Dynamic Sender-Receiver Games

Name: Atulya Jain
Advisor: Prof. Nicolas Vieille

September 9, 2018

Contents

Chapter 1

Dynamic Sender-Receiver Game

1.1 Introduction

A sender-receiver game is a special class of game of incomplete information. It models strategic interactions between an informed expert and an uninformed decision maker. It consists of a set of states S. A state s is chosen randomly by nature with a prior probability p. The sender is informed about the state s while the receiver is not. The sender then chooses a message $a \in A$ to be sent to the receiver. The receiver, who is unaware about the true state, after receiving the message a plays action $b \in B$. The sender and the receiver get payoff $u^1(s,b)$ and $u^2(s,b)$ respectively. Note that the payoff function does not depend on the message a sent by the sender.

The behaviour strategy σ of the sender is a mapping from the set of states S to the set of mixed messages ΔA ($\sigma : S \rightarrow \Delta A$). While the behaviour strategy τ of the receiver is a mapping from the set of messages A to the set of mixed actions ΔB ($\tau : A \rightarrow \Delta B$).

There always exists a *babbling equilibrium* where the sender transmits a constant message that is independent of the true state. The receiver's strategy is to play the action that maximizes the expected payoff given prior probability, irrespective of the sender's message. The babbling payoff v_2 is given by:

$$v_2 = \max_{b \in B} \sum_{s \in S} p(s) u(s,b)$$

Apart from this, there might also exist other equilibria. A *fully revealing equilibrium* corresponds to the sender sending a different signal m_s for each state s, i.e, $\sigma(\theta_s) = m_s$ and if $\theta_s \neq \theta_s^{'}$ then $m_s \neq m_s^{'}$. A *partially revealing equilibrium* corresponds to one where the sender discloses some information regarding the state. The transmission is not 1-1 like the fully revealing equilibrium but also not a constant mapping like the babbling equilibrium. Consider the following game :

		Receiver		
		L	M	R
	s_0	$1,1$	$0,0$	$0,0$
States	s_1	$0,0$	$1,0$	$0,1$
	s_2	$0,0$	$0,1$	$1,0$

The set of states is given by $S = \{s_0, s_1, s_2\}$ and each state has an equal probability of being chosen. The set of messages and actions is given by $A = \{\hat{s_0}, \hat{s_1}, \hat{s_2}\}$ and $B = \{L, M, R\}$

respectively. Consider the babbling equilibrium, the sender does not reveal anything about the states. All the actions L, M and R of the receiver achieve the same payoff. Thus, the babbling equilibrium payoff is given by $(\frac{1}{3}, \frac{1}{3})$. Next, consider the the fully revealing strategy of the sender. The sender sends the message $\hat{s_0}, \hat{s_1}$ and $\hat{s_2}$ at the states s_0, s_1 and s_2 respectively. The best reply of the player is to play L at $\hat{s_0}$, R at $\hat{s_1}$ and M at $\hat{s_2}$. Notice that this is not a equilibrium, as the sender strictly prefers sending message $\hat{s_0}$ as compared to $\hat{s_1}$ or $\hat{s_2}$. Hence, there is no fully revealing equilibrium. Finally, consider the strategy where the sender reveals if the state is s_0 or not. So, $\sigma(s_0) = m_1$ and $\sigma(s_1) = (s_2) = m_2$. The best response of the receiver is given by $\tau(m_1) = L$ and $\tau(m_2) = M$ or $\tau(m_2) = R$. The receiver plays the best action corresponding to his belief of the true state. The sender also is indifferent between sending the message m_1 or m_2. Hence, this is an equilibrium.

The first model of sender-receiver game was developed by Crawford and Sobel in [1]. The model described information transmission through a noisy channel. In [2], Golosov et al. consider repeated rounds of play. The sender sends messages repeated to the receiver who then takes actions. They consider the case when the state of the nature remains fixed. Renault, Solan and Vieille consider a dynamic version of this repeated play in [7]. In this model the state process is not constant but rather follows a Markov chain. They are able to characterize the limit set of equilibrium payoffs in the model. The equilibrium payoffs satisfy an individual rationality condition for the receiver, and an incentive compatibility condition for the sender.

In this report, we obtain the same characterization as in [7] but in the case of uniform equilibria. We also try to extend the model to the case of imperfect monitoring. In this model, the sender does not directly observe the action of the receiver but rather observes an action dependent signal.

1.2 Model

At each stage n, the game is in state s_n. The sender is aware of the true state and sends a message $a_n \in A$ to the receiver. The receiver, who does not know the true state, takes an action $b_n \in B$ based on the sender's message. The action b_n is publicly disclosed and players get stage payoff $u^1(s_n, b_n)$ and $u^2(s_n, b_n)$ respectively. The game then proceeds to the next stage s_{n+1} according to the transition matrix $P(s_{n+1} \mid s_n)$. The state process $\{s_n\}_{n\in\mathbb{N}}$ follows a Markov chain, which is irreducible and aperiodic. The unique invariant measure is given by m.

The game is determined by the following components:

- **S**: States of nature.
- **A**: action set for sender, **B**: action set for receiver.
- $\boldsymbol{u(s,b)} \in \mathbb{R}^2$, where $u^1(s,b)$ gives the stage payoff for the sender and $u^2(s,b)$ gives the stage payoff for the receiver.
- $\boldsymbol{P} \in \mathbb{R}^{S\times S}$: Transition matrix

The strategies of the sender and receiver map information they have to their possible actions. For the sender this is a mapping from the set of past and current states along with the past actions of both the players to possible messages: $\sigma_n : (S\times A\times B)^{n-1}\times S \to \Delta A$. For the receiver it is a function which maps past messages and actions to possible actions $\tau_n : (A\times B)^{n-1} \to \Delta B$.

Denote by $\sigma_1(a_1 \mid s_1)$ the probability that the sender chooses a_1 in the first stage when the state is s_1 and by $\sigma_n(a_n \mid s_1, a_1, b_1, ..., s_{n-1}.a_{n-1}, b_{n-1})$ the probability that the sender chooses a_n in the n^{th} stage conditioned that the previous states and actions were $s_1, a_1, b_1, ...s_{n-1}, a_{n-1}, b_{n-1}$. In case of the receiver, the action is conditioned only on past actions and not states. Using this, we define the probability measure over the n-period history H_n. The probability that the sequence of states and actions $h_n = s_1, a_1, b_1, ...s_n, a_n, b_n$ takes place is given by the following equation.

$$\mathbb{P}_{\sigma,\tau}(h_n) = P(s_1) \times \sigma_1(a_1 \mid s_1) \times \tau_1(b_1|a_1) \times .... \times P(s_n \mid s_{n-1}) \times \sigma_t(a_n \mid h_{n-1}, s_n) \times \tau_n(b_n \mid a_1, b_1, ..., a_n)$$

The probability measure for the the $T-$period repeated game forms a basis for the infinite game. Let us denote with H, the collection of all plays in the infinite repeated game. Each play is an infinite sequence of states and actions $\{s_n, a_n, b_n\}_{\{n\in\mathbb{N}\}}$. Let $(X_n)_{n\in\mathbb{N}}$ be a sequence of finite sets, and let $X^\infty = \prod_{n\in\mathbb{N}} X_n$. An element $Y \in X^\infty$ is called a cylinder set if there exists $N \in \mathbb{N}$ and $(Y_n)_{n=1}^N$ such that $Y_n \subset X_n$ and $B = (\prod_{n=1}^N Y_n) \times (\prod_{n=N+1}^\infty X^n)$. The σ-algebra $\mathcal{Y}$ generated by the cylinder sets forms the measurable space over X^∞. Using this $\sigma-$algebra there exists a unique extension extending the probability distribution $P^n_{\sigma,\tau}$ over finite history X_n to the probability distribution $\mathbb{P}_{\sigma,\tau}$ over the measurable space $(X^\infty, \mathcal{Y})$ such that

$$\mathbb{P}^n_{\sigma,\tau}(A) = \mathbb{P}_{\sigma,\tau}(A \times X_{n+1} \times ....) \quad \forall n \in \mathbb{N}, \forall A \in X^n$$

This defines a well defined probability measure over the entire play. In our model, $X_n = S_n \times A_n \times B_n$ and $X^\infty = \prod_{n\in\mathbb{N}}(S_n \times A_n \times B_n)$. The payoff is given by the expectation of the $\delta-$discounted payoff:

$$\gamma_\delta(\sigma, \tau) = \boldsymbol{E}_{\sigma,\tau}[(1-\delta)\sum_{n=1}^{\infty} \delta^{n-1} u(\boldsymbol{s}_n, \boldsymbol{b}_n)]$$

1.3 Results

Consider $\mathcal{M} \subset \Delta(S \times A)$ such that the marginal distributions over S and A are equal to m, the invariant measure of the Markov chain. Given a copula $\mu \in \mathcal{M}$ and a stationary strategy $y : A \to \Delta B$ we define:

$$U(\mu, y) := \sum_{s\in S, a\in A} \mu(s, a) u(s, y(\cdot \mid a)) \in \mathbb{R}^2 \tag{1.1}$$

This corresponds to the expected (undiscounted) payoff when the sender and receiver use stationary strategies given by $\mu(\cdot \mid s)$ and y respectively. And let us denote by v_2 the babbling equilibrium payoff for the receiver.

$$v_2 = \max_{b\in B} \sum_{s\in S} m(s) u^2(s, b) \tag{1.2}$$

Denote by $E(\mathcal{M})$, the payoff vectors $U(\mu_0, y)$ that satisfy

C1. $U^1(\mu_0, y) \geq U^1(\mu, y)$ for every $\mu \in \mathcal{M}$ (Incentive Compatibility)

C2. $U^2(\mu_0, y) \geq v^2$ (Individual Rationality)

A **uniform equilibrium** payoff $\gamma_* = (\gamma_*^1, \gamma_*^2)$ is a vector such that for every $\epsilon > 0$ there exists a δ_0 and strategy profile $(\sigma_\epsilon, \tau_\epsilon)$ such that $\forall \delta > \delta_0$

$$\gamma_\delta^1(\sigma_\epsilon, \tau_\epsilon) + \epsilon \geq \gamma_*^1 \geq \gamma_\delta^1(\sigma', \tau_\epsilon) - \epsilon \quad \forall \sigma' \tag{1.3}$$

$$\gamma_\delta^2(\sigma_\epsilon, \tau_\epsilon) + \epsilon \geq \gamma_*^2 \geq \gamma_\delta^2(\sigma_\epsilon, \tau') - \epsilon \quad \forall \tau' \tag{1.4}$$

Let us denoted by UEQ the set of all uniform equilibria of the game. In Theorem 1, we shall prove that the set of uniform equilibria UEQ contains $E(\mathcal{M})$. In Theorem 2, we shall show the converse holds under certain assumptions.

1.3.1 Example

We now present an example presented in [7] to demonstrate the results. Consider the following game:

		Receiver		
		L	M	R
	s_0	$1, 1$	$0, 0$	$0, 0$
States	s_1	$0, 0$	$1, 0$	$0, 1$
	s_2	$0, 0$	$0, 1$	$1, 0$

There are three states s_0, s_1 and s_2. We assume each successive stage is independent and identically distribution. Each state is equally likely to be picked. Thus, we have $P(s_{n+1} \mid s_n) = \frac{1}{3}$ for each $s_n, s_{n+1} \in S$ The initial state distribution is given by the invariant measure $m(s) = \frac{1}{3} \quad \forall s \in S$. To obtain the extreme points of set of feasible payoffs, we consider all the strategies where the sender tells the truth and the receiver plays a pure state dependent action.

$$y_0(s_0) = M, \quad y_0(s_1) = L, \quad y_0(s_2) = L \qquad U(\mu_0, y_0) = (0, 0)$$

$$y_1(s_0) = L, \quad y_1(s_1) = M, \quad y_1(s_2) = R \qquad U(\mu_0, y_1) = (1, \frac{1}{3})$$

$$y_2(s_0) = L, \quad y_2(s_1) = R, \quad y_2(s_2) = M \qquad U(\mu_0, y_2) = (\frac{1}{3}, 1)$$

$$y_3(s_0) = M, \quad y_3(s_1) = M, \quad y_3(s_2) = R \qquad U(\mu_0, y_3) = (\frac{2}{3}, 0)$$

$$y_4(s_0) = M, \quad y_4(s_1) = R, \quad y_4(s_2) = M \qquad U(\mu_0, y_4) = (0, \frac{2}{3})$$

The set of feasible payoff is given by conv(S) where $S = \{(0,0), (1, \frac{1}{3}), (\frac{1}{3}, 1), (\frac{2}{3}, 0), (0, \frac{2}{3})\}$. Now, we get the set of equilibrium payoffs $E(\mathcal{M})$. We first calculate the babbling payoff v_2. When the sender babbles, all the receiver's action results in the same payoff, which equals $\frac{1}{3}$. So, the condition **C2** becomes $\gamma^2 \geq \frac{1}{3}$. For the sender, we use the symmetry of the payoff matrix to obtain the condition **C1**. We first show that its necessary to have $\gamma^1 \geq \gamma^2$. We then show that this condition is sufficient. Consider the copula $\mu^{'}$ where the states s_1 and s_2 are exchanged. From the payoff matrix, we have $U^1(\mu^{'}, y) = U^2(\mu_0, y)$ and $U^1(\mu_0, y) = U^2(\mu^{'}, y)$.

The condition **C2** implies $\gamma^1 \geq \gamma^2$. For the sufficiency part, we show that the extreme points of the region satisfying $\gamma^2 \geq \frac{1}{3}$ and $\gamma^1 \geq \gamma^2$ indeed belong to $E(\mathcal{M})$. The extreme points is given by the set $C = \{(\frac{1}{3}, \frac{1}{3}), (1, \frac{1}{3}), (\frac{2}{3}, \frac{2}{3})\}$. Firstly, the payoff $(\frac{1}{3}, \frac{1}{3})$ corresponds to the babbling payoff and satisfies **C1** and **C2**. The payoff $(1, \frac{1}{3})$, obtained by playing y_1, is the maximum payoff the sender can achieve and so we have $U^1(\mu_0, y_1) \geq U^1(\mu, y_1) \quad \forall \mu \in \mathcal{M}$. The payoff $(\frac{2}{3}, \frac{2}{3})$ is achieved when the receiver plays L in the state s_0 and plays the mixed action $\frac{M+R}{2}$ in the states s_1 and s_2. The key point is that the sender cannot profit by exchanging s_0 with either s_1 or s_2. Coupled with the fact that the receiver uses the same strategy for s_1 and s_2, we see that $(\frac{2}{3}, \frac{2}{3})$ indeed satisfies condition **C1** and belongs to $E(\mathcal{M})$. The set of equilibrium payoffs $E(\mathcal{M})$ is given by the dark triangle in the figure below.

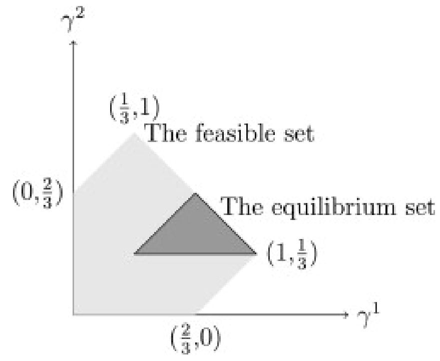

1.4 Proofs

1.4.1 Theorem 1

$$E(\mathcal{M}) \subset UEQ$$

Given a payoff $U(\mu_0, y)$ and $\epsilon > 0$ we will construct a strategy profile (σ, τ) such that for every $\delta \geq \delta_0$ it is an $\epsilon-$equilibrium in the $\delta-$discounted infinite game. The set of stages are divided into consecutive blocks of increasing lengths $B^1, ..., B^m$, where $m \in \mathbb{N}$. After each block, the players forget the past states and actions, and start fresh. We use block strategies and divide the play into "normal blocks" and "punishment blocks". Block strategies were first used to obtain folk theorem with uniform equilibria in [5] and [6] by Lehrer.

In the normal block, the strategy σ of the sender is to *speak the truth* and report the true state s. While the strategy τ of the receiver is to *listen* to the sender's announcement a in stage n and play the mixed action $y(\cdot \mid a)$, as long as no state is reported too often. As soon a state is reported too often, the receiver replaces the sender's report with a *fictitious report* θ and plays $y(\cdot \mid \theta)$.

At the end of the block the sender imposes a statistical test to check whether the receiver has deviated or not from his described strategy. If the receiver has not deviated the game proceeds to another normal block, else the game proceeds to a punishment phase. This phase consists of a series of punishment blocks. After the punishment phase, the game again returns to a normal block. The increase in block lengths allow the statistical test to get more accurate as the game progresses.

In the punishment phase, the sender does not provide any information about the state. The receiver resorts to playing the action corresponding to the babbling payoff v_2. If the receiver fails the test in block m, then he is punished for the next m^2 blocks.

We now define the construction of the fictitious states. Let $m_N \in \Delta S$ be the best "integral" approximation to the invariant measure m. By this we mean that out of all distributions $m' \in \Delta S$ such that $Nm'(s)$ is an integer $\forall s \in S$, m_N is the best approximation to m. For each $s \in S$ denote by $\boldsymbol{N}_n(s) = |\{k \leq n : a_k = s\}|$ the number of stages till stage n when state s was reported. The key idea is that the receiver keeps a check on the number of times a state has been reported. If a reported state crosses its quota, then the receiver substitutes the reported state with the fictitious state. This enforces that the distribution of announcements is always equal to m_N. The threshold for a $N-$stage block is defined as:

$$q := \min\{1 \leq n \leq N : \boldsymbol{N}_n(a_n) > N m_N(a_n)\}$$

The sequence θ_n satisfies:
F1. $\theta_n = a_n$ for $n < q$
F2. For each $s \in S$, the equality $|\{n \leq N : \theta_n = s\}| = N m_N(s)$ always hold.
F3. Conditional on $(a_1, ..a_q)$ the variables $(\theta_q, ..., \theta_N)$ are deterministic.

We will refer to θ_n as the announcement at stage n.

On the other hand, the sender at the end of the each normal block checks if the the distribution of the receiver's action is within an ϵ_m-neighbourhood of the *expected* distribution. The sender counts the number of times the receiver plays action b when the announcement was s. To do this, we define the random variable $\boldsymbol{N^s}(b) = |\{k \leq |B^m| : \theta_k = s, b_k = b\}|$, the number of stages in the block m where the action was b corresponding to the announcement s. The sender expects action b to be played $m(s)y(b|s)|B^m|$ times for the announcement s during the block. To avoid punishment, the receiver's strategy has to satisfy the following test:

$$|\frac{1}{|B^m|}N^s(b) - m(s)y(b \mid s)| \leq \epsilon_m \quad \forall s \in S, b \in B \tag{1.5}$$

We show below that by choosing appropriate block lengths and error margins, the strategy profile (σ, τ) ensures that the set of blocks when the play is not in normal blocks is almost surely finite. This will ensure that the strategy profile achieves the payoff $U(\mu_0, y)$ when $\delta \to 1$. We choose $|B^m| = m^{10}$ and $\epsilon_m = \frac{1}{m^2}$.

Lemma

$$\lim_{\delta \to 1} \gamma_\delta(\sigma, \tau) = U(\mu_0, y)$$

Proof For each block m we define the following event

$$\mathcal{D}_m = \{B^m \text{ is a normal block}\}$$

Fix a m where B^{m-1} is a normal block. We now calculate the conditional probability that the next block B^m is not a normal block.

$$\mathbb{P}_{\sigma,\tau}(\mathcal{D}_m^c \mid \mathcal{D}_{m-1}) = \mathbb{P}_{\sigma,\tau}(|\frac{1}{|B^m|}N^s(b) - m(s)y(b \mid s)| \geq \frac{1}{m^2} \mid \mathcal{D}_{m-1}) \text{ for some } s \in S,\, b \in B$$

We use the generalized Tchebychev's inequality proved in [5] to find an upper bound on this term:

Lemma Let $R_1, ..., R_n$ be a sequence of identically distributed Bernoulli random variables, with parameter p. Let $Y_1, ..., Y_n$ be a sequence of Bernoulli random variables such that for each $i < m < n$, R_m is independent of $R_1, ..., R_{m-1}, Y_1 ...., Y_m$. Then

$$\mathbb{P}(|\frac{R_1Y_1 + ... + R_nY_n}{n} - p\frac{Y_1 + ... + Y_N}{n}| \geq \epsilon) \leq \frac{1}{n\epsilon^2} \text{ for every } \epsilon > 0$$

Define $R_i = 1$ if the state is s in stage i else $R_i = 0$. As we have the initial distribution $s_0 = m(s)$, we have that $R_i = 1$ with probability $m(s)$ and is a Bernoulli random variable. Also, define $Y_i = 1$ if the receiver takes the action b given that the state is s in stage i. This again is a Bernoulli random variable with mean $y(b \mid s)$. Define x_i as $\frac{Y_i}{|B^m|}$. We thus have

$$\mathbb{P}_{\sigma,\tau}(|\frac{1}{|B^m|}N^s(b) - m(s)\sum_{n=1}^{|B^m|} x_n| \geq \frac{1}{m^2} \mid \mathcal{D}_{m-1}) \leq \frac{m^4}{|B^m|} = \frac{1}{m^6}$$

We can now use the Borel-Cantelli Theorem to show that number of punishment phases are finite almost surely and that there exists a random $m_1 \in \mathbb{N}$ such that for all $m \geq m_1$ the game is in normal blocks.

Lemma Suppose that $\{A_n : n \geq 1\}$ is a sequence of events in a probability space. If

$$\sum_{n=1}^{\infty} P(A_n) < \infty,$$

then $P(A(i.o.)) = 0$; only a finite number of the events occur, with probability 1.

$$\sum_{n=1}^{\infty} \mathbb{P}(\mathcal{D}_{n+1}^c \mid \mathcal{D}_n) = \sum_{n=1}^{\infty} \frac{1}{n^6} < \infty$$

So we get

$$\mathbb{P}(\lim_{n\to\infty} \mathcal{D}_{n+1}^c \cap \mathcal{D}_n) = 0$$

Hence, after a certain block m_1 there will be no punishment phase. Now, we show the strategy profile (σ, τ) leads to a payoff arbitrary close to $U(\mu_0, y)$ when δ is close to 1. The main idea is that the expected (undiscounted) payoff is equal to $U(\mu_0, y)$ in the normal blocks. Combining this with the fact that the the punishment blocks are almost surely finite, the discounted payoff converges to $U(\mu_0, y)$ as $\delta \to 1$.

We first show that the expected payoff in the normal blocks is in the ϵ-neighbourhood of $U(\mu_0, y)$ when δ tends to 1. Denote by $\mu_{\sigma,\tau}$ the expected (undiscounted) joint distribution of states and announcements in the B^m-stage game:

$$\mu^m_{\sigma,\tau}(s,a) = \boldsymbol{E}_{\sigma,\tau}[\frac{1}{|B^m|}\sum_{n=1}^{|B^m|} 1_{\{s_n=s,\theta_n=a\}}]$$

From the definition of the strategy profile (σ, τ), we get the expected (undiscounted) payoff in the normal block m is equal to $U(\mu^m_{\sigma,\tau}, y)$. When we consider infinite repetitions of the the block strategy, we get that $\gamma^N_\delta(\sigma, \tau)$ converges to $U(\mu_{\sigma,\tau}, y)$ as δ approaches 1, where $\mu_{\sigma,\tau} = \lim_{m\to\infty} \mu^m_{\sigma,\tau}(s, a)$ ($\gamma^N_\delta(\sigma, \tau)$ is the payoff obtained in the normal blocks). We show below that $\mu_{\sigma,\tau}$ is arbitrary close to μ_0.

Now, we combine the above result with the fact that the number of punishment blocks are finite almost surely. We show below that the total T-stage payoff converges to $U(\mu_{\sigma,\tau}, y)$ and use the lemma proved in [8] to show the $\delta-$discounted payoff converges to the same limit.

Lemma Given an arbitrary sequence of real numbers $(x^t)_{t\in\mathbb{N}}$, let $\bar{x}^T = \frac{1}{T}\sum_{t=1}^T x^t$ and $\bar{x}^\delta = (1-\delta)\sum_{n=1}^\infty \delta^{t-1}x^t$. Then

$$\limsup_{T\to\infty} \bar{x}^T \geq \limsup_{\delta\to 1} \bar{x}^\delta \geq \liminf_{\delta\to 1} \bar{x}^\delta \geq \liminf_{T\to\infty} \bar{x}^T \tag{1.6}$$

We have showed that after a certain random $m_1 \in \mathbb{N}$, only normal blocks remain. Thus, in the infinite set of blocks, the number of punishments blocks are finite. And the expected payoff during a normal block is equal to $U(\mu_{\sigma,\tau}, y)$. Hence, we can conclude that

$$\begin{aligned} &\lim_{T\to\infty} \bar{\gamma}^T(\sigma,\tau) = U(\mu_{\sigma,\tau}, y) \\ \Rightarrow\quad &\lim_{\delta\to 1} \bar{\gamma}^\delta(\sigma,\tau) = U(\mu_{\sigma,\tau}, y) \end{aligned}$$

We also know the reported and fictitious states match almost surely, i.e, $\|\mu_{\sigma,\tau} - \mu_0\| \leq \epsilon$. This follows from the *Ergodic theorem for strongly stationary processes* presented in [3].

Lemma Let $X = \{X_n : n \geq 1\}$ be a strongly stationary process such that $\mathbb{E}|X_1| \leq \infty$. There exists a random variable Y with the same mean as the X_n such that

$$\frac{1}{n}\sum_{j=1}^n X_j \to Y \, a.s \text{ and in mean} \tag{1.7}$$

The process $X = \{X(t) : t \in \mathbb{N}\}$ taking values in $\mathbb{R}$, is called *strongly stationary* if the families $\{X(t_1), X(t_2), ..., X(t_n)\}$ and $\{X(t_1+h), X(t_2+h), ...X(t_n+h)\}$ have the same joint distribution for all $t_1, t_2, ..., t_h$ and $h \in \mathbb{N}$. Under the assumption that $X_0 = m$, it can be shown that the Markov chain is strongly stationary and converges almost surely to the invariant measure m. Thus, we have that for any $\epsilon > 0$ and η, there exists N_0 such that $\forall N \geq N_0$.

$$\mathbb{P}(|\frac{1}{N}\sum_{n=1}^{N} 1_{\{s_n=s\}} - m(s)| > \epsilon) < \eta \quad \forall s \in S$$
$$\frac{1}{N}\sum_{n=1}^{N} 1_{\{s_n=s\}} \to m(s) \quad a.s$$

Now, let's calculate the probability when the joint distributions don't match, i.e, $\mathbb{P}(|\mu_{\sigma,\tau} - \mu_0| > \epsilon)$

$$\begin{aligned} & \mathbb{P}(|\frac{1}{|B^m|}\sum_{n=1}^{|B^m|}\sum_{a\in A} 1_{\{s_n=s,\theta_n=a\}} - \frac{1}{|B^m|}\sum_{n=1}^{|B^m|} 1_{\{s_n=s,\theta_n=s\}}| > \epsilon) \quad \forall s \in S \\ = \; & \mathbb{P}(\frac{1}{|B^m|}\sum_{n=1}^{|B^m|}\sum_{a\in A} 1_{\{s_n=s,\theta_n\neq s\}} > \epsilon) \quad \forall s \in S \\ \leq \; & \mathbb{P}(|\frac{1}{|B^m|}\sum_{n=1}^{|B^m|} 1_{\{s_n=s\}} - m(s)| > \frac{\epsilon}{M}) \leq \eta^{'} \quad \forall s \in S \end{aligned}$$

Hence, after a certain block $m_2 \in \mathbb{N}$, we have the true and reported states matching almost surely. This combined with the convergence of the payoff gives us the following relations.

$$\|\gamma_\delta(\sigma,\tau) - U(\mu_{\sigma,\tau}, y)\| \leq \epsilon_1 \quad \forall \delta \geq \delta_0 \tag{1.8}$$
$$\|U(\mu_0, y) - U(\mu_{\sigma,\tau}, y)\| \leq \epsilon_2 \tag{1.9}$$

Hence, using (1.8) and (1.9), we get that $\gamma_\delta(\sigma,\tau)$ can be brought arbitrary close to $U(\mu_0, y)$ when δ tends to 1.

Now that we have shown that the strategy profile (σ,τ) achieves the limiting payoff $U(\mu_0, y)$, we need to show that it is a uniform equilibrium. We first consider the deviations of the receiver. To do this, we first show that if any strategy of the receiver passes the test in block B^m, then the payoff in the block is close to the payoff obtained when the players use (σ,τ). We then show that if receiver gains in a block by failing the test, then the punishment is efficient in restricting the future payoff. In the case of the sender, we use condition **C1** to enforce that truth telling is indeed optimal.

Theorem The payoff vector $U(\mu_0, y)$ is a uniform equilibrium.

Proof We show that the receiver has no profitable deviation. First consider a strategy $\tau^{'}$ of the receiver. It can be complex and thus depending on the random sequence of reports, it may or may not pass the test. We show that if $\tau^{'}$ passes the test in a block B^m , then the expected gain in payoff cannot be much higher than when the receiver plays τ. We then show that if the receiver gains a lot in the block B^m by failing the test, the punishment is efficient to cut down on future payoffs.

We first show that the strategy $\tau^{'}$ can be approximated by a stationary strategy $\bar{y}$ to give a similar payoff. Assuming $\tau^{'}$ does not fail the statistical test in block B^m, we show that $\bar{y}$ is close to y. Assuming the strategy passes the test, we have the following relation

$$|\frac{1}{|B^m|}\sum_{n=1}^{|B^m|} 1_{\{s_n=s,b_n=b\}} - m(s)y(b \mid s)| \leq \frac{1}{m^2}$$

We define the stationary strategy $\bar{y}$ using the joint (undiscounted) distribution of states and actions:

$$\bar{y}(b \mid s) = \frac{1}{m(s)|B^m|}\sum_{n=1}^{|B^m|} 1_{\{s_n=s,b_n=b\}}$$
$$\|y(b \mid s) - \bar{y}(b \mid s)\| \leq \frac{1}{m^2 m(s)} \quad \forall b \in B, \forall s \in S$$

Now, let's calculate the difference in payoffs in the block B^m when the receiver uses $\bar{y}$ instead of y:

$$\begin{aligned}
=&(1-\delta)\sum_{n=1}^{m^{10}}\delta^{n-1}u(s_n,\bar{y}) - (1-\delta)\sum_{n=1}^{m^{10}}\delta^{n-1}u(s_n,y)\\
=&(1-\delta)\sum_{n=1}^{m^{10}}\delta^{n-1}[u(s_n,\bar{y}) - u(s_n,y)]\\
=&(1-\delta)\sum_{n=1}^{m^{10}}\delta^{n-1}\sum_{b\in B}u(s_n,b)[\bar{y}(b \mid s) - y(b \mid s)]\\
\leq&(1-\delta)\sum_{n=1}^{m^{10}}\delta^{n-1}\sum_{b\in B}u(s_n,b)\frac{1}{m^2 m(s_n)}\\
\leq&\frac{(1-\delta^{m^{10}})C}{m^2}
\end{aligned}$$

So, using this we can conclude that the receiver cannot gain much without being detected in a block. Next, we show that the punishment phase is efficient.

If the receiver fails the test in block B^m, he is punished for the next m^2 blocks. In each of the punishing block, the payoff cannot be greater than the babbling payoff v_2. Consider the blocks m,...., m^2+m, the limit of the stage payoffs when $\delta \to 1$ does not exceed $v_2+\epsilon_m$, where $\epsilon_m \to 0$ when $m \to \infty$. We also know that the block payoffs when the players use strategy profile (σ, τ) is approximately $U(\mu_0, y)$ during the normal blocks and v_2 during the punishment block. Using condition **C2**, which states that $U^2(\mu_0, y) \geq v_2$, it is easy to check that $\tau^{'}$ cannot be a profitable deviation against τ.

Now, let's consider the possible deviations by the sender. Assume the sender uses strategy $\sigma^{'}$. The receiver listens to the reports of the sender until the quota is reached. So, in each block B^m, we still have the announcements matching the quota limit. This ensures that the joint distribution μ in every block belongs to the set of copulas $\mathcal{M}$. So, we have

$$\mu^{\delta}_{\sigma',\tau}(s,a) = \boldsymbol{E}[(1-\delta)\sum_{n=1}^{\infty}\delta^{n-1}1_{\{s_n=s,\theta_n=a\}}] \tag{1.10}$$

$$\gamma_\delta(\sigma',\tau) = \boldsymbol{E}[(1-\delta)\sum_{n=1}^{\infty}\delta^{n-1}\sum_{s\in S,a\in A}1_{\{s_n=s,\theta_n=a\}}u(s,y(\cdot\mid a))] \tag{1.11}$$

$$\gamma_\delta(\sigma',\tau) = \sum_{s\in S,a\in A}\mu^{\delta}_{\sigma'.\tau}(s,a)u(s,y(\cdot\mid a)) \tag{1.12}$$

So, the expected payoff when the sender's report is drawn by $\mu^{\delta}_{\sigma',\tau}$ and the receiver plays y is equal to $U(\mu^{\delta}_{\sigma',\tau},y)$. But we know from (1.8) and (1.9) that if the players use strategy profile (σ,τ) then the expected payoff can be brought arbitrary close to $U(\mu_0,y)$ when $\delta \geq \delta_0$. But from condition **C1** we have $U^1(\mu_0,y) \geq U^1(\mu,y) \quad \forall\mu \in \mathcal{M}$. Thus, $U(\mu_0,y)$ is a uniform equilibrium.

1.4.2 Theorem 2

This theorem provides the converse inclusion to Theorem 1 under the following assumption.

Assumption A *There exist non-negative numbers* $(\alpha_s)_{s\in S}$*, with* $\sum_{s\in S\setminus\{\bar{s}\}}\alpha_s \leq 1$ *(for every* $\bar{s}\in S$*), such that* $p(s'|s)=\alpha_{s'}$ *whenever* $s'\neq s$*.*

Theorem Suppose that **Assumption A** holds. Then,

$$UEQ \subset E(\mathcal{M}) \tag{1.13}$$

Proof Consider a uniform equilibrium payoff (γ^1_*,γ^2_*) of the game. Choose a sequence $\epsilon_m > 0$ which converges to 0. We can define a sequence of strategies (σ_m,τ_m) and δ_m such that for $\delta \geq \delta_m$ the strategy (σ_m,τ_m) is ϵ_m-optimal. From the definition of uniform equilibrium we have $\lim\limits_{m\to\infty}\gamma_{\delta_m}(\sigma_m,\tau_m) = \gamma_*$. Let us define the stationary strategy of the receiver $y_m(b\mid s)$ as follows:

$$y_m(b\mid s) = \frac{1}{m(s)}\boldsymbol{E}_{\sigma_m,\tau_m}[\sum_{n=1}^{\infty}(1-\delta_m)\delta_m^{n-1}1_{\{\boldsymbol{s_n}=s,\boldsymbol{b_n}=b\}}] = \frac{\boldsymbol{E}_{\sigma_m,\tau_m}[\sum_{n=1}^{\infty}(1-\delta_m)\delta_m^{n-1}1_{\{\boldsymbol{s_n}=s,\boldsymbol{b_n}=b\}}]}{\boldsymbol{E}_{\sigma_m,\tau_m}[\sum_{n=1}^{\infty}(1-\delta_m)\delta_m^{n-1}1_{\{\boldsymbol{s_n}=s\}}]}$$

By construction, we have $\gamma_{\delta_m}(\sigma_m,\tau_m) = U(\mu_0,y_m)$.

$$\gamma_{\delta_m}(\sigma_m,\tau_m) = (1-\delta_m)\sum_{n=1}^{\infty}\delta_m^{n-1}\boldsymbol{E}_{\sigma_m,\tau_m}[u(\boldsymbol{s_n},\boldsymbol{b_n})] = (1-\delta_m)\sum_{n=1}^{\infty}\delta_m^{n-1}\sum_{s\in S,b\in B}\boldsymbol{E}_{\sigma_m,\tau_m}[1_{\{\boldsymbol{s_n}=s,\boldsymbol{b_n}=b\}}]u(s,b)$$
$$= \sum_{s\in S,b\in B}u(s,b)m(s)y_m(b\mid s) = U(\mu_0,y_m)$$

We define $y(b\mid s)$ as a limit point of $y_m(b\mid s)$ when $m\to\infty(\epsilon_m\to 0)$. This way we get $(\gamma_1^*,\gamma_2^*) = \lim\limits_{m\to\infty}\gamma_{\delta_m}(\sigma_m,\tau_m) = \lim\limits_{m\to\infty}U(\mu_0,y_m) = U(\mu_0,y)$.

Since the distribution of $\boldsymbol{s_n}$ is equal to $\mathbf{m}$ for each stage $n \in \mathbb{N}$, we have that $\gamma^2_{\delta_m} \geq v^2 - \epsilon_m$. Taking the limit we get $\gamma^2_* \geq v^2$. So, condition **C2** holds.

Now, for a given $\mu \in \mathcal{M}$, we define a sequence of strategies (σ'_m, τ_m) such that $\gamma_{\delta_m}(\sigma'_m, \tau_m) = U(\mu, y_m)$ $\forall m \in \mathbb{N}$. We can construct such a sequence using the idea of fictitious states. The sender generates a sequence $\boldsymbol{t}_n$ of states, which has the same distribution as the sequence $\boldsymbol{s}_n$. The sequences $\boldsymbol{s}_n$ and $\boldsymbol{t}_n$ are statistically indistinguishable and the joint distribution $(\boldsymbol{s}_n,\boldsymbol{t}_n)$ is given by μ. Replacing the states $\boldsymbol{s}_n$ with $\boldsymbol{t}_n$ in the sender's strategy σ gives us the new strategy σ'. The existence of such a sequence is given by the lemma below:

Lemma *Assume Assumption A and let $\mu \in \mathcal{M}$ be given. There exists S-valued process $\boldsymbol{t}_n$ such that following hold*

H1 Conditional on $\boldsymbol{s}_n$, the vector $(\boldsymbol{t}_1, ..., \boldsymbol{t_n})$ is independent of the future states $(\boldsymbol{s_{n+1}}, \boldsymbol{s_{n+2}}..)$.
H2 The law of the sequence $\boldsymbol{t_n}$ is the same as the law of the sequence $\boldsymbol{s}_n$.
H3 The law of the pair $(\boldsymbol{s}_n, \boldsymbol{t}_n)$ is μ for each stage $n \in \mathbb{N}$.
H4 The conditional law of $\boldsymbol{s}_n$ given $\boldsymbol{t}_1, ..., \boldsymbol{t}_n$ is $\mu(\cdot \mid \boldsymbol{t_n})$.

$$\gamma_\delta(\sigma', \tau) = \sum_{s,t,b} \mu(t \mid s)\mu(t)y(b \mid t)u(s, b) = U(\mu, y) \tag{1.14}$$

Now, using the result above, we get the following result:

$$\gamma^1_{\delta_m}(\sigma_m, \tau_m) \geq \gamma^1_{\delta_m}(\sigma'_m, \tau_m) - \epsilon_m \qquad \forall m \in \mathbb{N} \tag{1.15}$$
$$U^1(\mu_0, y_m) \geq U^1(\mu, y_m) - \epsilon_m \qquad \forall m \in \mathbb{N} \tag{1.16}$$
$$\lim_{m\to\infty} U^1(\mu_0, y_m) = U^1(\mu_0, y) \geq U^1(\mu, y) = \lim_{m\to\infty} U^1(\mu, y_m) \tag{1.17}$$

The inequality (1.15) follows from the definition of uniform equilibrium. Combining this inequality with (1.14) gives us (1.16). Taking the limit, we get (1.17). Thus, condition **C1** is also satisfied. Hence, we have proved that under Assumption A, $UEQ \subset E(\mathcal{M})$.

Chapter 2

Model with imperfect monitoring

In this chapter, we consider the model where the sender does not observe the action of the receiver but rather observes an action dependent signal. The mapping $\psi : B \to \Delta M$ maps each action b of the receiver to a mixed signal. The sender only observes the signal. As seen in the original model, the equilibrium strategies rely heavily on monitoring. We try to investigate the effect of imperfect monitoring and try to characterize the set of uniform equilibrium for this model.

At each stage n, the game is in the state s_n. The sender is aware of the true state and sends a message $a_n \in A$ to the receiver. The receiver, who does not know the true state, takes an action $b_n \in B$. The players get stage payoff $u^1(s_n, b_n)$ and $u^2(s_n, b_n)$ respectively. The sender does not observe the action of the receiver but a signal m_n is randomly chosen according to mapping ψ. The signal m_n is publicly observed. The game then moves to the next stage s_{n+1} according to the transition matrix $P(s_{n+1} \mid s_n)$. The state process again follows a Markov chain, which is irreducible and aperiodic, with unique invariant measure m.

2.1 Examples

In this section we describe examples highlighting the changes signalling brings. We give insights to why certain payoffs in $E(\mathcal{M})$ will not be a equilibrium in the signaling game.

Example:

		Receiver		
		L	M	R
	s_0	$1,1$	$0,0$	$0,0$
States	s_1	$0,0$	$1,0$	$0,1$
	s_2	$0,0$	$0,1$	$1,0$

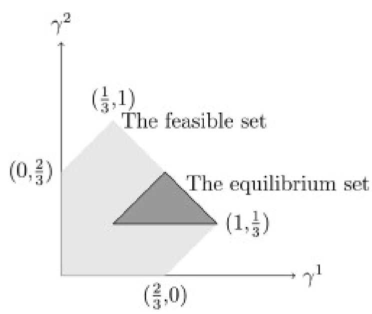

We use the payoff matrix from the game introduced in the previous section. Consider the *blind game*, where the sender receives no information on the receiver's action. We have $\psi(b) = m_0 \quad \forall b \in B$. Assume that all states are chosen independently with probability $\frac{1}{3}$ in every stage.

Consider the payoff $(1, \frac{1}{3})$, which is an equilibrium payoff in the model with perfect monitoring. This can only be achieved when the sender truthfully reports the states and receiver plays the strategy $y_1(s_0) = L, y_1(s_1) = M, y_1(s_2) = R$. Our claim is that this cannot be an equilibrium payoff in the blind game. As alternatively, the receiver can deviate and play R at s_1 and M at s_2 without being detected by the sender. This is a profitable deviation for the receiver. Hence, $(1, \frac{1}{3})$ is not an equilibrium payoff of the blind game. Notice that $(1, \frac{1}{3})$ will also not be an equilibrium payoff in the imperfect monitoring game where the signalling mechanism is given by $\psi(L) = l, \psi(M) = \psi(R) = \frac{m+r}{2}$.

Also, notice that unlike the equilibrium in the original game, where we could construct a strategy where the sender's best response was to speak the truth most of the time, it is not the case now. There may be scenarios where it is wise for the sender to only *partially reveal*. Consider the following stationary strategy: $\sigma(s_0) = \hat{s_0}, \sigma(s_1) = \sigma(s_2) = \frac{\hat{s_1}+\hat{s_2}}{2}$ and $y(s_0) = L, y(s_1) = y(s_2) = \frac{M+R}{2}$. This corresponds to the payoff $(\frac{2}{3}, \frac{2}{3})$. Here the sender only reveals whether the state is s_0 or not. Notice that this is an equilibrium in the blind game. It's not profitable for the sender to misreport s_0 as the receiver is playing the best response possible. And any exchange between s_1 and s_2 does not change the sender's payoff. This payoff belongs to the set of equilibrium payoffs $E(\mathcal{M})$ of the original model. But in the blind game, this cannot be achieved by the truth telling strategy of the sender. Because then the receiver could play his best response in each state to get the payoff $(\frac{1}{3}, 1)$ without being detected by the sender.

Using such "partial revelations" we can construct multiple equilibria where it is not beneficial for the sender to speak the truth most of the time. These equilibria payoff also belong to $E(\mathcal{M})$ but in the original model the payoff could be achieved by a truth telling strategy.

2.2 Characterization

Consider $\mathcal{M} \subset \Delta(S \times A)$ such that the marginal distributions S and A are equal to m, the invariant measure of the Markov chain. Given a copula $\mu \in \mathcal{M}$ and stationary strategy $y : A \to \Delta B$ we set:

$$U(\mu, y) = \sum_{s \in S, a \in A} \mu(s, a) u(s, y(\cdot \mid a)) \in \mathbb{R}^2$$

This corresponds to the expected payoff when the sender and receiver use stationary strategies given by μ and y respectively. We shall characterize the limit equilibrium payoffs using this function.

Denote by $S(\mathcal{M})$, the payoff vectors $U(\mu, y)$ that satisfy

D1. $U^1(\mu, y) \geq U^1(\mu', y) \quad \forall \mu' \in \mathcal{M}$

D2. $y = \underset{y'}{argmax}\{U^2(\mu, y')| \quad \psi(y \mid s) = \psi(y' \mid s) \quad \forall s \in S\}$

The equilibrium strategies in the original model, where the sender observes the receiver's action, were based on players able to monitor each others actions. The players could then check if the other player is deviating from his prescribed strategy. To devise optimal strategies of this model, we will have to use techniques for *imperfect* monitoring. The sender, instead of checking for the distribution of actions played in an interval, checks the distributions of signals. Depending on ψ, different actions could generate the same distribution of signals. We refer to actions inducing the same distribution of signals as being *equivalent*. So, if a player deviates to an equivalent action, it is undetectable. Extending the definition introduced in [4], we say two mixed strategies y and $y^{'}$ are equivalent if $\psi(y \mid s) = \psi(y^{'} \mid s) \quad \forall s \in S$. We need the equivalence across all states because the sender can check the distribution of signals for each state $s \in S$. Between two equivalent actions, a player chooses at equilibrium the one that yields the highest stage payoff. This is the basis for the condition **D2**.

In the original model, any equilibrium payoff could be achieved by a stationary strategy y of the receiver and the truth telling strategy of the sender. This was due to condition **C2** being lenient . But now due to condition **D2**, we have a rather restrictive set and have to look at a more generalized condition **D1**. We now check if the copula $\mu \in \mathcal{M}$, not necessarily μ_0, achieves the best payoff for the sender when the receiver uses the stationary strategy y.

2.3 Results

In this section, we consider the blind game, where the sender does not observe the actions of the receiver and has no way to check if the receiver is sticking to his prescribed strategy τ. We only informally prove one side of inclusion, i.e, $S(\mathcal{M}) \subset UEQ$. For the blind game, the condition **D1** and **D2** translate to:

G1 $U^1(\mu, y) \geq U^1(\mu', y) \quad \forall \mu' \in \mathcal{M}$

G2 $y = \underset{y'}{argmax} \quad \{U^2(\mu, y')\}$

We construct strategy profile (σ, τ) that achieves the payoff $U(\mu, y)$ and then show that it is a uniform equilibrium. The set of stages are divided into consecutive blocks of constant size N. The sender's strategy σ is to report the states according to $\mu(a \mid s)$. The receiver's strategy τ is to play the stationary strategy $y(b \mid a)$ as long as the sender does not cross the quota of announcements. We again use fictitious states θ so that in each block the joint distribution of states and announcements is a copula $\mu \in \mathcal{M}$. The sender cannot monitor the actions of the receiver and thus there is no punishment phase. As the receiver's strategy is stationary and the strategy profile is periodic after N stages, we have $\lim\limits_{\delta \to 1} \gamma_\delta(\sigma, \tau) = U(\mu, y)$.

Now, we prove there are no profitable deviations. Consider any strategy $\sigma^{'}$ of the sender. Using the definition of $\mu^{\delta}_{\sigma^{'}, \tau}$ from (1.10), we have $\gamma_\delta(\sigma^{'}, \tau) = U(\mu^{\delta}_{\sigma^{'}, \tau}, y)$. From condition **G1** we have $U^1(\mu, y) \geq U^1(\mu^{'}, y) \quad \forall \mu^{'} \in \mathcal{M}$. So, there is no profitable deviation of the sender. Now, consider the receiver's strategy $\tau^{'}$. we define y^δ as follows:

$$y^\delta(b \mid a) = \frac{1}{m(a)} \boldsymbol{E}_{\sigma, \tau'}[\sum_{n=1}^{\infty}(1-\delta)\delta^{n-1} 1_{\{\boldsymbol{\theta_n}=a, \boldsymbol{b_n}=b\}}]$$

Choosing a large N, we have the reported and fictitious states matching almost surely. So, we can then show that $\gamma_\delta(\sigma, \tau^{'}) = U(\mu, y^\delta)$. Using condition **G2**, it is easy to see that this will not be a profitable deviation. So, we have that $U(\mu, y)$ is an equilibrium payoff. Thus, $S(\mathcal{M}) \subset UEQ$ for the blind game.

2.4 Conclusion

In this report, we describe a dynamic version of the sender-receiver game. We characterize the equilibrium payoffs and construct optimal strategies that achieve the payoff vector. These strategies heavily rely on monitoring. We see this explicitly when we consider the model with imperfect monitoring. We show the impact of monitoring by considering the extreme case: the blind game, where the sender does not get any information on the receiver's action. We also provide insight into what changes are needed in the construction of the equilibrium strategies in this model and provide an informal proof to construct them.

Bibliography

[1] Vincent P. Crawford and Joel Sobel. "Strategic Information Transmission". In: *Econometrica* 50.6 (1982), pp. 1431–1451. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/1913390.

[2] Mikhail Golosov et al. "Dynamic strategic information transmission". In: *Journal of Economic Theory* 151.C (2014), pp. 304–341. URL: https://EconPapers.repec.org/RePEc:eee:jetheo:v:151:y:2014:i:c:p:304-341.

[3] Geoffrey Grimmett and David Stirzaker. *Probability and random processes*. Oxford; New York: Oxford University Press, 2001. ISBN: 0198572239 9780198572237 0198572220 9780198572220. URL: http://www.worldcat.org/search?qt=worldcat_org_all&q=9780198572220.

[4] E. Lehrer. "Lower equilibrium payoffs in two-player repeated games with non-observable actions". In: *International Journal of Game Theory* 18.1 (Mar. 1989), pp. 57–89. ISSN: 1432-1270. DOI: 10.1007/BF01248496. URL: https://doi.org/10.1007/BF01248496.

[5] E. Lehrer. "Nash equilibria of n-player repeated games with semi-standard information". In: *International Journal of Game Theory* 19.2 (June 1990), pp. 191–217. ISSN: 1432-1270. DOI: 10.1007/BF01761076. URL: https://doi.org/10.1007/BF01761076.

[6] E. Lehrer. "On the equilibrium payoffs set of two player repeated games with imperfect monitoring". In: *International Journal of Game Theory* 20.3 (Sept. 1992), pp. 211–226. ISSN: 1432-1270. DOI: 10.1007/BF01253776. URL: https://doi.org/10.1007/BF01253776.

[7] Jérôme Renault, Eilon Solan, and Nicolas Vieille. "Dynamic sender–receiver games". In: *Journal of Economic Theory* 148.2 (2013), pp. 502–534. ISSN: 0022-0531. DOI: https://doi.org/10.1016/j.jet.2012.07.006. URL: http://www.sciencedirect.com/science/article/pii/S0022053113000161.

[8] Jérôme Renault and Tristan Tomala. "General Properties of Long-Run Supergames". In: *Dynamic Games and Applications* 1.2 (May 2011), p. 319. ISSN: 2153-0793. DOI: 10.1007/s13235-011-0018-3. URL: https://doi.org/10.1007/s13235-011-0018-3.