

Zero-sum Revision Games*

Fabien Gensbittel[†], Stefano Lovo[‡], Jérôme Renault[§] and Tristan Tomala[¶]

September 1, 2016

Abstract

In a zero-sum asynchronous revision game, players can revise their actions only at exogenous random times. Players' revision times follow Poisson processes, independent across players. Payoffs are obtained only at the deadline by implementing the last prepared actions in the 'component game'. The value of this game is called *revision value*. We characterize it as the unique solution of an ordinary differential equation and show it is continuous in all parameters of the model. We show that, as the duration of the game increases, the *limit revision value* does not depend on the initial position and is included between the min-max and max-min of the component game. We fully characterize the equilibrium in 2×2 games. When the component game minmax and maxmin differ, the revision game equilibrium paths have a wait-and-wrestle structure: far from the deadline, players stay put at *sur-place* action profile, close to the deadline, they take best responses to the action of the opponent.

Key-words: Revision Games, Zero-sum Games, Deadline Effect.

*This work was supported by project PEPS MoMIS JDH funded by the CNRS.

[†]Toulouse School of Economics, University of Toulouse Capitole, Toulouse, France, fabien.gensbittel@tse-fr.eu

[‡]HEC Paris and GREGHEC, 78351 Jouy-en-Josas, France. lovo@hec.fr. Stefano Lovo acknowledges financial support from ANR/Investissements d'Avenir under grant ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047

[§]Toulouse School of Economics, University of Toulouse Capitole, Toulouse, France, jerome.renault@tse-fr.eu

[¶]HEC Paris and GREGHEC, 78351 Jouy-en-Josas, France. tomala@hec.fr. Tristan Tomala acknowledges financial support from the HEC foundation.

1 Introduction

A revision game models a situation where players prepare the actions they want to see actually implemented at a future predetermined deadline. Players enter the game at an initial time, choose actions and then wait for their actions to be implemented at a deadline and for payoffs to be distributed according to a known mapping from action profiles into payoffs (the component game). Between the initial time and the deadline, players observe the actions the other players have prepared and may have the opportunity to revise their action and reconsider their choices. Revision opportunities arrive at exogenous random times according to independent Poisson processes. As the deadline gets closer and closer, the probability to get another revision opportunity vanishes and the last chosen actions become more and more binding.

This model is a variant of usual stochastic games as defined by Shapley (1953). In this seminal paper, Shapley introduced discounted stochastic games as games with terminal payoffs and random duration, where the probability of termination of the game is stationary over time. By contrast, in revision games, the probability of termination increases with time. Lovo and Tomala (2015) study a general model of revision games with stochastic transitions.

Revision games can be used to model a number of economic and social relations. For example, in some stock markets, thirty minutes before the opening hour, traders can submit orders which can be withdrawn and changed until the opening time. Only the orders recorded in the system at the opening time are binding and payoff relevant. Another example are international negotiations, or negotiations about regulations between a government and the industry to be regulated. During an international negotiation process, parties prepare proposals and counterproposals but only the final draft of the agreement signed at the end of the summit matters. Regulation of industries are renewed at known dates and before the expiration of the current regulation, the regulator and the industry representatives negotiate on the next version. At the same time, companies prepare

their production technology to complain to the next standards in the regulation.¹ Sports also provide many applications for revision games. The composition of teams in football or rugby, needs to be made before the match, the choice of the tires and car setting before the care race starts, etc.

In such frameworks, it is common to empirically observe a *deadline effect*. In the stock markets, little happens at the beginning of the pre-opening period whereas most of the activity of submitting and revising orders concentrates in the last five minutes before the official opening hour (Biais et al., 1999). Negotiations often start with little progress in the early phase, whereas most activity is concentrated when the deadline approaches. These phenomena have been explained using war of attrition (Alesina and Drazen, 1991), or bargaining games with incomplete information (Spier, 1992, Vives and Medrano 2001), or optimism (Simsek and Yildiz, 2014). In a revision game, there is no waiting cost (or war of attrition), no incomplete information, and players are perfectly rational. Nevertheless, we show that if the component game does not have a pure strategy equilibrium, the deadline effect emerges in the revision game from the equilibrium behavior. Thus, we provide a new alternative explanation of the deadline effect in these economic frameworks.

Revision games have been introduced by Kamada and Kandori (2013), and further studied in Calcagno, Kamada, Lovo and Sugaya (2014). In particular, Calcagno et al. (2014) study revision games with 2×2 component games having two strict Nash equilibria.

In the present paper, we consider the case of two-player zero-sum component games. We are interested in the properties of the revision game equilibrium payoff, which we call *the revision value*, and its limit when the length of the revision phase goes to infinity, which we call *the limit revision value*.

¹Some example are: the banking sectors regulations Basel 1, 2, and 3 expire and are renewed known basis; Global warming forums set targets for a given period to be reviewed on regular basis; European Union emission standards Euro 4 effective January 1, 2008, Euro 5 effective January 1, 2010 etc.

The revision value exists (from Lovo and Tomala, 2015) and we characterize it as the unique solution of an ordinary differential equation. We prove that players have *pure* optimal strategies. Then we show that, in contrast with non-zero-sum revision games (see Calcagno et al., 2014), the equilibrium payoff (*i.e.* the value) is a continuous function of the parameters of the model: the payoff matrix of the component game, the relative frequency with which each player has a revision opportunity, and the duration of the revision game. We show that it is advantageous to be relatively fast, which is not the case for the games studied by Calcagno et al. (2014) where a player might gain from being forced to revise at a relatively low frequency.

Then, we analyze the limit revision value as the length of the duration phase (or equivalently, the frequency of revisions) goes to infinity. We show that the limit revision value does not depend on the initial action profile prepared at the beginning of the game and that, moving backward from the deadline, the revision value converges to the limit revision value exponentially fast. When the component game has a pure Nash equilibrium, the limit revision value coincides with the value of the component game. However, when the component game has no pure strategy equilibrium, the limit revision value generically differs from the component game value. In this case, the limit revision value to a player is strictly increasing in the relative frequency with which she can revise her action. The limit revision value converges to the player's min-max in the component game as this relative frequency goes to infinity.

Finally, we fully characterize the equilibrium strategies for 2×2 zero-sum component games. We show that if the component game has a pure Nash equilibrium, when the revision phase is sufficiently long, players prepare the component game equilibrium from the start. If the component game equilibrium is in mixed strategies, at the beginning of the revision phase, players prepare an action profile that they do not revise until the deadline is close enough. When the end of the revision phase approaches, players start changing their action whenever they have a revision opportunity. That is, the equilibrium generates a deadline effect: at the beginning, players settle on a *sur-place* action profile, which is not an equilibrium of the component game.

They wait for the deadline to be close enough and then all the activity of revision is concentrated at the last minute. This is consistent with robust empirical regularity observed in bargaining and negotiations.

The paper is organized as follows: The model is introduced in Section 2. Section 3 analyzes the equilibrium property of a general zero-sum revision game. The finite duration phase is studied in Section 3. Section 4 considers the limit case when the duration goes to infinity. In Section 5, we characterize the equilibrium of the revision game for 2×2 component games. The proofs for this case are relegated to the Appendix. Section 6 concludes.

2 Zero-sum revision games

2.1 Component Game

The first ingredient is a two-player zero-sum game called the *component game* throughout. Let X_1 and X_2 be finite and non-empty sets of actions for players 1 and 2 respectively, and denote $X := X_1 \times X_2$ the set of action profiles. Let $U : X \rightarrow \mathbb{R}$ be player 1's payoff function, the payoff function of player 2 is $-U$. We also denote by U the payoff matrix $(U(x))_{x \in X}$.

For any pure action $x_{-i} \in X_i$ of player $-i$, denote $BR_i^U(x_{-i}) \subseteq X_i$ the set of player i 's pure best responses to x_{-i} in the component game with payoff matrix U . Formally, for any x_1 in X_1 and x_2 in X_2 , we denote the best responses by:

$$BR_1^U(x_2) := \arg \max_{y_1 \in X_1} U(y_1, x_2), \quad BR_2^U(x_1) := \arg \min_{y_2 \in X_2} U(x_1, y_2)$$

We say that U is *regular* if for each player i and x_{-i} , $BR_i^U(x_{-i})$ is a singleton.

Let V be the minmax value of this game in mixed strategies, that is, V is player 1's unique Nash equilibrium payoff of the game U . Let S_i be player i 's Stackelberg payoff in the game where the first mover is player $i = 1, 2$. That is, S_1 (resp. S_2) is the maxmin (resp. minmax) value of

the game in pure strategies. One has,

$$\max_{x_1} \min_{x_2} U(x_1, x_2) = S_1 \leq V \leq S_2 = \min_{x_2} \max_{x_1} U(x_1, x_2)$$

and $S_1 = S_2$ if and only if the component game has a Nash equilibrium in pure strategies.

2.2 Revision Game

We describe now the revision game built on the component game. Time is continuous, $t \in [0, T]$ where T is a fixed deadline. At time 0, there is an initial “prepared” action profile $x(0) \in X$, given exogenously. For each player i , there is a Poisson process with intensity $\lambda_i > 0$, the two processes being independent. A player is able to change her prepared action only at an arrival of her Poisson process which we call a revision opportunity. At every time $t < T$, if the current prepared action profile is $x(t)$ and player i receives a revision opportunity, then she chooses $y_i \in X_i$ and the new prepared action profile is $(y_i, x_{-i}(t))$. At time $t = T$, the action profile $x(T)$ is actually played and each player receives the corresponding payoff, $U(x(T))$ for player 1 and $-U(x(T))$ for player 2. This is the only payoff that players receive in the revision game.

Revision opportunities and prepared actions are perfectly and publicly observed, the description and the parameters of the game are common knowledge. The revision game is called *asynchronous* since with probability one, only one player receives a revision opportunity at a time.

We will use the following parametrization. Denote $\lambda = \lambda_1 + \lambda_2$ and $q = \lambda_1/\lambda$. The revision process is equivalently described as follows. There is a common Poisson process with intensity λ , which we call the clock. At each arrival, or beat of the clock, a coin flip $(q, 1 - q)$ determines whether player 1 or player 2 gets the revision opportunity.

Next, we define strategies. A history of length n is a sequence $(x_0, \tau_1, i_1, x_1, \tau_2, \dots, \tau_n, i_n, x_n)$ in $H_n = X \times ([0, T] \times \{1, 2\} \times X)^n$, with $H_0 = X$ and where x_0 denotes the initial action profile, τ_m the time of the m -th beat of the clock, i_m the player who had the opportunity to revise at time τ_m and x_m the new action profile after the revision. A (behavior) strategy σ_i of player i is a measurable mapping from $\cup_{n \geq 0} (H_n \times [0, T])$ to $\Delta(X_i)$. A strategy is Markov if it depends on the time of revision and the other player's current action only. Formally, a Markov strategy σ_i for player i is a measurable mapping from $X_{-i} \times [0, T]$ to $\Delta(X_i)$.

A pair $\sigma = (\sigma_1, \sigma_2)$ induces a unique probability P_σ over the set $\cup_{n \geq 0} H_n$ of finite histories which are identical to plays of the revision games. If a play $(x, \tau_1, i_1, x_1, \tau_2, \dots, \tau_n, i_n, x_n)$ is realized, it means that τ_n was the last ring of the clock and the terminal action profile is $x(T) = x_n$. If the history $x_0 \in H_0 = X$ is realized, it means that the clock did not ring and $x(T) = x_0$. A play induces a unique right-continuous trajectory for the profile of actions $(x(t))_{t \in [0, T]}$ which jumps at revision times. This allows to define an expected payoff $\mathbb{E}_\sigma[U(x(T))]$.

We define the minmax and maxmin value of this game as,

$$\bar{u}(T, x) = \inf_{\sigma_2} \sup_{\sigma_1} \mathbb{E}_\sigma[U(x(T))], \quad \underline{u}(T, x) = \sup_{\sigma_1} \inf_{\sigma_2} \mathbb{E}_\sigma[U(x(T))].$$

If these two quantities are equal, the game is said to have a value $u(T, x) = \bar{u}(T, x) = \underline{u}(T, x)$.

In the sequel we denote by $t \geq 0$ **the remaining time to the deadline**, whereas T denotes the total length of the revision game, so that $t = T - \tau$ in the above description. The revision game over the interval $[0, T]$ with initial position x is denoted $\Gamma(T, x)$.

Notations: Given a real-valued function $u(t, x)$, we will denote by $u(t)$ the matrix $(u(t, x))_{x \in X}$, and $\|u(t)\| := \max_{x \in X} |u(t, x)|$. For $t \geq 0$ and $x = (x_1, x_2)$, we write:

$$u^-(t, x) := \min_{y_2 \in X_2} u(t, x_1, y_2) \text{ and } u^+(t, x) := \max_{y_1 \in X_1} u(t, y_1, x_2).$$

Clearly, $u^- \leq u \leq u^+$ pointwise.

3 Revision values

We start by analyzing the equilibrium and value for a revision game where T and λ are fixed. Let us first note that existence of equilibrium and value is guaranteed.

Lemma 1 *The revision game admits a Markov perfect equilibrium σ , the revision game with length t and initial position x has a value denoted by $u(t, x)$. This equilibrium and the value are such that for every $x \in X$ and almost every $t \in [0, T]$,*

$$\sigma_i(t, x_{-i}) \in BR_i^{u(t)}(x_{-i}). \quad (1)$$

Proof. The result follows immediately from Lemma 2.1 and Theorem 4.1 in Lovo and Tomala (2015). ■

From this result, a zero-sum revision game has an equilibrium and hence a value. Also there exist an equilibrium where each player's strategy only depends on the last action prepared by the other player and the remaining time to the end of the game. This Markov equilibrium strategies satisfy equation (1). To interpret this equation, suppose that at time t from the end, player i has a revision opportunity whereas player $-i$'s last prepared action is x_{-i} . After preparing an action $y_i \in X_i$, the players will be playing a revision game with length t and initial action profile $x(0) = (y_i, x_{-i})$, that is, the game $\Gamma(t, (y_i, x_{-i}))$. Thus, at a revision time, player i has the opportunity to choose which revision game to play from the set $\{\Gamma(t, (y_i, x_{-i}))\}_{y_i \in X_i}$. Since the equilibrium payoff associated to $\Gamma(t, (y_i, x_{-i}))$ is $u(t, (y_i, x_{-i}))$, it is optimal for player i to prepare an action $y_i \in BR_i^{u(t)}(x_{-i})$. That is, if a player has a revision opportunity at t , then he will prepare a best response to the other player's last prepared action in the component game with payoff $u(t)$. Thus, for any remaining time $t \geq 0$ before the deadline, the continuation payoff

is $u(t, x)$ and only depends on t and on the current state x .

From now on, we call *revision value* the mapping $(t, x) \mapsto u(t, x)$. The main result of this section is that the revision value is characterized as the solution of an ordinary differential equation, from which we derive regularity properties. We first prove that $u(t, x)$ is continuous, in fact Lipschitz, with respect to time. When t and t' are close, with high probability there will be no revision between t and t' , thus $u(t, x)$ and $u(t', x)$ are close. In other words, a small increase of the duration of revision game has a small effect on the equilibrium payoff. More formally:

Lemma 2 *For all $x \in X$, $u(\cdot, x)$ is $2\lambda\|U\|$ -Lipschitz.*

Proof. Consider two remaining times $t' > t$. Since histories in $\Gamma(t, x)$ are also histories in $\Gamma(t', x)$, any strategy σ'_i in $\Gamma(t', x)$ induces naturally (by restriction) a strategy in $\Gamma(t, x)$ (also denoted σ'_i for simplicity). On the other hand, any strategy σ_i in $\Gamma(t, x)$ may be seen as the restriction of some strategy in $\Gamma(t', x)$. For instance, one may define a strategy $\hat{\sigma}_i$ in $\Gamma(t', x)$ in which player i uses σ_i when the remaining time is less than t , and does not revise his action when the remaining time is between t and t' . As a consequence, $u(t, x) = \sup_{\sigma'_1} \inf_{\sigma'_2} \mathbb{E}_{\sigma} [u(x(t))]$ where the supremum and infimum range over strategies in $\Gamma(t', x)$.

For the strategy profile (σ'_1, σ'_2) , the probability that at least one revision time occurs in the interval $(t, t']$ is $(1 - e^{-\lambda|t-t'|})$, so that $P_{\sigma'}(x(t) = x(t')) \geq e^{-\lambda|t-t'|}$, which implies,

$$\mathbb{E}_{\sigma'}[|u(x(t)) - u(x(t'))|] \leq 2\|U\|(1 - e^{-\lambda|t-t'|}) \leq 2\|U\|\lambda|t - t'|.$$

The result follows. ■

Next, we show that the value function $u(t, x)$ can be obtained as the solution of a system of differential equations.

Proposition 1 *For all $t \geq 0$ and all $x \in X$, the value $u(t, x)$ of the revision game $\Gamma(t, x)$ is the unique solution of the following system of differential equations:*

Evolution equations: $\forall x \in X, \forall t \geq 0$,

$$\frac{\partial u(t, x)}{\partial t} = \lambda_1(u^+(t, x) - u(t, x)) + \lambda_2(u^-(t, x) - u(t, x)), \quad u(0, x) = U(x). \quad (2)$$

Proof. Let us first observe that the value satisfies the following Dynamic Programming Principle. For $s \geq 0$, denote $u^s(t, x)$ the value of the revision game with length t , initial position x , and component game $u(s)$. Then $u^0(t, x) = u(t, x)$, $u^t(0, x) = u(t, x)$. More generally:

Dynamic Programming Principle:

$$\forall x \in X, \forall t \geq 0, \forall s \in [0, t], \quad u^s(t - s, x) = u(t, x). \quad (3)$$

Consider now the game with length t , and let s be the realization of the waiting time before the first revision opportunity, drawn from the exponential distribution. If $s \geq t$, no revision occurs and the payoff is $u(x)$. If $s \leq t$ and player 1 has the opportunity to revise, the new payoff is $u(t - s, \sigma_1(s, x), x_2) = u^+(t - s, x)$, and if player 2 has the opportunity to revise, the new payoff is $u(x_1, \sigma_2(s, x), t - s) = u^-(t - s, x)$. So that:

$$u(t, x) = U(x)e^{-\lambda t} + \int_{s=0}^t \lambda e^{-\lambda s} (qu^+(t - s, x) + (1 - q)u^-(t - s, x)) ds. \quad (4)$$

The change of variable ($s \mapsto t - s$) yields the following **Dynamic programming equation:**

$$u(t, x) = U(x)e^{-\lambda t} + \int_{s=0}^t e^{-\lambda(t-s)} (\lambda_1 u^+(s, x) + \lambda_2 u^-(s, x)) ds. \quad (5)$$

Since u^+ and u^- are continuous, u is continuously differentiable in t , and the law of motion of the continuation values is given by the system:

$$\frac{\partial u(t, x)}{\partial t} = \lambda_1(u^+(t, x) - u(t, x)) + \lambda_2(u^-(t, x) - u(t, x)), \quad u(0, x) = U(x).$$

The mapping $(u \mapsto \lambda_1 u^+ + \lambda_2 u^- - \lambda u)$ is Lipschitz continuous, so by the Cauchy-Lipschitz Theorem, this ordinary differential equation has a unique solution. ■

From this characterization, we deduce the following.

Corollary 1 *The revision game has a Markov perfect equilibrium in pure strategies.*

Proof. As argued before, the revision game $\Gamma(t, x)$ has an equilibrium, or equivalently, players have optimal strategies guaranteeing the value, but these are possibly mixed. We prove now that there exist also *pure* optimal strategies. The main argument is that a player who is randomizing can select any of the pure actions in the mixing support without affecting the continuation payoffs. The zero-sum aspect is crucial here, in a non-zero-sum game, it is possible to change the continuation payoff of one player without affecting the continuation payoff of the other.

We now detail the construction, note that we have to make sure that the strategy is measurable w.r.t. time. To this end, we identify the action set X_i with $\{1, \dots, |X_i|\}$ according to some fixed ordering. This allows us to break ties in a measurable way.

Define a strategy $\sigma(1)$ of player 1 in the game $\Gamma(t, x)$ as follows:

- If the first beat of the clock occurs at time s and if player 1 is called to play, then she chooses the action,

$$x'_1 = \min \left(\operatorname{argmax}_{y_1 \in X_1} u^{t-s}(y_1, x_2) \right).$$

That is, instead of playing the original Markov equilibrium strategy (which is possibly mixed), player 1 plays the action in the support of the mixed strategy which is minimal in the given ordering.

- Otherwise, player 1 plays some (possibly mixed) optimal strategy given by some fixed Markov perfect equilibrium.

Against any strategy of player 2, this strategy $\sigma(1)$ guarantees the payoff:

$$u(x)e^{-\lambda t} + \int_{s=0}^t \lambda e^{-\lambda s} (qu^+(t-s, x) + (1-q)u^-(t-s, x)) ds.$$

From equation (4), this is precisely the value $u(t, x)$, and therefore $\sigma(1)$ is optimal in the game $\Gamma(t, x)$.

So in the game $\Gamma(t, x)$, player 1 has an optimal strategy which is pure at the first beat of the clock. Iteratively, define for every $n \geq 1$, the following strategy $\sigma(n)$ of player 1:

- If for some $m \leq n$, player 1 receives a revision opportunity at the m -th beat of the clock occurring at time s , player 1 chooses the action $x'_1 = \min \left(\operatorname{argmax}_{y_1 \in X_1} u^{t-s}(y_1, x_2(s)) \right)$, where $x_2(s)$ denotes the current action of player 2 at time s .

- Otherwise, player 1 plays some (possibly mixed) optimal strategy given by some fixed Markov perfect equilibrium.

A simple induction argument shows that $\sigma(n)$ is optimal for player 1 in the game $\Gamma(t, x)$. Then, for any length t , the clock will almost surely beat finitely many times. Thus, letting n go to infinity shows that the strategy $\sigma(\infty)$ which plays at time s the pure action $x'_1 = \min \left(\operatorname{argmax}_{y_1 \in X_1} u^{t-s}(y_1, x_2(s)) \right)$, is a pure optimal Markov strategy of player 1.

By symmetry, the same holds for player 2, which concludes the proof. ■

Our next result states that the revision value $u(t, x)$ is continuous with respect to the parameters of the model: the payoff matrix U and the revision parameters λ and q .

Lemma 3 *The revision value $u(t, x)$ is 1-Lipschitz in U , Lipschitz-continuous in $\lambda \in (0, \infty)$, continuous and non-decreasing in $q \in (0, 1)$.*

This result is in contrast with properties of non zero-sum revision games, see Calcagno et al. (2014). In a Battle of the Sexes game, a slight change in the relative speed of players, measured

by q , may drastically change the equilibrium payoff. Further, it might be advantageous to be relatively slower than the opponent. This cannot be the case in a zero-sum game.

Proof.

i) $u(t, x)$ is 1-Lipschitz in U . Consider first two payoff matrices U and U' . For any $T \geq 0$ and x in X , the strategy spaces of the players in the revision games $\Gamma(T, x)$ and $\Gamma'(T, x)$ with respective component games U and U' are the same. The payoffs of a given strategy profile in the two games can be written $\mathbb{E}_P(U)$ and $\mathbb{E}_P(U')$, with P the same probability distribution over X , hence they can not differ by more than $\|U - U'\| := \max_x |U(x) - U'(x)|$. Consider now the equilibrium strategy in the revision game with component game U . Player 1's equilibrium strategy yields a payoff at least $u(T, x)$, no matter the strategy of player 2. If player 1 uses the same strategy in the revision game with component game U' , she cannot lose more than $\|U - U'\|$. Thus, we have $|u(T, x) - u'(T, x)| \leq \|U - U'\|$ for all T and x .

ii) Lipschitz-continuity in λ . For fixed q , modifying λ just changes the frequency of beats of the Poisson clock, but not the relative frequency of beats of player 1 and player 2. In other words, a change in λ is equivalent to a change of unit of time. Namely, denoting $u(t, x)|_\lambda$ the value of $\Gamma(t, x)$ with parameter λ , we have $u(t, x)|_{\lambda'} = u(\frac{\lambda'}{\lambda}t, x)|_\lambda$. Thus, $u(t)$ is Lipschitz in λ since it is Lipschitz in t .

iii) Continuity in q . Let us prove the continuity of $u(t, x)$ in the parameter $q \in [0, 1]$, λ and U being fixed. Given a strategy profile σ , denote \mathbb{E}_σ the expectation with respect to P_σ when the parameter is q , and \mathbb{E}'_σ the expectation with respect to P'_σ when the parameter is q' . We have the following uniform continuity property:

$$\forall \varepsilon > 0, \exists \alpha > 0, |q - q'| \leq \alpha \Rightarrow \forall \sigma, |\mathbb{E}_\sigma(U) - \mathbb{E}'_\sigma(U)| \leq \varepsilon$$

To see this, fix a strategy σ and compare $\mathbb{E}_\sigma(U)$ with $\mathbb{E}'_\sigma(U)$. To evaluate these expectations, what matters is the distribution of the total number of beats, and of the alternation between

players. We claim that this distribution is continuous with respect to q . For any $\varepsilon > 0$, there is an integer M large enough such that, with probability at least $1 - \varepsilon$, the total number of beats is no more than M . Then, conditional on a number of beats $m \leq M$, finitely many configurations of alternations of players are possible. For given m , the conditional probability of these configurations is continuous with respect to q (for each configuration, this can be written $q^k(1 - q)^{m-k}$, with k the number of revisions for player 1 in the configuration). So, for q' close enough to q , the differences in probabilities are bounded by ε .

The payoff function of the revision game $\Gamma(t, x)$ with parameter q , is thus close to the payoff function of the revision game $\Gamma'(t, x)$ with parameter q' , uniformly over strategies. The values of the two games are therefore also close to each other.

iv) Monotonicity in q . Take $q' < q$ and let $u(t, x)|_q$ be the revision game value for a given q . We prove that $u(t, x)|_{q'} \leq u(t, x)|_q$ for all t and x . This is a direct consequence of the Gronwall inequality. At $t = 0$, we have $u(0)|_{q'} = u(0)|_q = U$ and from equation (2) we have $\frac{\partial u(0, x)|_{q'}}{\partial t} \leq \frac{\partial u(0, x)|_q}{\partial t}$ for all x . From Gronwall, the inequality is maintained over time and $u(t, x)|_{q'} \leq u(t, x)|_q$ for all t and x . ■

4 Limit revision value

In this section we analyse the asymptotic behavior of the revision value as the duration T of the revision phase goes to infinity, or equivalently, as the intensity of arrivals of revisions λ goes to infinity (for T fixed).

We denote by $\underline{R}(x)$ and $\overline{R}(x)$ the lim inf and lim sup of the revision value $u(t, x)$ as t goes to infinity. If $\underline{R}(x) = \overline{R}(x) = R$ is independent of x , then we call R the *limit revision value*. We will prove the existence of the limit revision value and compare it with the static solutions of the game: V , S_1 and S_2 .

Proposition 2 For each component game U and revision parameters $\lambda > 0$ and $q \in (0, 1)$, there exists a number $R \in [S_1, S_2]$, such that for all $t \geq 0$ and all $x \in X$,

$$|u(t, x) - R| \leq 2\|U\| \exp(-\lambda(1 - \max\{q, 1 - q\})t). \quad (6)$$

It follows directly that the limit revision value $R = \lim_t u(t, x)$ exists and is independent of the initial position x . The number R satisfies $S_1 \leq R \leq S_2$, thus if the component game has a pure Nash equilibrium, then $V = S_1 = S_2 = R$.

Remark that the condition $0 < q < 1$ is important. If $q = 1$, only player 1 can change her action and the value depends on the initial action of player 2.

Proof. Consider the maximum $M(t) = \max_{x \in X} u(x, t)$ of the revision value in the game of length t . If $t > s$, by the dynamic programming principle (3), we have for all x , $u(t, x) = u^s(t - s, x) \leq M(s)$, so $M(t) \leq M(s)$. $M(t)$ is non-increasing with time and thus converges to some M_∞ . Similarly, $m(t) = \min_{x \in X} u(t, x)$ is non-decreasing and converges to some m_∞ , and we have,

$$m = m(0) \leq m(t) \leq m_\infty \leq M_\infty \leq M(t) \leq M = M(0).$$

For a given $t \geq 0$, let y be such that $m(t) = u(t, y)$ and x be such that $M(t) = u(t, x)$. From the dynamic programming equation in $\Gamma(t, x)$, we have:

$$M(t) = u(t, x) \leq \|U\|e^{-\lambda t} + \int_{s=0}^t \lambda e^{-\lambda(t-s)}(qM(s) + (1 - q)u(s, x_1, y_2))ds$$

since $u^+(s, x) \leq M(s)$ and $u^-(s, x) \leq u(s, x_1, y_2)$. Similarly, considering $\Gamma(t, y)$, we have:

$$m(t) = u(t, y) \geq -e^{-\lambda t}\|U\| + \int_{s=0}^t \lambda e^{-\lambda(t-s)}((1 - q)m(s) + qu(s, x_1, y_2))ds.$$

Define $w(t) = M(t) - m(t)$. If $q \geq 1/2$, we write $qM(s) - (1 - q)m(s) = qw(s) + (2q - 1)m(s)$ and deduce:

$$\begin{aligned} w(t) &\leq 2\|U\|e^{-\lambda t} + q \int_0^t \lambda e^{-\lambda(t-s)} w(s) ds + \int_0^t \lambda e^{-\lambda(t-s)} (2q - 1)(m(s) - u(s, x_1, y_2)) ds \\ &\leq 2\|U\|e^{-\lambda t} + q \int_0^t \lambda e^{-\lambda(t-s)} w(s) ds, \end{aligned}$$

where the second inequality follows from $(2q - 1)(m(s) - u(s, x_1, y_2)) \leq 0$.

Similarly, if $q \leq 1/2$, we write $qM(s) - (1 - q)m(s) = (1 - q)w(s) + (2q - 1)M(s)$ and get,

$$\begin{aligned} w(t) &\leq 2\|U\|e^{-\lambda t} + (1 - q) \int_0^t \lambda e^{-\lambda(t-s)} w(s) ds + \int_0^t \lambda e^{-\lambda(t-s)} (2q - 1)(M(s) - u(s, x_1, y_2)) ds \\ &\leq 2\|U\|e^{-\lambda t} + (1 - q) \int_0^t \lambda e^{-\lambda(t-s)} w(s) ds. \end{aligned}$$

We deduce that for each q ,

$$w(t) \leq 2\|U\|e^{-\lambda t} + \max\{q, 1 - q\} \int_0^t \lambda e^{-\lambda(t-s)} w(s) ds.$$

Denoting $g(t) = e^{\lambda t} w(t)$, we have,

$$g(t) \leq 2\|U\| + \max\{q, 1 - q\} \int_0^t \lambda g(s) ds.$$

Applying the Gronwall inequality to the function g we obtain,

$$g(t) \leq 2\|U\|e^{\lambda t \max\{q, 1 - q\}}.$$

We deduce finally that,

$$w(t) \leq 2\|U\|e^{\lambda t(\max\{q, 1 - q\} - 1)},$$

which concludes the proof: this implies $M_\infty = m_\infty = R$ and $|u(t, x) - R| \leq w(t)$.

To show $S_1 \leq R \leq S_2$, it is enough to prove that,

$$S_1 \leq \liminf_t u(t, x) \leq \limsup_t u(t, x) \leq S_2.$$

Consider the strategy of player 1 which always prepares a maxmin action, *i.e.*, an optimal action in the Stackelberg game where player 1 moves first. Since the length of the revision phase is arbitrary long and $\lambda q > 0$, the probability that player 1 gets at least one revision opportunity, is arbitrarily close to one. Hence, this strategy guarantees player 1 get approximately S_1 when T is large and $\liminf_t u(t, x) \geq S_1$. Dually, $\limsup_t u(t, x) \leq S_2$. ■

From this result, we can define for the matrix game U , the limit revision value $R(U) = \lim_t u(t, x)$, independently of x . Notice that the dynamic programming principle (3) gives that:

$$\forall s \geq 0, \quad R(u(s)) = R(U).$$

As for the fixed duration revision value, the limit revision value R is continuous in the parameters of the model, as stated in the following lemma.

Lemma 4 *The revision game value R is 1-Lipschitz in U , independent of $\lambda \in (0, \infty)$, continuous and non-decreasing in $q \in (0, 1)$. It converges to the pure minmax S_2 when q goes to 1, and to the pure maxmin S_1 when q goes to 0.*

Proof.

i) R is 1-Lipschitz in U . Take two payoff matrices U and U' . From Lemma 3, for all t and x , $|u(t, x) - u'(t, x)| \leq \|U - U'\|$. Letting $t \rightarrow \infty$ gives $|R(U) - R(U')| \leq \|U - U'\|$.

ii) R does not depend on λ . We have seen in the proof of Lemma 3 that a change in λ corresponds to a change in the unit of time. Taking the limit when t goes to infinity removes the dependency on λ , and the revision value R depends only on q and U .

iii) R is continuous in $q \in (0, 1)$. Fix $q \in (0, 1)$ and take a sequence $q_n \rightarrow q$. For any $\varepsilon > 0$, there exists t large enough such that $u(t, x) \in [R - \varepsilon, R + \varepsilon]$, for all x in X . From Lemma 3, $u(t, x)$ is continuous in q , and for n large enough, the revision value $u_n(t, x)$ of the game with parameter q_n belongs to $[R - 2\varepsilon, R + 2\varepsilon]$, for all x . Thus, for n large enough, the limit revision value of the game with parameter q_n belongs to $[R - 2\varepsilon, R + 2\varepsilon]$, concluding the proof.

iv) R is non-decreasing in q . From Lemma 3, $u(t, x)$ is continuous in t and non-decreasing in q . Thus, $R = \lim_{t \rightarrow \infty} u(t, x)$ is also non-decreasing in q .

v) R converges to S_2 when q goes to 1. Take t large enough such that in the game with duration t , there is at least one revision with probability at least $1 - \varepsilon$. Whenever called to play, Player 1 can choose a pure best reply to the current position in the component game U . With probability at least $(1 - \varepsilon)q$, there is at least one revision opportunity and player 1 is called to play at the last clock beat. Thus for all x :

$$u(t, x) \geq (1 - \varepsilon)qS_2 - (1 - (1 - \varepsilon)q)\|U\|.$$

By letting t go to infinity and ε go to 0, we get that the revision value satisfies $R \geq qS_2 - (1 - q)\|U\|$. Since $R \leq S_2$, R converges to S_2 when q goes to 1.

Symmetrically, R converges to S_1 when q goes to 0. ■

5 2×2 zero-sum revision games

In this section, we focus on 2×2 component games. For this class, we characterize the equilibrium strategies and the limit revision value R .

Consider a 2×2 zero-sum game and denote $X_i = \{\alpha, \beta\}$ the actions available to player i . We focus on the generic case where u is regular. First, we remark that, for regular 2×2 zero-sum games, there are only three possible scenarios, describing the players' best responses *in the component game*.

Definition 1 1. *Scenario DD: Each player has a dominant action.*

2. *Scenario DN: One player has a dominant action, the other player's best response varies with the opponent's action.*

3. *Scenario NN: Each player's best response varies with his opponent's action.*

The game has a pure Nash equilibrium for the scenarios DD and DN, and a mixed one for the scenario NN.

Without loss of generality, let us normalize payoffs by setting $U(\alpha, \alpha) = 0$. Observe then that if $U(\alpha, \beta) + U(\beta, \alpha) - U(\alpha, \alpha) - U(\beta, \beta) = 0$, each player has a dominant action: Player 1's (Player 2's) dominant action is α if $U(\beta, \alpha) \leq 0$ (resp. $U(\alpha, \beta) \geq 0$) and β if $U(\beta, \alpha) \geq 0$ (resp. $U(\alpha, \beta) \leq 0$).

Now generically, $U(\alpha, \beta) + U(\beta, \alpha) - U(\alpha, \alpha) - U(\beta, \beta) \neq 0$, and it is possible to further normalize payoffs by setting $U(\alpha, \beta) + U(\beta, \alpha) - U(\alpha, \alpha) - U(\beta, \beta) = 1$. Thus, we adopt the following specification of U :

	α	β
α	0	b
β	c	$b + c - 1$

Figure 1

This parametrization is compatible with each of the three scenarios as illustrated in the following example.

Example 1 1. If $c, b > 1$, the component game scenario is *DD*: the dominant actions are β for Player 1 and α for Player 2.

2. If $0 < b < 1$ and $c > 1$, the component game scenario is *DN*: the dominant action for Player 1 is β , whereas for Player 2, $BR_2^U(\alpha) = \alpha$, $BR_2^U(\beta) = \beta$.

3. If $0 < b < 1$ and $0 < c < 1$, the component game scenario is *NN*: for Player 1, $BR_1^U(\alpha) = \beta$, $BR_1^U(\beta) = \alpha$, whereas for Player 2, $BR_2^U(\alpha) = \alpha$, $BR_2^U(\beta) = \beta$.

Consider a Markov equilibrium of the revision game. For any $t \geq 0$, we denote the continuation payoff matrix as follows:

	α	β
α	$u(t, \alpha, \alpha)$	$u(t, \alpha, \beta)$
β	$u(t, \beta, \alpha)$	$u(t, \beta, \beta)$

where $u(0, x) = U(x)$. The next Proposition shows that studying the equilibrium strategies for each of the three scenarios, fully characterizes the equilibrium play in the revision game.

Proposition 3 Consider a regular 2×2 payoff matrix U . The revision game equilibrium σ is as follows.

1. If the component game scenario is *DD*, then each player prepares her dominant action in the component game, for all $t \geq 0$. It results that $R = V$.

2. If the component game scenario is *DN*, then there is a player i who has a dominant action \hat{x}_i and prepares it at all $t \geq 0$:

$$\sigma_i(t, x_{-i}) = \hat{x}_i, \quad \forall t, x_{-i}.$$

There exists $t^* > 0$ such that:

- for $t < t^*$, Player $-i$ prepares her best response in the component game to the current action of Player i :

$$\sigma_{-i}(t, x_i) = B_{-i}^U(x_i);$$

- for $t > t^*$, Player $-i$ prepares her component game best response to Player i 's dominant action in the component game:

$$\sigma_{-i}(t, x_i) = B_{-i}^U(\hat{x}_i), \forall x_i.$$

It results that $R = V$.

3. If the component game scenario is NN, then there are $t_2^* > t_1^* > 0$ and a player i , such that:

(a) Player i prepares her component game best response to Player $-i$'s current action when $t < t_1^*$, and prepares a fixed action \hat{x}_i when $t > t_1^*$, i.e.,

$$\sigma_i(t, x_{-i}) = \begin{cases} BR_i^U(x_{-i}) & \text{for } t < t_1^* \\ \hat{x}_i & \text{for } t \geq t_1^*, \forall x_{-i} \end{cases}$$

(b) Player $-i$ prepares her component game best response to Player i 's current action when $t < t_2^*$. When $t > t_2^*$, Player $-i$ prepares her component game best response to the fixed action that Player i prepares when $t > t_1^*$, i.e.,

$$\sigma_{-i}(t, x_i) = \begin{cases} BR_{-i}^U(x_i) & \text{for } t < t_2^* \\ BR_{-i}^U(\hat{x}_i) & \text{for } t \geq t_2^*, \forall x_i \end{cases}$$

It results that $R \neq V$ for generic NN component games. In this scenario, we define $\hat{x}_{-i} := BR_{-i}^U(\hat{x}_i)$ and call $\hat{x} := (\hat{x}_i, \hat{x}_{-i})$ the sur-place action profile.

We provide here the intuition for the proof, the details are in the appendix. The proof is divided in three steps. First, if at some t from the deadline, the continuation payoff matrix $u(t)$ is in scenario DD, then at any $t' > t$ the payoff matrix $u(t')$ is also in scenario DD, *i.e.* the players' dominant actions are the same in t and in t' . Thus, for a component game in scenario DD, at all revision opportunities, players prepare the same dominant actions which correspond to the pure Nash equilibrium of the component game.

Second, we consider scenario DN. We show that if at some t , the matrix $u(t)$ is in scenario DN, then there is $t^* > t$, such that for $t' < t^*$ the continuation payoff matrix $u(t')$ is in the same scenario DN, whereas for $t' > t^*$ it is in scenario DD. Further, the scenario DD in the early stages of the revision game ($t' > t^*$) is such that players prepare the pure Nash equilibrium of the component game. Once the prepared actions reach this action profile, players do not revise anymore.

Third, we analyze scenario NN, the most interesting case. We show that if U is in scenario NN and payoffs are generic, then there is a finite $t_1^* > 0$ such that for $t < t_1^*$ the continuation payoff matrix $u(t)$ is in the same scenario NN (with the same best responses). For $t > t_1^*$ and close to t_1^* , the payoff matrix $u(t)$ is in scenario DN, and one can apply the analysis of the former two steps.

It results that in the early stages of the revision game, players prepare a *sur-place* action profile, and do not revise until t_1^* from the deadline. Players start “wrestling” and change their prepared action at any revision opportunity, only when the remaining time is less than t_1^* . This is illustrated in Figure 1.

In Figure 1, the component game scenario is NN: player 2 prefers to play the same action as player 1, who prefers to play the opposite. In red is represented the equilibrium behavior when the deadline is close $t < t_1^*$: a revising player prepares the best response to the other player current action. In blue is represented the equilibrium behavior when the remaining time




Figure 1: Evolution of scenarios

is between t_2^* and t_1^* . The continuation payoff function $u(t)$ is in scenario DN: α is dominant for Player 1. If the prepared action profile is (α, α) , players do not change their actions. Otherwise, at least one player changes his action. Compared to the strategy used for $t < t_1^*$, only Player 1's strategy changes: she now prepares α when player 2 prepares α (orange vertical arrow). In black is represented the equilibrium behavior for $t > t_2^*$. Both players' dominant action is α . Compared to the equilibrium strategy for $t \in (t_1^*, t_2^*)$, the only difference is that player 2 prefers α also when player 1 is playing β (orange horizontal arrow).

The sur-place action profile depends on the specific values of b , c and the relative speed of revision $(1 - q)/q$. Note that U being in scenario NN requires that $0 < b, c < 1$.² Without loss of generality, we assume $0 < b, c < 1$ from now on.

Denote $\lambda_1 = \lambda q$, $\lambda_2 = \lambda(1 - q)$, define $\theta := |\lambda_1^2 - 6\lambda_1\lambda_2 + \lambda_2^2|^{\frac{1}{2}}$ and let $\hat{t}(A, B, \lambda_1, \lambda_2)$ be the smallest positive t such that,

$$e^{-\frac{\lambda_1 + \lambda_2}{2}t} + A \cos\left(\frac{\theta t}{2}\right) + \frac{(\lambda_2 - \lambda_1)A + 2\lambda_2 B}{\theta} \sin\left(\frac{\theta t}{2}\right) = 0. \quad (7)$$

²Namely player 2 has a dominant action if $b > 1$ or $b < 0$ whereas player 1 has a dominant action if $c > 1$ or $c < 0$.

If this equation has no solution in $t \geq 0$, then $\hat{t}(A, B, \lambda_1, \lambda_2)$ is set to infinity. Let,

$$t_{\alpha,\alpha} = \hat{t}(2c - 1, 2b - 1, \lambda_1, \lambda_1)$$

$$t_{\alpha,\beta} = \hat{t}(2b - 1, 1 - 2c, \lambda_2, \lambda_1)$$

$$t_{\beta,\alpha} = \hat{t}(1 - 2b, 2c - 1, \lambda_2, \lambda_1)$$

$$t_{\beta,\beta} = \hat{t}(1 - 2c, 1 - 2b, \lambda_1, \lambda_1)$$

The following result holds.

Lemma 5 *Consider a 2×2 component game in scenario NN. The sur-place action profile is,*

$$\hat{x} = \arg \min_{y \in X} t_y.$$

Then $t_1^ = t_{\hat{x}}$, \hat{x} is a pure Nash equilibrium of $u(t_{\hat{x}})$ and the limit revision value is $R = u(t_{\hat{x}}, \hat{x})$.*

The proof is in the Appendix.

We detail now two specific cases. First, we consider the revision game where U is the matching pennies game. This is a non-generic case where in equilibrium, each player prepares her best response in the component game at each revision opportunity. The limit revision value R is then a simple function of q .

Example 2 *For $b = c = 1/2$, the component game is the matching pennies game:*

	α	β
α	0	1/2
β	1/2	0

In this case, $\hat{t}_x = \infty$ for all $x \in X$ and therefore, no matter the remaining time, at any revision opportunity, each player prepares his component game best response to the other player's current action. This implies that the prepared action profile will keep cycling counterclockwise. As t goes to infinity, the equilibrium payoff converges to,

$$R = \frac{q}{2}.$$

Note that in this case $R > V = 1/4$ if and only if $q > 1 - q$, *i.e.*, player 1 gains more in a revision game than in the one-shot game, only if he is faster than the other player. When q goes to 1 or to 0, the fastest player gets approximately the Stackelberg second mover's payoff.

As a second specific case, we consider games where players are equally fast. We thus set $q = 1/2$ and analyze how R varies with b and c . In this case, R is trigonometric function of b , c and t_1^* , the latter being defined as the smallest solution of an equation similar to equation (7).

Example 3 *Fix $q = 1/2$, the sur-place action depends on the parameters b and c in $[0, 1]$. Figure 2 represents the regions for b and c corresponding to different surplace actions.*

Figure 2 outlines an interesting property of zero-sum revision games. For $b \simeq c \simeq \frac{1}{2}$, a small change in the payoff b can change the sur-place action to any other action. In other words, there are regions for b and c in which the equilibrium strategy is “very” discontinuous in the parameters. Nevertheless, the resulting equilibrium payoff change smoothly (see Lemma 3).

As comparative statics exercise, let's ask whether player 1 prefers to play the one-shot component game, or the revision game with long duration. This is equivalent to comparing the component game value V with the limit revision value R . We already know that when the component game equilibrium is in mixed strategies, then player 1 prefers to play the revision game

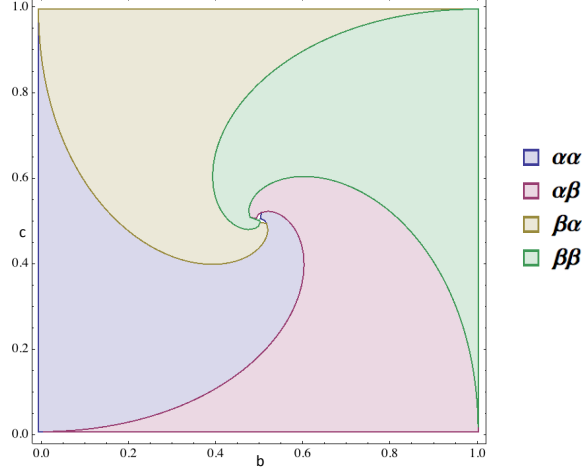


Figure 2: Regions in the square $b \times c$ leading to each of the four action profile as the sur-place action. For $b = c$ we get the matching pennies case of Example 2.

rather than the component game if q is large (Proposition 2). If one keeps $q = 1/2$, it is not a-priori obvious how R compares to V . The next lemma provides a more explicit formula for R in the case $q = 1/2$.

Lemma 6 *If $b, c \in [0, 1] \times [0, 1]$ and $q = 1/2$, the limit revision value is given by:*

$$R(b, c) = \frac{1}{4}(2b + 2c - 1) + \frac{1}{2}(b + c - 1)(b - c) \sin(2t_1^*) + \frac{1}{4}(2c - 1)(2b - 1) \cos(2t_1^*), \quad (8)$$

where t_1^* is the smallest t in \mathbb{R}_+ satisfying:

$$e^{-t} = \max\{(1 - 2c) \cos(t) + (1 - 2b) \sin(t), (1 - 2b) \cos(t) - (1 - 2c) \sin(t), \\ -(1 - 2b) \cos(t) + (1 - 2c) \sin(t), -(1 - 2c) \cos(t) - (1 - 2b) \sin(t)\}.$$

The value of the one-shot component game is $V = bc$. It is then possible to represent the values for b and c for which $R(b, c) > V$. This is illustrated panel (a) of Figure 3.

From Figures 2 and 3, one might read that Player 1 prefers the component game to the

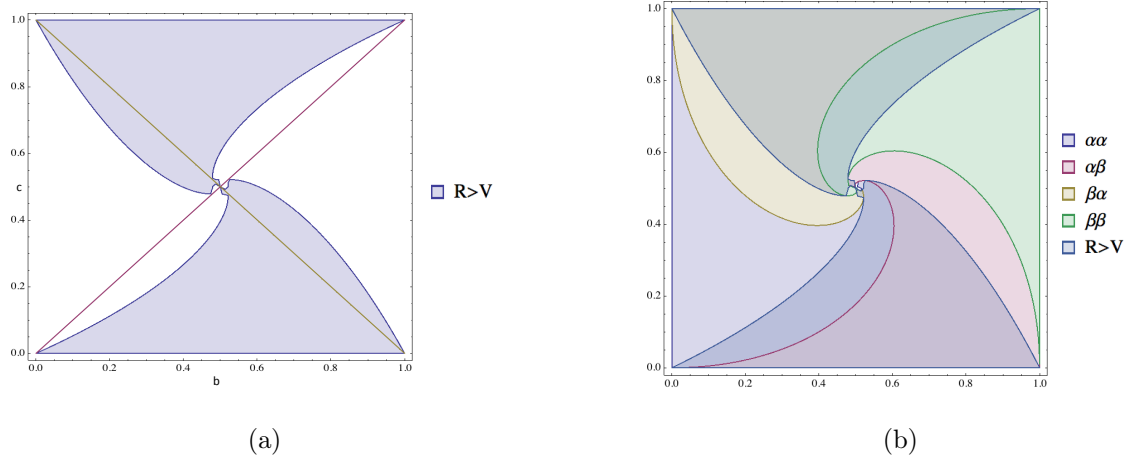


Figure 3: Panel (a): The shaded area represents the region of parameters (b, c) for which $R > V$. The lines are $c = b$ and $c = 1 - b$. Panel (b): the shaded area represents the region of parameters (b, c) for which $R > V$ and is superposed over the regions corresponding to the four sur-place actions.

revision game when the sur-place action is (α, α) or (β, β) . However, this is not necessarily the case as illustrated by panel (b) of Figure 3 that is obtained by superimposing panel (a) over Figure 2, there is no one-to-one relation between the preference of revision over component game, and the equilibrium strategy played in the revision game.

6 Conclusion

We have analyzed asynchronous zero-sum revision games and shown that the properties of equilibrium payoff and strategies qualitatively differ from those in non-zero-sum revision games studied in the literature. First, in zero-sum revision games the equilibrium payoff is continuous in all the parameters of the model and a player cannot lose from being faster than his opponent in revising. This applies both to the equilibrium payoff finite duration games, as well as for the limit of the equilibrium payoff as the duration goes to infinity. The opposite occurs for non-zero-sum revision games. Second, whereas in the games analyzed in Calcagno et al. (2014), all the activity occurs at the beginning of the revision phase, this is not necessarily the case for the

zero-sum case. If the component game equilibrium is mixed, then the revision game equilibrium displays a deadline effect where all the revision activity is concentrated toward the end of the revision phase.

7 Appendix

7.1 Proof of Proposition 3

We start with a preliminary result which shows the local uniqueness of equilibrium when the matrix of continuation payoff is generic. Let $u(t)$ be some 2×2 continuation payoff matrix. We have:

Lemma 7 *If $u(t)$ is regular, then there is $\varepsilon > 0$ such that in all equilibria, and for all $t' \in [t+\varepsilon, t]$,*

$$\sigma_i(x, t') = BR_i^{u(t)}(x_{-i})$$

Proof. The result comes from the fact that $u(t)$ is a continuous function in time. Since $u(t)$ is regular, for any x_{-i} there exists x_i^* such that $BR_i^{u(t)}(x_{-i}) = \{x_i^*\}$. Hence, there is $\delta > 0$ such that for all i , all x_i and all $y \neq x_i^*$, $u(t, x_i^*, x_{-i}) > u(t, y, x_{-i}) + \delta$. Fix $\varepsilon < \frac{\delta}{4\lambda\|U\|}$. Since $t \rightarrow u(t)$ is $2\lambda\|U\|$ -Lipschitz,

$$\begin{aligned} u(t', x_i^*, x_{-i}) &\geq u(t, x_i^*, x_{-i}) - 2\lambda\|U\|\varepsilon \geq u(t, y, x_{-i}) - 2\lambda\|U\|\varepsilon + \delta \\ &\geq u(t', y, x_{-i}) - 4\lambda\|U\|\varepsilon + \delta > u(t', y, x_{-i}), \end{aligned}$$

thus, x^* is the unique best response at t' . ■

Scenario DD: Without loss of generality, we can focus on the case $c, b > 1$ implying that the dominant actions in the component game are β and α for Player 1 and Player 2, respectively.

From Lemma 7, when the remaining time t is close to 0, these are also the actions prepared by the players. Thus, using backward induction, as the remaining time t increases, players continuation payoffs evolve according to the following differential system:

$$\begin{aligned}
u'(t, \alpha, \alpha) &= \lambda_1(u(t, \beta, \alpha) - u(t, \alpha, \alpha)) \\
u'(t, \alpha, \beta) &= \lambda_1(u(t, \alpha, \alpha) - u(t, \alpha, \beta)) + \lambda_2(u(t, \beta, \beta) - u(t, \alpha, \beta)) \\
u'(t, \beta, \alpha) &= 0 \\
u'(t, \beta, \beta) &= \lambda_2(u(t, \beta, \alpha) - u(t, \beta, \beta))
\end{aligned}$$

where derivative is taken with respect to the remaining time t . Solving this system with terminal condition $u(0, x) = U(x)$ for all $x \in X$, one finds that for any $t \geq 0$,

$$\begin{aligned}
u(t, \beta, \alpha) - u(t, \alpha, \alpha) &= ce^{-\lambda_1 t} > 0 \\
u(t, \beta, \beta) - u(t, \alpha, \beta) &= e^{-(\lambda_1 + \lambda_2)t}(ce^{\lambda_2 t} - 1) > 0 \\
u(t, \alpha, \beta) - u(t, \beta, \alpha) &= e^{-(\lambda_1 + \lambda_2)t}(1 + (b - 1)e^{\lambda_1 t}) > 0 \\
u(t, \beta, \beta) - u(t, \beta, \alpha) &= (b - 1)e^{-\lambda_1 t} > 0.
\end{aligned}$$

The first two inequalities, together with the equilibrium condition (1), guarantee that at any revision time t and no matter the last prepared action profile, player 1 prepares β . The latter two inequalities guarantee that at any revision time t and no matter the last prepared action profile, player 2 prepares α .

Scenario DN: Without loss of generality, we focus on the case $0 < b < 1$ and $c < 0$. This implies that for t close to 0, in equilibrium Player 1 prepares α , independently of the last prepared action, and Player 2 prepares α if and only if the last prepared action of Player 1 is α . Thus, for t close to 0 the differential system describing the evolution of continuation payoffs is,

$$\begin{aligned}
u'(t, \alpha, \alpha) &= 0 \\
u'(t, \alpha, \beta) &= \lambda_2(u(t, \alpha, \alpha) - u(t, \alpha, \beta)) \\
u'(t, \beta, \alpha) &= \lambda_1(u(t, \alpha, \alpha) - u(t, \beta, \alpha)) + \lambda_2(u(t, \beta, \beta) - u(t, \beta, \alpha)) \\
u'(t, \beta, \beta) &= \lambda_1(u(t, \alpha, \beta) - u(t, \beta, \beta)).
\end{aligned}$$

If $\lambda_1 \neq \lambda_2$, solving this system with terminal condition $u(0, x) = U(x)$, one finds that for any $t \geq 0$,

$$\begin{aligned}
u(t, \alpha, \alpha) - u(t, \beta, \alpha) &= -e^{-(\lambda_1 + \lambda_2)t} \\
&\quad + e^{-\lambda_1 t} \left(1 - c + b \frac{\lambda_2}{\lambda_1 - \lambda_2} (1 - e^{(\lambda_1 - \lambda_2)t}) \right) \\
u(t, \alpha, \beta) - u(t, \beta, \beta) &= e^{-\lambda_1 t} \left(1 - c + b \frac{\lambda_2}{\lambda_1 - \lambda_2} (1 - e^{(\lambda_1 - \lambda_2)t}) \right) \\
u(t, \alpha, \beta) - u(t, \alpha, \alpha) &= b e^{-\lambda_2 t} > 0 \\
u(t, \beta, \alpha) - u(t, \beta, \beta) &= e^{-(\lambda_1 + \lambda_2)t} (1 - b e^{\lambda_1 t}).
\end{aligned}$$

All these expressions are strictly positive for t close enough to 0. We are interested in the smallest t such that at least one of these expressions is nil. We want to show that such t is the one for which $u(t, \beta, \alpha) - u(t, \beta, \beta) = 0$.

Note first, that $u(t, \beta, \alpha) - u(t, \beta, \beta) = 0$ for $t = t^* := \frac{1}{\lambda_1} \ln \left(\frac{1}{b} \right) > 0$. Second, $u(t, \alpha, \beta) - u(t, \alpha, \alpha) > 0$ for all t . Third, note that $u(t, \alpha, \alpha) - u(t, \beta, \alpha) = u(t, \alpha, \beta) - u(t, \beta, \beta) - e^{-(\lambda_1 + \lambda_2)t}$, which implies $u(t, \alpha, \alpha) - u(t, \beta, \alpha) < u(t, \alpha, \beta) - u(t, \beta, \beta)$ for all t . Hence, it is sufficient to show that for $t \leq t^*$ one has $u(t, \alpha, \alpha) - u(t, \beta, \alpha) > 0$. This implies that when the remaining time is less than t^* , the revision strategies coincide with the component game best responses.

Define $f(t) := (u(t, \alpha, \alpha) - u(t, \beta, \alpha))e^{(\lambda_1 + \lambda_2)t}$, which has the same sign as $u(t, \alpha, \alpha) - u(t, \beta, \alpha)$.

One has,

$$f(t) = (1 - b - c)e^{\lambda_2 t} + b \frac{\lambda_1 e^{\lambda_2 t} - \lambda_2 e^{\lambda_1 t}}{\lambda_1 - \lambda_2} - 1.$$

Note that $f(0) = -c > 0$ and $f'(0) = \lambda_2(1 - b - c) > 0$. Note also that $f'(t)$ has the same sign as $\frac{1}{\lambda_2} f'(t) e^{-\lambda_2 t} = b \frac{\lambda_1}{\lambda_1 - \lambda_2} (1 - e^{(\lambda_1 - \lambda_2)t}) + 1 - b - c$ which is positive for all $t > 0$ if $\lambda_2 \geq \lambda_1 \frac{1-c}{1-c-b}$. For $\lambda_2 < \lambda_1 \frac{1-c}{1-c-b}$, $\frac{1}{\lambda_2} f'(t) e^{-\lambda_2 t}$ is positive for $t < \log\left(\frac{\lambda_1(1-c) - \lambda_2(1-c-b)}{\lambda_1 b}\right) / (\lambda_1 - \lambda_2)$ and negative otherwise. In summary, the function $f(t)$ has the same sign as $u(t, \alpha, \alpha) - u(t, \beta, \alpha)$. This latter expression is positive at $t = 0$, and either increasing in t (for $\lambda_2 \geq \lambda_1 \frac{1-c}{1-c-b}$), or hump-shaped (for $\lambda_2 < \lambda_1 \frac{1-c}{1-c-b}$). Thus, it is sufficient to show that $u(t^*, \alpha, \alpha) - u(t^*, \beta, \alpha) > 0$. Note that,

$$\begin{aligned} u(t^*, \alpha, \alpha) - u(t^*, \beta, \alpha) &= b \left(\frac{\lambda_1}{\lambda_2 - \lambda_1} (b^{\lambda_2/\lambda_1} - b) + 1 - b - c \right) \\ &> b \left(\frac{\lambda_1}{\lambda_2 - \lambda_1} (b^{\lambda_2/\lambda_1} - b) + 1 - b \right), \end{aligned}$$

where the first inequality follows from $c < 0$. The last expression is positive if $\lambda_2 > \lambda_1$ and has the same sign as $\phi(\lambda_2/\lambda_1) := 1 - b^{\lambda_2/\lambda_1} - \frac{\lambda_2}{\lambda_1}(1 - b)$ which is strictly positive for $\lambda_2/\lambda_1 \in (0, 1)$. In fact, this function is zero for $\lambda_2 = 0$ and $\lambda_2 = \lambda_1$, and strictly concave because $\phi''(\lambda_2/\lambda_1) = -b^{\lambda_2/\lambda_1} \log(b)^2 < 0$.

If $\lambda_1 = \lambda_2 = \lambda/2$, solving this system with the terminal condition $u(0, x) = U(x)$, one finds that for any $t \geq 0$,

$$\begin{aligned} u(t, \alpha, \alpha) - u(t, \beta, \alpha) &= -e^{-\lambda t} + e^{-\lambda t/2} \left(1 - c - b \frac{\lambda t}{2}\right) \\ u(t, \alpha, \beta) - u(t, \beta, \beta) &= e^{-\lambda t/2} \left(1 - c - b \frac{\lambda t}{2}\right) \\ u(t, \alpha, \beta) - u(t, \alpha, \alpha) &= b e^{-\lambda t/2} > 0 \\ u(t, \beta, \alpha) - u(t, \beta, \beta) &= e^{-\lambda t} (1 - b e^{\lambda t/2}). \end{aligned}$$

Applying exactly the same method, define $f(t) := (u(t, \alpha, \alpha) - u(t, \beta, \alpha))e^{\lambda t}$, which has the same sign as $u(t, \alpha, \alpha) - u(t, \beta, \alpha)$. One has,

$$f(t) = e^{\lambda t/2} \left(1 - c - b \frac{\lambda t}{2}\right) - 1.$$

Note that $f(0) = -c > 0$ and,

$$f'(t) = \frac{\lambda}{2} e^{\lambda t/2} \left(1 - c - b - b \frac{\lambda t}{2}\right).$$

f' is positive for $t < \frac{2(1-c-b)}{b\lambda}$ and negative otherwise. From this, we deduce that $u(t, \alpha, \alpha) - u(t, \beta, \alpha)$ is hump-shaped. Thus, it is sufficient to show that $u(t^*, \alpha, \alpha) - u(t^*, \beta, \alpha) > 0$. Note that,

$$\begin{aligned} u(t^*, \alpha, \alpha) - u(t^*, \beta, \alpha) &= b(1 - c + b \ln(b) - b) \\ &> b(1 + b \ln(b) - b) > 0 \end{aligned}$$

where the first inequality follows from $c < 0$. To prove the last inequality, note that $1 + b \ln(b) - b$ is decreasing and equals 0 for $b = 1$.

For all pairs (λ_1, λ_2) , $u(\cdot, \cdot)$ solves (2) and is continuously differentiable. One may therefore compute the derivative $u'(t^*, x)$ and we have:

$$u(t^*, \beta, \alpha) - u(t^*, \beta, \beta) = 0,$$

$$u'(t^*, \beta, \alpha) - u'(t^*, \beta, \beta) = \lambda_1(u(t^*, \alpha, \alpha) - u(t^*, \beta, \alpha) - u(t^*, \alpha, \beta) - u(t^*, \beta, \beta)) = -\lambda_1 e^{-(\lambda_1 + \lambda_2)t^*} < 0.$$

This implies that there exists δ such that for $t \in (t^*, t^* + \delta)$, we have,

$$u(t, \alpha, \alpha) - u(t, \beta, \alpha) > 0$$

$$u(t, \alpha, \beta) - u(t, \beta, \beta) > 0$$

$$u(t, \alpha, \beta) - u(t, \alpha, \alpha) > 0$$

$$u(t, \beta, \alpha) - u(t, \beta, \beta) < 0.$$

We deduce that $u(t)$ is in scenario DD and we may apply the preceding analysis to conclude that,

$$R = u(t, \alpha, \alpha) = u(t^*, \alpha, \alpha) = u(\alpha, \alpha) = V,$$

where the second equality follows from the continuity of u , the third from the differential system satisfied by u on $[0, t^*]$.

Scenario NN: Without loss of generality we focus on the case $0 < b, c < 1$. This implies that for t close to 0, in equilibrium, Player 1 prepares an action different from Player 2's last prepared action, whereas Player 2 prepares the same action as the last prepared by Player 1. Thus, for t close to 0, the differential system describing the evolution of continuation payoffs is,

$$u'(t, \alpha, \alpha) = \lambda_1(u(t, \beta, \alpha) - u(t, \alpha, \alpha))$$

$$u'(t, \alpha, \beta) = \lambda_2(u(t, \alpha, \alpha) - u(t, \alpha, \beta))$$

$$u'(t, \beta, \alpha) = \lambda_2(u(t, \beta, \beta) - u(t, \beta, \alpha))$$

$$u'(t, \beta, \beta) = \lambda_1(u(t, \alpha, \beta) - u(t, \beta, \beta)).$$

If $\theta := |\lambda_1^2 - 6\lambda_1\lambda_2 + \lambda_2^2|^{1/2} \neq 0$, solving this system with terminal condition $u(0, x) = U(x)$, one finds that for t close to 0:

$$u(t, \alpha, \alpha) = \frac{1}{2} \left(b + c - \frac{\lambda_1 e^{-(\lambda_1 + \lambda_2)t} + \lambda_2}{\lambda_1 + \lambda_2} - \eta(b + c - 1, b - c, \lambda_1, \lambda_2) \right) \quad (9)$$

$$u(t, \alpha, \beta) = \frac{1}{2} \left(b + c - \frac{\lambda_2(1 - e^{-(\lambda_1 + \lambda_2)t})}{\lambda_1 + \lambda_2} + \eta(b - c, b + c - 1, \lambda_2, \lambda_1) \right) \quad (10)$$

$$u(t, \beta, \alpha) = \frac{1}{2} \left(b + c - \frac{\lambda_2(1 - e^{-(\lambda_1 + \lambda_2)t})}{\lambda_1 + \lambda_2} - \eta(b - c, b + c - 1, \lambda_2, \lambda_1) \right) \quad (11)$$

$$u(t, \beta, \beta) = \frac{1}{2} \left(b + c - \frac{\lambda_1 e^{-(\lambda_1 + \lambda_2)t} + \lambda_2}{\lambda_1 + \lambda_2} + \eta(b + c - 1, b - c, \lambda_1, \lambda_2) \right), \quad (12)$$

with,

$$\eta(A, B, \lambda_1, \lambda_2, t) := e^{-\frac{\lambda_1 + \lambda_2}{2}t} \left(A \cos\left(\frac{\theta t}{2}\right) + \frac{(\lambda_2 - \lambda_1)A + 2\lambda_1 B}{\theta} \sin\left(\frac{\theta t}{2}\right) \right).$$

This implies that,

$$u(t, \beta, \alpha) - u(t, \alpha, \alpha) = \psi(2c - 1, 2b - 1, \lambda_1, \lambda_2, t) \quad (13)$$

$$u(t, \alpha, \beta) - u(t, \beta, \beta) = \psi(1 - 2c, 1 - 2b, \lambda_1, \lambda_2, t) \quad (14)$$

$$u(t, \alpha, \beta) - u(t, \alpha, \alpha) = \psi(2b - 1, 1 - 2c, \lambda_2, \lambda_1, t) \quad (15)$$

$$u(t, \beta, \alpha) - u(t, \beta, \beta) = \psi(1 - 2b, 2c - 1, \lambda_2, \lambda_1, t), \quad (16)$$

with,

$$\psi(A, B, \lambda_1, \lambda_2, t) := \frac{e^{-\frac{\lambda_1 + \lambda_2}{2}t}}{2} \left(e^{-\frac{\lambda_1 + \lambda_2}{2}t} + A \cos\left(\frac{\theta t}{2}\right) + \frac{(\lambda_2 - \lambda_1)A + 2\lambda_2 B}{\theta} \sin\left(\frac{\theta t}{2}\right) \right).$$

If $\theta = 0$ (*i.e.* $\lambda_1 = (3 \pm \sqrt{8})\lambda_2$), then one has to replace the above expression by,

$$\psi(A, B, \lambda_1, \lambda_2, t) := \frac{e^{-\frac{\lambda_1 + \lambda_2}{2}t}}{2} \left(e^{-\frac{\lambda_1 + \lambda_2}{2}t} + A + \frac{(\lambda_2 - \lambda_1)A + 2\lambda_2 B}{2}t \right).$$

Let t_1^* be the smallest $t > 0$ such that at least one of expressions (13)-(16) is nil.

For generic values of b, c , the time t_1^* is finite and only one of expression (13)-(16) is nil at t_1^* . Without loss of generality suppose that $u(t_1^*, \beta, \alpha) - u(t_1^*, \beta, \beta) = 0$. Then, for a remaining time $t > t_1^*$ and close to t_1^* , the continuation payoff matrix is in the scenario DN analyzed above and the same conclusion apply. That implies that there is $t_2^* > t_1^*$ such that for $t > t_2^*$, Players 1 and 2 prepare β and α , respectively. For $t_1^* < t < t_2^*$, Player 1 prepares the action that differs from Player 2's current action, whereas Player 2 prepares α , no matter Player 1's current action. When the remaining time is less than t_1^* , the revision strategies coincide with the component game best responses. This ends the proof.

7.2 Proof of Lemma 5

The first part of the Lemma follows from the proof for scenario NN in Proposition 3. To see that $R = u(t_1^*, \hat{x})$, observe that for a remaining time t between t_2^* and t_1^* , the continuation payoff matrix $u(t)$ is in scenario DN so that, as long as the prepared action profile is \hat{x} , players do not revise their actions. For a remaining time t larger than t_2^* , $u(t)$ is in scenario DD where the dominant actions are \hat{x} . Hence, for any $t > t_1^*$ one has $u(t, \hat{x}) = u(t_1^*, \hat{x})$. Observe that no matter the starting action profile, as t goes to infinity, the probability that the prepared action profile becomes \hat{x} before t_2^* , goes to 1. Hence $\lim_{t \rightarrow \infty} u(t, x) = u(t_1^*, \hat{x})$, for all $x \in X$.

7.3 Proof of Lemma 6

Observe that (c, b) are such that the component game is in scenario *NN*. For t close to 0, continuation payoffs are given by expressions (9)-(12). For $\lambda_1 = \lambda_2 = 1$, we have $\theta = 2$, and

hence:

$$\begin{aligned}
u(t, \alpha, \alpha) &= \frac{1}{2} \left(b + c - \frac{e^{-2t} + 1}{2} - e^{-t} ((b + c - 1) \cos(t) + (b - c) \sin(t)) \right), \\
u(t, \alpha, \beta) &= \frac{1}{2} \left(b + c - \frac{1 - e^{-2t}}{2} + e^{-t} ((b - c) \cos(t) + (b + c - 1) \sin(t)) \right), \\
u(t, \beta, \alpha) &= \frac{1}{2} \left(b + c - \frac{1 - e^{-2t}}{2} - e^{-t} ((b - c) \cos(t) + (b + c - 1) \sin(t)) \right), \\
u(t, \beta, \beta) &= \frac{1}{2} \left(b + c - \frac{e^{-2t} + 1}{2} + e^{-t} ((b + c - 1) \cos(t) + (b - c) \sin(t)) \right).
\end{aligned}$$

Time t_1^* is the smallest $t > 0$ such that at least one of expressions (13)-(16) is nil. Denote,

$$\begin{aligned}
h_{\alpha, \alpha}(t) &:= (1 - 2c) \cos(t) + (1 - 2b) \sin(t) \\
h_{\alpha, \beta}(t) &:= (1 - 2b) \cos(t) - (1 - 2c) \sin(t) \\
h_{\beta, \alpha}(t) &:= -(1 - 2b) \cos(t) + (1 - 2c) \sin(t) \\
h_{\beta, \beta}(t) &:= -(1 - 2c) \cos(t) - (1 - 2b) \sin(t).
\end{aligned}$$

Then, for $\lambda_1 = \lambda_2 = 1$, finding t_1^* boils down to finding the smallest t in \mathbb{R}_+ satisfying:

$$e^{-t} = \max\{h_{\alpha, \alpha}(t), h_{\alpha, \beta}(t), h_{\beta, \alpha}(t), h_{\beta, \beta}(t)\}. \quad (17)$$

Whether the sur-place action profile \hat{x} is (α, α) , (α, β) , (β, α) or (β, β) , depends on whether at t_1^* , the l.h.s. of (17) equals the first, the second, the third or the fourth argument of the max operator on the r.h.s. In the non-generic case where the argmax of the l.h.s. is not a singleton, then $\hat{x} \in A(b, c) := \arg \max_{x \in X} h_x(t_1^*)$. Figure 2 is obtained by letting (b, c) vary in $[0, 1] \times [0, 1]$, numerically solving for t_1^* , and hence finding \hat{x} . The frontier between two regions corresponds to the non-generic cases where two sur-place actions (and hence two payoff equivalent equilibria) are possible.

Let us assume that $\hat{x} = (\alpha, \alpha)$, so that $R = u(t_1^*, \alpha, \alpha)$. Substituting $e^{-t_1^*} = (1 - 2c) \cos(t_1^*) + (1 - 2b) \sin(t_1^*)$, into the above expression of $u(t, \alpha, \alpha)$ we obtain,

$$\begin{aligned} R &= \frac{1}{2} \left[b + c - \frac{1}{2} + \frac{1}{2}(1 - 2c)(1 - 2b) \cos^2(t_1^*) - \frac{1}{2}(1 - 2c)(1 - 2b) \sin^2(t_1^*) \right. \\ &\quad \left. + 2(b + c - 1)(b - c) \sin(t_1^*) \cos(t_1^*) \right] \\ &= \frac{1}{2} \left[b + c - \frac{1}{2} + \frac{1}{2}(1 - 2c)(1 - 2b) \cos(2t_1^*) + (b + c - 1)(b - c) \sin(2t_1^*) \right]. \end{aligned}$$

Doing the same computation for the other profiles $\hat{x} = (\alpha, \beta), (\beta, \alpha), (\beta, \beta)$ leads to the same result which proves the formula.

So far, we have proved the formula only for generic payoff matrices. In order to complete the proof and extend it to all matrices, it is enough to prove that r.h.s. of equality (8) is continuous in (b, c) , since we know that $R(b, c)$ is continuous. Denote $t_1^*(b, c)$ the smallest t in \mathbb{R}_+ satisfying $e^{-t} = \max_{x \in X} h_x(t)$. We argue now that $t_1^*(b, c)$ is continuous in (b, c) .

The function $t_1^*(b, c)$ is lower semi-continuous (as the minimum of finitely many lower semi-continuous functions). Let us prove that it is upper semi-continuous on $[0, 1]^2 \setminus \{(1/2, 1/2)\}$.

Claim 1 For $(b, c) \neq (1/2, 1/2)$, $t_1^* < \infty$ and there exists $x \in A(b, c)$ such that $-e^{-t_1^*} < h'_x(t_1^*)$.

Proof. Take $x \in A(b, c)$. The map $g(t) = e^{-t} - h_x(t)$ is positive on $[0, t_1^*)$ and equal to zero at t_1^* . Thus, its derivative at t_1^* is non-positive and $-e^{-t_1^*} \leq h'_x(t_1^*)$. If $g'(t_1^*) = 0$, then $-e^{-t_1^*} = h'_x(t_1^*) = h'_y(t_1^*)$ for some $y \neq x$. This implies that $y \in A(b, c)$ and $-e^{-t_1^*} < h'_y(t_1^*)$ which proves the claim. ■

Now, if $x \in A(b, c)$ is such that $-e^{-t_1^*} < h'_x(t_1^*)$, then using this claim, for all $\varepsilon > 0$ sufficiently small, we have $e^{-(t_1^* - \varepsilon)} - h_x(t_1^* - \varepsilon) < 0$ and $e^{-(t_1^* + \varepsilon)} - h_x(t_1^* + \varepsilon) > 0$. Since h_x is continuous with

respect to (b, c) , we deduce that:

$$\limsup_{(b', c') \rightarrow (b, c)} t_1^*(b', c') \leq t_1^* + \varepsilon,$$

which proves that t_1^* is upper semi-continuous, by sending ε to zero.

Clearly, $t_1^*(b, c)$ goes to $+\infty$ when (b, c) goes to $(1/2, 1/2)$ and we deduce that the r.h.s. of (8) is continuous on $[0, 1]^2$ as a function of (b, c) .

References

- [1] Alesina A. and A. Drazen “Why are Stabilizations Delayed?”, *The American Economic Review*, 1991, 81(5), 1170–118.
- [2] Biais, Bruno, Pierre Hillion, and Chester Spatt, “Price Discovery and Learning during the Preopening Period in the Paris Bourse”, *The Journal of Political Economy*, 1999, 107(6), 1218–1248.
- [3] Calcagno, R., S. Lovo, Y. Kamada and T. Sugaya, “Asynchronicity and Coordination in Common and Opposing Interest Games”, *Theoretical Economics*, 2014, 9(2), 409–434.
- [4] Kamada, Y. and M. Kandori, “Revision Games”, working paper, 2013.
- [5] Lovo, S. and T. Tomala, “Markov Perfect Equilibria in Stochastic Revision Games”, 2015, Available at SSRN: <http://ssrn.com/abstract=2622084> or <http://dx.doi.org/10.2139/ssrn.2622084>.
- [6] Simsek and A. and M. Yildiz, “Durability, Deadline and Election Effects in Bargaining”, working paper, 2014.

- [7] Shapley, L. S., “Stochastic games”, Proceedings of the national academy of sciences, 1953, 39(10), 1095–1100.
- [8] Spier K. “The Dynamics of Pretrial Negotiation”, The Review of Economic Studies, 1992, 59(1), 93–108.
- [9] Vives X. and L.A. Medrano, “Strategic Behavior and Price Discovery”, The RAND Journal of Economics, 2001, 32(2), 221–248.