follows: for every $x\in X$, for every $a\in A$, and for every $b\in B$,

$$q'(x,a,b)=q(x,H(a,b)_{t(a,b)SU}),$$

where $H(a,b)_{t(a,b)SU}$ is the truncation of $H(a,b)$ up to stage $t(a,b)SU$ (action included).

5.3 Definition of $\widetilde{\Gamma}_\varepsilon$

The construction of $\widetilde{\Gamma}_\varepsilon=(X_f,A,B,\widetilde{q},\widetilde{g})$ is similar to the construction of Γ'_ε but with an additional step of projection onto a finite state space X_f. As in the previous section, we omit the dependence on ε for A, B, $\widetilde{q}$, and $\widetilde{g}$. The finite set X_f will be constructed starting from two finite auxiliary sets: a first set that is an ε-dense subset of the full space X and a second set that is based on the stable vectors of opinions for the Markov chain M. We then project each communicating class differently depending on the actions of the players.

The sets of actions and the payoff functions in $\widetilde{\Gamma}_\varepsilon$ are the same as in the game Γ'_ε. Hence, we have $A = I^\Omega$ and $B = J^\Omega$, where $\Omega = \cup_{l=1}^{(N+T)SU}(I \times J)^{l-1}$. For every $x \in X_f$, for every $m \in \{1,2\}$, for every $a \in A$, and for every $b \in B$, we have

$$\widetilde{g}^m(x,a,b) = (g^m)'(x,a,b).$$

We now define the set of states X_f. We start with some notation. Given a vector $x \in [0,1]^K$ and $E \subseteq \{1,...,K\}$, we define x_E to be the $|E|$-vector comprising only the coordinates of x in E. Given a subset O of $[0,1]^K$, we define the trace of O along the coordinates in E by

$$O_E = \{x_E, x \in O\},$$

and

$$\|x\|_{E,\infty} = \sup_{k \in E} |x_k|.$$

When $E = \{k\}$ is a singleton, the set $O_{\{k\}}$ will be denoted by O_k. Let G be an ε-dense finite subset of $[0,1]^K$ for the infinity norm:

$$\forall x \in [0,1]^K,\ \exists z \in G,\ \|x - z\|_\infty \leq \varepsilon.$$

Let $X_\infty = \{B_0 x, x \in X\}$ be the set of stable vectors of opinions for M^U. By construction, X_∞ is a closed subspace of a compact set, and hence compact. There exists an ε-dense finite subset F of X_∞ for the infinity norm:

$$\forall x \in X_\infty,\ \exists z \in F,\ \|x - z\|_\infty \leq \varepsilon.$$

We define $X_f \subset [0,1]^K$ by

$$X_f = \left\{(x_1, x_2, \ldots, x_{|K|}), \forall k \in K, x_k \in G_k \cup F_k\right\}.$$

By construction, X_f is finite.

We now define the auxiliary transition by cleverly projecting the transition q' onto the set X_f. Given $E \subseteq \{1,...,K\}$ and $x \in X$, we define the projection $\Pi_E(x)$ from X to X_f such that

- $(\Pi_E(x))_E \in G_E$,
- $(\Pi_E(x))_{K \setminus E} \in F_{K \setminus E}$,
- the distance between x and $\Pi_E(x)$ is minimal.

There may be several elements satisfying these three conditions. Whenever this is the case, we select one of them. We define the transition $\widetilde{q}$ for all $x \in X_f$, for all $a \in A$, and for all $b \in B$ by

$$\widetilde{q}(x,a,b) = \Pi_{K(a,b) \cup K^0}\left(q'(x,a,b)\right).$$

Hence, the transition is defined in two steps: the transition q' and a projection onto the finite set X_f. The coordinates of nodes in $K(a,b)$ and in K^0 are projected onto the grid G while the coordinates of nodes not in $K(a,b)$ nor in K^0 are projected onto the grid F, hence onto an invariant vector of opinions for M (restricted to these communicating classes). The mapping $\widetilde{q}$ is well defined from X_f to X_f by construction.

5.4 Comparison between the auxiliary game and the original game

We now prove several results regarding these three games. In the first section, we prove that plays in the games Γ'_ε and $\widetilde{\Gamma}_\varepsilon$ stay close, whence the existence of an equilibrium in $\widetilde{\Gamma}_\varepsilon$ implies the existence of an approximate equilibrium in Γ'_ε. In the second section, we show that averages in Γ'_ε are almost averages in Γ.

5.4.1 Proximity of the transitions

Let us prove that the sequence of states along a play in $\widetilde{\Gamma}_\varepsilon$ stay close to the sequence of states along the same play in Γ'_ε. We will first study each absorbing communicating class of M'' separately. Then, we will focus on transient communicating classes.

First we describe the key elements of the proof informally. Fix a communicating class K^d, a play h (in the auxiliary games), and an initial vector of opinions x_1. The play generates two distinct trajectories in $\widetilde{\Gamma}_\varepsilon$ and Γ'_ε. Let $\widetilde{x}$ and x' denote the vectors of opinions along these two trajectories at the beginning of the mth large block. What is the difference in behavior between the two sequences of vectors of opinions restricted to K^d on this large block? It will depend on two properties: does any player target the communicating class K^d in block m and did any player target the communicating class K^d in block $m-1$(the previous one, before reaching $\widetilde{x}$ and x'). We distinguish three cases:

- First, if the class is influenced by any player in block m, then the original vectors of opinions $\widetilde{x}$ and x' only impact for a fraction $1-\theta(d)$ the opinions at the beginning of the next block. Hence, before the projection, the two trajectories got strictly closer due to external influence. Then, an additional margin of error needs to be taken into account due to the projection that is only carried out along one of the trajectories (in game $\widetilde{\Gamma}_\varepsilon$).
- Second, if the class is not influenced by one of the players in block m and in block $m-1$, then the two trajectories stay at the same distance and no margin of error needs to be added when projecting. Indeed, by definition, the restriction of $\widetilde{x}$ to K^d is a stable vector of opinions for M. Since the opinions propagate according to M, this vector of opinions do not change and therefore the projection does not introduce an error.
- Lastly, if the class is not under influence in block m, but was under influence in block $m-1$, then we only know that the trajectories are not moving apart and that we need to add a margin of error associated with the projection.

Clearly, the second case is nice since there is no error. The first is also a nice case due to the contraction property. Assume that there are only blocks of this type. Then on each block, the error accumulated up to stage t is shrunk by $1-\theta(d)$ and a margin of error ε is added. Summing over all stages, we see that the total error is, therefore, bounded from above by $\varepsilon/\theta(d) < +\infty$. The third case is the problematic one since several blocks of this type could yield an accumulation of errors that diverges. Indeed, we need to add a margin of ε on each block. We notice that it is impossible to have only blocks of this form: between two blocks of the third type, at least one block of the first type is needed. Hence, the distance between the two trajectories is bounded above by a converging series and it can then be controlled.

Recall, that $K(a,b)$ is a union of communicating classes. Therefore, given a communicating absorbing class K^d and a pair of actions (a,b), we only have two cases: either $K^d \subseteq K(a,b)$ or $K^d \cap K(a,b) = \varnothing$. We first prove the following lemma, in which we focus on the transition q' on one block.

Lemma 5.3 *Let K^d be a communicating class. There exists $\theta(d) \in [0,1[$ such that, for every $a \in A$ and $b \in B$ such that $K^d \subseteq K(a,b)$, we have*

$$\forall x, y \in X,\ \|q'(x,a,b) - q'(y,a,b)\|_{K^d,\infty} \leq \theta(d)\|x-y\|_{K^d,\infty}.$$

Proof. Given a sequence of n pairs of actions, $(i_1, j_1, \ldots, i_n, j_n) \in H_n$, we define

$$\overleftarrow{\prod_{l=1}^{n}} B(i_l, j_l) := B(i_n, j_n)B(i_{n-1}, j_{n-1})...B(i_2, j_2)B(i_1, j_1).$$

With this notation, one can rewrite the transition $q'(x,a,b)$ as follows

$$\begin{pmatrix} 1 \\ 0 \\ q'(x,a,b) \end{pmatrix} = \begin{pmatrix} 1 & 0 & \mathbf{0} \\ 0 & 1 & \mathbf{0} \\ A_1(a,b) & A_0(a,b) & B(a,b) \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ x \end{pmatrix}$$

where

$$B(a,b) = \overleftarrow{\prod_{t=1}^{t(a,b)SU}} B(i_t, j_t),$$

$$A_1(a,b) = \sum_{u=1}^{t(a,b)SU} \left(\overleftarrow{\prod_{t=u+1}^{t(a,b)SU}} B(i_t, j_t) \right) A_1(i_u, j_u),$$

and

$$A_0(a,b) = \sum_{u=1}^{t(a,b)SU} \left(\overleftarrow{\prod_{t=u+1}^{t(a,b)SU}} B(i_t, j_t) \right) A_0(i_u, j_u).$$

Let k be in K^d and therefore in $K(a,b)$. By definition of $K(a,b)$, there exists $l \in K$ and $u \geq 1$ such that:

(i) $A_1(i_u, j_u)_l > 0$ or $A_0(i_u, j_u)_l > 0$,

(ii) there exists a sequence of nodes $(k_t)_{t=u+1,\dots,t(a,b)SU+1}$ such that $k_u = l$, $k_{t(a,b)SU+1} = k$ and for all $t \in \{u+1, \dots, t(a,b)SU\}$, $M(k_{t+1}, k_t) > 0$,

(iii) $u \leq (t(a,b)-1)SU$.

Let us consider the shortest sequence satisfying (i) and (ii), and denote the parameters by l^*, u^* and $(k_t^*)_{t=u^*+1,\dots,t(a,b)SU+1}$, respectively. We also assume without loss of generality that $A_1(i_{u^*}, j_{u^*})_k > 0$.

Since it is the shortest sequence satisfying (i) and (ii), it follows that, for $t \in \{u^*+1, \dots, t(a,b)SU+1\}$, $A_1(i_t, j_t)_{k_t^*} = 0$ and $A_2(i_t, j_t)_{k_t^*} = 0$, otherwise u^* could be replaced by a larger natural number contradicting the minimality. We deduce that

$$\forall t \in \{u^*+1, \dots, t(a,b)SU\},\ B(i_t, j_t)_{k_t^* k_{t+1}^*} = M_{k_{t+1}^*, k_t^*} > 0,$$

and therefore

$$\left(\overleftarrow{\prod_{t=u^*+1}^{t(a,b)SU}} B(i_t, j_t) \right)_{kl^*} > 0.$$

Finally, we have

$$A_1(a,b)_k > \left(\overleftarrow{\prod_{t=u^*+1}^{t(a,b)SU}} B(i_t, j_t) \right) A_1(i_{u^*}, j_{u^*})_{l^*} > 0.$$

Hence, if $K^d \subseteq K(a,b)$, for every $k \in K^d$, the kth coordinate of either $A_1(a,b)$ or $A_0(a,b)$ is strictly positive. Since $M(a,b)$ is row-stochastic, being a product of raw-stochastic matrices, there exists $\theta(a,b) \in (0,1)$ such that, for every $k \in K^d$,

$$0 \leq \sum_{l \in K^d} B(a,b)_{kl} \leq \theta(a,b).$$

Let $\theta(d) = \max_{(a,b) \in A \times B,\ \text{s.t. } K^d \subseteq K(a,b)} \theta(a,b)$. Since $A \times B$ is finite and not empty, $\theta(d)$ is well defined and $\theta(d) < 1$.

Since the class is supposed to be communicating and absorbing, which implies

$$\forall k \in K^d,\ \forall l \notin K^d, \forall n \in \mathbb{N},\ (M^n)_{kl} = 0,$$

it follows that, for all $x, y \in X_f$ and for all $k \in K^d$,

$$|q'(x,a,b)(k) - q'(y,a,b)(k)| = \left| \sum_{l \in K^d} B(a,b)_{kl} x(l) - \sum_{l \in K^d} B(a,b)_{kl} y(l) \right|,$$
$$\leq \left| \sum_{l \in K^d} B(a,b)_{kl} \right| \|x - y\|_{K^d,\infty}.$$

Hence,

$$\|q'(x,a,b) - q'(y,a,b)\|_{K^d,\infty} = \|q'(x,a,b) - q'(y,a,b)\|_{K^d,\infty}, \tag{5.3}$$
$$\leq \theta(a,b)\|x - y\|_{K^d,\infty}, \tag{5.4}$$
$$\leq \theta(d)\|x - y\|_{K^d,\infty}. \tag{5.5}$$

The players have a strictly positive influence on the class K^d and therefore the original vector of opinions weighs strictly less than 1 in the final opinions.

■

We can then deduce the following lemma which formalizes our introductory discussion on transitions. We extend both $q'$and $\widetilde{q}$ to finite sequences of pairs of actions similarly to the extension of q. For clarity, we also simplify the notation $q(x,(a_1,b_1,a_2,b_2))$ to $q(x,a_1,b_1,a_2,b_2)$.

Lemma 5.4 *For $d \in \{1,...,D\}$, $x, y \in X_f$, $(a_1,b_1) \in A \times B$, and $(a_2,b_2) \in A \times B$, there exists $\theta(d)$ such that:*

- *if $K^d \subseteq K(a_2,b_2)$, then*

$$\|q'(x,a_1,b_1,a_2,b_2) - \widetilde{q}(y,a_1,b_1,a_2,b_2)\|_{K^d,\infty} \leq \theta(d)\|q'(x,a_1,b_1) - \widetilde{q}(y,a_1,b_1)\|_{K^d,\infty} + \varepsilon,$$

- *if $K^d \cap K(a_2,b_2) = \varnothing$ and $K^d \cap K(a_2,b_2) = \varnothing$, then*

$$\|q'(x,a_1,b_1,a_2,b_2) - \widetilde{q}(y,a_1,b_1,a_2,b_2)\|_{K^d,\infty} \leq \|q'(x,a_1,b_1) - \widetilde{q}(y,a_1,b_1)\|_{K^d,\infty},$$

- *if $K^d \cap K(a_2,b_2) = \varnothing$ and $K^d \subseteq K(a_1,b_1)$, then*

$$\|q'(x,a_1,b_1,a_2,b_2) - \widetilde{q}(y,a_1,b_1,a_2,b_2)\|_{K^d,\infty} \leq \|q'(x,a_1,b_1) - \widetilde{q}(y,a_1,b_1)\|_{K^d,\infty} + 2\varepsilon.$$

Proof.

We first assume that $K^d \subset K(a_2,b_2)$. By Lemma 5.3 applied to $q'(x,a_1,b_1)$, $\widetilde{q}(y,a_1,b_1)$, and the pair of actions (a_2,b_2), we know that

$$\|q'(x,a_1,b_1,a_2,b_2) - q'(\widetilde{q}(y,a_1,b_1),a_2,b_2)\|_{K^d,\infty} = \|q'(q'(x,a_1,b_1),a_2,b_2) - q'(\widetilde{q}(y,a_1,b_1),a_2,b_2)\|_{K^d,\infty}, \tag{5.6}$$
$$\leq \theta(d)\|q'(x,a_1,b_1) - \widetilde{q}(y,a_1,b_1)\|_{K^d,\infty}. \tag{5.7}$$

The players have a strictly positive influence on the class K^d and therefore the original vector of opinions weights strictly less than 1 in the final opinions. Moreover, we know that

$$\widetilde{q}(y,a_1,b_1,a_2,b_2) = \Pi_{K(a_2,b_2)}(q'(\widetilde{q}(y,a_1,b_1),a_2,b_2)),$$

so, by projecting, an error ε is made:

$$\|\widetilde{q}(y,a_1,b_1,a_2,b_2) - q'(\widetilde{q}(y,a_1,b_1),a_2,b_2)\|_{K^d,\infty} \leq \varepsilon. \tag{5.8}$$

Inequalities 5.7 and 5.8 imply the first result by the triangle inequality:

$$\|q'(x,a_1,b_1,a_2,b_2) - \widetilde{q}(y,a_1,b_1,a_2,b_2)\|_{K^d,\infty} \leq \theta(d)\|q'(x,a_1,b_1) - \widetilde{q}(y,a_1,b_1)\|_{K^d,\infty} + \varepsilon.$$

Secondly, we assume that $K^d \cap K(a_2,b_2) = \varnothing$ and $K^d \cap K(a_1,b_1) = \varnothing$. After playing (a_1,b_1), one projects onto the stable vectors of opinions, and the vectors of opinions evolve according to M. It follows that, on the K^d-coordinates, we have

$$q'(\widetilde{q}(y,a_1,b_1),a_2,b_2)(k) = \widetilde{q}(y,a_1,b_1)(k) = \widetilde{q}(y,a_1,b_1,a_2,b_2)(k).$$

Therefore,

$$\begin{aligned}\|q'(x,a_1,b_1,a_2,b_2) - \widetilde{q}(y,a_1,b_1,a_2,b_2)\|_{K^d,\infty} &= \|q'(x,a_1,b_1,a_2,b_2) - q'(\widetilde{q}(y,a_1,b_1),a_2,b_2)\|_{K^d,\infty},\\ &\leq \|q'(x,a_1,b_1) - \widetilde{q}(y,a_1,b_1)\|_{K^d,\infty}.\end{aligned}$$

This is the second result.

Finally, we assume that $K^d \cap K(a_2,b_2) = \varnothing$ and $K^d \subset K(a_1,b_1)$. It follows that, restricting to the coordinates in K^d,

$$\forall k \in K^d,\ q'(\widetilde{q}(y,a_1,b_1),a_2,b_2)(k) = e_k M^{t(a_2,b_2)SU}\widetilde{q}(y,a_1,b_1),$$

where e_k is the unitary vector corresponding to the kth coordinate. Informally, the opinions evolve according to M^{SU} without influence from the actions played by the players. By assumption $t(a,b) \geq N \geq R$, we thus know that $\|(M'')^{t(a_2,b_2)} - B_0\|_\infty \leq \varepsilon$, and the distance between a state y and its projection is less than ε. Therefore, we have

$$\|\widetilde{q}(y,a_1,b_1,a_2,b_2) - q'(\widetilde{q}(y,a_1,b_1),a_2,b_2)\|_{K^d,\infty} \leq 2\varepsilon.$$

Moreover, the mapping q' is 1-Lipschitz and one can isolate what is happening inside K^d, since K^d is a communicating absorbing class. Hence,

$$\|q'(x,a_1,b_1,a_2,b_2) - q'(\widetilde{q}(y,a_1,b_1),a_2,b_2)\|_{K^d,\infty} \leq \|q'(x,a_1,b_1) - \widetilde{q}(y,a_1,b_1)\|_{K^d,\infty}.$$

By the triangle inequality, we obtain

$$\|q'(x,a_1,b_1,a_2,b_2) - \widetilde{q}(y,a_1,b_1,a_2,b_2)\|_{K^d,\infty} \leq \|q'(x,a_1,b_1) - \widetilde{q}(y,a_1,b_1)\|_{K^d,\infty} + 2\varepsilon.$$

■

It follows that we can control the distance between the two transitions on absorbing communicating classes.

Lemma 5.5 *There exists $\theta < 1$ such that, for every $x_1 \in X$, for every $h' \in H'$, and for all $n \geq 1$,*

$$\forall d \in \{1,...,D\},\ \|q'(x_1,h') - \widetilde{q}(x_1,h')\|_{K^d,\infty} \leq \frac{3\varepsilon}{1-\theta},$$

and therefore,

$$\|q'(x_1,h') - \widetilde{q}(x_1,h')\|_{K\setminus K^0,\infty} \leq \frac{3\varepsilon}{1-\theta}.$$

Proof. Let $\theta = \max_{d\in\{1,...,D\}}\theta(d)$, let $h' = (a_1,b_1,\ldots,a_n,b_n) \in H'$, and let $d \in \{1,...,D\}$. We define the sequence of stages $(n_s)_{s\geq 1}$ as follows:

- $n_1 = 1$,
- for all $s \geq 1$, $n_{s+1} = \inf\{n \geq n_s + 1, K^d \subset K(a_n,b_n)\}$.

For every $s \geq 1$ and $n \in (n_s, n_{s+1})$, $K^d \cap K(a_n,b_n) = \varnothing$. We now use Lemma 5.4. We apply the first case once, then the third case $n_{s+1} - n_s - 2$ times, and finally the second case once:

$$\begin{aligned}\|q'(x_1,h'_{n_{s+1}}) - \widetilde{q}(x_1,h'_{n_{s+1}})\|_{K^d,\infty} &\leq \theta\|q'(x_1,h'_{n_{s+1}-1}) - \widetilde{q}(x_1,h'_{n_{s+1}-1})\|_{K^d,\infty} + \varepsilon,\\ &\leq \theta\|q'(x_1,h'_{n_s+1}) - \widetilde{q}(x_1,h'_{n_s+1})\|_{K^d,\infty} + \varepsilon,\\ &\leq \theta\|q'(x_1,h'_{n_s}) - \widetilde{q}(x_1,h'_{n_s})\|_{K^d,\infty} + \theta 2\varepsilon + \varepsilon,\\ &\leq \theta\|q'(x_1,h'_{n_s}) - \widetilde{q}(x_1,h'_{n_s})\|_{K^d,\infty} + 3\varepsilon.\end{aligned}$$

By induction, we deduce that, for every $s \geq 1$,

$$\|q'(x_1, h'_{n_s}) - \widetilde{q}(x_1, h'_{n_s})\|_{K^d,\infty} \leq \frac{1-\theta^s}{1-\theta} 3\varepsilon.$$

Since the transition is 1-Lipchitz, we obtain an upper bound at every stage:

$$\forall n \geq 1, \|q'(x_1, h'_n) - \widetilde{q}(x_1, h'_n)\|_{K^d,\infty} \leq \frac{3\varepsilon}{1-\theta}.$$

■

We now use the previous result to identify a bound on the transient states for one-stage transitions.

Lemma 5.6 *There exists $\eta \in (0,1)$ such that, for every $x, y \in X_f$,*

$$\|q'(x, a_1, b_1) - \widetilde{q}(y, a_1, b_1)\|_{K^0,\infty} \leq \eta \|x - y\|_{K^0,\infty} + \|x - y\|_{K \setminus K^0,\infty} + \varepsilon.$$

Proof. By definition of S and the transience of the states in K^0, there exists η such that

$$\|q'(x, a_1, b_1) - q'(y, a_1, b_1)\|_{K^0,\infty} \leq \eta \|x - y\|_{K^0,\infty} + \|x - y\|_{K \setminus K^0,\infty}.$$

By the discretization that we have chosen, we have

$$\|q'(y, a_1, b_1) - \widetilde{q}(y, a_1, b_1)\|_{K_0,\infty} \leq \varepsilon.$$

The triangle inequality yields the result. ■

We can now establish the bound along a play during the whole game.

Lemma 5.7 *Let $x_1 \in X_f$ and $h' \in H'_f$ then*

$$\|q'(x_1, h') - \widetilde{q}(x_1, h')\|_\infty \leq \xi(\varepsilon),$$

with $\xi(\varepsilon) = \frac{4\varepsilon}{(1-\eta)(1-\theta)}$.

Proof. Let $n \geq 1$ and $h' = (a_1, b_1, ..., a_n, b_n) \in H'_n$ and $(a_{n+1}, b_{n+1}) \in (A \times B)$. By Lemma 5.6, we know that

$$\begin{aligned} &\|q'(x_1, (h', a_{n+1}, b_{n+1})) - \widetilde{q}(x_1, (h', a_{n+1}, b_{n+1}))\|_{K_0,\infty} \\ \leq \eta &\|q'(x_1, h') - \widetilde{q}(x_1, h')\|_{K^0,\infty} + \|q'(x_1, h') - \widetilde{q}(x_1, h')\|_{K \setminus K^0,\infty} + \varepsilon. \end{aligned}$$

Hence, by Lemma 5.5, we obtain

$$\begin{aligned} \|q'(x_1, (h', a_{n+1}, b_{n+1})) - \widetilde{q}(x_1, (h', a_{n+1}, b_{n+1}))\|_{K_0,\infty} &\leq \eta \|q'(x_1, h') - \widetilde{q}(x_1, h')\|_{K^0,\infty} + \frac{3\varepsilon}{1-\theta} + \varepsilon, \\ &\leq \eta \|q'(x_1, h') - \widetilde{q}(x_1, h')\|_{K^0,\infty} + \frac{4\varepsilon}{1-\theta}. \end{aligned}$$

Similarly to the proof of Lemma 5.5, we prove by induction that, for every $h' \in H'$,

$$\|q'(x_1, h') - \widetilde{q}(x_1, h')\|_{K_0,\infty} \leq \frac{1}{1-\eta} \frac{4\varepsilon}{1-\theta}.$$

Since $\eta < 1$, this is also an upper bound outside of K^0 and we obtain the bound on the whole space. ■

In Lemma 5.7, we proved that the trajectories obtained from the same starting point by the same play in the game Γ' and $\widetilde{\Gamma}$ remain at a maximum distance of $\xi(\varepsilon)$. Since, for every $(i, j) \in I \times J$, the payoff functions $g^1(., i, j)$ and $g^2(., i, j)$ are continuous and the space X is compact, each function is uniformly

continuous on X. Moreover, since there are a finite number of such functions, they form a uniformly equicontinuous family: there exists a function ω from $\mathbb{R}_+$ to $\mathbb{R}_+$ continuous at 0 such that $\omega(0) = 0$ and

$$\forall (i,j) \in I \times J,\ \forall x, x' \in X,\ |g^1(x,i,j) - g^1(x',i,j)| \leq \omega(\|x - x'\|_\infty)$$

and

$$\forall (i,j) \in I \times J,\ \forall x, x' \in X,\ |g^2(x,i,j) - g^2(x',i,j)| \leq \omega(\|x - x'\|_\infty).$$

We now have the following proposition.

Proposition 5.8 *Let $x_1 \in X$, $n \geq 1$, $\sigma' \in \Sigma'$, and $\tau' \in \mathcal{T}'$. Then we have*

$$|\gamma'^m_n(x_1,\sigma',\tau') - \widetilde{\gamma}^m_n(x_1,\sigma',\tau')| \leq \omega(\xi(\varepsilon)),\ \textit{for all } m \in \{1,2\}.$$

5.4.2 Proximity of the time-average payoffs

We now show that the average payoffs in the game Γ'_ε (without projection) and in the game Γ_ε are close. There are two problems. First, there is a normalization problem since, in the game Γ'_ε, we consider blocks of different sizes. Hence, when considering the average, the weight of the stages are no longer uniform. Second, the set of strategies is different so we need to relate strategies in the two games. We first focus on the averaging problem and then on the strategies.

Lemma 5.9 *Let $x_1 \in X$ and $h' \in H_f$ be a sequence of pairs of actions in the auxiliary game Γ'_ε. Let h be the sequence of actions associated with h' in the original game. We have*

$$|\gamma'^m_n(x_1,h') - \gamma^m_{nNSU}(x_1,h)| \leq 2\varepsilon,\ \textit{for } m \in \{1,2\}.$$

Proof. Let $x \in X$ and let $h' = (a_1,b_,...,a_n,b_n)$ be a sequence of pair of actions in the auxiliary game Γ'_ε and $m \in \{1,2\}$. They generate:

- a finite sequence $h' = (x'_1,a_1,b_1,...,x'_n,a_n,b_n) \in (X \times A \times B)^n$ in Γ'_ε,
- a finite sequence $h = (x_1,i_1,j_1,...,x_L,i_L,j_L)$ in $(X \times I \times J)^L$ in Γ_ε, where $L = \sum_{t=1}^n t(a_t,b_t) \in [nNSU, n(N+T)SU]$.

Note that $(x'_1,...,x'_n)$ is a subsequence of $(x_1,...,x_L)$. By definition of the payoffs in Γ'_ε (with a slight abuse of notation in the second line), we have

$$\begin{aligned}
(\gamma')^m_n(x_1,h') &= \frac{1}{n}\sum_{u=1}^n (g')^m(x'_u,a_u,b_u),\\
&= \frac{1}{n}\sum_{u=1}^n \left(\sum_{t'=1}^{t(a_u,b_u)SU} \frac{1}{t(a_u,b_u)SU} g^m(x^u_{t'},i^u_{t'},j^u_{t'}) \right),\\
&= \sum_{t=1}^L \mu(t) g^m(x_t,i_t,j_t),
\end{aligned}$$

where $\mu(t) = \frac{1}{nSUt(a_u,b_u)}$ for some $u \in \{N,...,N+T\}$. Hence, for every $t \leq R$, we have $\mu(t) \in \left[\frac{1}{n(N+T)SU}, \frac{1}{nNSU}\right]$. Therefore, for every $t \in [1,L]$, we have

$$\left|\mu(t) - \frac{1}{L}\right| \leq \left|\frac{1}{nNSU} - \frac{1}{n(N+T)SU}\right| \leq \left|\frac{T}{nN(N+T)SU}\right|$$

and

$$\begin{aligned}\left|(\gamma')_n^m(x_1,h') - \frac{1}{L}\sum_{t=1}^{L} g^m(x_t,i_t,j_t)\right| &\leq \sum_{t=1}^{L}\left|\mu(t)-\frac{1}{L}\right|,\\ &\leq L\frac{T}{nN(N+T)SU},\\ &\leq \frac{T}{N}\frac{L}{n(N+T)SU},\\ &\leq \frac{T}{N},\\ &\leq \varepsilon,\end{aligned}$$

by the choice of T and N. In the previous computation, L depends on the sequence of actions. In order to obtain the result with $nNSU$, we use the standard result

$$|\gamma_{nNSU}^m(x_1,h) - \gamma_L^m(x_1,h)| \leq \frac{|L-nNSU|}{n(N+T)SU} = \frac{nTSU}{nNSU} = \frac{T}{N} \leq \varepsilon,$$

and the triangle inequality. ■

Proposition 5.10 *Let $x_1 \in X$ and $n \geq 1$.*

- *If $\sigma' \in \Sigma'$, there exists $\sigma \in \Sigma$ such that*

$$\forall \tau \in \mathcal{T},\ \exists \tau' \in \mathcal{T}',\ \text{such that } |(\gamma^1)'_n(x_1,\sigma',\tau') - \gamma^1_{nNSU}(x_1,\sigma,\tau)| \leq 2\varepsilon.$$

- *If $\tau' \in \mathcal{T}'$, there exists $\tau \in \mathcal{T}$ such that*

$$\forall \sigma \in \Sigma,\ \exists \sigma' \in \Sigma',\ \text{such that } |(\gamma^2)'_n(x_1,\sigma',\tau') - \gamma^2_{nNSU}(x_1,\sigma,\tau)| \leq 2\varepsilon.$$

Proof. We only prove the result for Player 1. The proof is similar for Player 2. In Lemma 5.9, we compared the payoffs along a sequence of pairs of actions. We now extend the result to strategies. By Kuhn's theorem, a strategy is equivalent to a mixed distribution over pure strategies, so it suffices to establish the result for pure strategies of the second player: for all $\sigma' \in \Sigma'$, there exists $\sigma \in \Sigma$ such that, for every pure strategy τ of Player 2, there exists τ', a pure strategy of Player 2, such that

$$|(\gamma^1)'_n(x,\sigma',\tau') - \gamma^1_{nNSU}(x,\sigma,\tau)| \leq 2\varepsilon.$$

Fix σ', a strategy of Player 1 in the auxiliary game. The definition of σ is the following:

- Start by following $a_1 = \sigma'(\varnothing)$.
- During the game, compute the number $t(h)$ of the current history up to some stage n such that $t(h)$ is finite. We have shown that it happens before stage $(N+T)SU$. Let t_1 be this stage.
- Construct a fictitious action b_1 of Player 2 that chooses to play, at stage $t \leq t_1$, the action j_t that has been observed for any partial history and plays some fixed actions between t_1 and $(N+T)SU$.
- It follows that $a_2 = \sigma'(a_1,b_1)$.
- Construct a fictitious action b_2, and so on.

By construction, we do indeed have the same distribution over actions in both games. By Lemma 5.9, it follows that

$$|(\gamma^1)'_n(x_1,\sigma',\tau') - \gamma^1_{nNSU}(x_1,\sigma,\tau)| \leq 2\varepsilon.$$

■

5.5 Conclusion

We now conclude the proof by constructing a uniform Nash equilibrium payoff in the game $\Gamma(x_1)$.

First, the game $\widetilde{\Gamma}_\varepsilon(x_1)$ is equivalent to a game with finite sets of actions and a finite set of states, $\{x_1\} \cup X_f$. Applying the result due to Vieille [22, 23], there exists at least one uniform Nash equilibrium payoff which we denote by $\gamma_\varepsilon(x_1)$. Moreover, each player can guarantee the equilibrium payoff with a strategy depending only on the states and the payoffs. We consider an accumulation point of the family $(\gamma_\varepsilon(x_1))_{\varepsilon>0}$ when ε goes to 0, denoted by $\gamma_*(x_1)$. Let us prove that $\gamma_*(x_1)$ is a uniform Nash-equilibrium payoff of $\Gamma(x_1)$, i.e., for every $\eta > 0$, there exists (σ_*, τ_*) and $N \geq 1$ such that, for every $n \geq N$,

$$\forall \tau \in \mathcal{T},\ \gamma_n^2(x_1, \sigma_*, \tau) \leq \gamma_*^2(x_1) + 6\eta,$$
$$\forall \sigma \in \Sigma,\ \gamma_n^1(x_1, \sigma, \tau_*) \leq \gamma_*^1(x_1) + 6\eta,$$

and

$$\|\gamma_n(x_1, \sigma_*, \tau_*) - \gamma_*(x_1)\|_\infty \leq 6\eta.$$

Let $\eta > 0$. By definition of ω and ξ, there exists ε_0 such that, for every $\varepsilon \leq \varepsilon_0$, $\omega(\xi(\varepsilon)) \leq \eta$. Fix $\varepsilon > 0$ such that $\varepsilon \leq \min(\eta, \varepsilon_0)$ and

$$\|\gamma_*(x_1) - \gamma_\varepsilon(x_1)\| \leq \eta. \tag{5.9}$$

By construction, γ_ε is a uniform Nash-equilibrium payoff in the auxiliary game $\widetilde{\Gamma}_\varepsilon(x_1)$. Hence, there exists a pair of strategies $(\sigma'_*, \tau'_*) \in \Sigma' \times \mathcal{T}'$ and $L \geq 1$ such that, for every $l \geq L$,

$$\forall \tau' \in \mathcal{T}',\ \widetilde{\gamma}_l^2(x_1, \sigma'_*, \tau') \leq \gamma_\varepsilon^2(x_1) + \eta,$$
$$\forall \sigma \in \Sigma',\ \widetilde{\gamma}_l^1(x_1, \sigma', \tau'_*) \leq \gamma_\varepsilon^1(x_1) + \eta,$$

and

$$\|\widetilde{\gamma}_l(x_1, \sigma'_*, \tau'_*) - \gamma_\varepsilon(x_1)\|_\infty \leq \eta.$$

By Proposition 5.10, we know that there exists σ_* such that, for all $l \geq 1$,

$$\forall \tau \in \mathcal{T},\ \exists \tau' \in \mathcal{T}', \forall m \in \{1, 2\},\ |(\gamma^m)'_l(x_1, \sigma'_*, \tau') - \gamma^m_{lNSU}(x_1, \sigma_*, \tau)| \leq 2\varepsilon,$$

and, by a similar result for Player 2, there exists τ_* such that, for all $l \geq 1$,

$$\forall \sigma \in \Sigma,\ \exists \sigma' \in \Sigma', \forall m \in \{1, 2\},\ |\gamma'_l(x_1, \sigma', \tau'_*) - \gamma_{lNSU}(x_1, \sigma, \tau_*)| \leq 2\varepsilon.$$

Let us focus on the deviation of Player 2. If $\tau \in \mathcal{T}$ and $n \geq N_0 \coloneqq \max(\frac{NSU}{\eta}, LNSU)$, then

$$\begin{aligned}\gamma_n^2(x_1, \sigma_*, \tau) &\leq \gamma^2_{\lfloor \frac{n}{NSU} \rfloor NSU}(x_1, \sigma, \tau) + \frac{NSU}{n}, \\ &\leq (\gamma^2)'_{\lfloor \frac{n}{NSU} \rfloor}(x_1, \sigma'_*, \tau') + \frac{NSU}{n} + 2\varepsilon.\end{aligned}$$

By Proposition 5.8 and the definition of (σ'_*, τ'_*), we obtain

$$\begin{aligned}\gamma_n^2(x_1, \sigma_*, \tau) &\leq \widetilde{\gamma}^2_{\lfloor \frac{n}{NSU} \rfloor}(x_1, \sigma'_*, \tau')| + \frac{NSU}{n} + 2\varepsilon + \omega(\xi(\varepsilon)), \\ &\leq \gamma_\varepsilon^2(x_1) + \frac{NSU}{n} + 2\varepsilon + \omega(\xi(\varepsilon)) + \eta.\end{aligned}$$

Finally, using Equation 5.9, we have

$$\begin{aligned}\gamma_n^2(x_1, \sigma_*, \tau) &\leq \gamma_*^2(x_1) + \frac{NSU}{n} + 2\varepsilon + \omega(\xi(\varepsilon)) + 2\eta, \\ &\leq \gamma_*^2(x_1) + 6\eta.\end{aligned}$$

We have obtained the first result. The second result can be obtained by exchanging the roles of Players 1 and 2. To obtain the last inequality, we have once again the same computation in which we invert the inequalities. This concludes the proof.

6 Proof of Proposition 3.6

First, let us introduce an example where the sequence of discounted Nash equilibrium sets converges to a set disjoint from the set of uniform Nash equilibrium. It is based on a non-zero-sum version of the Big Match introduced by Sorin [18] to prove the corresponding result for stochastic games with a finite set of states and finite sets of actions.

	L	R
T	$(1,0)^*$	$(0,2)^*$
B	$(0,1)$	$(1,0)$

k_1

This stochastic game is an absorbing game. There are three states: two absorbing states represented by the stars and one non-absorbing state k_1 represented by the preceding matrix. The game is interpreted as follows. When Player 1 plays B, he obtains a bad payoff against R and a good payoff against L. Then, the two players play again. If Player 1 plays T at stage t, the game is absorbed and the payoff is fixed forever depending on the action of Player 2 at stage t. The vector of payoffs is $(1,0)$ if Player 2 was playing L and $(0,2)$ if Player 2 was playing R. Sorin [18] showed that, for every $\lambda \in (0,1]$, $E_\lambda(k_1) = \{(1/2, 2/3)\}$, whereas $E_\infty(k_1) = \{(a, 2(1-a); 1/2 \leq a \leq 2.3\}$. Therefore the sequence of discounted Nash equilibrium payoffs converges to a set disjoint from $E_\infty(k_1)$.

We introduce a variant of this game where we slightly change the absorbing states to make it easier to encode into a dynamic opinion game. The arrows indicate the deterministic transition corresponding to the pair of actions played. When the state does not change, we omit the arrow.

	L	R
T	$(1,0)_{\to k_2}$	$(0,2)_{\to k_3}$
B	$(0,1)$	$(1,0)$

k_1

	L	R
T	$(1,0)$	$(1,0)_{\to k_4}$
B	$(1,0)$	$(1,0)$

k_2

	L	R
T	$(0,2)_{\to k_4}$	$(0,2)$
B	$(0,2)$	$(0,2)$

k_3

	L	R
T	$(0,0)$	$(0,0)$
B	$(0,0)$	$(0,0)$

k_4

The states k_2 and k_3 are no longer absorbing. Nevertheless, there exists a unique Nash equilibrium payoff equal to $(1,0)$ from k_2 and a unique Nash equilibrium payoff equal to $(0,2)$ from k_3. Indeed, 1 is the maximal payoff for Player 1 and he can guarantee to stay in k_2 by playing B. The situation is similar for Player 2 in k_3. Hence, the set of Nash equilibrium payoffs is the same in the original game and in the variant.

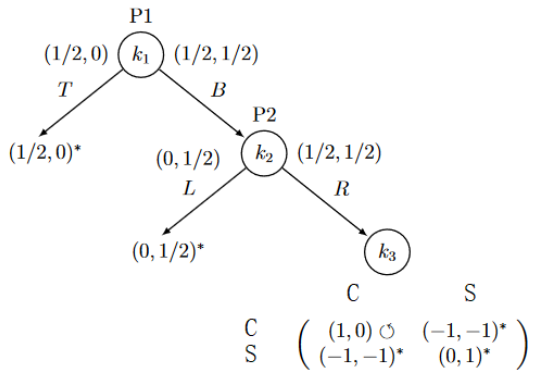
We now introduce a dynamic opinion game with limited lobbies that is equivalent to this game. We consider $K = \{T, B, L, R\}$ and assume that all nodes are isolated. Let $I = \{T, B\}$ and $J = \{L, R\}$. We assume that ρ and μ are given as follows:

- $\rho(.,.,.)$ is constant equal to 0,
- $\mu(B,.,.)$ is identically equal to 0,
- $\forall k \in K, \mu(T, L, k) = \delta_L(k)$ and $\mu(T, R, k) = \delta_R(k)$.

Hence, if Player 1 plays B, the vector of opinions does not change, whereas if Player 1 plays T, then the action played at this stage by player 2 is encoded in the state.

Let us introduce some special vectors of opinions: let $k_1 = (1,1,1,1)$, $k_2 = (1,1,0,1)$, $k_3 = (1,1,1,0)$, and $k_4 = (1,1,0,0)$ and define the payoff function to be equal to the payoffs of the variant stochastic game. It is easy to see that the transitions in the dynamic games are the same as in the variant. This is formally the same game, so we obtain a dynamic opinion game with limited lobbies such that the sequence of the discounted Nash equilibrium set does not converge.

We now introduce an example where the sequence of discounted Nash equilibrium sets does not converge. The example is obtained by encoding a variation of an example introduced by Renault and Ziliotto [16]. The original example is the following.



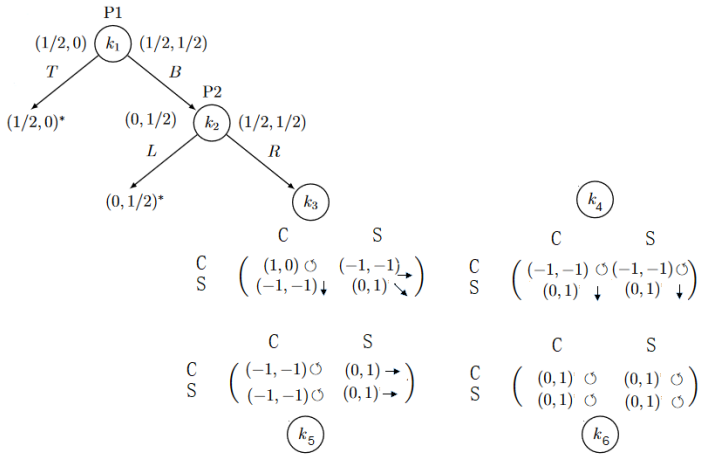
In state k_1, Player 1 is playing alone and has two actions. The first action T gives the vector of payoffs $(1/2, 0)$ and leads to an absorbing state with payoff $(1/2, 0)$. The second action B gives the vector of payoffs $(1/2, 1/2)$ and leads to k_2. In k_2, Player 2 is playing alone and has two actions leading respectively to an absorbing state with payoff $(0, 1/2)$ or to k_3. At k_3, both players are playing. If both players play C, they get $(1, 0)$ and they play again. Otherwise, the game is absorbed in a state depending on the subset of players that played Q and they obtain the corresponding payoff.

Let Λ_1 be the set of discount factors of the form $\lambda = 1 - \left(\frac{1}{2}\right)^{1/N}$ where N is a positive integer and $\Lambda_2 = (0, 1] \setminus \Lambda_1$. The authors show that

- for every $\lambda \in \Lambda_2$, $E_\lambda(k_1) = \{(1/2, 0)\}$,
- for every $\lambda \in \Lambda_1$, $E_\lambda(k_1) = 1/2 \times [0, 1/2]$.

Therefore, the sequence of discounted equilibria sets does not converge. The idea is that each player can guarantee a payoff of $1/2$ before reaching state k_3. Indeed, Player 1 can play T in k_1 and Player 2 can play L in k_2. Therefore, they accept to play respectively B and R only if each of them is expecting a payoff greater than $1/2$, which can only be possible if the cumulative weight on stages where they play (C, C) and the cumulative weight on stages after playing (S, S) are equal to $1/2$. It is only possible if the discount factor is in Δ_1.

Let us introduce a variant of this game denoted by Γ where we change the transitions in state k_3. The transitions are described by the following picture:



The difference with the example of Renault and Ziliotto [16] is the following. If only one of the players plays S, then the state does not go to an absorbing state but moves to an intermediate state k_4 or k_5. Conditionally on reaching k_3, if neither player plays S for the first time at the same stage, there is a loss that makes it impossible to satisfy both players by giving them at least $1/2$. Hence, we still have

- for every $\lambda \in \Lambda_2$, $E_\lambda(k_1) = \{(1/2, 0)\}$,
- for every $\lambda \in \Lambda_1$, $E_\lambda(k_1) = 1/2 \times [0, 1/2]$.

The key idea to translate this game into a dynamic opinion game is that the payoff on each finite history is uniquely determined by the actions that have been played previously by each player. It does not depend on the order or on how many times each action has been played. Hence, it is sufficient to encode this information into a vector of opinions.

We now consider the following dynamic opinion game that replicates this game. The state space is $K = \{T, BC, BS, L, RC, RS\}$ and every node is isolated. Player 1 can target $I = \{T, BC, BS\}$ and Player 2 can target $J = \{L, RC, RS\}$. We assume that ρ and μ are designed such that a player persuades the node that he targets. We will start with the initial vector of opinions $(0, 0, 0, 1, 1, 1)$. Since, the players targets disjoint sets, the opinions are always in $\{0, 1\}$. Moreover, the opinions on I can only increase, whereas opinions on J can only decrease.

The payoff function depends only on the opinions on I and J as follows:

I\J	$(0,1,1)$	$(1,0,1)$	$(1,1,0)$	$(1,0,0)$	$(0,0,1)$	$(0,1,0)$	$(0,0,0)$
$(1,0,0)$	$(1/2,0)$	$(1/2,0)$	$(1/2,0)$	$(1/2,0)$	$(1/2,-10)$	$(1/2,-10)$	$(1/2,-10)$
$(0,1,0)$	$(0,1/2)$	$(1,0)$	$(-1,-1)$	$(-1,-1)$	$(0,-10)$	$(0,-10)$	$(0,-10)$
$(0,0,1)$	$(0,1/2)$	$(-1,-1)$	$(0,1)$	$(0,1)$	$(0,-10)$	$(0,-10)$	$(0,-10)$
$(0,1,1)$	$(0,1/2)$	$(-1,-1)$	$(0,1)$	$(0,1)$	$(0,-10)$	$(0,-10)$	$(0,-10)$
$(1,1,0)$	$(-10,1/2)$	$(-10,0)$	$(-10,0)$	$(-10,0)$	$(-10,-10)$	$(-10,-10)$	$(-10,-10)$
$(1,0,1)$	$(-10,1/2)$	$(-10,0)$	$(-10,0)$	$(-10,0)$	$(-10,-10)$	$(-10,-10)$	$(-10,-10)$
$(1,1,1)$	$(-10,1/2)$	$(-10,0)$	$(-10,0)$	$(-10,0)$	$(-10,-10)$	$(-10,-10)$	$(-10,-10)$

Let us simplify this matrix and introduce some notations to make the analysis easier. First the opinions on I and J evolve independently with some easy constraints. For example, restricted to I, there exists at least one agent with opinion 1, opinions can only increase and only one node at a time. Similarly, restricted to J, there exists at least one agent with opinion 0, opinions can only decrease and only one node at a time. We see immediately that, for each player, there is some bad state (bottom side and left side of the table) that each player can avoid by never playing an action leading to it. Moreover, let us give some names to some configurations of opinions:

- On I: $(1,0,0) = T*$, $(0,1,0) = BC^\infty$, $(0,0,1) = BS^\infty$ and $(0,1,1) = B(S \cup C)^\infty$.
- On J: $(0,1,1) = L*$, $(1,0,1) = RC^\infty$, $(1,1,0) = RS^\infty$ and $(1,0,0) = R(S \cup C)^\infty$.

The interesting states reduce to the following matrix:

	$L*$	RC^∞	RS^∞	$R(S \cup C)^\infty$
$T*$	$(1/2, 0)$	$(1/2, 0)$	$(1/2, 0)$	$(1/2, 0)$
BC^∞	$(0, 1/2)$	$(1, 0)$	$(-1, -1)$	$(-1, -1)$
BS^∞	$(0, 1/2)$	$(-1, -1)$	$(0, 1)$	$(0, 1)$
$B(S \cup C)^\infty$	$(0, 1/2)$	$(-1, -1)$	$(0, 1)$	$(0, 1)$

The names are chosen to relate each vector of opinions to a set of histories in the game Γ. Restricted to I, $(1,0,0)$ represents histories where Player 1 plays T, $(0,1,0)$ represents histories where Player 1 plays B and only C up to now, $(0,0,1)$ represents histories where Player 1 plays B and only S up to now, $(0,1,1)$ represents histories where Player 1 plays B and at least once S and once C up to now. Similar interpretations lead to the corresponding notations for configurations on J.

We observe that the natural domination that we noticed before implies that, if Player 1 (resp. 2) targets T (resp. L), he continues forever. Moreover, we have:

- If Player 1 targets T, then the payoff is $(1/2, 0)$, independently of the choice of the second player.
- Conditionally on Player 1 not targeting T, if Player 2 plays L, the payoff is $(0, 1/2)$.
- If Player 1 does not target T whereas Player 2 does not target L, then the sequence of payoffs has to be of the following form: $(1, 0)$ for some stages, then $(-1, -1)$ for some stages, and finally $(0, 1)$ forever. The change depends on how many players decided to target S.

Hence, we obtain the same payoffs as in the variant stochastic game. The game Γ and the dynamic opinion game have the same sets of discounted Nash equilibrium payoffs. This concludes the proof.

References

[1] Coralio Ballester, Antoni Calvó-Armengol, and Yves Zenou. Who's who in networks. wanted: The key player. *Econometrica*, 74(5):1403–1417, 2006.

[2] Coralio Ballester, Yves Zenou, and Antoni Calvó-Armengol. Delinquent networks. *Journal of the European Economic Association*, 8(1):34–61, 2010.

[3] A. Banerjee, A. Chandrasekhar, E. Duflo, and M. Jackson. Diffusion of microfinance. *Science*, 341, DOI: 10.1126/science.1236498, 2013.

[4] Kostas Bimpikis, Asuman Ozdaglar, and Ercan Yildiz. Competitive targeted advertising over networks. *Operations Research*, 64(3):705–720, 2016.

[5] Phillip Bonacich. Factoring and weighting approaches to status scores and clique identification. *Journal of Mathematical Sociology*, 2(1):113–120, 1972.

[6] Stephen P Borgatti. Identifying sets of key players in a social network. *Computational & Mathematical Organization Theory*, 12(1):21–34, 2006.

[7] Derrick L Cogburn and Fatima K Espinoza-Vasquez. From networked nominee to networked nation: Examining the impact of web 2.0 and social media on political participation and civic engagement in the 2008 obama campaign. *Journal of Political Marketing*, 10(1-2):189–213, 2011.

[8] Morris H DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

[9] Linton C Freeman. A set of measures of centrality based on betweenness. *Sociometry*, pages 35–41, 1977.

[10] Sanjeev Goyal, Hoda Heidari, and Michael Kearns. Competitive contagion in networks. *Games and Economic Behavior*, 2015.

[11] Michel Grabisch, Antoine Mandel, Agnieszka Rusinowska, and Emily Tanimura. Strategic influence in social networks. *Mathematics of Operations Research*, forthcoming, 2017.

[12] Jason Hartline, Vahab Mirrokni, and Mukund Sundararajan. Optimal marketing strategies over social networks. In *Proceedings of the 17th international conference on World Wide Web*, pages 189–198. ACM, 2008.

[13] C. R. Lever. *Strategic Competitions over Networks*. Ph.D. dissertation, Stanford University, 2010.

[14] Antoine Mandel and Xavier Venel. Dynamic competition over social networks. *European Journal of Operational Research*, 2019.

[15] J-F Mertens and Abraham Neyman. Stochastic games. *International Journal of Game Theory*, 10(2):53–66, 1981.

[16] Jérôme Renault and Bruno Ziliotto. Limit equilibrium payoffs in stochastic games. *Mathematics of Operations Research*, 2020.

[17] Dinah Rosenberg, Eilon Solan, and Nicolas Vieille. Blackwell optimality in markov decision processes with partial observation. *Annals of statistics*, pages 1178–1193, 2002.

[18] S. Sorin. Asymptotic properties of a non-zero sum stochastic game. *International Journal of Game Theory*, 15(2):101–107, Jun 1986.

[19] Thomas W Valente. Network interventions. *Science*, 337(6090):49–53, 2012.

[20] Xavier Venel. Commutative stochastic games. *Mathematics of Operations Research*, 40(2):403–428, 2014.

[21] Xavier Venel and Bruno Ziliotto. Strong uniform value in gambling houses and partially observable markov decision processes. *SIAM Journal on Control and Optimization*, 54(4):1983–2008, 2016.

[22] Nicolas Vieille. Two-player stochastic games i: A reduction. *Israel Journal of Mathematics*, 119(1):55–91, 2000.

[23] Nicolas Vieille. Two-player stochastic games ii: The case of recursive games. *Israel Journal of Mathematics*, 119(1):93–126, 2000.

[24] B. Ziliotto. Zero-sum repeated games: Counterexamples to the existence of the asymptotic value and the conjecture $\max\min = \lim v_n$. *The Annals of Probability*, 44(2):1107–1133, 2016.