Learning in Games Where Agents Sample

Eduardo García-Echeverri*

Abstract

This paper proposes an equilibrium concept –the Sampling Bayesian Equilibrium– for games in which players observe the actions of only a small random sample of other players. I show the existence of these equilibrium points for *i)* a class of coordination of global games and for *ii)* general static games in normal form. For the first, I further show the existence of a unique interior Sampling Bayesian Equilibrium, easing comparative statics over the set of equilibria. Using asymptotic Bayesian analysis, in particular Bernstein- von Mises theorem, I show that most equilibrium points in the complete information games (where agents have perfect foresight over the actions of all other players) can be obtained as limits of pure-strategy Sampling Bayesian Equilibria of the perturbed games, as agents learn and sample sizes tend to infinity. These purification results are robust to a wide class of prior distributions over strategy profiles and are consistent with Nash's 'mass-action' interpretation of mixed strategies.

JEL Classification: C11, C72, D74, D82 and D83.

Keywords: *Sampling Bayesian Equilibrium; Global Games; Endogenous Information; Purification.*

*This Version: July 28, 2021. García-Echeverri: Department of Economics, University of Rochester, Harkness Hall, P.O. Box: 270156, Rochester, NY 14627, Email: egarc12@ur.rochester.edu. I would like to thank Paulo Barelli, Neşe Yildiz, Yu Awaya, Srihari Govindan, Narayana Kocherlakota, Yiran Wang, Richard Grice, Roman Merga, Stepan Gordeev, Paulo Lins, Marcos Ceron, Kaixin Liu, Jose Villegas, Chen Kan, and the audience at the Student Seminar for their suggestions and useful comments.

1. Introduction

An implicit assumption in most solution concepts in game theory –the Nash equilibrium included– is that each player has perfect foresight on the equilibrium behaviour of all other players and chooses her strategy accordingly. The process by which agents learn and calibrate their beliefs on the actions of others is typically left unspecified. In many contexts, such as games with an uncountable infinite set of players, the assumption of perfect foresight is debatable. The aim of this paper is to propose a solution concept – the Sampling Bayesian Equilibrium– for games in which players are assumed rational but have limited information, only getting to know the actions of a finite sample of other players.

Analysing this class of games with limited information provides an interpretation for mixed strategy equilibria, as will be shown below. Two aspects of this type of equilibria have puzzled game theorists since the origins of the field (see von Neumann and Morgenstern (1947); and Nash (1950)). First, players rarely appear to be randomizing in practise. Second, it is not clear why they would choose precisely the actions that make other players indifferent, given that they can deviate without penalty from their equilibrium strategies. By modelling the learning process as sample sizes growing to infinity, I show that most equilibrium points in the complete information games (in pure or mixed strategies) can be obtained as limits of pure-strategy Sampling Bayesian Equilibria (SBE herein thereafter) of the games with limited information. I do so in the context of *i)* coordination global games –specifically a regime switch model a la Morris and Shin (1999)– and *ii)* general static games in normal form played by a continuum of agents.

In both of these classes of games, agents are interested in learning (by means of statistical inference) about the proportion of agents in the continuum playing each of the different actions. To do so, they use the empirical frequencies of the actions observed in their samples to update their beliefs –via Bayesian Updating– and then choose their strategies to maximize their expected payoffs. For each individual player, I'll refer to the action profile observed in her sample as her *private signal*. Many economic strategic interactions in which agents have incentives to coordinate –such as bank runs, attacks on currency pegs, or models of investment with complementarities– can be modelled under this endogenous information structure in which agent's actions affect the information (signals) that agents receive.

The equilibrium concept for this class of games has to be adapted in order to account for the dual causality between the strategy profiles and the distribution of the *private signals.* I impose a *coherency* requirement in the solution of the game: given a strategy profile there should be a particular distribution of the private signals, and in turn, given this distribution of signals, the strategy profile should be consistent with agents optimizing their expected payoffs. I define a strategy profile in which these two requirements are met a Sampling Bayesian Equilibrium. I show the existence of such equilibria for *i)* a regime switch model a la Morris and Shin (1999), and *ii)* general static games in normal form played by a continuum of agents.

Finally, to model the learning process of the agents, I analyze the asymptotic behaviour of SBE as the sample size of the *private signals* tends to infinity. This gives rise to the two purfication results presented in this paper. First, for the regime switch model, I show that the distribution of actions in the unique interior SBE converges to that of the mixed- strategy Nash equilibria of the game with complete information. Moreover, in this unique interior SBE, all players are using pure strategies. Second, I show that any equilibria (in mixed or pure strategies) of any static 2×2 game, in which no player uses a weakly dominated strategy, can be obtained as the limit of a sequence of SBE, as the sample size tends to infinity and common knowledge is asymptotically restored.

1.0.1. Related Literature: This paper is related to four streams of literature. First, it contributes to the literature on global games with endogenous information structures, as in Angeletos et al. (2006), Tarashev (2003); Tsyvinski et al. (2004); and Morris and Shin (2006). These authors have focused on endogenizing public signals, making their distributions depend on the actions of players. To the best of my knowledge, such attempts have not been tried with private signals. Endogenizing such signals via the sampling procedure introduced, gives a theoretical foundation to the assumption of a normal distribution, very common in the literature.

Second, it contributes to the literature on purification of mixed strategies that began with Harsanyi (1973) and then further developed with Radner and Rosenthal (1982), Aumann et al. (1983), Bhaskar et al. (2008), Govindan et al. (2003), and Barelli and Duggan (2015), among many others. It does so by introducing techniques of Bayesian statistics, specifically Bernstein-von Mises theorem, that allow to study the asymptotic distribution of posterior beliefs derived from Bayesian updating. This result is then used to derive the purification results in the context of an endogenous information structure in which agent's actions affect the distribution of types.

Third, this paper is related to the literature of learning in games– thoroughly reviewed in Fudenberg and Levine (2009). In particular, it is related to the literature studying the convergence of (stochastic) fictitious play and calibrated algorithms to Nash or correlated equilibria (Fudenberg and Kreps (1993); Foster and Vohra (1997)). In these models, convergence (or divergence) is analyzed by letting the number of times the stage game is repeated grow to infinity. Sample size is kept constant in each period and typically equals one. In contrast, convergence to Nash equilibria in my model is not longitudinal (as $T \to \infty$) but cross-sectional (as sample size $N \to \infty$). This means that I study how individuals learn about the actions of the whole continuum as they observe more and more different individuals in their sample.

Fourth and finally, this paper contributes to the literature of games with sampling procedures, developed by Osborne and Rubinstein (2003) for agents with bounded rationality. The solution concept proposed in this paper –the SBE– borrows the *coherency* requirement of Osborne and Rubinstein's Sampling Equilibrium. In both solution concepts, a given strategy profile induces a random sample result, and given this random sample result, the strategy profile should be consistent with agents maximizing their expected payoffs. In a SBE, however, agents are fully rational and thus aware of the approximation error in their samples. They updadte their priors via Bayesian updating.

The rest of this paper is organized as follows. Section 2 presents a coordination global game with endogenous private signals and studies convergence of SBE in this context. Specifically, Section 2.1 presents the setup of the regime switch model (as in Angeletos and Werning (2006)), the endogenous information structure, and the signal generating process. Sections 2.2 and 2.3 define strategy profiles and SBE, and proof the existence of the latter for this class of games. The asymptotic analysis of SBE can be found in Section 2.4. Section 3 defines SBE for static games played by a continuum of agents, shows their existence for this class of games, and studies their convergence in the context of 2×2 games. Finally, the proofs of all theorems and propositions can be found in Appendix A.

2. A Coordination Global Game where Agents Sample

As in the regime switch model in Morris and Shin (1999), suppose there is a status-quo and a continuum of agents with measure unity. Agents are indexed by $i \in [0, 1]$ and have a binary action set: to attack the status-quo, $a_i = 1$, or not to attack it, $a_i = 0$. For simplicity, payoff of not attacking is normalized to zero. Let $C \in (0, 1)$ be the utility cost of attacking, θ the exogenous strength of the status-quo, and A the mass of agents attacking. Payoff from attacking is $1-C$ if the regime is overthrown, i.e. $A > \theta$, and $-C$ otherwise. Thus, the payoff of agent i can be expressed as:

$$U(a_i, A, \theta) = a_i(1\{A > \theta\} - C), \tag{2.1}$$

where $1\{A > \theta\}$ is an indicator function that equals 1 under the event that $A > \theta$, and equals 0 otherwise. The weak monotonicity of $U(a_i, A, \theta)$ with respect to A incentivizes agents to coordinate and gives rise to multiple Nash equilibria.

Many strategic interactions of economic interest can be modelled as regime switch models. In a bank run, for instance, as more people withdraw their deposits the risk of default raises, pushing even more people to withdraw their funds. Eventually, reserves may be insufficient to cover all withdrawals and the bank enters into default, triggering a regime change (Goldstein and Pauzner (2005), Rochet and Vives (2004), and Cañón and Margaretic (2014)). In a similar fashion, in models of self-fulfilling currency crises, central banks are forced to abandon their pegs when a sufficiently large number of speculators attack the currency (Obstfeld (1986), Obstfeld (1996), and Morris and Shin (1998)). Other applications include investments with complementarities, where the project's profitability depends on the number of investors (Dasgupta (2007)), debt crises (Cole and Kehoe (2000)) and political revolutions against autocratic regimes (Edmond (2005)).

In the complete information benchmark, when θ is common knowledge and agents have perfect foresight on the action profile of the entire continuum, then both $A = 0$ and $A = 1$ arise as the only Nash equilibria of the game whenever $\theta \in (0, 1)$. When knowing that all the rest of the agents are attacking, the unique best response for agent i is $a_i = 1$. Similarly, when knowing no other agent is attacking the opposite happens, and the unique best response for agent i is $a_i = 0$.

2.1. Information Structure: Information on the regime's strength θ is assumed imperfect. At the beginning of the game, nature draws θ from an exogenous publicly-known distribution P_θ, which constitutes the agents' common prior. Suppose that this prior satisfies the support condition $[0, 1] \subset \text{supp}[P_\theta]$. After θ is realized, agents receive a noisy public signal $Y = \theta + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon)$ is a random error term independent of θ. According to this signal, agents update their beliefs following Bayes' rule and derive their (common) posterior beliefs on the regime's strength $G_{\theta|Y}$.[1]

Suppose also that agents don't have perfect foresight on the action profile of the entire continuum, but instead only observe the actions of a finite sample of other players. Suppose all agents observe exactly n other agents, chosen randomly from the continuum. Let X_i be the number of attackers observed in agent i's sample. We will refer to X_i as agent's i private signal. A strategy is thus defined as a mapping $s : \{0, 1 \cdots n\} \to \{0, 1\}$ that indicates the agent what action to choose for each possible realization of her private signal. A strategy profile is a collection $\{s_i\}_{i \in [0,1]}$ of such mappings.

2.1.1. Timing of the Game For each individual player, taking the actions of all other players as given, the timing of the game is the following:

i) Exogenous public signal Y and cost C are observed. The posterior belief on θ, $G_{\theta|Y}$, is formed.

ii) Based on C and Y, agent forms a prior belief on the size of the attack, $p_0(A|Y, C)$.[2]

iii) The agent's sample of other players is randomly chosen and private signal X_i is observed. Posterior beliefs on the size of the attack, $p_i(A|Y, C, X_i)$, are derived via Bayesian updating.

iv) Based on her public signal X_i, the agent decides whether to attack or not.

[1] In the global games literature, it is common to take the improper uniform distribution over the entire real line as the prior for θ. Under this assumption, the posterior belief $G_{\theta|Y}$ has density $\frac{1}{\sigma_\varepsilon}\phi\left(\frac{\theta - Y}{\sigma_\varepsilon}\right)$, where $\phi(\cdot)$ is the probability density function of the standard normal distribution.

[2] Suppose that the prior $p_0(A|Y, C)$ is non-degenerate and absolutely continuous with respect to the Lebesgue measure.

2.2. Optimal Threshold Strategies: Under the assumptions on the information structure, we will proceed to show that if agents are maximizing their expected payoffs, then the strategy profile is symmetric (i.e. $s_i = s_j$ for all $i, j \in [0,1]$) and it follows a threshold criterion. Agents attack if an only if their private signal is realized above a threshold that depends on public signal Y, cost C, and sample size n.

If observations in the sample are independently chosen and noting that the probability of observing an attacker from a single observation is A, then private signal X_i follows a binomial distribution with n trials and a success probability of A, i.e. $X_i \sim B(n, A)$. According to Bayes' rule, an agent with private signal X_i has a posterior belief:

$$p_i(A|Y,C,X_i) = \frac{\Pr[X_i|A] \cdot p_0(A|Y,C)}{\int_0^1 \Pr[X_i|A] \cdot p_0(A|Y,C)\,\mathrm{d}A} = \frac{A^k(1-A)^{n-k} \cdot p_0(A|Y,C)}{\int_0^1 A^k(1-A)^{n-k} \cdot p_0(A|Y,C)\,\mathrm{d}A}.$$

Thus, the subjective belief of the status-quo being overthrown, for an agent with private signal X_i, is given by:

$$\Pr[A > \theta|Y,C,X_i] = \int\int_\theta^1 p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y}.$$

As shown in Proposition 1, observing a higher number of attackers shifts the posterior distribution over A to the right.

Proposition 1: If $X_i > \tilde{X}_i$, then $p_i(A|Y,C,X_i) \succsim_{\text{FSD}} p_i(A|Y,C,\tilde{X}_i)$.[3]

Consequently, the expected payoff of attacking, $\Pr[A > \theta|Y,C,X_i] - C$, is strictly increasing in X_i. Since an agent attacks if and only if this expected payoff is strictly positive, then it follows that rational agents must use a threshold strategy.

Proposition 2: $a_i = 1 \Leftrightarrow X_i > \bar{X}_n(Y,C) \equiv \sup\big\{\bar{X} \in \{0,1,...n\}|\Pr[A > \theta|Y,C,\bar{X}] \leq C\big\}$.

2.3. Sampling Bayesian Equilibrium: Endogenous private signals in the coordination global game introduce circularity between action profiles and beliefs. An aggregate attack, A, induces a particular binomial distribution of private signals, $X_i \sim B(n, A)$. These in turn, according to the threshold strategy derived in Proposition 2, determine the proportion of agents attacking.

[3] Where $\succsim_{\text{FSD}}$ represents First Order Stochastic Dominance. I.e. let $f_1(x)$ and $f_2(x)$ be two probability densities, then $f_1(x) \succsim_{\text{FSD}} f_2(x)$ if and only if for all α, $\int_{-\infty}^{\alpha} f_1(x)\,\mathrm{d}x \leq \int_{-\infty}^{\alpha} f_2(x)\,\mathrm{d}x$, with strict inequality for at least one α.

As part of a solution concept for this type of games, I impose a requirement of *coherency*[4] between the aggregate attack and the information structure that arises from it. Specifically, in equilibrium, the aggregate attack must equal the probability that the private signal realizes above the threshold $\bar{X}_n(Y, C)$, given that $X_i \sim B(n, A)$, i.e.:

$$A = \Pr[X_i > \bar{X}_n(Y, C)] \tag{2.2}$$

Equation 2.2 guarantees that the size of the attack generates the specific distribution of private signals that generates exactly that same size of attack. I call a strategy profile and an aggregate attack that satisfies this *coherency* requirement a Sampling Bayesian Equilibrium (abbreviated SBE herein thereafter). Formally,

Definition [Sampling Bayesian Equilibrium]: Let $s \equiv \big\{s_i : \{0, 1 \cdots n\} \to \{0, 1\}\big\}_{i \in [0,1]}$ be a strategy profile and let $A \in [0, 1]$ be an aggregate size of attack. Given public signal Y and cost C, (s, A) is a SBE if and only if:

1. For all i, $s_i(x_i) = 1$ if and only if $x_i > \bar{X}_n(Y, C)$;

2. $A = \Pr[X_i > \bar{X}_n(Y, C)]$, where $X_i \sim B(n, A)$.

The first condition is simply a requirement that agents behave optimally by maximizing their subjective expected payoffs after observing their respective private signal. The second condition imposes *coherency* in the system.

To find all the SBE we must find the aggregate attacks that satisfy Equation 2.2. This equation can be re-expressed replacing the cumulative binomial distribution in the right hand side by a Regularized Incomplete Beta Function, as shown in Equation 2.3.

$$A = 1 - \Pr[X_i \leq \bar{X}_n(Y, C)] = 1 - (n - \bar{X}_n)\binom{n}{\bar{X}_n} \int_0^{1-A} t^{n-\bar{X}_n-1}(1-t)^{\bar{X}_n}\, \mathrm{d}t \tag{2.3}$$

This equation characterizes all aggregate attacks in which there is *coherency* between the signals and the action profile. That is, it characterizes all SBE. As shown in Lemma 1 (Appendix A), the threshold strategy satisfies $1 \leq \bar{X}_n(Y, C) \leq n - 2$ for all sufficiently large n. Under this condition,

[4] This is a common requirement in equilibrium concepts of games where agents sample. See Osborne and Rubinstein (2003) and Osborne and Rubinstein (1998) for further reference.

$A = 0$ and $A = 1$ solve the Equation 2.3. Besides these two equilibria, it can be shown that there exists a unique interior equilibrium with aggregate attack in the interval $(0, 1)$, that satisfies Equation 2.3.

Proposition 3: For sufficiently large n, there exists a unique $A_n^* \in (0, 1)$ that satisfies Equation 2.3.

This limited multiplicity eases comparative statics over the set of equilibria, as will be shown in Sections 2.4 and 2.5. Figure 1 illustrates the solutions to Equation 2.2 for a particular coordination global game with a sample size of $n = 6$ and a threshold strategy of $\bar{X}_n(Y, C) = 2$. Under this specification, the unique aggregate attack is $A_6^* \approx 0.347$ and consequently the regime survives under this SBE.

Figure 1: Solutions to Equation 2.2

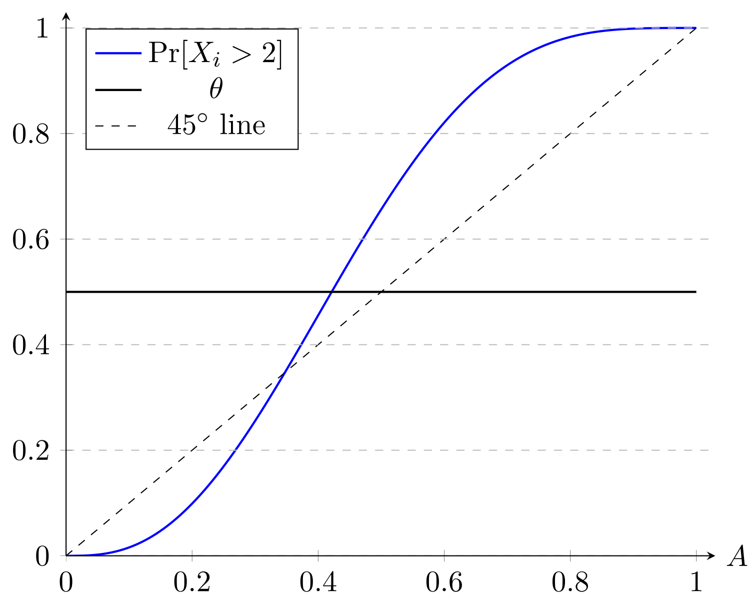

Coordination global game with $n = 6$, $\theta = 0.5$, $Y = 0.2$, $C = 0.3$, and uniform prior $p_0(A|Y, C)$.

2.4. Asymptotic Behavior of SBE: This section studies the convergence of SBE to Nash equilibria of the complete information game (where agents have perfect foresight on the action profile of the entire continuum), as sample size n tends to infinity. I model the learning process of agents as them acquiring information about more and more different players, and thus increasing their sample sizes. I show that asymptotically, as common knowledge is restored, the distribution of actions of any mixed strategy equilibria of the complete information game can be approximated arbitrarily well by a SBE, given some regularity conditions on prior $p_0(A|Y, C)$. The strategy profile in the SBE is entirely composed of pure strategies and involves no randomization from the agents. The mixed strategy equilibrium is thus *purified*.

This result uses Bayesian asymptotic analysis, in particular Bernstein-von Mises theorem.[5] This theorem studies the asymptotic behaviour of posterior beliefs derived from Bayesian updating. As data becomes more and more abundant and sample size tends to infinity, three results stem from the theorem. First, the distribution of the posterior beliefs is asymptotically independent of the prior. Second, the posterior distribution is asymptotically normal and centered around the Maximum Likelihood Estimator (MLE). Third and finally, the asymptotic variance is given by Cramér-Rao's lower bound.

In the regime switch model studied here, the MLE for the size of the attack is the sample average $n^{-1}X_i$. Its asymptotic variance is given by Cramér-Rao's lower bound: $n^{-2}A(1-A)$. Bernstein-von Mises' theorem thus implies that, given some regularity on the prior,

$$\big|\big|p_i(A|Y, C, X_i) - \mathcal{N}\big(n^{-1}X_i,\, n^{-2}A(1-A)\big)\big|\big|_{TV} \xrightarrow{\text{P}} 0. \tag{2.4}$$

Where $||P - Q||_{TV}$ is the total variation distance between measures P and Q on the Borel σ-algebra of $\mathbb{R}$, i.e.:

$$||P - Q||_{TV} = \sup_{B \in \mathcal{B}(\mathbb{R})} |P(B) - Q(B)|.$$

Expression 2.4 has two important implications. First, as a consequence of the Law of Large Numbers, the means of the posterior distributions start clustering in probability around the true size of the attack. Second, as sample sizes tend to infinity, the variance of these posterior distributions becomes vanishes as the tails of the distribution become thinner. These two effects mean that agents start agreeing on the true size of the attack and become more confident about their estimates.

Given the normality assumption on ε and the conditions on the prior P_θ, we have that $[0, 1] \subseteq$

[5]The precise formulation of Bernstein-von Mises theorem used in this paper can be found in Nickl (2013), as Theorem 5.

supp$[\theta|Y]$ for any given realization of public signal Y. This large support condition guarantees that there is a unique aggregate attack A^* that satisfies the indifference condition $\int_0^{A^*} g_{\theta|Y}\,\mathrm{d}\theta = C$, required for players to randomize in the complete information game. This means that, in every Nash equilibrium of the complete information game in which at least one player is playing a non-degenerate mixed strategy, the aggregate attack must be exactly A^*. Using this notation, Theorem 1 states formally the main purification result.

Theorem 1: If prior $p_0(A|Y,C)$ is continuous and strictly positive in $[0,1]$, then

$$\lim_{n\to\infty} A_n^* = A^*,$$

where $A_n^* \in (0,1)$ denotes the unique interior SBE of the game with sample size n.

It should be noted that convergence is not limited to the unique interior SBE only. For large enough sample size n, $A=0$ and $A=1$ are also SBE of the coordination global game, as they satisfy Equation 2.2. Given the conditions imposed on prior P_θ and cost C, action profile $a_i=0$ for all i constitutes a Nash equilibrium of the complete information game. Similarly, action profile $a_i=1$ for all i also constitutes a Nash equilibrium. This means that the SBE $A=0$ and $A=1$ trivially converge to equilibrium points of the complete information game. Figure 2 shows the convergence of the set of SBE to the set of Nash equilibria for a particular specification of the coordination game.

This purification result is attained under an endogenous information structure in which action profiles affect the distribution of types/signals and vice-versa. This is done while keeping the *coherency* between the strategy profile and the fully endogenous private signals. In traditional purification results– see for instance Harsanyi (1973) or Radner and Rosenthal (1982)– strategy profiles don't affect the distribution of types and the information structure is assumed exogenous. This assumption is also present in the purification literature that developed after.[6]

2.5. Comparative Statics on the set of SBE: This section presents the comparative statics of the model with respect to the value of public signal Y and cost C, under a fixed and finite sample

[6]See Aumann et al. (1983), Bhaskar et al. (2008), and Barelli and Duggan (2015) for further reference.

Figure 2: Asymptotic Behavior of SBE

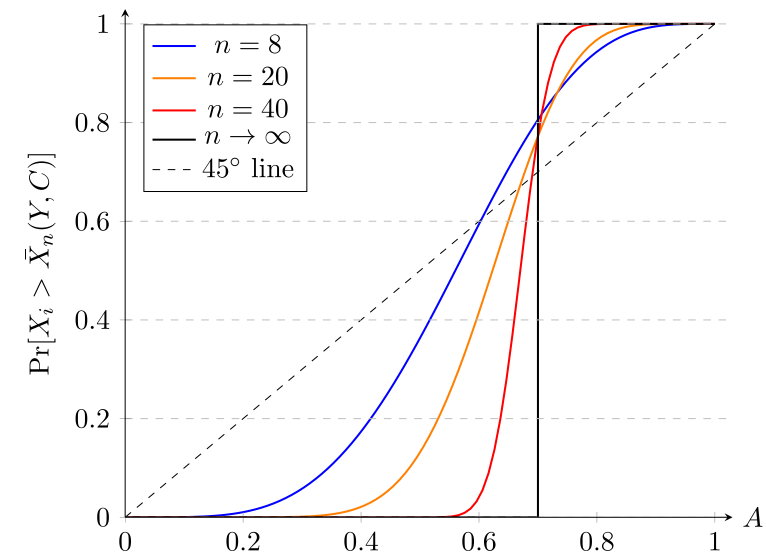

Coordination global game with $\theta = 0.5$, $Y = 0.7$, $C = 0.5$, and uniform prior $p_0(A|Y,C)$.

size. In order to do this, some assumptions on the behaviour of the prior need to be established first. How do agents' prior beliefs on the aggregate size of the attack A change with respect to Y and C? I assume that agents believe that a regime that is perceived as stronger or a higher cost of attacking make the rest of the agents less likely to attack. This could be modelled with the following monotonicity assumptions on the prior,

$$Y' > Y \Rightarrow p_0(A|Y,C) \succsim_{\text{FSD}} p_0(A|Y',C),$$

$$C' > C \Rightarrow p_0(A|Y,C) \succsim_{\text{FSD}} p_0(A|Y,C').$$

However, First Order Stochastic Dominance is not preserved under Bayesian updating, as shown by Klemens (2007), which makes this stochastic order problematic to handle in the context of this model. These conditions must be re-stated in terms of a stronger stochastic order like the Monotone Likelihood Ratio (MLR) property,

$$Y' > Y \Rightarrow p_0(A|Y,C) \succsim_{\text{MLR}} p_0(A|Y',C) \iff \frac{p_0(A|Y,C)}{p_0(A|Y',C)} \text{ is increasing in } A,$$

$$C' > C \Rightarrow p_0(A|Y,C) \succsim_{\text{MLR}} p_0(A|Y,C') \iff \frac{p_0(A|Y,C)}{p_0(A|Y,C')} \text{ is increasing in } A.$$

The MLR property has the advantage that it implies FSD and is preserved under Bayesian updating (Klemens (2007)). This is useful, as it allows to derive monotonicity conditions for the posterior belief distributions from the assumptions on the prior. Thus, for every X_i,

$$Y' > Y \Rightarrow p_0(A|Y,C) \succsim_{\text{MLR}} p(A|Y',C) \Rightarrow p_i(A|Y,C,X_i) \succsim_{\text{FSD}} p_i(A|Y',C,X_i),$$

$$C' > C \Rightarrow p_0(A|Y,C) \succsim_{\text{MLR}} p(A|Y,C') \Rightarrow p_i(A|Y,C,X_i) \succsim_{\text{FSD}} p_i(A|Y,C',X_i).$$

Consequently, the subjective probability of the attack succeeding,

$$\Pr[A > \theta|Y,C,X_i] = \int\int_{\theta}^{1} p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y}$$

is decreasing in C. It is also decreasing in Y because: **i)** if $Y' > Y$ then $g_{\theta|Y'} \succsim_{\text{FSD}} g_{\theta|Y}$, by construction; and **ii)** $\int_{\theta}^{1} p_i(A|Y,C,X_i)\,\mathrm{d}A$ is a decreasing function of θ. This implies that the threshold $\bar{X}_n(Y,C)$ is increasing in Y and C. The stronger a regime is perceived or the higher the cost of attacking, the more attackers an agents needs to see in her sample in order to attack.

This implies that the aggregate size of attack in the unique interior SBE is monotonically increasing in both Y and C. More attackers are needed to sustain an equilibrium with strictly positive aggregate attack if the cost is larger or the regime is perceived as stronger. Figure 3 shows the set of SBE in a particular specification of the coordination global game for two different values of public signal Y.

3. Sampling Bayesian Equilibria in Static Games Where Agents Sample

To study agents' learning process and the convergence of strategy profiles to Nash equilibria in a more general class of games, this section extends the definition of SBE to static games played by a continuum of agents. As before, players are assumed rational but having limited information, only getting to know the actions of a finite sample of other players. Agent's learn by increasing the size of their samples. Under certain regularity conditions, I will show that SBE exist and converge to the set of Nash equilibria of the game with complete information, as sample sizes tend to infinity

Figure 3: Comparative Statics of SBE

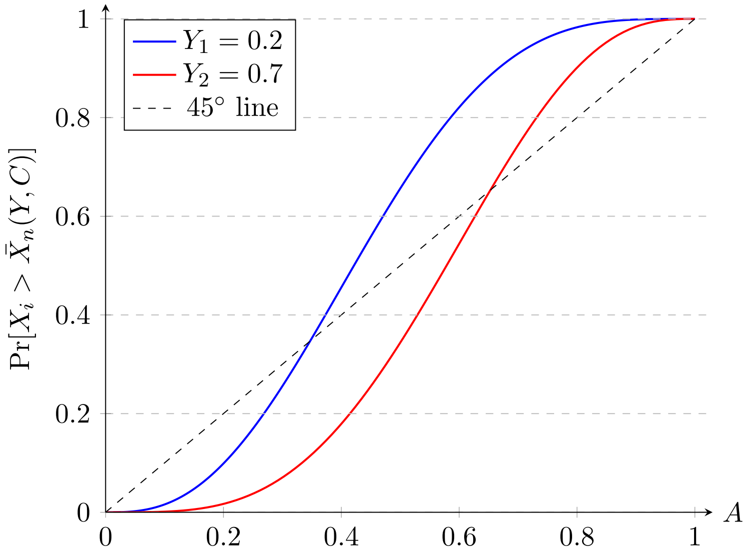

Coordination global game with $n = 6$, $\theta = 0.5$, $C = 0.3$, and triangular prior $p_0(A|Y,C)$.[7]

and common knowledge is restored.

Let the tuple $\Gamma = \left(A_i, \pi_i\right)_{i=1}^{N}$ – indexed by a set $\{1, 2 \cdots N\}$ of players– denote a finite static game in normal form, where A_i is the finite set of pure actions for player i and $\pi_i : \times_{i=1}^{N} A_i \to \mathbb{R}$ its respective payoff function.

To define a SBE for this class of games, let's first introduce the notion of a *sampling game* with sample size n, represented herein thereafter as Γ_n. Drawing an analogy with the purification result in Harsanyi (1973), these games represent perturbed versions of the complete information game Γ. In a *sampling game*, index i denotes classes of players each composed by an infinite unit-mass population of identical agents (i.e. agents with an identical action set and payoff function). Each individual player is grouped randomly with players of the remaining $N - 1$ classes. Then, each of them simultaneously picks an action and payoffs are realized, in accordance to the functions π_i.

For each player class $\{1, 2 \cdots N\}$ and each action $a_i \in A_i$, define:

i) m_i as the cardinality of set A_i;

ii) $\eta^i(a_i)$ as the proportion of players in class i playing strategy a_i;

[7] Specifically, let $\omega = \frac{\exp\{1-Y-C\}}{1+\exp\{1-Y-C\}}$ and $p_0(A|Y,C) = \begin{cases} \frac{2}{\omega} & 0 \le A < \omega, \\ \frac{2}{1-\omega}(1-A) & \omega \le A \le 1. \end{cases}$

iii) Vector $\eta^i = \big(\eta^i(a_1), \eta^i(a_2), \cdots \eta^i(a_{m_i})\big) \in \Delta^{m_i - 1}$ as the aggregate strategy of class i;[8]

iv) Vector $\eta^{-i} = \big(\eta^1, \cdots \eta^{i-1}, \eta^{i+1} \cdots \eta^N\big) \in \times_{j \neq i} \Delta^{m_j - 1}$ the aggregate strategy of all player classes except i.

3.1. Information Structure: Suppose that agents don't have perfect foresight on the action profile of the entire population, but instead only observe the actions of finite samples of the $(N-1)$ remaining classes. For simplicity, assume all these samples have the same size, n, for all agents and all classes. Agents use the observed action profile in their samples, X^i, to derive posterior beliefs on η^{-i} via Bayesian updating. To keep the parallel with the coordination global games studied above, we'll refer to X^i as the private signal for a player in class i.

Thus, strategies in *sampling games* are mappings from private signals to actions, $s : \mathbb{X}^i \to A_i$, where $\mathbb{X}^i$ is the support of random vector X^i. A strategy profile is a set of maps $s_{i,t} : \mathbb{X}^i \to A_i$, indexed by $(i,t) \in \{1, \cdots N\} \times [0,1]$, describing for each individual in each player class the action to take after observing the realization of her private signal. We call a strategy profile symmetric if for every $t \neq t'$ and every i, $s_{i,t} = s_{i,t'}$.

3.1.1. Distribution of the Private Signals: To study the distribution of the private signals, define the following random variables and vectors:

i) $X^i_j(a_j)$ the number of class j players playing action $a_j \in A_j$, that a player of class i observes in her sample;

ii) $X^i_j = \big(X^i_j(a_1), X^i_j(a_2), \cdots X^i_j(a_{m_j})\big)$ the empirical frequency of actions for class j, that a player of class i observes in her sample;

iii) $X^i = \big(X^i_1, \cdots X^i_{i-1}, X^i_{i+1} \cdots X^i_N\big)$ the empirical frequency of actions for all the other $N-1$ classes of players, that a player of class i observes in her sample.

Random variable $X^i_j(a_j)$ follows a binomial distribution with n trials and success probability of

[8] Where Δ^n is the unit n-simplex. I.e.: $\Delta^n = \big\{(t_0, t_1, \cdots t_n) \in \mathbb{R}^{n+1} : \sum_{i=1}^n t_i = 1 \,\text{and}\, t_i \geq 0 \text{ for all } i\big\}$.

$\eta^j(a_j)$. Thus, vectors X^i_j follow a multinomial distribution of the form:

$$X^i_j \sim \text{Multi}\big(n; \eta^j(a_1), \eta^j(a_2) \cdots \eta^j(a_{m_j})\big),$$

Assuming that samples of different classes of players are mutually independent, we have that the private signals follow the distribution:

$$X^i \sim \otimes_{j\neq i}\text{Multi}\big(n; \eta^j(a_1), \eta^j(a_2) \cdots \eta^j(a_{m_j})\big).$$

3.1.2. Timing of the Game: For an individual player in class i, taking the actions of all other players as given, the timing of the game is the following:

i) The agent observes the payoff structure in Γ and then forms a prior $p^0_i(\eta^{-i}|\,\Gamma)$ on the aggregate actions of the remaining $N-1$ classes.[9]

ii) The agent's samples of the $(N-1)$ remaining classes are randomly chosen and private signal X^i is observed. Posterior beliefs on the aggregate actions of the remaining $N-1$ classes are derived:

$$p_i(\eta^{-i}|X^i = x^i,\, \Gamma) = \frac{p(X^i = x^i|\eta^{-i})p^0_i(\eta^{-i}|\,\Gamma)}{\int p(X^i = x^i|\eta^{-i})p^0_i(\eta^{-i}|\,\Gamma)\,\mathrm{d}\eta^{-i}}.$$

iii) The agent is randomly coupled with players of the remaining $N-1$ classes and chooses her action in A_i, based on her private signal X^i.

3.2. Expected Payoffs and Optimal Actions: Given the posterior beliefs on η^{-i}, the expected payoff for an agent in class i, playing action $a_i \in A_i$, after observing private signal X^i is given by:[10]

$$U_i(a_i|X^i) = \int \Big(\sum_{a_{-i}\in A_{-i}} \eta^{-i}(a_{-i})\pi_i(a_i, a_{-i})\Big) p_i(\eta^{-i}|X^i,\, \Gamma)\,\mathrm{d}\eta^{-i}.$$

[9] Suppose that this prior is absolutely continuous with respect to the Lebesgue measure in $\times_{j\neq i}\Delta^{m_j-1}$ and therefore follows a probability density function.

[10] Where $A_{-i} = \times_{j\neq i}A_i$ and for $a_{-i} = (a_{-i,1}, \cdots a_{-i,i-1}, a_{-i,i+1} \cdots a_{-i,N}) \in A_{-i}$ let $\eta^{-i}(a_{-i}) = \prod_{j\neq i}\eta^j(a_{-i,j})$

3.2.1. Optimal Actions: For a particular realization of the private signal $x^i \in \mathbb{X}^i$, define the set of *optimal actions* as the set of actions in A_i such that:

$$\mathrm{OA}(x^i) = \arg\max_{a_i \in A_i} U_i(a_i|x^i)$$

3.3. Sampling Bayesian Equilibium: Sampling games have endogenous information structures as players' actions alter the distribution of private signals. In turn, the distribution of these signals shapes actions, as agents try to respond optimally given their posterior beliefs. To deal with this circularity, I impose a *coherency* requirement in the solution of the game: given a strategy profile there should be a particular distribution of the private signals, and given this distribution of signals the strategy profile should be consistent with agents optimizing their expected payoff. I define a strategy profile and a share vector of aggregate actions, η, in which these two requirements are met a Sampling Bayesian Equilibrium of sampling game Γ_n. For simplicity, we restrict our definition to SBE with symmetric strategy profiles.

Definition: [Sampling Bayesian Equilibrium of Γ_n in Symmetric Strategies]
Let $s \equiv \{s_i : \mathbb{X}^i \to A_i\}_{i=1}^N$ be a symmetric strategy profile and let $\eta \equiv \left(\eta^i\right)_{i=1}^N$ be a collection of share vectors (such that, for every i, $\eta^i \in \Delta^{m_i-1}$).
(s, η) is a SBE in symmetric strategies if and only if for all i, all $a_i \in A_i$, and all $x_i \in \mathbb{X}^i$:

1. $s_i(x^i) \in \mathrm{OA}(x^i)$;

2. $\eta^i(a_i) = \Pr\left[X^i \in s_i^{-1}(a_i)\right]$

Where $X^i \sim \otimes_{j \neq i}\mathrm{Multi}\big(n; \eta^j(a_1), \eta^j(a_2) \cdots \eta^j(a_{m_j})\big)$.

The first condition is a requirement that agents behave optimally by maximizing their subjective expected payoffs after observing their respective private signal. The second condition imposes *coherency* in the share vectors: Given the proportion of agents playing each of the actions, the distribution of private signals thus induced coupled with agent's optimal behaviour should generate these same proportions of agents playing each action.
A natural question that follows is if for every static game Γ and any finite sample size n, one can find a SBE in symmetric strategies for the corresponding sampling game Γ_n. As stated in Theorem

2, this is always the case. This result rests on Brouwer's fixed-point theorem, which is used to show the existence of the share vectors satisfying requirement 2. The existence of a symmetric strategy profile satisfying condition 1 is immediate from the extreme value theorem and the finiteness of set $\mathbb{X}^i$.[11]

Theorem 2: Every *sampling game* Γ_n has a SBE in symmetric strategies.

3.4. Convergence of SBE to Nash Equilibria: This section studies the asymptotic behaviour of SBE in symmetric strategies as the sample size n tends to infinity. In particular, it focuses on the approachability –a la Harsanyi (1973)- of the Nash equilibria in 2×2 games by sequences of SBE of the respective perturbed sampling games. As will be shown below, the payoff structure in 2×2 games gives strategies in the corresponding sampling games a monotonicity property over the signals. This implies that optimal strategies must follow a threshold condition, as happened in the coordination global game studied above. This similarity allows us to extend the previous purification result into this context.

Suppose game Γ has the following payoff matrix:

		Player B: 1	Player B: 2
Player A	1	(a_{11}, b_{11})	(a_{12}, b_{12})
	2	(a_{21}, b_{21})	(a_{22}, b_{22})

In line with the notation introduced above, let $X_B^A(1)$ denote the number of class A players playing action 1, observed by a player of class B in her sample. Let $X_A^B(1)$ be defined analogously. Since action spaces are binary, we have that $X_B^A(2) = n - X_B^A(1)$, and therefore, to simplify notation, we can consider $X_B^A(1)$ by itself as the private signal for a player in class A.

The expected payoff difference between actions 1 and 2 for a player in class A –after observing a

[11] The complete proof of Theorem 2 can be found in the mathematical appendix.

private signal $X_B^A(1)$– can be expressed as:

$$
\begin{aligned}
&U_A(1|X_B^A(1),\,\Gamma) - U_A(2|X_B^A(1),\,\Gamma) = \\
&\qquad\int_0^1 \Big[\eta_1^B(a_{11}-a_{21}) + (1-\eta_1^B)(a_{12}-a_{22})\Big]\cdot p_i(\eta_1^B|X_B^A(1),\,\Gamma)\,\mathrm{d}\eta_1^B \\
&\qquad\qquad = a_{12}-a_{22} + [a_{11}-a_{21}+a_{22}-a_{12}]\int_0^1 \eta_1^B\cdot p_i(\eta_1^B|X_B^A(1),\,\Gamma)\,\mathrm{d}\eta_1^B
\end{aligned}
\tag{3.1}
$$

Where $p_i(\eta_1^B|X_B^A(1),\,\Gamma)$ is the posterior belief of the unknown share η_1^B that a class A player has after observing private signal $X_B^A(1)$.

From Proposition 1, we know that if $X_B^A(1) > \tilde{X}_B^A(1)$, then we have that $p_i(\eta_1^B|X_B^A(1),\,\Gamma) \succsim_{\text{FOSD}} p_i(\eta_1^B|\tilde{X}_B^A(1),\,\Gamma)$ and thus $\int \eta_1^B\cdot p_i(\eta_1^B|X_B^A(1),\,\Gamma)\,\mathrm{d}\eta_1^B > \int \eta_1^B\cdot p_i(\eta_1^B|\tilde{X}_B^A(1),\,\Gamma)\,\mathrm{d}\eta_1^B$. Therefore, Equation 3.1 implies that the optimal threshold strategy depends on the sign of $[a_{11}-a_{21}+a_{22}-a_{12}]$. If $\text{sign}[a_{11}-a_{21}+a_{22}-a_{12}] \geq 0$, then there exists $\bar{X}\in\{0,1,\cdots n\}$ such that $s_i(x_B^A(1)) = 1$ if and only if $x_B^A(1) > \bar{X}$. Similarly, if $\text{sign}[a_{11}-a_{21}+a_{22}-a_{12}] \leq 0$, then there exists $\bar{X}\in\{0,1,\cdots n\}$ such that $s_i(x_B^A(1)) = 0$ if and only if $x_B^A(1) > \bar{X}$.

To analyze the convergence of SBE, we must find the asymptotic distribution of the expected ratio of players B playing action 1, $\int_0^1 \eta_1^B\cdot p_i(\eta_1^B|X_B^A(1),\,\Gamma)\,\mathrm{d}\eta_1^B$. We rely again on Bernstein-von Mises theorem.[12] This theorem guarantees that, given some regularity on the prior distribution, the posterior distribution is asymptotically independent of the prior and normally distributed, centered around the MLE. For the sampling game studied, the MLE is given by the sample average, i.e. $\hat{\eta}_{1,MLE}^B = X_B^A(1)/n$. Bernstein -von Mises theorem also implies that the variance of the posterior distribution attains Cramér-Rao's lower bound. In the present context, this lower bound is given by $n^{-2}\eta_1^B(1-\eta_1^B)$.

Thus, the theorem states that if the prior has a Lebesgue density that is continuous and positive in a neighborhood of η_1^B, then

$$
\Big|\Big|p_i(\eta_1^B|X_B^A(1),\,\Gamma) - \mathcal{N}\Big(n^{-1}X_B^A(1),\,n^{-2}\eta_1^B(1-\eta_1^B)\Big)\Big|\Big|_{TV} \xrightarrow{\text{P}} 0. \tag{3.2}
$$

Where $||P-Q||_{TV}$ is the total variation distance between measures P and Q on the Borel σ-algebra

[12]For the precise formulation of Bernstein- von Mises' Theorem please refer to Nickl (2013), specifically Theorem 5 in page 37.

of $\mathbb{R}$, i.e.:

$$||P-Q||_{TV} = \sup_{B\in\mathcal{B}(\mathbb{R})} |P(B)-Q|.$$

The convergence of SBE to Nash equilibria is driven by two forces. First, as $n\to\infty$, the means of the posterior distributions cluster in probability around the true ratio η_1^B. Second, the variances of these distributions tend to 0. This means that, as agents learn and sample sizes tend to infinity, more and more agents of class A are more and more sure about the true size of η_1^B. Common knowledge on the aggregate action profile is asymptotically restored. Theorem 3 states formally the purification result of this section.

Theorem 3:[13] Let $\sigma\in\Delta^1\times\Delta^1$ be an equilibrium of the 2×2 game Γ in which no player uses a weakly dominated strategy. Then, there exists a sequence of SBE in symmetric strategies $(s_n,\eta_n)_{n=1}^{\infty}$ of the sampling games $(\Gamma_n)_{n=1}^{\infty}$, such that for $i\in\ \{A,B\}$ and $j\in\{1,2\}$:

$$\lim_{n\to\infty}\eta_n^i(j)=\sigma_i(j).$$

This result shares two similarities with the classic rationale of mixed strategies developed in Harsanyi (1973). First, both purification results give conditions for the approachability of equilibria in the complete information game by sequences of pure strategy equilibria of the perturbed games. This contrasts with most of the purification literature that has been developed in more recent years, after Radner and Rosenthal (1982). This literature does not aim for the approximation of mixed strategy equilibria, but on finding conditions on the information structure of an incomplete information game that guarantee the existence of an outcome equivalent pure strategy equilibrium for every (possibly mixed) equilibrium point (Morris and Shin (2006)).

The second similarity relies on the type of equilibria that can be purified, namely equilibria that do not involve weakly dominated strategies. In the context of sampling games, this condition is also necessary for purification.[14] For any realization of the private signal, it would be sub-optimal to play a weakly dominated strategy as the posterior on parameter η_1^B assigns positive probability

[13]Where $\sigma_i(j)$ is the probability that player i assigns to strategy j.

[14]This condition is not a very restrictive one because, as shown in Theorem 2.6.1 in Damme (1991), all equilibria of almost all games do not involve weakly dominated strategies when payoffs are chosen according to any Lebesgue-absolutely continuous measure.

to the open interval $(0, 1)$. Hence, approachability fails as agents in a SBE behave optimally.

The main difference between both purification results rests on the assumptions on the information structure. In Harsanyi (1973), convergence of the perturbed games to the complete information benchmark occurs as the variance of the random payoffs, or types, is exogenously shrunk to zero. The distribution of types is not affected by the actions of players. In contrast, in sampling games, the information structure is endogenous. The distribution of signals/types is affected by agents' actions, and vice-versa. This is the main contribution of sampling games and studying the convergence of SBE in this context: exploring the purification of mixed strategy equilibria under endogenous information structures.

3.5. Example –Convergence of SBE in the Battle of the Sexes: This section illustrates the convergence of SBE to Nash equilibria– proved in Theorem 3– for the game Battle of the Sexes. Suppose the complete information game Γ has the following payoff matrix:

		Player B: 1	Player B: 2
Player A	1	$(2, 1)$	$(0, 0)$
Player A	2	$(0, 0)$	$(1, 2)$

This game has two Nash equilibria in pure strategies–both players coordinate on strategy 1 or they both coordinate on strategy 2– and a unique equilibrium in mixed strategies. The latter corresponds to Player A playing strategy 1 with probability 2/3 and Player B playing strategy 1 with probability 1/3. Notice that in these three Nash equilibria no player is using a weakly dominated strategy.

Figure 4 shows the convergence of the SBE in the corresponding sampling games to the set of Nash equilibria of game Γ, as sample size grows to infinity. The continuous lines in the figure show the proportion of agents in class B that would play strategy 1 given the distribution of signals generated by η_1^A. Symmetrically, the dashed lines show the proportion of agents in class A that would play strategy 1 given the distribution of signals generated by η_1^B. Thus, the SBE are found in the intersections of the continuous lines with the dashed lines of their corresponding color, as these are the points were the share vectors satisfy *coherency*. Notice how the interior SBE approach the unique mixed strategy equilibria of the Battle of the Sexes, as sample size grows. Similarly, the two fixed SBE, $(\eta^A(1), \eta^B(1)) = (0, 0)$ and $(\eta^A(1), \eta^B(1)) = (1, 1)$, also (trivially) converge to the two pure strategy Nash equilibria of game Γ.

Figure 4: SBE Convergence to Nash Equilibria

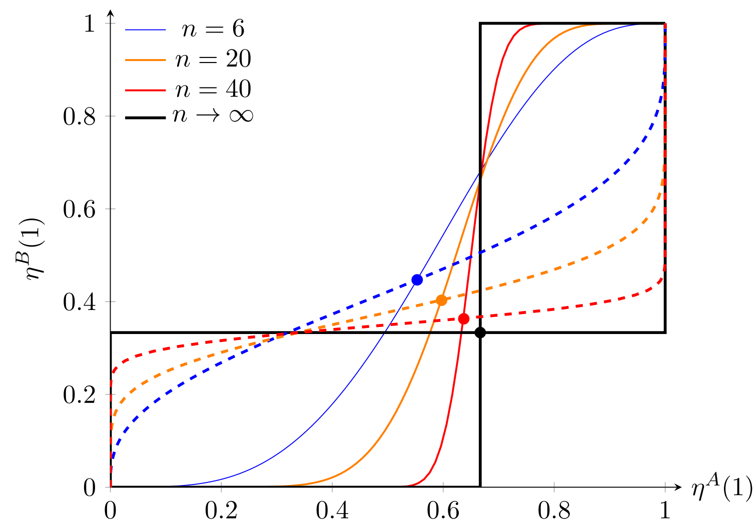

Sampling game of Battle of the Sexes with uniform priors $p_1^0(\eta^{-1}|\,\Gamma)$ and $p_2^0(\eta^{-2}|\,\Gamma)$.

4. Concluding Remarks

This paper defines an equilibrium concept–the Sampling Bayesian Equilibrium (SBE)– for games in which players are assumed rational but have limited information, only getting to know the actions of a finite sample of other players. This framework is useful to study many economic strategic interactions –such as bank runs or attacks on currency pegs– in which the assumption that each player has perfect foresight on the equilibrium behaviour of all other players is debatable. I show the existence of these equilibrium points for a coordination global game – specifically a regime switch model (Morris and Shin (1999))– and for finite static games played by a continuum of agents. I model the agents' learning process as them acquiring information about more and more different players, and thus increasing the size of their samples. As the latter tend to infinity, I show conditions under which the set of SBE converges to the set of Nash equilibria of the complete information games (where agents observe the actions of the entire continuum).

The convergence results proved in this paper are reminiscent of Nash's 'mass-action' interpretation of equilibrium points, in particular those composed of mixed strategies. In his doctoral dissertation, he argued that the approximation of mixed strategy equilibria could arise from games with a population of participants (in the sense of statistics), where players try to estimate the average frequency of each pure strategy (Nash (1950)).[15] The sampling games introduced in this paper and the asymptotic analysis of SBE provide a possible mathematical foundation to Nash's 'mass-action' interpretation of mixed strategy equilibria.

[15] "We shall now take up the 'mass-action' interpretation of equilibrium points [...] [P]articipants are supposed to accumulate empirical information on the relative advantages of the various pure strategies at their disposal.

To be more detailed, we assume that there is a population (in the sense of statistics) of participants for each position of the game. Let us also assume that the 'average playing' of the game involves N participants selected at random from the N populations, and that there is a stable average frequency with which each pure strategy is employed by the 'average member' of the appropriate population [...] Thus the assumptions we made in this 'mass-action' interpretation lead to the conclusion that the mixed strategies representing the average behaviour in each of the populations form an equilibrium point[...].Actually, of course, we can only expect some sort of approximate equilibrium, since the information, its utilization, and the stability of the average frequencies will be imperfect." [Nash (1950); pages 21–23.]

A. Appendix: Proofs of Theorems and Propositions

Propositions 1 and 2: Notice that,

$$\frac{p_i(A|Y,C,X_i=k)}{p_i(A|Y,C,X_i=k+1)} = \frac{1-A}{A}\frac{\int_0^1 A^{k+1}(1-A)^{n-k-1}p_0(A|Y,C)\,\mathrm{d}A}{\int_0^1 A^{k}(1-A)^{n-k}p_0(A|Y,C)\,\mathrm{d}A} \equiv \frac{1-A}{A}M.$$

Therefore, for any $A \in (0,1)$,

$$p_i(A|Y,C,X_i=k) \lesseqqgtr p_i(A|Y,C,X_i=k+1) \Leftrightarrow A \gtreqqless \frac{M}{1+M}$$

Consequently, for any $\lambda \leq \frac{M}{1+M}$,

$$\int_0^\lambda p_i(A|Y,C,X_i=k+1)\,\mathrm{d}A < \int_0^\lambda p_i(A|Y,C,X_i=k)\,\mathrm{d}A$$

And for any $\lambda > \frac{M}{1+M}$,

$$\int_0^\lambda p_i(A|Y,C,X_i=k+1)\,\mathrm{d}A = 1 - \int_\lambda^1 p_i(A|Y,C,X_i=k+1)\,\mathrm{d}A$$
$$< 1 - \int_\lambda^1 p_i(A|Y,C,X_i=k)\,\mathrm{d}A = \int_0^\lambda p_i(A|Y,C,X_i=k)\,\mathrm{d}A$$

Thus, $p_i(Y,C,A|X_i=k+1) \succsim_{\text{FSD}} p_i(A|Y,C,X_i=k)$. Finally, conclude by noticing that stochastic dominance is a transitive binary relation.

□

Proposition 3:

[Existence]: Let $Q(A) = 1-(n-\bar{X}_n)\binom{n}{\bar{X}_n}\int_0^{1-A} t^{n-\bar{X}_n-1}(1-t)^{\bar{X}_n}\,\mathrm{d}t$ and $R(A) = Q(A) - A$.

$$\frac{\partial Q(A)}{\partial A} = (n-\bar{X}_n)\binom{n}{\bar{X}_n}(1-A)^{n-\bar{X}_n-1}A^{\bar{X}_n} \Rightarrow \frac{\partial Q(A)}{\partial A}\bigg|_{A=0,1} = 0 \Rightarrow \frac{\partial R(A)}{\partial A}\bigg|_{A=0,1} = -1$$

Since $R(A)$ is C^∞ in the interval $(0,1)$, then $\exists A_1 \in (0,\varepsilon)$ and $\exists A_2 \in (1-\varepsilon,1)$ s.t. $R(A_1) < 0$ and $R(A_2) > 0$, with $\varepsilon < 1/2$. By a direct result of the Intermediate Value Theorem, $\exists A \in (A_1, A_2)$ s.t. $R(A) = 0$.

[Uniqueness]:Notice that,

$$\frac{\partial^2 R(A)}{\partial A^2} = \frac{\partial^2 Q(A)}{\partial A^2} = (n-\bar{X}_n)\binom{n}{\bar{X}_n}A^{\bar{X}_n-1}(1-A)^{n-\bar{X}_n-1}\big(\bar{X}_n-(n-1)A\big)$$

Therefore, $A = \frac{\bar{X}_n}{n-1}$ is the only inflection point of $R(A)$ in the interval $(0,1)$, for large enough n.[16] Suppose by means of contradiction that $\exists A_1, A_2 \in (0,1)$ s.t. $R(0) = R(A_1) = R(A_2) = R(1) = 0$.

[16] In the proof of Theorem 1, Lemma 1 shows that there exists N such that for all $n \geq N$, $1 \leq \bar{X}_n(Y,C) \leq n-2$.

Applying the Mean Value Theorem: $\exists c_1, c_2, c_3 \in (0,1)$ s.t. $0 < c_1 < A_1 < c_2 < A_2 < c_3 < 1$ and $\frac{\partial R(A)}{\partial A}\Big|_{A=c_1,c_2,c_3} = 0$. By a second application of the Mean Value Theorem: $\exists d_1, d_2 \in (0,1)$ s.t. $c_1 < d_1 < c_2 < d_2 < c_3$ and $\frac{\partial^2 R(A)}{\partial A^2}\Big|_{A=d_1,d_2} = 0$.$\Rightarrow\Leftarrow$

□

Proof of Theorem 1:

Let $V \subset \mathbb{R}$ be the the neighbourhood around A^* where the prior $p_0(A|Y,C)$ is continuous and strictly positive. Pick $\delta > 0$ such that $(A^* - \delta, A^* + \delta) \subset V$. Pick $\underline{A} \in (A^* - \delta, A^*)$ and $\bar{A} \in (A^*, A^* + \delta)$. Notice that we can rewrite Equation 2.3 as:

$$A = \Pr\left[\iint_\theta^1 p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right]$$

We will show that,

1. If $X_i \sim \text{Bin}(n, \underline{A})$, then $\lim_{n\to\infty} \Pr\left[\iint_\theta^1 p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] = 0$;
2. If $X_i \sim \text{Bin}(n, \bar{A})$, then $\lim_{n\to\infty} \Pr\left[\iint_\theta^1 p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] = 1$.

Therefore, as δ can be made arbitrarily small and $\Pr\left[\iint_\theta^1 p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right]$ is continuous with respect to A, we would have the desired result. First, we proceed to prove Lemma 1 with the aid of these two limits. This lemma guarantees that for large enough n, A_n^* exists.

Lemma 1: There exists N such that for all $n \geq N$, we have $1 \leq \bar{X}_n(Y,C) \leq n-2$.

[Proof:]

Suppose, for the sake of contradiction, that for every N there exists $n \geq N$ such that $\bar{X}_n(Y,C) = 0$. Then, when $X_i \sim \text{Bin}(n, \underline{A})$, we have that $\Pr(X_i > \bar{X}_n(Y,C)) = 1 - (1 - \underline{A})^n$ for arbitrarily large n. Therefore, when $X_i \sim \text{Bin}(n, \underline{A})$, $\lim_{n\to\infty} \Pr\left[\iint_\theta^1 p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] \neq 0$ which contradicts the first limit. Thus, we conclude that $1 \leq \bar{X}_n(Y,C)$. For the second inequality, a similar contradiction argument holds involving the second limit.

□

We proceed to prove the first limit. Notice that,

$$\Pr\left[\iint_\theta^1 p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] = \Pr\left[\iint_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] + \Pr\left[\iint_\theta^1 p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] - \Pr\left[\iint_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right]$$

Where $\phi(\mu,\,\sigma^2)$ denotes the density of a normal distribution with mean μ and variance σ^2. We divide the proof into two steps.

Step 1: Show that, if $X_i \sim \mathrm{Bin}(n, \underline{\mathrm{A}})$,

$$\lim_{n\to\infty} \Pr\left[\iint_\theta^1 p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] - \Pr\left[\iint_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] = 0.$$

The MLE is given by $n^{-1}X_i$ and its asymptotic variance is given by the Cramér-Rao lower bound $n^{-2}\underline{A}(1-\underline{A})$. Bernstein-von Mises' theorem, which specific formulation can be found in Nickl (2013) as Theorem 5, states that, given the technical conditions imposed on the prior,

$$\left|\left|p_i(A|Y,C,X_i) - \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\right|\right|_{TV} \xrightarrow{\mathrm{P}} 0.$$

Where, $||P-Q||_{TV}$ is the total variation distance between measures P and Q on the Borel σ-algebra of $\mathbb{R}$, i.e.: $||P-Q||_{TV} = \sup_{B\in\mathcal{B}(\mathbb{R})} |P(B)-Q(B)|$. We use this theorem to prove Lemma 2.

Lemma 2:

$$\iint_\theta^1 p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} - \iint_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \xrightarrow{\mathrm{p}} 0$$

[Proof]:

$$\begin{aligned}
&\left|\iint_\theta^1 p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} - \iint_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y}\right| \leq \\
&\int\left|\int_\theta^1 p_i(A|Y,C,X_i)\,\mathrm{d}A - \int_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\right|\mathrm{d}G_{\theta|Y} \leq \\
&\int \sup_{\mathcal{A}\in\mathcal{B}(\mathbb{R})}\left|\int_{\mathcal{A}} p_i(A|Y,C,X_i)\,\mathrm{d}A - \int_{\mathcal{A}} \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\right|\mathrm{d}G_{\theta|Y} = \\
&\sup_{\mathcal{A}\in\mathcal{B}(\mathbb{R})}\left|\int_{\mathcal{A}} p_i(A|Y,C,X_i)\,\mathrm{d}A - \int_{\mathcal{A}} \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\right| \xrightarrow{\mathrm{p}} 0
\end{aligned}$$

Where $\mathcal{B}(\mathbb{R})$ denotes the Borel σ-algebra of $\mathbb{R}$. The last convergence in probability comes from directly applying Bernstein-von Mises' theorem.

$\square$

Applying Lemma 2, we obtain, $\iint_\theta^1 p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \xrightarrow{\mathrm{d}} \iint_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y}$. Therefore, we conclude Step 1:

$$\lim_{n\to\infty} \Pr\left[\iint_\theta^1 p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] - \Pr\left[\iint_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] = 0.$$

Step 2: Show that, if $X_i \sim \mathrm{Bin}(n, \underline{\mathrm{A}})$,

$$\lim_{n\to\infty} \Pr\left[\iint_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] = 0.$$

Pick $\nu > 0$ such that $\underline{A} + \nu < A^*$. Notice that,

$$\begin{aligned}
&\Pr\left[\iint_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] = \\
&\quad \Pr\left[n^{-1}X_i \leq \underline{A}+\nu\right]\Pr\left[\iint_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\,\Big|\, n^{-1}X_i \leq \underline{A}+\nu\right] + \\
&\quad\quad \Pr\left[n^{-1}X_i > \underline{A}+\nu\right]\Pr\left[\iint_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\,\Big|\, n^{-1}X_i > \underline{A}+\nu\right]
\end{aligned}$$

The Law of Large Numbers implies that $\lim_{n\to\infty} \Pr\big[n^{-1}X_i > \underline{A}+\nu\big] = 0$. Therefore, we have that:

$$\begin{aligned}
&\lim_{n\to\infty} \Pr\left[\iint_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] \leq \\
&\quad\quad \lim_{n\to\infty} \Pr\left[\iint_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\,\Big|\, n^{-1}X_i \leq \underline{A}+\nu\right]
\end{aligned}$$

Conditional on $n^{-1}X_i \leq \underline{A}+\nu$, we have that $\phi\big(\underline{A}+\nu,\, n^{-2}\underline{A}(1-\underline{A})\big) \succsim_{\mathrm{FSD}} \phi\big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\big)$. Thus, conditional on $n^{-1}X_i \leq \underline{A}+\nu$, we have that: $\iint_\theta^1 \phi\Big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\Big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \leq \iint_\theta^1 \phi\big(\underline{A}+\nu,\, n^{-2}\underline{A}(1-\underline{A})\big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y}$. Pick $\tilde{A} \in (\underline{A}+\nu, A^*)$ and notice that,

$$\begin{aligned}
&\iint_\theta^1 \phi\big(\underline{A}+\nu,\, n^{-2}\underline{A}(1-\underline{A})\big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} = \\
&\quad \int_{-\infty}^{\tilde{A}}\int_\theta^1 \phi\big(\underline{A}+\nu,\, n^{-2}\underline{A}(1-\underline{A})\big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} + \int_{\tilde{A}}^{\infty}\int_\theta^1 \phi\big(\underline{A}+\nu,\, n^{-2}\underline{A}(1-\underline{A})\big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \leq \\
&\quad\quad \int_{-\infty}^{\tilde{A}} 1\,\mathrm{d}G_{\theta|Y} + \int_{\tilde{A}}^{\infty}\int_{\tilde{A}}^1 \phi\big(\underline{A}+\nu,\, n^{-2}\underline{A}(1-\underline{A})\big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y}
\end{aligned}$$

Since $\underline{A}+\nu < \tilde{A}$, then as $n \to \infty$, we have by the Dominated Convergence Theorem that $\int_{\tilde{A}}^{\infty}\int_{\tilde{A}}^1 \phi(\underline{A}+$

$\nu,\, n^{-2}\underline{A}(1-\underline{A})\big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \to 0$. Therefore, as $n\to\infty$, we get,

$$\iint_{\theta}^{1} \phi\big(\underline{A}+\nu,\, n^{-2}\underline{A}(1-\underline{A})\big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \to \int_{-\infty}^{\tilde{A}} 1\,\mathrm{d}G_{\theta|Y} < \int_{-\infty}^{A^*} 1\,\mathrm{d}G_{\theta|Y} = C$$

Consequently,

$$\lim_{n\to\infty} \Pr\left[\iint_{\theta}^{1} \phi\big(n^{-1}X_i,\, n^{-2}\underline{A}(1-\underline{A})\big)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] = 0,$$

and we have concluded Step 2. Combining these two results, we have that,

$$\lim_{n\to\infty} \Pr\left[\iint_{\theta}^{1} p_i(A|Y,C,X_i)\,\mathrm{d}A\,\mathrm{d}G_{\theta|Y} \geq C\right] = 0.$$

The limit second limit can be proved with a very similar argument as the one exposed above.

□

Proof of Theorem 2: Notice that, for any fixed $x^i \in \mathbb{X}^i$, function $U_i(a_i|X^i = x^i)$ can only take take finitely many values. Thus, as a direct consequence of the extreme value theorem, there exists at least one strategy profile such that $s_i(x^i) \in \mathrm{OA}(x^i)$ for every i and every signal $x^i \in \mathbb{X}^i$.

For any one of such profiles define the function $H : \times_{j=1}^{N}\Delta^{m_j-1} \to \times_{j=1}^{N}\Delta^{m_j-1}$ as:

$$H(\eta) = \begin{bmatrix} \Pr\big[X^1 \in s_1^{-1}(a_1)\big] \\ \Pr\big[X^1 \in s_1^{-1}(a_2)\big] \\ \vdots \\ \Pr\big[X^N \in s_N^{-1}(a_{m_N})\big] \end{bmatrix}$$

Note that:

1. $\times_{j=1}^{N}\Delta^{m_j-1}$ is a compact convex subset of $\mathbb{R}^{\sum m_j}$.
2. Since $X^i \sim \otimes_{j\neq i}\mathrm{Multi}\big(n; \eta^j(a_1), \eta^j(a_2)\cdots\eta^j(a_{m_j})\big)$, then for all i and $a_i \in A_i$, $\Pr\big[X^i \in s_i^{-1}(a_i)\big]$ is a continuous function.

Therefore, as a consequence of Brouwer's fixed-point theorem, there exists $\eta_* \in \times_{j=1}^{N}\Delta^{m_j-1}$ such that $H(\eta_*) = \eta_*$.

□

Proof of Theorem 3: Let $\sigma \in \Delta^1 \times \Delta^1$ be an equilibrium of game Γ in which no player uses a weakly dominated strategy. Under this condition, it can be easily shown that σ is either totally mixed or in pure strategies. Thus, we separate the proof for these two cases.

Case 1: [Totally mixed equilibrium] Let $\sigma = (\alpha, 1-\alpha;\ \beta, 1-\beta)$ with $\alpha, \beta \in (0,1)$ represent a totally mixed equilibrium of game Γ. Since no player uses a weakly dominated strategy, the payoffs of player 1 must satisfy either: **i)** $a_{11} - a_{21} > 0$ and $a_{12} - a_{22} > 0$, or **ii)** $a_{11} - a_{21} > 0$ and $a_{12} - a_{22} < 0$. Similarly, the payoffs of player 2 must satisfy either: **i)** $b_{11} - b_{21} > 0$ and $b_{12} - b_{22} > 0$, or **ii)** $b_{11} - b_{21} > 0$ and $b_{12} - b_{22} < 0$.. Without loss of generality, assume condition 1 holds for the payoffs of both players. The proofs in the remaining four cases are identical.

Let $\delta > 0$ be arbitrarily small and pick $\underline{\eta}^B \in (\beta - \delta,\ \beta)$, $\bar{\eta}^B \in (\beta,\ \beta + \delta)$, $\underline{\eta}^A \in (\alpha - \delta,\ \alpha)$, and $\bar{\eta}^A \in (\alpha,\ \alpha + \delta)$. We want to show:

1. If $X_B^A(1) \sim \text{Bin}(n, \underline{\eta}^B)$, then $\lim_{n\to\infty} \Pr\big[U_A(1|X_B^A(1),\, \Gamma) - U_A(2|X_B^A(1),\, \Gamma) \geq 0\big] = 0$;
2. If $X_B^A(1) \sim \text{Bin}(n, \bar{\eta}^B)$, then $\lim_{n\to\infty} \Pr\big[U_A(1|X_B^A(1),\, \Gamma) - U_A(2|X_B^A(1),\, \Gamma) \geq 0\big] = 1$;
3. If $X_A^B(1) \sim \text{Bin}(n, \underline{\eta}^A)$, then $\lim_{n\to\infty} \Pr\big[U_B(1|X_A^B(1),\, \Gamma) - U_B(2|X_A^B(1),\, \Gamma) \geq 0\big] = 0$;
4. If $X_A^B(1) \sim \text{Bin}(n, \bar{\eta}^A)$, then $\lim_{n\to\infty} \Pr\big[U_B(1|X_A^B(1),\, \Gamma) - U_B(2|X_A^B(1),\, \Gamma) \geq 0\big] = 1$.

Given 1-4, the continuity of $\Pr\big[U_A(1|X_B^A(1),\, \Gamma) - U_A(2|X_B^A(1),\, \Gamma) \geq 0\big]$ with respect to $\eta^B(1)$, and the continuity of $\Pr\big[U_B(1|X_A^B(1),\, \Gamma) - U_B(2|X_A^B(1),\, \Gamma) \geq 0\big]$ with respect to $\eta^A(1)$, the result follows as $\delta > 0$ can be picked arbitrarily small.

We'll start by proving 1. Suppose $X_B^A(1) \sim \text{Bin}(n, \underline{\eta}^B)$ and notice that:

$$
\begin{aligned}
&U_A(1|X_B^A(1),\, \Gamma) - U_A(2|X_B^A(1),\, \Gamma) = \\
&\qquad \int_0^1 \Big[\eta_1^B(a_{11} - a_{21}) + (1 - \eta_1^B)(a_{12} - a_{22})\Big] \cdot p_i(\eta_1^B|X_B^A(1),\, \Gamma)\, \mathrm{d}\eta_1^B \\
&\qquad = a_{12} - a_{22} + [a_{11} - a_{21} + a_{22} - a_{12}] \int_0^1 \eta_1^B \cdot p_i(\eta_1^B|X_B^A(1),\, \Gamma)\, \mathrm{d}\eta_1^B
\end{aligned}
$$

We proceed to show that $\int \eta_1^B \cdot p_i(\eta_1^B|X_B^A(1),\, \Gamma)\, \mathrm{d}\eta_1^B \xrightarrow{\text{p}} \underline{\eta}^B$. Let $\phi(\cdot)$ denote the pdf of a normal distribution with mean $n^{-1}X_B^A(1)$ and variance $n^{-2}\underline{\eta}^B(1 - \underline{\eta}^B)$. Notice that:

$$
\begin{aligned}
&\int_0^1 \eta_1^B \cdot p_i(\eta_1^B|X_B^A(1),\, \Gamma)\, \mathrm{d}\eta_1^B = \\
&\qquad \int_0^1 \eta_1^B \cdot \phi(\eta_1^B)\, \mathrm{d}\eta_1^B + \int_0^1 \eta_1^B \cdot \Big[p_i(\eta_1^B|X_B^A(1),\, \Gamma) - \phi(\eta_1^B)\Big]\, \mathrm{d}\eta_1^B \equiv A + B
\end{aligned}
$$

Let $Z = \Phi(1) - \Phi(0)$ where $\Phi(\cdot)$ is the cdf of a normal distribution with mean $n^{-1}X_B^A(1)$ and variance $n^{-2}\underline{\eta}^B(1 - \underline{\eta}^B)$. Thus,

$$
A = Z \int_0^1 \eta_1^B \cdot \frac{\phi(\eta_1^B)}{Z}\, \mathrm{d}\eta_1^B
$$

Since $\frac{\phi(\eta_1^B)}{Z}$ is the pdf of a truncated normal distribution on the interval $[0,1]$, we have that:

$$A = Z \cdot \left[n^{-1} X_B^A(1) + \frac{\phi(0) - \phi(1)}{Z} \cdot n^{-1} \sqrt{\underline{\eta}^B (1 - \underline{\eta}^B)} \right]$$

By a very similar argument as the one in the proof of Theorem 1, it can be shown that $Z \xrightarrow{\text{p}} 1$ and therefore, as a consequence of the Law of Large Numbers and Slutsky's Theorem, we have that $A \xrightarrow{\text{p}} \underline{\eta}^B$. For the case of B, notice that:

$$\begin{aligned} |B| &\leq \int_0^1 \eta_1^B \cdot \left| p_i(\eta_1^B | X_B^A(1), \Gamma) - \phi(\eta_1^B) \right| \mathrm{d}\eta_1^B \leq \int_0^1 \left| p_i(\eta_1^B | X_B^A(1), \Gamma) - \phi(\eta_1^B) \right| \mathrm{d}\eta_1^B \\ &= \int_{-\infty}^{\infty} \left| p_i(\eta_1^B | X_B^A(1), \Gamma) - \phi(\eta_1^B) \right| \mathrm{d}\eta_1^B - \int_{-\infty}^{0} \phi(\eta_1^B)\,\mathrm{d}\eta_1^B - \int_1^{\infty} \phi(\eta_1^B)\,\mathrm{d}\eta_1^B \\ &= 2\left|\left| p_i(\eta_1^B | X_B^A(1), \Gamma) - \phi(\eta_1^B) \right|\right|_{TV} - (1 - Z) \end{aligned}$$

And therefore, as a consequence of Bernstein-von Mises Theorem, we have that $|B| \xrightarrow{\text{p}} 0$ and consequently $B \xrightarrow{\text{p}} 0$. Finally, notice that:

$$\begin{aligned} U_A(1|X_B^A(1), \Gamma) - U_A(2|X_B^A(1), \Gamma) &\xrightarrow{\text{p}} a_{12} - a_{22} + [a_{11} - a_{21} + a_{22} - a_{12}]\underline{\eta}^B \\ &< a_{12} - a_{22} + [a_{11} - a_{21} + a_{22} - a_{12}] \cdot \beta = 0 \end{aligned}$$

And therefore, when $X_B^A(1) \sim \text{Bin}(n, \underline{\eta}^B)$, then $\lim_{n\to\infty} \Pr\left[U_A(1|X_B^A(1), \Gamma) - U_A(2|X_B^A(1), \Gamma) \geq 0\right] = 0$. The three remaining statements, 2-4, are proved in an analogous way.

Case 2: [Pure Equilibrium] Now let $\sigma = (\alpha, 1-\alpha; \beta, 1-\beta)$ represent an equilibrium in pure strategies of game Γ. Without loss of generality, assume that $\alpha = \beta = 1$. We proceed to prove that, for large enough n, the share vectors $\left(\eta^A(1), \eta^A(2)\right) = \left(\eta^B(1), \eta^B(2)\right) = (1, 0)$ form part of a SBE in symmetric strategies of game Γ_n. Since $\sigma = (1,0;1,0)$ is an equilibrium of Γ, we have two possible cases for the payoffs of player 1: **i)** $a_{11} - a_{21} > 0$ and $a_{12} - a_{22} > 0$, or **ii)** $a_{11} - a_{21} > 0$ and $a_{12} - a_{22} < 0$. We proceed separately for each case:

Case 2.i: Player 1 has a dominant strategy and therefore, for every realization of signal $X_B^A(1)$, $U_A(1|X_B^A(1), \Gamma) - U_A(2|X_B^A(1), \Gamma) > 0$. Applying the definition of an SBE, we have that $\eta^A(1) = \Pr[U_A(1|X_B^A(1), \Gamma) - U_A(2|X_B^A(1), \Gamma) > 0] = 1$.

Case 2.ii: In a SBE, the strategy for player 1 follows a threshold, i.e. $s_i(x_B^A(1)) = 1$ if and only if $x_B^A(1) \geq \bar{X}$. Using the limits found in 2 and by means of a similar contradiction argument as the one employed in the proof of Theorem 1, it can be shown that: for large enough n, $\bar{X} \leq n$. Thus, when $\eta^B(1) = 1$ and $X_B^A(1) \sim \text{Bin}(n, \underline{\eta}^B)$, we have that $\eta^A(1) = \Pr[U_A(1|X_B^A(1), \Gamma) - U_A(2|X_B^A(1), \Gamma) > 0] \geq \Pr[X_B^A(1) \geq n] = 1$.

Therefore, in either of the two cases, the share vector $\left(\eta^A(1), \eta^A(2)\right) = (1, 0)$ is part of an SBE. The analysis for player 2 and share vector $\left(\eta^B(1), \eta^B(2)\right)$ is completely symmetrical.

□

References

Angeletos, G., Hellwig, C., and Pavan, A. (2006). Signaling in a global game: Coordination and policy traps. *Journal of Political Economy*, 114(3):452–484.

Angeletos, G.-M. and Werning, I. (2006). Crises and prices: Information aggregation, multiplicity and volatility. *The American Economic Review*, 96(5):1720–1736.

Aumann, R. J., Katznelson, Y., Radner, R., Rosenthal, R. W., and Weiss, B. (1983). Approximate purification of mixed strategies. *Mathematics of operations research*, 8(3):327–341.

Barelli, P. and Duggan, J. (2015). Purification of bayes nash equilibrium with correlated types and interdependent payoffs. *Games and Economic Behavior*, 94:1–14.

Bhaskar, V., Mailath, G. J., and Morris, S. (2008). Purification in the infinitely-repeated prisoners' dilemma. *Review of Economic Dynamics*, 11(3):515–528.

Cañón, C. and Margaretic, P. (2014). Correlated bank runs, interbank markets and reserve requirements. *Journal of Banking & Finance*, 49(C):515–533.

Cole, H. L. and Kehoe, T. J. (2000). Self-fulfilling debt crises. *The Review of Economic Studies*, 67(1):91–116.

Damme, E. v. (1991). *Stability and Perfection of Nash Equilibria*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2. edition.

Dasgupta, A. (2007). Coordination and delay in global games. *Journal of Economic Theory*, 134:195–225.

Edmond, C. (2005). Information and the limits to autocracy.

Foster, D. P. and Vohra, R. V. (1997). Calibrated learning and correlated equilibrium. *Games and economic behavior*, 21(1-2):40–55.

Fudenberg, D. and Kreps, D. M. (1993). Learning mixed equilibria. *Games and economic behavior*, 5(3):320–367.

Fudenberg, D. and Levine, D. K. (2009). Learning and equilibrium. *Annual review of economics*, 1(1):385–419.

Goldstein, I. and Pauzner, A. (2005). Demand–Deposit Contracts and the Probability of Bank Runs. *Journal of Finance*, 60(3):1293–1327.

Govindan, S., Reny, P. J., and Robson, A. J. (2003). A short proof of harsanyi's purification theorem. *Games and economic behavior*, 45(2):369–374.

Harsanyi, J. C. (1973). Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, 2:1–23.

Klemens, B. (2007). When do ordered prior distributions induce ordered posterior distributions? *SSRN Electronic Journal*.

Morris, S. and Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 88(3):587–97.

Morris, S. and Shin, H. S. (1999). Private versus public information in coordination problems. *Unplubished*.

Morris, S. and Shin, H. S. (2006). Endogenous public signals and coordination. *Working Paper*.

Nash, J. F. (1950). Non-cooperative games. *PhD Thesis, Princeton University*.

Nickl, R. (2013). Statistical theory. *Statistical Laboratory, Department of Pure Mathematics and Mathematical Statistics, University of Cambridge*.

Obstfeld, M. (1986). Rational and self-fulfilling balance-of-payments crises. *American Economic Review*, 76(1):72–81.

Obstfeld, M. (1996). Models of currency crises with self-fulfilling features. *European Economic Review*, 40(3-5):1037–1047.

Osborne, M. J. and Rubinstein, A. (1998). Games with procedurally rational players. *The American economic review*, 88(4):834–847.

Osborne, M. J. and Rubinstein, A. (2003). Sampling equilibrium, with an application to strategic voting. *Games and economic behavior*, 45(2):434–441.

Radner, R. and Rosenthal, R. W. (1982). Private information and pure-strategy equilibria. *Mathematics of Operations Research*, 7(3):401–409.

Rochet, J.-C. and Vives, X. (2004). Coordination Failures and the Lender of Last Resort : Was Bagehot Right After All? IDEI Working Papers 294, Institut d'Économie Industrielle (IDEI), Toulouse.

Tarashev, N. (2003). Currency crises and the informational role of interest rates. *Working Paper, Bank of International Settlements.*

Tsyvinski, A., Hellwig, C., and Mukherji, A. (2004). Coordination failures and asset prices. *Working Paper.*

von Neumann, J. and Morgenstern, O. (1947). *Theory of games and economic behavior.* Princeton University Press.