



HAL
open science

Using or Hiding Private Information? An Experimental Study of Zero-Sum Repeated Games with Incomplete Information

Nicolas Jacquemet, Frédéric Koessler

► **To cite this version:**

Nicolas Jacquemet, Frédéric Koessler. Using or Hiding Private Information? An Experimental Study of Zero-Sum Repeated Games with Incomplete Information. *Games and Economic Behavior*, 2013, 78, pp.103-120. 10.1016/j.geb.2012.12.002 . halshs-00773412

HAL Id: halshs-00773412

<https://shs.hal.science/halshs-00773412>

Submitted on 14 Jan 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Using or Hiding Private Information?

An Experimental Study of Zero-Sum Repeated Games with Incomplete Information *

Nicolas JACQUEMET[†]

Frédéric KOESSLER[‡]

November 2012

Abstract

This paper studies the value of private information in strictly competitive interactions in which there is a trade-off between (i) the short-run gain of using information, and (ii) the long-run gain of concealing it. We implement simple examples from the class of zero-sum repeated games with incomplete information. While the empirical value of information does not always coincide with the theoretical prediction, the qualitative properties of the value of information are satisfied in the laboratory: (i) it is never negative, (ii) it decreases with the number of repetitions, (iii) it is bounded below by the value of the infinitely repeated game, and (iv) it is bounded above by the value of the one-shot game. In line with the theory, the empirical use of private information is almost complete when it should be, and decreases in longer interactions.

KEYWORDS: Concealment of information; laboratory experiments; value of information; zero-sum repeated games.

JEL CLASSIFICATION: C72; D82.

*This paper is a revised version of CES WP n°2011-02. We are very grateful to Kene Boun My for his assistance in designing the experiments and for his comments, and to Thibault Baron, Marie Coleman, Vincent Montamat, Marie-Laure Nauleau and Ivan Ouss for their research assistance. We thank Nick Feltovich for his comments and for sharing his data with us. We also thank the editor, Thomas Palfrey, the associate editor, two anonymous referees, Alessandra Casella, Vincent Crawford, David Ettinger, Françoise Forges, Guillaume Fréchette, Olivier Gossner, Jeanne Hagenbach, Philippe Jehiel, Rida Laraki, Paul Pezanis-Christou, Dinah Rosenberg, Larry Samuelson, Tristan Tomala, and Shmuel Zamir for comments and discussions on the theory and/or experimental protocol. Financial support in the form of an ACI grant from the French Ministry of Research is gratefully acknowledged. Jacquemet has greatly benefited from the support of the *Institut Universitaire de France*.

[†]Paris School of Economics and University Paris I Panthéon–Sorbonne. Centre d'Économie de la Sorbonne, 106 Bd. de l'Hôpital, 75013 Paris, France. Nicolas.Jacquemet@univ-paris1.fr.

[‡]Paris School of Economics – CNRS. 48, Boulevard Jourdan, 75014 Paris, France. koessler@pse.ens.fr.

1 Introduction

It is well known that private information may not always be valuable in strategic contexts. In other situations, although it may be valuable, information should not be used, or should only be partially used. A typical example is the use of balanced strategies in poker: a good poker player should sometimes resist the temptation to bet or raise the best hand in the current round of betting to hide his information from his opponents, even though naive reasoning would suggest that he raise his hand in order to increase the pot. This is best illustrated by the following advice from the famous professional poker player Dan Harrington:

“No matter what style you finally adopt as your own, you’ll have to learn to play what I call a balanced strategy. Simply put, this means that you have to learn to vary both your raises and calls, as well as the actual size of your bets, to avoid giving your opponents a read on your style. You’ll have to do this even when you believe that a certain bet is clearly correct. What you sacrifice in terms of making a slightly incorrect bet on a given occasion will be recovered later, when your opponents have to guess at what you’re really doing, and they guess wrong.”¹

Harrington and Robertie (2007, p. 52)

Choosing payoff dominated actions today in order to get greater benefit from information in the future is also a feature of well-known deceptive military strategies. This includes the WWII example of British military intelligence. When they were able to secretly read radio communications of the Axis powers enciphered using Enigma machines, they sent spotter submarines and aircraft to search for Axis ships just to disguise the source of the intelligence behind the Allied attacks (Hinsley, 1993). The Axis forces who observed these spotters and their radio transmissions then concluded that their ships were being located by conventional reconnaissance. Similar strategies were used to disguise the intelligence source from Allied crews themselves, by sending them on useless search missions.

This trade-off between *(i)* using information to get higher information rents today, but at the cost of losing an informational advantage tomorrow and *(ii)* choosing non-informed, and therefore more costly decisions today, in order to keep an informational advantage in the future, can also be found in strategic economic interactions. Consider for example a repeated common value auction in which identical objects are successively put up for sale and bidders receive private and independent signals about the true value. If bids are observed after each period, a well-informed bidder faces the same trade-off: *“if a bidder’s proprietary information indicates that the units are all of high quality, he would like to increase his chance of winning the object by bidding more aggressively, but*

¹He further explains: *“Here’s a simple example. Suppose you believe that when you hold aces in first or second position, the “right” play is to open with a raise [...] If you always make this play with aces [...] they will know that when you call, you don’t have aces. This is dangerous information to be giving away, so you need to take some countermeasures. The simplest countermeasure is to vary your play at random, giving a higher probability to the play you think is correct, but mixing in other plays frequently enough so that your opponents can’t put you on a hand easily”*.

doing so may prove costly later on as other bidders may compete away the value of the information released through the bid” (Hörner and Jamison, 2008, p. 476). As in the previous examples, this leads the informed bidder to delay information revelation by acting as an uninformed bidder, so as to make the other bidders more cautious and thus win with a lower bid.

This paper studies experimentally how subjects react to the above-mentioned trade-off between the short-run cost and long-run benefit of concealment. We raise the following general questions. When it is optimal to do so in the long-run, is there empirical evidence in a controlled environment that subjects only partially use their information, or delay information revelation, even though it is costly in the short-run? That is, do informed subjects refrain from using a naive, fully-revealing strategy, in order to benefit from an informational advantage in subsequent periods? Do these same subjects fully use their information when there is no value for information concealment, for example in late periods or in one-shot interactions? Do they suffer from a curse of knowledge, *i.e.*, could the empirical value of information be negative in some instances? Do uninformed subjects extract information from informed subjects’ behavior?

We address these questions by studying three examples from the class of zero-sum repeated games with incomplete information on one side and perfect monitoring, drawn from Aumann and Maschler (1966, 1967) and Stearns (1967). In this class of games, repetition is the channel through which information is transmitted from one stage to another. Because of perfect monitoring, information is (at least partially) revealed by the informed player’s action whenever he decides to use it. Exactly as in the strategic applications mentioned above, the basic problem for the informed player is to find the optimal balance between using information as much as possible and revealing as little as possible to his opponent. On the other side, the uninformed player tries to find out the actual information of his opponent and to minimize its value. In the three examples under study, the stage games are trivial, *i.e.*, the informed player has a different dominant strategy in either state. However, if he simply plays this naive strategy, the uninformed player will learn the state and subsequently choose actions that result in a very low payoff for the informed player.

As an illustration, Figure 1 presents the payoff matrices of the leading repeated games we implement in the laboratory.² At the beginning of a repeated game, one of the two payoff matrices, A^1 or A^2 , is drawn at random with the same probability. Player 1 (the row player) is privately informed about the state of nature (matrix A^1 or A^2), and has a dominant action in each state: Top in A^1 , Bottom in A^2 . This is clearly the best strategy for him in the one-shot game: he should completely use his private information. Player 2 (the column player) would like to play Right in A^1 and Left in A^2 but, being uninformed, he has to choose the same action in both states, yielding an expected payoff equal to 5 for both players. When the game is repeated, the past decisions of the informed player become a signal about his information: if this private information is fully used, then player 2 becomes aware of the actual state after the first stage, and plays his perfectly informed decision forever: Right following Top (*i.e.*, in A^1) and Left following Bottom (*i.e.*, in A^2). In this case, the payoff of the informed player is 0 in each subsequent stage, yielding an average

²In the game presented in Figure 1, as in all instances we implement in the laboratory, the sum of players’ utilities is constant across actions profiles. They are thus strategically equivalent to zero-sum games.

Figure 1: Payoff matrices in the NR games.

	Left	Right	
Top	10, 0	0, 10	
Bottom	0, 10	0, 10	
	A^1		


	Left	Right
Top	0, 10	0, 10
Bottom	0, 10	10, 0
	A^2	

payoff of $5/n$ in the n -stage game. Alternatively, the informed player can keep his information private by using a pooling strategy which is independent of the actual payoff matrix, for example mixing his play uniformly between Top and Bottom – just as if he was uninformed. The cost of this non-revealing strategy is the loss of the rent derived from information in the first stage. The benefit is that it provides an expected payoff equal to $5/2$ in each stage, whatever the length of the game. So, whenever the game is not one-shot, not using information at all is better than using the previous naive, fully-revealing strategy. How this trade-off is solved more generally determines the extent to which information should be fully used, partially used or ignored by the informed player.

The theory of zero-sum repeated games with incomplete information provides precise predictions on players' optimal strategies and the value (expected payoff) that they can provide in the long-run as a function of the payoff matrices and the prior beliefs of the uninformed player. From a behavioral point of view, the class of zero-sum repeated games is probably one of the cleanest environments for studying the use, revelation and value of information in strategic contexts. In particular, since equilibrium strategies are equivalent to max-min strategies, they are interchangeable and there is no equilibrium selection problem or miscoordination issue between players. This strategic behavior coincides with what players play not only when they knew their opponent's strategy perfectly, but also when they have vague beliefs about it (Marinacci, 2000). The induced expected payoff for each player is always unique, which defines an unambiguous value for each player in the game. This value is also the only rational expectation a player might have in such a game (Aumann and Dreze, 2008), and simple learning procedures lead to this value (see, e.g., Hart and Mas-Colell, 2000). Finally, and again because payoffs are zero-sum, cooperation is not an issue even if the (same or similar) repeated games are played several times between the same players (as will be the case in our experiments).

We consider three types of repeated games, implemented as separate experiments. The games differ in that, in equilibrium, information should be either fully revealed (the FR games), partially revealed (the PR games) or almost never revealed (the NR games based on our introductory example). To study the impact of the length of the game (*i.e.*, the number of repetitions), each type of game is repeated from 1 to 5 stages in each experimental treatment. To generate an empirical benchmark for the value of information, we also consider two types of information structures in the NR games: no information (in which no player knows the actual payoff matrix) and incomplete information on one side (only one player knows the payoff matrix).

Figure 2: The value of the NR game for the informed player.



Legend. - · - · - Theoretical value ; — Value of the one-shot game ; - - - Value of the pooling strategy ; - · - · - Value of the naive strategy ; ■ Mean empirical value observed in Feltovich (1999, 2000).

Note. The figure plots the theoretical value, the theoretical value of the one-shot game, the value when player 1 does not use his information (*i.e.*, plays Top and Bottom uniformly), and the value stemming from the naive (fully revealing) strategy for each length $n \in \{1, \dots, 5\}$ in abscissa of the NR game presented in Figure 1. We also report the average empirical value observed by Feltovich for $n = 2$.

Feltovich (1999, 2000) is, to the best of our knowledge, the only experimental analysis to use games from this class.³ Although the focus of his work is different from ours,⁴ his lab implementation of a 2-stage zero-sum game concludes that informed players use their information too much in the first stage of the game, even when they can play the same repeated game for 40 rounds. This results in a low average payoff for informed subjects compared to what they could get with an optimal strategy. This might suggest that a curse of knowledge may occur in longer repeated games from this class, in the sense that the actual value of information may become negative when information should, optimally, be concealed.⁵

The empirical challenge is described in more detail in Figure 2. For every game length (from $n = 1$ to $n = 5$ stages) of the repeated simultaneous-move game of Figure 1, the value of the game is decreasing in the length of the game (plotted as the theoretical value). But it is always strictly higher than the value of the pooling strategy of the game without private information, consisting for player 1 in playing Top and Bottom with probability 1/2 in each stage of the repeated game whatever the payoff matrix (value of the pooling strategy). By contrast, the value of the naive and fully-revealing strategy of the informed player, consisting for player 1 in playing the stage-dominant action in each stage of the repeated game, reaches this lower bound already for $n = 2$, and then

³Non zero-sum repeated experimental games with incomplete information have also been studied in the literature (see, e.g., McKelvey and Palfrey, 1995) but they have different strategic features from our games.

⁴It compares different models of learning (belief-based learning vs. reinforcement learning).

⁵This conjecture is also supported by Chaudhuri (1998), who found in a 2-stage principal-agent laboratory experiment that informed subjects played naively (*i.e.*, revealed their type in the first stage) even though this is sub-optimal in the dynamic game.

goes beyond as the length increases (value of the naive strategy). As shown in Figure 2, the average empirical value of the 2-stage game observed in Feltovich is significantly lower than the theoretical value but remains above this lower bound. The question is thus whether the empirical value might cross over the value of the pooling strategy in longer games, in which case informed subjects would suffer severe losses compared to the situation in which they are completely uninformed.

According to our experiments, this conjecture about a curse of knowledge in longer dynamic interactions turns out to be false: the empirical value of information is always positive in all versions of the repeated games, whatever their length. The average empirical payoff of the informed player is actually always bounded below by the value of the infinitely repeated game. This means that the value of information is not only positive, but strictly positive when predicted to be so; hence informed subjects use their information at least fairly efficiently. This is particularly remarkable when the values of both the naive strategy and the pooling strategy of the informed player are strictly lower than this benchmark. We also show that in each type of n -stage repeated game with $n \geq 2$, the average empirical payoff of the informed player is always strictly below the value of the one-shot game, and is decreasing in the length of the game. This means that uninformed subjects manage to correctly extract (at least some) relevant informational rent by observing informed subjects' behavior.

The analysis of empirical strategies confirms these insights, and reveals additional interesting features of subjects' behavior. One of our most interesting findings is the strategic sophistication of informed subjects' behavior in the n -stage repeated NR games of Figure 1: the empirical use of information, at any stage $t < n$, is strictly decreasing with the length n of the game. This is to be contrasted with the informed subjects' behavior in all stages of the FR games and in the last stage ($t = n$) of the NR games in which information is used, as predicted by the theory, almost 100% of the time. This proves that informed subjects' behavior reacts strongly to the trade-off between the short-run cost and the long-run gain of concealment. The fact that subjects understand this trade-off correctly is confirmed by comparing the empirical correlation of the uninformed subjects' actions and the informed subjects' type: it is significantly higher in the FR games than in the NR games, especially in longer versions of the games.

Not surprisingly (given the complexity and the stochastic nature of the optimal strategies in some finite versions of the repeated games), in some situations the strategies supporting these results do not coincide perfectly with theoretical predictions.⁶ In particular, while experimental subjects use private information with very high accuracy when it is worth using, they are unable to completely hide it when it is optimal to do so.

In Section 2 we present the basic model and some general theoretical predictions. The three examples on which our experiment is based on and more specific theoretical predictions are provided in Section 3. Then, in Section 4 we present our research hypotheses and the experimental design implemented in the laboratory. In Section 5 and 6, we analyze subjects' performances and behavior in the different experimental treatments. Section 7 concludes.

⁶This is neither new nor surprising, and has already been observed and discussed in simpler zero-sum laboratory and field experiments (see, e.g., Palacios-Huerta, 2003, Wooders, 2010).

2 The theory of zero-sum repeated games

The theory of zero-sum repeated games with incomplete information first appeared in the period 1966–1968 in technical reports of the United States Arms Control and Disarmament Agency (see Aumann and Maschler, 1966, 1967 and Stearns, 1967). In the simplest model of zero-sum repeated game, one of two finite zero-sum two-person payoff matrices A^1 or A^2 is played repeatedly.⁷ The payoff matrix A^1 (A^2 , respectively) is chosen once and for all according to the common prior probability $p^1 = p \in [0, 1]$ ($p^2 = 1 - p$, respectively). When the matrix is A^k , $k \in \{1, 2\}$ is also called the type of player 1, or state. The real $A^k(i, j)$ denotes player 1's payoff when the state is $k \in \{1, 2\}$, player 1's action is $i \in I$, and player 2's action is $j \in J$. Player 2's payoff is a constant minus $A^k(i, j)$, so player 2 is the minimizer and player 1 is the maximizer. The value (or, equivalently, Nash equilibrium payoff for player 1) of the complete information game A^k is denoted by $w^k = \max_{x \in \Delta(I)} \min_{y \in \Delta(J)} A^k(x, y)$ (payoffs are extended to mixed actions in the usual way).

The n -stage repeated game with incomplete information is denoted $G_n(p)$. Only player 1 knows the state. That is, he knows the payoff matrix from the outset, but player 2 does not. Both players publicly observe past actions (perfect monitoring) but they do *not* observe their past payoffs before the end of stage n (however, player 1 can obviously deduce his past payoffs from his knowledge of the actual payoff matrix and past actions). Player 1's *average payoff* is $\frac{1}{n} \sum_{m=1}^n A^k(i_m, j_m)$, when the state is k and the final history of play is $((i_1, j_1), (i_2, j_2), \dots, (i_n, j_n))$. The *value* of $G_n(p)$ is denoted by $v_n(p)$.⁸

Let $u(p)$ be the value of the average game $\sum_k p^k A^k$, *i.e.*, the value of the game in which player 1 is also uninformed about the payoff matrix. The following proposition states the positive value of information for player 1 whatever the length of the repeated game.

Proposition 1 (Positive value of information) *For all n and p , $v_n(p) \geq u(p)$.*

Albeit positive, the value of the repeated game is decreasing in the number of repetitions, as stated in the following proposition. In particular, the value of the one-shot game, $v_1(p)$, is an upper bound of the value of the repeated game, whatever its length. The intuition is that when n increases, the amount of information revealed by player 1 to player 2 is weakly increasing, so the value for player 1 should decrease.

Proposition 2 (Decreasing value of information) *For all p , $v_n(p)$ is weakly decreasing in n .*

Interestingly, the value of the repeated game does not necessarily decrease to the value of the average game, even in the long run. The bound below $v_n(p)$ is in fact very easy to characterize directly from the value of the average game, and is given by the *concavification* of u , denoted by $\text{cav } u$, which is the smallest concave function which is higher than u . The intuition of this property

⁷See, e.g., Zamir (1992) for a thorough treatment of this class of games.

⁸Since all games under study are zero-sum, we will use the term “optimal strategy” for a max-min strategy, remembering that max-min and equilibrium strategies are equivalent; in addition, since associated expected payoffs are uniquely defined, the “value” of the game denotes the equilibrium expected payoff for the informed player, an expected payoff that each player can guarantee using an optimal strategy whatever his opponent's strategy.

Figure 3: Payoff matrices in the FR games.

	Left	Right
Top	6, 4	4, 6
Bottom	0, 10	0, 10

A^1

	Left	Right
Top	0, 10	0, 10
Bottom	4, 6	6, 4

A^2

relies on the “splitting procedure”: starting from a prior belief p of the uninformed player about the informed player’s type, the informed player can generate two new beliefs $p_1 < p$ and $p_2 > p$ and new expected payoffs z_1 and z_2 by revealing information through appropriate (mixed) type-dependent actions. This procedure moves the informed player’s payoff from z to a convex combination of z_1 and z_2 , and can be beneficial for him only when $u(\cdot)$ is not concave.

Proposition 3 (Value in the long repeated games) $v_\infty(p) = \lim_{n \rightarrow \infty} v_n(p) = \text{cav } u(p)$.

The theory of zero-sum repeated games with incomplete information provides simple lower bounds, upper bounds, and comparative statics for the value of information as a function of the length and payoff matrices of the repeated game. We now turn to the application of those results to examples that lead to very different uses of information in equilibrium.


3 Analysis of the experimental games

Our experiment implements three modified examples from the literature in which information should be fully used (FR games), should almost never be used (NR games), or should be partially used (PR games) by the informed player. The examples have been designed to ensure that each subject gets a positive payoff in all experiments, equilibrium payoffs are similar across treatments, and there is a different stage-dominant action for the informed player in each state (*i.e.*, a different dominant action in each state of the static game). All games are implemented as between-subject treatments, with a common prior of $p = 1/2$ for each payoff matrix. Moreover, we consider 5 different lengths ($n = 1, \dots, 5$), implemented as within-subject treatment variables.

3.1 Full revelation of information: the FR games

The payoff matrices presented in Figure 3 are a modified (but not strategically equivalent) version of the second example studied by Aumann and Maschler (1995). They have a stage-dominant action but, contrary to the original example, the informed player (player 1) cannot guarantee himself the maximum payoff of the game, so the uninformed player (player 2) has an incentive to use the information revealed by the informed player. The values under complete information are $w^1 = w^2 = 4$. The stage-dominant (and fully revealing) strategy Top (T) in A^1 and Bottom (B) in A^2 is clearly the unique optimal strategy for player 1 in $G_n(p)$ for all n and p . Hence, player 2 plays any strategy in the first stage, and plays Right (R) after T and Left (L) after B . This yields

Figure 4: Concavification in the FR game.



Note. Plot of the functions $u(p)$ (dotted lines) and $\text{cav } u(p)$ (plain lines) against $p \in [0, 1]$ for the FR games in Figure 3.

$v_n(p) = \frac{1}{n}(5 + 4 + 4 + \dots) = \frac{1+4n}{n}$ for all n and p , which tends to 4 when n tends to infinity. On the contrary, by playing a non-revealing strategy, *i.e.*, not using his private information at all, player 1 gets the value of the *average game*

$$pA^1 + (1-p)A^2 = \begin{pmatrix} 6p & 4p \\ 4(1-p) & 6(1-p) \end{pmatrix}, \quad \text{i.e., } u(p) = \begin{cases} 4(1-p) & \text{if } p \leq 2/5 \\ 10p(1-p) & \text{if } p \in [2/5, 3/5] \\ 4p & \text{if } p \geq 3/5, \end{cases}$$


which is always strictly smaller than $\frac{1+4n}{n}$. Given the value of the average game, $u(p)$, Figure 4 shows the concavification in this game, which leads to $v_\infty(p) = \text{cav } u(p) = 4 > u(p)$ for all $p \in (0, 1)$.

3.2 No revelation of information: the NR games

Figure 1 in the Introduction presents payoff matrices which are strategically equivalent to those of the first example in Aumann and Maschler (1995) and those studied experimentally by Feltovich (1999, 2000). We consider two treatments. In the NR-NoInfo treatment, no player is informed about the actual payoff matrix. This provides a control of the actual behavior of uninformed players. This version of the game is strategically equivalent to matching pennies, so the value is trivial: whatever the length n of the game, every player plays each action with probability one half, leading to the average payoff $u(1/2) = 2.5$ for player 1. In the NR treatment, we implement the same payoff matrices but with information asymmetries: player 1 is informed about the actual payoff matrix, player 2 is not.

The values under complete information are $w^1 = w^2 = 0$, and in the one-shot incomplete information game $G_1(1/2)$ the optimal strategy of player 1 is T in A^1 and B in A^2 . Hence, $v_1(1/2) = 5$ (any strategy of player 2 is optimal). Next, consider the n -stage incomplete information game. If player 1 uses the previous stage-dominant strategy then player 2 learns the state in the second stage and will therefore play R in A^1 and L in A^2 . Hence, player 1's average payoff will be $\frac{1}{n}(5 + 0 + 0 + \dots) = \frac{5}{n}$, which decreases to 0 when n increases. An alternative strategy for player 1 is to play a non-revealing strategy, which is equivalent to playing the repeated average

Figure 5: Concavification in the NR game.



Note. Plot of the functions $u(p)$ (dotted lines) and $\text{cav } u(p)$ (plain lines) against $p \in [0, 1]$ for the NR game in Figure 1.

game $pA^1 + (1 - p)A^2$, the value of which is $u(p) = 10p(1 - p)$, so $u(1/2) = 2.5$ (player 1 plays each action with probability $1/2$ independently of his information). This is clearly better than the previous fully revealing strategy whenever $n \geq 2$, since $2.5 \geq 5/n$ for $n \geq 2$. As shown in Figure 5, 2.5 is in fact the maximum payoff that player 1 can guarantee himself in the long run, because $\lim_{n \rightarrow \infty} v_n(1/2) = \text{cav } u(1/2) = u(1/2) = 2.5$, so the optimal strategy of player 1 in the infinitely repeated game consists in not using his information.

The value and optimal strategies of the finitely-repeated games are much more difficult to calculate. However, one can start by noting that we necessarily have the simple property $v_n(1/2) \geq \frac{5(n+1)}{2n}$ for all n because player 1 can guarantee himself the average payoff $\frac{1}{n}[(n-1)\frac{5}{2} + 5]$ by playing a non-revealing strategy during the first $n - 1$ stages, and the stage-dominant strategy in stage n . But $\frac{5(n+1)}{2n}$ is not exactly the value of $G_n(1/2)$ for all n , only for $n = 1, 2, 3$ and $n \rightarrow \infty$. To see this, we compute the value $v_n(p)$ for all $p \in [0, 1]$ and $n \in \{1, \dots, 5\}$ using the recursive formula (see, e.g., Zamir, 1992, p. 126).⁹ Table 1 on page 13 represents these values for $p = 1/2$ and for $n = 1, \dots, 5$. Appendix A describes players' optimal strategies and the posterior beliefs of player 2 after every possible history of actions from player 1. The optimal strategy of player 1 is unique. It consists in playing non-informatively (Top and Bottom with probability $1/2$ whatever the state and the history) in all but the last stage of the 2-stage and 3-stage repeated games, and almost non-informatively in all but the last stage of the 4-stage and 5-stage repeated games (so that player 2's posterior belief about the state is $1/2$ or close to $1/2$ most of the time). The optimal strategy of player 2 is not unique in games $G_1(1/2)$, $G_2(1/2)$ and $G_5(1/2)$ in which a single free parameter remains (which can take any value in a continuous interval). In all those situations, we present in the tables of Appendix A the strategy of player 2 that reacts symmetrically against Top and Bottom in the history of play of player 1, yielding a symmetric payoff for player 1 (the same expected payoff in A^1 and A^2).

⁹The detailed calculations and explicit values of the function $v_n(\cdot)$ are available from the authors upon request. Computations are not difficult because the value function is a stepwise continuous and linear function, but they are tedious because the number of intervals on which v_n is defined rapidly increases with n .

Figure 6: Payoff matrices in the PR games.

	Left	Center	Right
Top	9, 0	3, 6	6, 3
Bottom	9, 0	3, 6	0, 9

A^1

	Left	Center	Right
Top	3, 6	9, 0	0, 9
Bottom	3, 6	9, 0	6, 3

A^2

3.3 Partial revelation of information: the PR games

The game presented in Figure 6 is an intermediary case, which is strategically equivalent to the third example in Zamir (1992). The values under complete information are $w^1 = w^2 = 3$. In $G_n(1/2)$, the fully-revealing, stage-dominant strategy (T in A^1 , B in A^2) guarantees player 1 the average payoff $\frac{1}{n}(6 + 3 + 3 + \dots) = \frac{3(1+n)}{n}$, which decreases to 3 when n increases. On the other hand, by using a non-revealing strategy, player 1 can obtain the value of the average game $pA^1 + (1-p)A^2$. The optimal strategy of player 1 in the average game is T if $p \geq 1/2$ and B if $p \leq 1/2$. Hence, player 2's optimal strategy is L if $p \leq 1/4$, R if $1/4 \leq p \leq 3/4$, and C if $p \geq 3/4$. The value is


$$u(p) = \begin{cases} 3 + 6p & \text{if } p \leq 1/4, \\ 6(1 - p) & \text{if } 1/4 \leq p \leq 1/2, \\ 6p & \text{if } 1/2 \leq p \leq 3/4, \\ 9 - 6p & \text{if } p \geq 3/4. \end{cases}$$

Thus, when $p = 1/2$, player 1 can only guarantee himself 3 by using a non-revealing strategy. Now, we show that when $p = 1/2$ player 1 can guarantee himself exactly 4.5 whatever $n \geq 2$ by using a *partially-revealing* (PR) strategy. If player 1 always plays T with probability $3/4$ and B with probability $1/4$ in A^1 , and always plays T with probability $1/4$ and B with probability $3/4$ in A^2 , then the posterior beliefs are $\Pr(k = 1 \mid i = T) = 3/4$ and $\Pr(k = 1 \mid i = B) = 1/4$. So, when he plays T , player 1's conditional expected payoff is $\frac{3}{4}(9, 3, 6) + \frac{1}{4}(3, 9, 0) = (7.5, 4.5, 4.5)$, and when he plays B his conditional expected payoff is $\frac{1}{4}(9, 3, 0) + \frac{3}{4}(3, 9, 6) = (4.5, 7.5, 4.5)$. In both situations, whatever the strategy of player 2, the expected payoff of player 1 is at least 4.5, so $v_n(1/2) \geq 4.5$ for all n . As shown in Figure 7, this is in fact the maximum payoff player 1 can guarantee himself in the repeated game since

$$\text{cav } u(p) = \begin{cases} 3 + 6p & \text{if } p \leq 1/4 \\ 4.5 & \text{if } p \in [1/4, 3/4], \\ 9 - 6p & \text{if } p \geq 3/4, \end{cases}$$

so $\text{cav } u(p) > u(p)$ for all $p \in (1/4, 3/4)$. Contrary to the previous games, the value of the n -stage

Figure 7: Concavification in the PR game.



Note. Plot of the functions $u(p)$ (dotted lines) and $\text{cav } u(p)$ (plain lines) against $p \in [0, 1]$ for the PR game in Figure 6.

PR game is the same as in the infinitely repeated game whenever $n \geq 2$. Indeed, it is easy to verify that the following strategy profile constitutes a Nash equilibrium of the 2-stage game, with an expected payoff of 4.5 for player 1: in stage 1 player 1 plays T in A^1 and B in A^2 with probability $x \geq 3/4$, and player 2 plays R with probability 1; in stage 2 player 1 plays T in A^1 and B in A^2 with probability 1, and player 2 plays C after T and L after B . Since Nash equilibrium payoffs are unique in zero-sum games and since $v_\infty(1/2) = 4.5$, Propositions 2 and 3 imply that $v_n(1/2) = 4.5$ for every $n \geq 2$.

4 Empirical approach

Our experiment aims to assess the empirical content of these predictions as regards the trade-off between hiding and using one's own information. To that end, we implement the three games analysed in the previous section with a uniform prior distribution over the set of states ($p = 1/2$). One advantage of these games is that they have the same stage-dominant strategy for the informed player (*i.e.*, player 1 has a dominant action, Top in one state and Bottom in the other state, in all static versions of the games).

4.1 Testable predictions from the theory

The theory translates into particular levels of the value function of the game, as summarized in Table 1. Although we will compare the data with the point predictions provided by the theory, our main focus is on the qualitative predictions about the shape of the value function both across lengths and across games. In this respect, we assess the following hypothesis.

Hypothesis 1 (Shape of the value) *In all n -stage FR, PR and NR repeated games, the empirical value is: (i) above the value of the corresponding repeated games in which player 1 is uninformed, (ii) strictly below the value of the corresponding one-shot game, (iii) above the value of the infinite game and, (iv) (weakly) decreasing in the length n .*

Table 1: Theoretical properties of the experimental games.

	Value of the game, $v_n(1/2)$						Value of the average game, u	Optimal use of information
	1	2	n			∞		
FR game	5.00	4.50	4.33	4.25	4.20	4	2.5	Fully Revealing
NR game	5.00	3.75	3.33	3.21	3.07	2.50	2.5	Non Revealing
PR game	6.00	4.50	4.50	4.50	4.50	4.50	3	Partially Revealing

While the shape of the value summarizes the theoretical predictions, it might not be uniquely related to optimal strategies. At the individual level, testing the theory means analyzing more closely the flow of information between players. For the informed player, the optimal behavior is to use information in the last stage of every repeated game, but to use it according to the labels of the game (fully in FR, partially in PR, etc) in the previous stages. The trade-off between using and hiding information which underlies these predictions relies on two main features: the straightforward benefit of the use of information is the increase in the expected payoff of the current stage game, while the cost of using information to a given extent is higher the more stages remain to be played.

Hypothesis 2 (Use of private information by informed subjects) *Informed subjects weight their use of information in line with the trade-off between short-run cost and long-run benefit of concealment: (i) The revelation of private information is higher in the last stage than in the previous stages of the NR repeated games, and higher in the one-shot than in the repeated NR and PR games; (ii) the revelation of private information is higher in the FR repeated games than in the PR repeated games, and higher in the PR repeated games than in the NR repeated games.*

The cost of using information is conditional on the equilibrium reaction of the uninformed player: as long as the use of information is not accounted for, the stage-dominant action can be played at no cost. In other words, the revelation strategy offers the uninformed player the opportunity to extract part of the informational rent from the informed player. This happens if the uninformed player favors his best action in the state of nature signaled by the informed player's history of actions.

Hypothesis 3 (Rent extraction and the reaction of uninformed subjects) *The correlation between the uninformed player's action and the state is positive in the FR and PR games, and higher in the FR than in the NR games. The correlation between the uninformed player's action and the history of play of player 1 is positive in the FR and NR games, and higher in the FR than in the NR games.*

These three hypotheses summarize the treatment effects predicted by the optimal trade-off between using and hiding one's own information on the outcomes of the games over two dimensions:

Table 2: Experimental Design.

Treatment	NR	NR-NoInfo	FR	PR
Payoff matrices	Figure 1	Figure 1	Figure 3	Figure 6
Player 1 informed	Yes	No	Yes	Yes
Player 2 informed	No	No	No	No
<i>Ex-ante</i> belief (p)	1/2	1/2	1/2	1/2
Number of sessions	3	2	1	1
Number of subjects per session	22	22–24	24	22
Total number of pairs	33	23	12	11

Note. All treatments are implemented according to a between-subject design. In all treatments, subjects interact in fixed pairs. The number of stages per game is $n \in \{1, \dots, 5\}$, each repeated 4 times. This results in 20 games, and 60 decisions per subject.

the content of the stage game’s payoffs and the length of the game. Our experiment implements exogenous variations over these two parameters.

4.2 Design of the experiment

To simplify subjects’ comprehension, we implement only one type of game in each session. We thus consider three (between-subject) treatments, in which the FR, NR or PR games are played during a whole experimental session. To obtain benchmark observations on the behavior of an actually uninformed player 1, we introduce a fourth treatment, NR-NoInfo, which is exactly the same as the NR treatment except that neither player 2 nor player 1 is informed about the actual payoff matrix (and this is common knowledge among the subjects). We use the length of the game as a (within-subject) treatment parameter and consider five games in each class $G_n(1/2)$, $n = 1, \dots, 5$.

An important concern in building this kind of design is to allow subjects to get enough familiarity with the functioning of the game.¹⁰ We thus implement several repetitions of each game-length. To allow a clean comparison of our results in the NR treatment when $n = 2$ with Feltovich (1999, 2000), we use a partner matching design – subjects remain in the same role and play against the same opponent during the whole session. The statistical benefit of this choice is that it generates independent data across pairs of subjects. The statistical cost is that we produce correlated data between one repetition and the next at the pair level. This will be accounted for in the statistical analysis.

Each pair of subjects plays 20 repeated games in the session. The number of stages of each repeated game changed after each play of a repeated game and the sequence was the same for all pairs in all sessions. More precisely, we divided a session into four “phases” consisting of five repeated games, and for every $n \in \{1, \dots, 5\}$ the repeated game $G_n(1/2)$ was played once in each phase. We mixed the sequence of lengths between phases to avoid systematic order effects, but kept it constant

¹⁰See, *e.g.*, Binmore (1999) for an extensive discussion of the importance of training subjects to accurately test theories in experiments.

across sessions to reduce unwanted noise. The precise ordering is 32145/15432/21543/32514 in all sessions and all treatments. Both players are provided with full feedback (realized payoff and state) at the end of each game. Table 2 summarizes the experimental design.¹¹

Subjects received the average points they earned in 3 repeated games out of the 20, randomly chosen at the end of the experiment, with a conversion rate of one euro for one point. A participation fee of five euros was also added to the NR and NR-NoInfo treatments because the probability that a subject gets zero payoff is relatively high compared to the two other treatments. Subjects were instructed on the rules of the game and the use of the computer program with written instructions which were read aloud before the start of the experiment. This was followed by a short questionnaire to assess the subjects' understanding of the instructions and one dry run. Afterwards, the twenty repeated games that constituted the experimental treatment took place. Communication between subjects was not allowed. Each session lasted between 45 and 60 minutes.

5 Results

5.1 Subjects' performance: empirical value


The experiment provides an empirical measure of the value of the games, v_n , as the average payoff earned by subjects in the role of player 1, \hat{v}_n . Figure 8 describes the empirical values we observe in each treatment and for each length of the game, $n = 1, \dots, 5$. We also draw the main theoretical benchmarks. We plot the theoretical value, v_n , based on the computations summarized in Table 1. In all experimental games in which player 1 is informed (NR, PR and FR treatments), the empirical value should be bounded above by the theoretical value of the one-shot game, v_1 , and bounded below by the theoretical value of the infinite game, $\text{cav } u$. The empirical value of information is derived from differences from the value of the average game (*i.e.*, the game in which player 1 is uninformed), u .

Table 3 provides a first overview of the comparison between empirical values and theoretical predictions. For each treatment and each length, we report the average payoff earned by informed subjects and the standard error between subjects along with a recall of theoretical values and the p -values of t -tests on the distance between the two.¹²

¹¹All sessions took place in the laboratory of experimental economics at the University of Strasbourg (France) in June 2007. The recruitment of subjects was managed using ORSEE (Greiner, 2004). The translated instructions and questionnaire are provided as supplementary material, Appendices D and E.

¹²In our data, each game-length combination is played 4 times by the same subjects. Observations from a given pair are thus correlated. In what follows, we perform statistical analyses on pooled data at the pairs level, by considering averages over the four repetitions of the same game (of a given length, n). In working at the pair-length level, we disregard the variability across repetitions of the same game. The tests are thus conservative when rejecting the null amounts to reject the theory – *i.e.*, we reject too often observed behavior coinciding with what the theory predicts. The tests are liberal when theory predicts rejection of the null. An alternative empirical approach would be to work at the pair-round level (considering each play of a game as an observation) and to estimate clustered errors at the pair level to account for correlation across games of a given length. This approach is valid only if the number of clusters is high enough. Otherwise, the standard errors must be corrected through bootstrap procedures. We have implemented both solutions, and found very few variations in the results. Between the two available solutions, we chose the one that relies on fewer statistical assumptions and provides fewer chances to conclude that observed behavior matches predicted behavior. This last reason also lead us to perform all statistical tests at the 10% level.

Figure 8: Mean empirical values and some theoretical benchmarks.



Legend. - - - Mean empirical value \hat{v}_n -■- Theoretical value v_n ; - - v_1 ; - - $\text{cav } u$; - - u
Note. For each game, the figures plot the mean empirical value, the theoretical value and the lower and upper bounds of the value functions for each length $n \in \{1, \dots, 5\}$ in abscissa. For the NR game with $n = 2$ stages, we also plot the average value observed in Feltovich (2000) (appears as a \times).

Although empirical values are close to the predictions, especially in the FR games, the dispersion in the FR games is much lower than in the others. This results in a few rejections of the equality test, even in the FR games: equality is rejected for $n = 1$ and $n = 4$. In all treatments, the empirical value is non-increasing in the length of the game, with the exception of $n = 3$ in the NR games. This is due to an empirical value of the NR games with $n = 2$ that is significantly lower than the theoretical prediction (at the 5% level). As can be seen in the upper-right graph of Figure 8, our results in this respect are very close to those obtained by Feltovich (2000). But in contrast to what might have been expected on that basis, the empirical value thereafter smoothly decreases in the NR treatment, though quicker than predicted. This results in a value significantly lower than the prediction for games repeated over $n = 5$ stages. In the PR games, the value of the short games ($n = 1, 2$) differs from the theory. The value of the one-shot game is significantly lower than it should be. The value then decreases, but in a way that is smoother than expected: the value of the 2-stage repeated game remains significantly higher than its theoretical level; it then stabilizes at its theoretical level.

Table 3: Empirical values by games against theoretical levels.

n		1	2	3	4	5
FR	Average empirical value, \hat{v}_n^{FR}	4.71	4.65	4.40	4.36	4.21
	Std. Error	0.48	0.35	0.22	0.20	0.16
	Theoretical value, v_n^{FR}	5.00	4.50	4.33	4.25	4.20
	p-value of $H_0 : \bar{v}_n^{FR} - v_n^{FR} = 0$	0.067	0.189	0.313	0.085	0.862
NR	Average empirical value, \hat{v}_n^{NR}	4.62	3.14	3.53	3.09	2.65
	Std. Error	2.56	1.61	1.61	1.61	1.24
	Theoretical value, v_n^{NR}	5.00	3.75	3.33	3.21	3.07
	p-value of $H_0 : \bar{v}_n^{NR} - v_n^{NR} = 0$	0.406	0.040	0.483	0.668	0.065
PR	Average empirical value, \hat{v}_n^{PR}	5.52	5.22	4.45	4.60	4.75
	Std. Error	0.59	0.68	0.53	0.53	0.59
	Theoretical value, v_n^{PR}	6.00	4.50	4.50	4.50	4.50
	p-value of $H_0 : \bar{v}_n^{PR} - v_n^{PR} = 0$	0.026	0.007	0.788	0.548	0.213

Note. For each game, the first row presents the average payoff of informed subjects (in the role of player 1) across all stages of the four repetitions of the game length in column. The second row presents the between-subject standard error, *i.e.*, deviations computed on these averages. The third row recalls the theoretical levels described in Table 1. The last row gives *p*-values of the test of equality between the two. The number of independent observations in each cell is the number of pairs in each treatment: $N = 33$ in NR, $N = 12$ in FR, $N = 11$ in PR.

Result 1 *The point predictions of the theoretical value of the game are not rejected by the data, except for $n = 1, 4$ in the FR games, $n = 2, 5$ in the NR games, and $n = 1, 2$ in the PR games. The value of the game is (weakly) decreasing in the length in all games except the short ($n \leq 3$) NR games (due to an empirical value much lower than predicted in $n = 2$).*

Support. The empirical value of the game is measured as the average payoff of informed players in each treatment $Tr \in \{NR, FR, PR\}$, \hat{v}_n^{Tr} . Denoting \bar{v}_n^{Tr} the true mean of this distribution, we statistically test whether these point observations are different from their theoretical counterparts v_n^{Tr} using t-tests of $H_0 : \bar{v}_n^{Tr} - v_n^{Tr} = 0$ against $H_1 : \bar{v}_n^{Tr} - v_n^{Tr} \neq 0$. As explained above (see footnote 12) we pool observations for all four repetitions of each game and rely on the individual estimates $\hat{v}_{i,n}^{Tr} = \sum_{r=1}^4 \hat{v}_{i,r,n}^{Tr}/4$. The p-values of the corresponding test statistics are presented in Table 3. The first part of the result relies on rather conservative tests, since we run all comparisons at the 10% level. That the value of information is (weakly) decreasing in the length of the game is straightforwardly deduced from the fact that the empirical value either decreases from one length to the other or is statistically equal to the theoretical value whenever that is not the case (in long PR games, $n \geq 3$). The only exception is \hat{v}_2^{NR} , which both significantly differs from prediction and is lower than \hat{v}_3^{NR} . ■

We now check whether the empirical value satisfies the main qualitative features of the theoretical predictions. First, in zero-sum games, it should always be worthwhile being informed about the actual payoff matrix. In addition, the theory predicts that the value of information should be

strictly positive (*i.e.*, strictly higher than the value of the average game, u) in the FR and PR games, whatever the number of repetitions, but should become negligible in the NR games when the number of repetitions increases. The empirical value should also be higher in the NR treatment than in the NR-NoInfo treatment (in which subjects in the role of player 1 are actually uninformed). Those properties are strongly supported by the data.

Result 2 *The empirical value of information is positive in all games, for all lengths.*

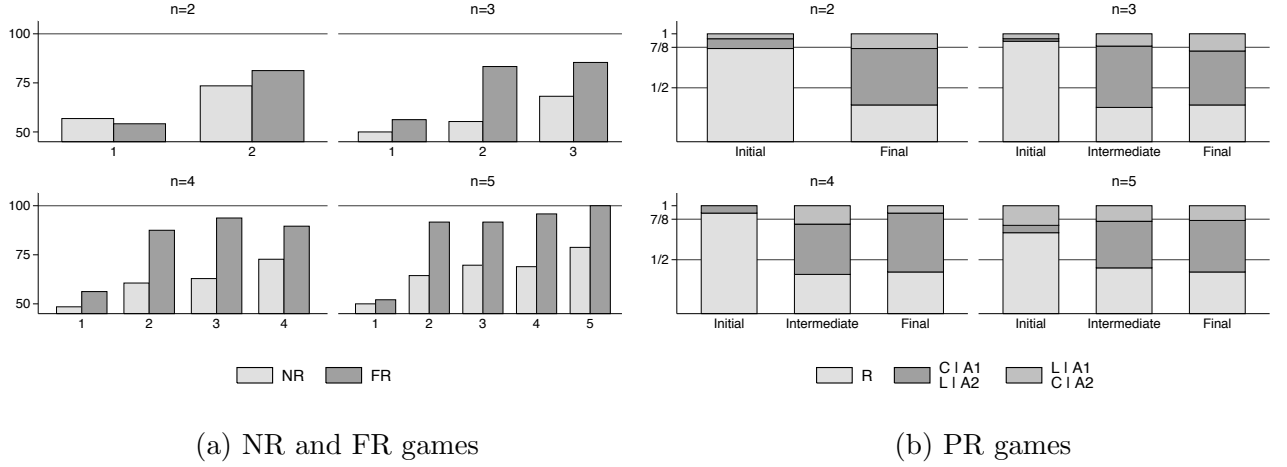
- *In the FR and PR games the empirical value is strictly higher than the theoretical value of the average game for all n ;*
- *In the NR games the empirical value is strictly higher than the theoretical value of the average game for $n \leq 4$, and always strictly higher than the empirical value without information, $\bar{v}_n^{NR-NoInfo}$.*

Support. The value of information is measured as the difference between the average payoff of the informed player, \hat{v}_n^{Tr} , and the value of the average game, u^{Tr} . As shown in Table 1, the values of the average games are: $u^{NR} = u^{FR} = 2.5$ and $u^{PR} = 3$. We statistically test whether these differences are significant for each length, in each treatment, using unilateral t-tests of $H_0 : \bar{v}_n^{Tr} - u^{Tr} \leq 0$ against $H_1 : \bar{v}_n^{Tr} - u^{Tr} > 0$. The null is rejected at the 5% level in all games, except for NR with $n = 5$ (p-value=.25). For the NR games, we also observe the empirical value of the average game, *i.e.*, the average payoff of uninformed players 1: $\hat{v}_n^{NR-NoInfo}$. For this difference, we apply a t-test to unpaired observations (averaged over the four repetitions of the same game at the individual level) from both treatments. The difference from the empirical value in the NR games is significantly positive at the 10% level for $n = 2$ and $n = 4$ and at the 5% level for all other lengths. ■

The theory imposes other simple bounds on expected payoffs the informed player can get from the game thanks to information. First, this expected payoff in the repeated game is strictly lower than the value of the one-shot game, since the informed player can only do worse when the uninformed player observes his past behavior. Second, the informational rent that the uninformed player can extract from observed decisions is higher the longer the game. The payoff of the informed player must thus stay above the value of the infinite game. The plots for all games provided in Figure 8 show that the empirical values always range between these two thresholds: the value moves away from the upper bound in all repeated games ($n > 1$) and reaches the lower bound when the horizon is long enough. As predicted, for the PR treatment we observe a quicker decrease of the empirical value towards the lower bound.

Result 3 *In all games, the empirical value is strictly lower than the value of the one-shot game for $n \geq 2$; it is weakly higher than the value of the infinite game, and strictly higher than the value of the infinite game in the FR games for all n , the NR games for $n \leq 4$, and the PR games for $n = 1, 2$.*

Figure 9: Empirical correlation between the uninformed player's actions and the state.



Note. The left-hand side figure presents results from the FR and NR games. In each graph, the bars plot the proportion of decisions from the uninformed player which are informed decisions (i.e., $(R | A^1, L | A^2)$) for each stage of the game (in abscissa). The right-hand side figure presents results from the PR games. In each graph, the bars plot the proportion of decisions from the uninformed player that correspond to the ones described in the legend, for the first stage ($t = 1$), the intermediate stages ($t < n$), and the final stage ($t = n$). In all figures, games are split by length in each sub-graph.

Support. We apply the same testing procedure as above to the distance between the empirical value and each of the two bounds presented in Table 1. The difference $\bar{v}_n^{Tr} - v_1^{Tr}$ is significant at the 1% level in all treatments, for $n = 2, \dots, 5$. The difference $\bar{v}_n^{Tr} - \text{cav } u^{Tr}$ is significant at the 1% level for all lengths of the FR games, for $n = 1, \dots, 3$ in the NR game and for $n = 1, 2$ in the PR game. It is significant with p-value=.02 for $n = 4$ in the NR game. The empirical value cannot be distinguished from the lower bound for $n = 5$ in the NR game (p-value=.25) and for all $n > 2$ in the PR games (p-value=.61 for $n = 3$; p-value=.27 for $n = 4$; p-value=.11 for $n = 5$.) ■

The experiment provides strong support in favor of the theoretical properties of the value function. Qualitative predictions are fulfilled in the laboratory in accordance to Hypothesis 1, and quantitative predictions are fulfilled in most instances; the main differences arise in the 2-stage NR and PR games.

5.2 Empirical flow of information

To get a flavor of the flows of information in the experiments, we look at the correlation between the uninformed player's actions and the informed player's private information. Such statistics aggregate the use of information on both sides, since the correlation results from how the uninformed player accounts for the informational content of the informed player's decisions. The results for each treatment are presented in Figure 9.

Consider first the FR and NR games in which the theoretical correlation is very simple. In both games, the uninformed player can only play his informed decisions $(R | A^1, L | A^2)$ with probability $1/2$ in the first stage (his action cannot depend on the state, and each state has probability $1/2$).

In equilibrium of the FR games, his action should be perfectly correlated with the payoff matrix in all subsequent stages, playing R in A^1 and L in A^2 . On the contrary, his informed decision should theoretically be played with probability $1/2$ whatever the payoff matrix in all stages of the NR game.¹³ Figure 9.a provides the proportion of informed decisions ($R | A^1, L | A^2$) from the uninformed player in both games, and for each length of the repeated games. While the empirical correlation is too low in most instances of the FR games and, to a larger extent, too high in the NR games, it reacts to changes in the game in keeping with our research Hypothesis 3. The correlation in the FR games always dominates the one observed in the NR games, showing that subjects do use information more in the contexts where they should do so. Moreover, the correlation is increasing in the number of stages in all versions of the repeated games, and is slightly increasing in the length for a given stage inside the game.

Result 4 *The correlation of the uninformed player's decisions with the actual state of nature is strictly higher in the FR games than in the NR games. However, in most instances, this correlation is strictly lower than predicted in the FR games and strictly higher than predicted in the NR games.*

Support. The above described correlation pattern provides predictions for the average value of the dummy variable $BR = \mathbf{1}[R|A^1, L|A^2]$: this should equal 0.5 in the NR games and at the first stage of the FR games, and equal 1 in any repetition of the FR games: $BR_{t,n}^{Tr} = 1$ if $\{Tr = FR, t > 1\}$; 0.5 otherwise. The distance of this variable from 0.5 thus measures the extent to which the decision of the uninformed player is in fact informed. Again, we disregard intra-pair variability and apply a t-test of $H_0 : \overline{BR}_{t,n}^{Tr} = BR_{t,n}^{Tr}$ on averages over the four occurrences of each game played by a given pair of subjects. The differences are significant at the 10% level in almost all repetitions of the games. The exceptions are: the NR game with $n = 3, t = 2$ and the FR game with $n = 5, t = \{4, 5\}$. The FR and NR games are implemented as separate treatments. Observations are thus independent between treatments. Applying the t-test procedure for unpaired observations of $H_0 : \overline{BR}_{t,n}^{FR} = \overline{BR}_{t,n}^{NR}$, the difference is highly significant ($p < 0.01$) for any repetition of the games. ■

In the PR games, if the informed player randomizes between the sequence of stage-dominant actions ($T \dots T$ in A^1 and $B \dots B$ in A^2) with probability $3/4$, and the sequence of stage-dominated actions ($B \dots B$ in A^1 and $T \dots T$ in A^2) with probability $1/4$, then the uninformed player plays R in stage 1, he plays R and C (L , respectively) with the same probability in stages $t = 2$ to $t = n - 1$ after a history of T (B , respectively), and in the last stage he plays C (L , respectively) after a history of T (B , respectively). Hence, along the equilibrium path of the PR games, the uninformed player plays on average

- R in the *Initial stage* – i.e., stage 1 for all $n > 1$;
- $(\frac{1}{2}R + \frac{3}{8}C + \frac{1}{8}L)$ in A^1 and $(\frac{1}{2}R + \frac{1}{8}C + \frac{3}{8}L)$ in A^2 in *Intermediate stages* – i.e., stages $2, \dots, n - 1$;

¹³This can be shown from the optimal strategies in the NR games presented in Appendix A, and is true more generally irrespective of the selected optimal strategy of player 2.

– $(\frac{3}{4}C + \frac{1}{4}L)$ in A^1 and $(\frac{1}{4}C + \frac{3}{4}L)$ in A^2 in the *Final stage* – stage n .

Figure 9.b displays the share of decisions of the uninformed player in each of those three sub-parts of the PR games. While we do observe some decisions other than R at the beginning of the PR games, the predicted decision is by far the most frequent in stage 1. This share decreases in intermediate stages, and stabilizes around the one-half share predicted by theory. This decrease is offset by very asymmetric increases in the two other kinds of decisions: the share of actions C in A^1 and L in A^2 rises dramatically, while the share of L in A^1 and C in A^2 only slightly increases. In all games, the observed shares of each pair of contingent decisions are very close to the $3/8 - 1/8$ levels predicted by theory. By contrast, experimental subjects do not adjust their strategies at the final stages of the PR games: decisions essentially remain the same as during the intermediate stages.

Result 5 *In all but the last stage of the PR games, the correlation of the uninformed player's decisions with the actual state of nature matches the theoretical predictions.*

Support. In the PR games, the uninformed player chooses between three actions. We statistically test the changes in probability inside each game through linear probability models. Denoting $I[R]$ the binary variable associated with decision R from the uninformed player, and $I[t = 1]$, $I[1 < t < n]$ and $I[t = n]$ the binary variables associated with the first stage, the intermediate stages and the final stage of the n -stage repeated game, we thus estimate the unknown parameters b_k in the model: $I[R] = b_1I[t = 1] + b_2I[1 < t < n] + b_3I[t = n] + \varepsilon$. To estimate the model, we consider all plays of a given game n, Tr by a given pair of subjects as one observation. Because the model is linear, the parameters measure the change in probability of the dependent variable induced by the explanatory variables. A well-known drawback of this specification is that errors are heteroscedastic; we therefore use robust standard errors. The same estimation procedure is applied to the decision to play C in A^1 and L in A^2 , $I[C|A^1, L|A^2]$. The results of separate estimation for each length are presented in Table 4, along with confidence intervals at the 95% level. The second column summarizes the theoretical predictions discussed above.

The coefficients and confidence intervals clearly confirm the theoretical predictions on behavior at the initial stage. The probability that the informed player plays R at the beginning of the game is higher than 75% whatever the length of the games, with upper bounds of the 95% confidence intervals close to the theoretical 100% level for $n = 2, 5$ and higher for $n = 3, 4$. The probability that C is played in A^1 and L is played in A^2 , by contrast, is statistically equal to 0 for all lengths – confidence intervals on the effect of $I[t = 1]$ in the bottom part of the Table always include 0.

In intermediate stages, the share of decisions R significantly decreases compared with the rate observed at the initial stage: confidence intervals associated with $I[t = 1]$ and $I[1 < t < n]$ never overlap. The estimated coefficients are generally lower than the 0.5 equilibrium share, although upper bounds of confidence intervals are again very close to this threshold, if not above ($n = 5$). Similarly, the share of decisions C in A^1 and L in A^2 significantly increases compared with the initial stage. The estimated share is generally higher than the predicted level (equal to 0.375 in this case) but the lower bound of confidence intervals are lower than this threshold for $n > 3$.

Table 4: Regressions on the uninformed player’s decisions: PR game

Stage dummies	Expected Share	Model			
		$n = 2$	$n = 3$	$n = 4$	$n = 5$
Dependent variable: Uninformed player chooses R					
$I[t = 1]$	1.000	0.864 [0.726, 1.002]	0.932 [0.863, 1.001]	0.932 [0.864, 1.000]	0.750 [0.605, 0.895]
$I[1 < t < n]$	0.500	– –	0.318 [0.185, 0.452]	0.364 [0.222, 0.505]	0.424 [0.322, 0.527]
$I[t = n]$	0.000	0.341 [0.163, 0.519]	0.341 [0.175, 0.506]	0.386 [0.209, 0.563]	0.386 [0.223, 0.550]
Dependent variable: Uninformed player chooses (C A¹, L A²)					
$I[t = 1]$	0.000	0.091 [–0.009, 0.191]	0.023 [–0.022, 0.067]	0.068 [0.000, 0.136]	0.068 [0.001, 0.136]
$I[1 < t < n]$	0.375	– –	0.568 [0.405, 0.731]	0.466 [0.323, 0.609]	0.432 [0.332, 0.531]
$I[t = n]$	0.750	0.523 [0.309, 0.736]	0.500 [0.313, 0.687]	0.545 [0.321, 0.770]	0.477 [0.278, 0.676]

Note. Results from linear probability models on the decision of the uninformed player in PR games. The upper part of each row presents estimated coefficients; the bottom part presents 95% confidence intervals computed from robust standard errors.

Observed decisions at the final stage clearly depart from predicted behavior. The shares of both decision R and either C in A^1 or L in A^2 essentially remain the same as during intermediate stages – confidence intervals on $I[1 < t < n]$ and $I[t = n]$ largely overlap. ■


We now turn to a less aggregated analysis of behavior underlying such correlation patterns: when and to what extent does the informed player reveal information and the uninformed player react to observed decisions?

6 Players’ behavior

6.1 Informed player’s behavior

The three experimental treatments are labelled according to the expected amount of information the informed player should use. The resulting local strategies at each stage may still finely depend on the history of the game, especially in the NR games (see Appendix A). Given the range of lengths we study, the set of possible histories is large, resulting in an untractable set of possible strategies. We circumvent this issue by focusing on the properties of informed players’ decisions at each stage independently of the history. We define $D_{i,g,t}$ as the dummy variable indicating whether the decision of the informed player i , in stage $t = 1, \dots, n$ of game $g = 1, \dots, 20$, is the stage-dominant decision: $D_{i,g,t} = \mathbf{1}[T|A^1, B|A^2]$. Thus, $\bar{D}_{\dots,t} = 1$ for the fully-revealing (stage-dominant) strategy, *i.e.*, the optimal strategy in the static game, and $\bar{D}_{\dots,t} = 1/2$ for a non-revealing strategy (pure randomization between the two actions).

Figure 10: Relative frequency of the stage-dominant action by informed subjects.




Note. For each treatment and each length, the figures display the mean share of the informed player's stage-dominant actions, in the final stage (*left-hand side*) and in intermediate stages (*right-hand side*).

Figure 10 provides a rough summary of how information is used in each treatment. Remember that all treatments have one prediction in common: the informed player has nothing to lose in using his private information (i.e., playing the stage-dominant action) at the last stage of the game. We thus separate the figures according to the stage inside each game: the last stage of all games is reported on the left-hand side, intermediate stages of all repeated games (in stages $t = 1$ to $t = n - 1$ for all $n > 1$) are reported on the right-hand side. From both the left-hand side figure and the frequency of the stage-dominant action observed in the FR and PR games, experimental subjects unambiguously use information whenever it is worthwhile to do so. The relative frequency of the stage-dominant action in the FR games ranges from 94% (first stage of the 5-stage repeated game) to 100% (in a vast majority of the FR games, with the exception of $(n = 1, t = 1)$, $(n = 3, t = 1)$ and $(n = 3, t = 2)$). In the PR games, information is used to its expected extent: in long games ($n > 2$) the relative frequency of the stage-dominant action always remains very close to the 75% theoretical level in intermediate stages (from 65% in $(n = 5, t = 4)$ to 88% in $(n = 4, t = 1)$). In the two-stage PR game this relative frequency is much higher at every stage of the game, which is compatible with the (multiple) equilibrium predictions when $n = 2$ (according to which the stage-dominant action is played with probability between 75% and 100% in the first stage; see Section 3.3). In the last stage of the NR games, the stage-dominant action is played more than 95% of the time. Thus, experimental subjects adjust their use of information not only as a reaction to experimental treatments, but also along the path of decisions over stages of a given game.

The ability of experimental subjects to optimally ignore information is much weaker. As an empirical benchmark, Figure 10 presents observations from the NR-NoInfo treatment, in which neither player 1 nor player 2 is informed about the actual payoff matrix. In the NR games, informed subjects should theoretically behave (almost) as in the NR-NoInfo treatment, i.e., they should ignore their private information and play each action with probability close to 50%. Contrasting

Figure 11: Relative frequency of the stage-dominant action by informed subjects, by stage.



Note. For each stage in *abscissa*, the dots give the mean share of informed player's decisions that are the stage-dominant action of the game actually played. All stages from a given length (see the legend) are connected.

these frequencies between the two treatments unambiguously suggests decisions are over-correlated with information: the relative frequencies of the stage-dominant actions we elicit in intermediate stages of the NR games dominate those observed in NR-NoInfo, and are very often similar to those observed in the PR games. Still, these frequencies remain lower than in the FR games.

While experimental subjects overuse their information in the NR games, their strategy reacts in a sophisticated manner to changes in the environment. Figure 11 disaggregates the relative frequency of the stage-dominant action according to each stage inside the n -stage repeated games. The cost of excessive use of the stage-dominant strategy is that the uninformed player becomes more and more able to play his informed decision ($R | A^1, L | A^2$) and thus to get a higher share of the total payoff in subsequent repetitions of the game. Clearly, the cost of revealing information is increasing in the number of stages towards the end of the current game; that is, it is higher if information is revealed earlier in the game for a given length, and it is higher if information is revealed in longer games for a given stage. Figure 11 confirms that informed subjects adjust their use of information according to those parameters.

Result 6 *The strategic use of information by the informed player has the following properties:*

- *The stage-dominant action is almost always chosen in all instances of the FR games and in the last stage of all games;*
- *The stage-dominant action is chosen too often in all other instances of the NR games; however, the relative frequency of the stage-dominant action in stage t is decreasing in the total number of stages, n , for every stage t , and increasing in t for every length, n , of the game;*
- *The relative frequency of the stage-dominant action is close to the theoretical level and slightly decreasing in t in intermediate stages of all games in the PR games.*

Table 5: Decision of the informed player to play the stage-dominant action

Stage (t)	$n = 1$	$n = 2$		$n = 3$			$n = 4$				$n = 5$				
	1	1	2	1	2	3	1	2	3	4	1	2	3	4	5
NR-NoInfo	0.51	0.57	0.55	0.56	0.53	0.53	0.57	0.53	0.51	0.49	0.49	0.52	0.45	0.43	0.57
Lower B.	0.41	0.46	0.44	0.46	0.44	0.44	0.48	0.44	0.42	0.39	0.38	0.41	0.33	0.31	0.49
Upper B.	0.61	0.68	0.67	0.66	0.62	0.63	0.67	0.62	0.60	0.59	0.60	0.63	0.57	0.55	0.65
NR	0.97	0.89	0.98	0.79	0.87	0.95	0.73	0.73	0.89	0.98	0.58	0.71	0.74	0.85	0.98
Lower B.	0.94	0.82	0.94	0.71	0.81	0.92	0.63	0.65	0.84	0.95	0.48	0.61	0.66	0.78	0.95
Upper B.	1.00	0.95	1.01	0.87	0.93	0.99	0.84	0.82	0.95	1.00	0.67	0.81	0.82	0.92	1.00
PR	0.95	0.93	0.89	0.77	0.75	0.93	0.89	0.80	0.73	0.82	0.80	0.75	0.73	0.66	0.91
Lower B.	0.89	0.86	0.75	0.65	0.60	0.86	0.81	0.67	0.56	0.63	0.68	0.62	0.62	0.51	0.83
Upper B.	1.01	1.00	1.02	0.90	0.90	1.00	0.96	0.92	0.90	1.01	0.91	0.88	0.83	0.81	0.98
FR	0.98	1.00	1.00	0.98	0.96	1.00	1.00	1.00	1.00	1.00	0.94	1.00	1.00	1.00	1.00
Lower B.	0.94	1.00	1.00	0.94	0.90	1.00	1.00	1.00	1.00	1.00	0.85	1.00	1.00	1.00	1.00
Upper B.	1.02	1.00	1.00	1.02	1.01	1.00	1.00	1.00	1.00	1.00	1.03	1.00	1.00	1.00	1.00

Note. OLS estimations on the decision of the informed player to play the stage-dominant action. An observation is the average for each pair over the four repetitions of the same game. For each treatment, the upper part of the row presents estimated coefficients; the two subsequent rows give the lower and upper bounds of 95% confidence intervals computed from robust standard errors.

Support. To obtain confidence intervals on the share of decisions by the informed player which are the stage-dominant action in each stage, we specify linear probability models on $D = \mathbf{1}[T|A^1, B|A^2]$. We estimate separate OLS regressions for each treatment and each length of the model: $D = b_1\mathbf{1}[t = 1] + \dots + b_n\mathbf{1}[t = n] + \varepsilon$, on data averaged at the pair-length level. We estimate robust standard errors to account for the induced heteroscedasticity. Results are presented in Table 5.

In instances where the informed player should fully use private information (FR games and last stage in all treatment) the 95% confidence intervals are close to the expected 100% share of stage-dominant actions. There is much less dispersion, though, in the FR games – in which the optimal share always remains the same – than in the last stages of other games. In intermediate stages of long PR games (*i.e.*, $n > 2$ and $t < n$), the confidence intervals contain percentages that are always higher than 50% and lower than 100%; in most cases, the prediction that 75% of decisions are the stage-dominant actions lies inside the confidence interval.

The statistical support for the qualitative variations in the use of information comes from regressions on the probability that the informed player uses the stage-dominant action. The detailed results are presented in the Supplementary Material, Section B. ■

To sum up, information is accurately used when it should be by informed players, and the use of information qualitatively reacts to changes in the environment according to our research Hypothesis 2. But subjects experience difficulties in behaving as if they were uninformed (*i.e.*, in the shorter versions of the NR games). We now turn to the ability of uninformed subjects to account for such revelation patterns.

6.2 Uninformed player’s behavior

In zero-sum games such as ours, the overuse of information from informed subjects can either improve or harm their payoff depending on how uninformed players accommodate this deviation (by accounting for the information contained in informed subjects’ actions). Analysis of the empirical values of the games has already shown that the expected payoffs match with the theory in many instances; so it is already clear that the uninformed player does not take full advantage of the overuse of information. The fact that we do observe a decrease, but not a drastic drop, in the values as the number of stages increases means that the uninformed players do account for the signal, but less than they could. This was further confirmed by the empirical correlation of uninformed subjects’ actions with the state of nature (Section 5.2).

While uninformed players do appear to account for the information they receive, this does not mean they do so in the expected way. To further explore this dimension, we now turn to the questions of how and to what extent uninformed players react to the signal, *i.e.*, the history of play of player 1.¹⁴ To that end we compare, for each possible history of play from player 1, empirical actions to theoretical mixed actions supporting the equilibrium. Evidence is presented in Tables 6 and 7 in the Supplementary Material, Section C.1. The first conclusion one can draw from the tables is that the empirical relative frequency of action L follows the same trend as theoretical ones: the relative frequency of action L is increasing in each stage of the NR and FR games in the difference between the number of actions B and the number of actions T in the history. In addition, and again in line with our research Hypothesis 3, this correlation is higher in the FR than in the NR games.¹⁵

7 Conclusion

This paper investigates the empirical content of the theoretical predictions associated with the class of zero-sum repeated games with incomplete information drawn from Aumann and Maschler (1966, 1967, 1995). We study three payoff structures with the same trivial optimal strategies in the one-shot versions of the games, but which differ according to the amount of information the informed player should exploit in the repeated versions of the games. The empirical value of information is in keeping with the qualitative predictions of the theory, and often with the quantitative predictions too. We reject, in particular, that the value of information becomes negative in long games in which information should be disregarded.

In line with our research hypotheses, we found that experimental subjects react to the length and payoff structure of the games according to the trade-off between the short-term benefit and

¹⁴In this section, we only focus on the NR and FR games. One reason is that they are not directly comparable with the PR games as regards the behavior of player 2: at each information set there are three possible actions in the PR games but only two possible actions in the NR and FR games. Another reason is that we only have one session of the PR games (11 pairs of subjects), which means that we do not have a lot of data to analyze player 2’s behavior for each possible length and history.

¹⁵The statistical significance of the effects commented on in this paragraph is derived from OLS regressions of the uninformed player’s decisions on the content of the history. This is presented as supplementary material, Appendix C.2.

long-term cost of using private information. In particular, even though optimal strategies of the informed player are equivalent in all static versions of the games, subjects behave very differently in the three versions of the repeated games. Uninformed subjects react to the history of actions of informed subjects, so the values of the games for informed subjects decrease in the length of the game. But since the use of information by informed subjects correctly reacts to the type and length of the game, this value not only remains positive, but never falls below the theoretical value of the infinitely repeated game. An interesting further avenue of research would be to study how this trade-off is managed between subjects who are trained differently to similar strategic situations, for example using professional poker players, or matching professional poker players with students.

Appendix

A Equilibrium strategies for the NR games

Denote p_t player 2's posterior belief about the state in stage t , $t = 1, \dots, n$, $y_t(h_{t-1})$ the strategy (probability of playing L) of player 2 in stage t given the history of actions $h_{t-1} \in \{T, B\}^{t-1}$ of player 1, and $x_t^1(p_t)$ and $x_t^2(p_t)$ the strategy (probabilities of playing T in A^1 and A^2) of player 1 in stage t given the posterior p_t of player 2 in stage t . In the following tables, we describe optimal strategies $((x_t^1, x_t^2), y_t)$ and player 2's posteriors p_t for the NR games. We use the following recursive formula (see, e.g., Zamir, 1992, p. 126), which considerably helps to find the value of any (finitely) repeated game and the associated optimal strategies:

$$v_{n+1}(p) = \frac{1}{n+1} \max_{x \in [\Delta(T)]^2} \min_{y \in \Delta(J)} \left(\sum_k p^k A^k(x^k, y) + n \sum_{i \in I} \left(\sum_k p^k x^k(i) \right) v_n(\hat{p}(x, i)) \right), \quad (1)$$

where $\hat{p}(x, i) = \left(\frac{p^k x^k(i)}{\sum_k p^k x^k(i)} \right)_{k \in \{1, 2\}}$ is the posterior probability over $\{1, 2\}$ given action i and strategy x of player 1. In the NR games, the recursive formula simplifies to

$$\begin{aligned} v_{n+1}(p) &= \frac{10}{n+1} \max_{(x^1, x^2) \in [0, 1]^2} \min_{y \in [0, 1]} \\ &\quad \left(px^1 y + (1-p)(1-x^2)(1-y) + n \left(\alpha v_n \left(\frac{px^1}{\alpha} \right) + (1-\alpha) v_n \left(\frac{p(1-x^1)}{1-\alpha} \right) \right) \right), \\ &= \frac{10}{n+1} \max_{(x^1, x^2) \in [0, 1]^2} \\ &\quad \left(\begin{cases} px^1 & \text{if } px^1 \leq (1-p)(1-x^2) \\ (1-p)(1-x^2) & \text{if } px^1 \geq (1-p)(1-x^2) \end{cases} + n \left(\alpha v_n \left(\frac{px^1}{\alpha} \right) + (1-\alpha) v_n \left(\frac{p(1-x^1)}{1-\alpha} \right) \right) \right), \end{aligned}$$

where $\alpha = px^1 + (1-p)x^2$ is the probability that player 1 plays T . The optimal strategy of player 1 depending on player 2's posterior belief is deduced from the max-minimization program of the recursive formula.¹⁶ Next, from the equilibrium conditions of player 1 (indifference when he uses a strictly positive mixed strategy) we deduce the equilibrium strategies of player 2. We obtain:

¹⁶As already mentioned in Section 3.2, the equilibrium strategy of player 2 is not always unique in some NR games (when $n = 1$, $n = 2$ and $n = 5$). In these instances, we select the one in which the uninformed player reacts symmetrically to the history – i.e., $y_t(h_{t-1}) = 1 - y_t(\bar{h}_{t-1})$ when \bar{h}_{t-1} is the history obtained from h_{t-1} by substituting actions B with actions T , e.g., $y_5(TBTT) = 1 - y_5(BTBB)$.

$G_1(1/2)$

t	h_{t-1}	x_t^1	x_t^2	p_t	y_t
1	\emptyset	1	0	1/2	1/2

 $G_2(1/2)$

1	\emptyset	1/2	1/2	1/2	1/2
2	T	1	0	1/2	1/4
	B				3/4

 $G_3(1/2)$

1	\emptyset	1/2	1/2	1/2	1/2
2	T	1/2	1/2	1/2	1/2
	B				
3	TT	1	0	1/2	0
	TB				1/2
	BT				
	BB				1

 $G_4(1/2)$

1	\emptyset	4/7	3/7	1/2	1/2	
2	T	1/2	1/3	4/7	3/7	
	B				2/3	4/7
3	TT	1/2	0	2/3	5/14	
	TB				1/2	1/2
	BT					
	BB				1	1/3
4	TTT	1	0	1/2	0	
	BTT				3/14	
	TBT				2/7	
	TTB				5/14	
	BBT				9/14	
	BTB				5/7	
	TBB				11/14	
	BBB			0	1	

 $G_5(1/2)$

t	h_{t-1}	x_t^1	x_t^2	p_t	y_t			
1	\emptyset	1/2	1/2	1/2	1/2			
2	T	4/7	3/7	1/2	7/16			
	B				9/16			
3	TT	1/2	1/3	4/7	41/112			
	BT				55/112			
	TB				2/3	1/2	3/7	57/112
	BB							71/112
4	TTT	1/2	0	2/3	13/56			
	BTT				27/56			
	TTB	1/2	1/2	1/2	1/2			
	TBT							
	BTB							
	BBT							
	TBB	1	1/2	1/3	29/56			
	BBB				43/56			
5	$TTTT$	1	0	1/2	0			
	$BTTT$							
	$TTTB$				13/56			
	$TTBT$				11/112			
	$TTBB$				67/112			
	$TBTT$				3/112			
	$TBTB$				59/112			
	$TBBT$				29/56			
	$BTTB$				27/56			
	$BTBT$				53/112			
	$BTBB$				109/112			
	$BBTT$				45/112			
	$BBTB$				101/112			
	$BBBT$			43/56				
$TBBB$	0	1						
$BBBB$								

References

- AUMANN, R., AND J. DREZE (2008): “Rational Expectations in Games,” *American Economic Review*, 98(1), 72–86.
- AUMANN, R. J., AND M. MASCHLER (1966): “Game Theoretic Aspects of Gradual Disarmament,” Report of the U.S. Arms Control and Disarmament Agency, ST-80, Chapter V, pp. 1–55.
- (1967): “Repeated Games with Incomplete Information: A Survey of Recent Results,” Report of the U.S. Arms Control and Disarmament Agency, ST-116, Chapter III, pp. 287–403.
- AUMANN, R. J., AND M. B. MASCHLER (1995): *Repeated Games of Incomplete Information*. MIT Press, Cambridge, Massachusetts.
- BINMORE, K. (1999): “Why Experiment in Economics?,” *Economic Journal*, 109(453), F16–F24.
- CHAUDHURI, A. (1998): “The Ratchet Principle in a Principal Agent Game with Unknown Costs: An Experimental Analysis,” *Journal of economic behavior & organization*, 37(3), 291–304.
- FELTOVICH, N. (1999): “Equilibrium and Reinforcement Learning in Private-Information Games: An Experimental Study,” *Journal of Economic Dynamics and Control*, 23(9-10), 1605–1632.
- (2000): “Reinforcement-Based vs. Belief-Based Learning Models in Experimental Asymmetric-Information Games,” *Econometrica*, 68(3), 605–641.
- GREINER, B. (2004): “An Online Recruitment System for Economic Experiments,” in *Forschung und wissenschaftliches Rechnen 2003. GWDG Bericht 63*, ed. by K. Kremer, and V. Macho, pp. 79–93. Gesellschaft für Wiss. Datenverarbeitung, Göttingen.
- HARRINGTON, D., AND B. ROBERTIE (2007): *Harrington on Hold'em: Expert Strategy for No-limit Tournaments. Volume I: Strategic Play*. Two Plus Two Pub.
- HART, S., AND A. MAS-COLELL (2000): “A Simple Adaptive Procedure Leading to Correlated Equilibrium,” *Econometrica*, 68(5), 1127–1150.
- HINSLEY, F. (1993): *British Intelligence in the Second World War*. Cambridge University Press.
- HÖRNER, J., AND J. JAMISON (2008): “Sequential Common-Value Auctions with Asymmetrically Informed Bidders,” *Review of Economic Studies*, 75(2), 475–498.
- MARINACCI, M. (2000): “Ambiguous Games,” *Games and Economic Behavior*, 31(2), 191–219.
- MCKELVEY, R., AND T. PALFREY (1995): “The holdout game: An experimental study of an infinitely repeated game with two-sided incomplete information,” in *Social choice, welfare, and ethics: Proceedings of the Eighth International Symposium in Economic Theory and Econometrics*, pp. 321–49. New York: Cambridge University Press.
- PALACIOS-HUERTA, I. (2003): “Professionals Play Minimax,” *Review of Economic Studies*, 70(2), 395–415.
- STEARNS, R. (1967): “A Formal Information Concept for Games with Incomplete Information,” Report of the U.S. Arms Control and Disarmament Agency, ST-116, Chapter IV, pp. 405–433.

WOODERS, J. (2010): “Does Experience Teach? Professionals and Minimax Play in the Lab,” *Econometrica*, 78(3), 1143–1154.

ZAMIR, S. (1992): “Repeated Games of Incomplete Information: Zero-Sum,” in *Handbook of Game Theory*, ed. by R. J. Aumann, and S. Hart, vol. 1, chap. 5, pp. 109–154. Elsevier Science B. V.

Supplementary material to: Jacquemet N., Koessler F. “Using or Hiding Private Information? An Experimental Study of Zero-Sum Repeated Games with Incomplete Information

B Statistical analysis of informed subjects’ reaction to the length of the game

The conclusions presented in result 6 regarding the sensitivity of informed player’s behavior to the length, n , and the stage, t , comes from regressions on the probability that the informed player uses the stage-dominant action. This led us to work at the game level, considering each play of a given length as one observation. To account for multiple observations from the same pair of subjects, the regressions include pair dummies and round dummies – which identify the order of the game among the 20 played by each pair. Since we do not need to interpret the marginal effects, and only signs are of interest, we use probit models. For both the NR games and the PR games, we estimate two specifications: the first model includes only the effect of the length (n) and the stage (t), the second isolates the last stage of the game.

	NR treatment				PR treatment			
	Coef. ⁽¹⁾	p-value	Coef. ⁽²⁾	p-value	Coef. ⁽¹⁾	p-value	Coef. ⁽²⁾	p-value
Constant	1.36	0.007	1.15	0.026	1.92	0.017	1.12	0.175
n	-0.47	0.000	-0.38	0.001	-0.27	0.351	0.05	0.870
t	0.42	0.000	0.31	0.000	0.03	0.629	-0.15	0.055
Final	–	–	0.57	0.000	–	–	0.67	0.001
Nb Obs.	1798				638			

Note. Probit regressions of the probability that the informed player uses the stage-dominant action on the stage (t) and the length (n) of the game (Models 1) and the dummy variable $I[t = n]$ isolating the last stage of the game (Models 2). All models include pair dummies and round dummies. P-values are computed according to robust standard errors.

For the NR games, the use of information is significantly decreasing in n and increasing in t even when the last stage of each game is separated from intermediate stages (model 2). For the PR games, the use of information is constant across different lengths. Once the final stage is identified separately, the use of information appears to decrease with the stage. This supports the last two items of the Result.

C Statistical analysis of uninformed subjects’ reaction to history

C.1 Descriptive statistics

In support to the discussion of the uninformed player’s behavior provided in Section 6.2, Tables 6 and 7 present the relative frequency of action L , denoted $\hat{y}_t(h_{t-1})$, in the FR and NR games for each stage/length combination and according to the actually observed sequence of decisions from player 1 (h_{t-1}). We also present a mixed action – $y_t(h_{t-1})$ – supporting the equilibrium of the corresponding games for each possible sequence of observed decisions. The variable $Use(h_{t-1}) = |\hat{y}_t(h_{t-1}) - 0.5| - |y_t(h_{t-1}) - 0.5|$ measures the distance between observed and theoretical frequencies of action L as a function of the history h_{t-1} of actions of player 1. By construction, it is positive (negative) when uninformed subjects over-react (under-react) to the history as compared with the theoretical prediction. Based on the average values provided in each cell,

Table 6: Uninformed player: Equilibrium and observed strategies in the FR games

			$n = 5$		$n = 4$		$n = 3$		$n = 2$	
t	h_{t-1}	y_t	\hat{y}_t	Use	\hat{y}_t	Use	\hat{y}_t	Use	\hat{y}_t	Use
2	B	1.000	0.870	-0.130	0.875	-0.125	1.000	0.000	0.750	-0.250
	T	0.000	0.000	0.000	0.125	-0.125	0.280	-0.280	0.125	-0.125
	Average	-	-	-0.063	-	-0.125	-	-0.146	-	-0.188
3	BB	1.000	0.909	-0.091	0.917	-0.083	1.000	0.000		
	TB	1.000	0.000	0.000	-	-	1.000	0.000		
	BT	0.000	0.000	0.000	-	-	-			
	TT	0.000	0.000	0.000	0.042	-0.042	0.227	-0.227		
	Average	-	-	-0.042	-	-0.062	-	-0.104		
4	BBB	1.000	0.909	-0.091	0.917	-0.083				
	TBB	1.000	1.000	0.000	-	-				
	BTT	0.000	0.000	0.000	-	-				
	TTT	0.000	0.000	0.000	0.125	-0.125				
	Average	-	-	-0.042	-	-0.104				
5	BBBB	1.000	1.000	0.000						
	TBBB	1.000	1.000	0.000						
	BTTT	0.000	0.000	0.000						
	TTTT	0.000	0.000	0.000						
	Average			0.000						
Average				-0.036		-0.097		-0.125		-0.188

Note. For each game length, n , and each stage, t , the table presents the empirical (\hat{y}_t) mixed strategies (in terms of the relative frequency of playing L) against all possible histories of play by the informed player (h_{t-1}) in the FR games. The third column presents the theoretical mixed strategy, y_t . The Use variable is an index of the distance between observed and theoretical use of the history, constructed as: $Use_t = |\hat{y}_t - 0.5| - |y_t - 0.5|$. The *Average* of the Use variable for each stage-length combination and for each length is computed using the weights defined by the frequency of each history in the data.

uninformed subjects appear to react less than predicted when the informed player's decision is/should be fully revealing, and more otherwise.¹⁷ In both tables, the histories are ordered according to the number of actions B in the history (from the highest to the smallest), in such a way that the theoretical prediction, $y_t(h_{t-1})$, is increasing in the number of actions B .

The analysis of uninformed subjects' strategies in the NR games sheds new light on the patterns observed in the previous sections. In particular, we observed in Figure 8 and Result 1 that the empirical value for informed subjects is higher than the theoretical value only in the 3-stage game, and it is significantly smaller than the theoretical value for $n = 2$ and $n = 5$. The difference between the 2-stage and the 3-stage games could partly be explained by the fact that informed subjects use less informative strategies in longer games (see Figure 11 (a) and Result 6), but this cannot explain the difference between the 3-stage and the 5-stage game. Similarly, we observed in Figure 9 that the correlation at a given period between uninformed subjects' actions and the state could be higher in longer games (e.g., in $(n = 5, t = 2)$) than in shorter games (e.g., in $(n = 3, t = 2)$), which seems inconsistent with the fact that informed subjects' actions in period $t = 1$ are less correlated with the state in longer games. These patterns can be explained by adding uninformed subjects' strategies into the picture: it is precisely in the 3-stage game that uninformed subjects react less to the history than they should according to the theory. For example, the correlation of uninformed subjects'

¹⁷Note that over- and under- reaction here refers to equilibrium predictions, not to the best reply to the actually observed behavior of the informed player.

Table 7: Uninformed player: Equilibrium and observed strategies in the NR games

t	h_{t-1}	$n = 5$			$n = 4$			$n = 3$			$n = 2$		
		y_t	\hat{y}_t	Use	y_t	\hat{y}_t	Use	y_t	\hat{y}_t	Use	y_t	\hat{y}_t	Use
2	B	0.562	0.833	0.271	0.571	0.620	0.049	0.500	0.636	0.136	0.750	0.753	0.003
	T	0.438	0.264	0.174	0.429	0.293	0.136	0.500	0.351	0.149	0.250	0.064	0.186
	Average			0.218			0.103			0.144			0.068
3	BB	0.634	0.926	0.292	0.643	0.727	0.084	1.000	0.743	-0.257			
	TB	0.491	0.484	0.007	0.500	0.500	0.000	0.500	0.667	0.167			
	BT	0.509	0.424	0.067	0.500	0.321	0.179	0.500	0.500	0.000			
	TT	0.366	0.049	0.317	0.357	0.154	0.203	0.000	0.125	-0.125			
	Average			0.177			0.132			-0.095			
4	BBB	0.768	1.000	0.232	1.000	0.857	-0.143						
	TBB	0.518	0.786	0.268	0.786	0.800	0.014						
	BTB	0.500	0.733	0.233	0.643	0.667	0.024						
	BBT	0.500	0.600	0.100	0.714	0.625	-0.089						
	TTB	0.500	0.429	0.071	0.286	0.250	0.036						
	TBT	0.500	0.353	0.147	0.357	0.450	-0.093						
	BTT	0.482	0.389	0.093	0.214	0.316	-0.102						
	TTT	0.232	0.111	0.121	0.000	0.114	-0.114						
Average			0.156			-0.082							
5	BBBB	1.000	0.938	-0.062									
	TBBB	1.000	0.750	-0.250									
	BTBB	0.973	0.714	-0.259									
	BBTB	0.902	0.875	-0.027									
	BBBT	0.768	1.000	0.232									
	TTBB	0.598	0.750	0.152									
	TBTB	0.527	0.857	0.330									
	TBBT	0.518	0.667	0.149									
	BTTB	0.482	0.500	-0.018									
	BTBT	0.473	0.500	-0.027									
	BBTT	0.402	0.000	0.402									
	TTTB	0.232	0.167	0.065									
	TTBT	0.098	0.500	-0.402									
	TBTT	0.027	0.200	-0.173									
BTTT	0.000	0.500	-0.500										
TTTT	0.000	0.095	-0.095										
Average			-0.064										
Average			0.129			0.051			0.025			0.068	

Note. For each game length, n , and each stage, t , the table presents the theoretical (y_t) and empirical (\hat{y}_t) mixed strategies (in terms of the probability/frequency of playing L) against all possible histories of play from the informed player (h_{t-1}) in the NR games. The *Average* for each stage-length combination is computed using the weights defined by the frequency of each history in the data. The *Use* variable is an index of the distance between observed and theoretical use of the history, constructed as: $Use_t = |\hat{y}_t - 0.5| - |y_t - 0.5|$.

actions with the state is much higher in ($n = 3, t = 2$) than in ($n = 5, t = 2$).

C.2 Statistical support

Herein we provide statistical evidence supporting that (i) the relative frequency of action L is increasing in each stage of the NR and FR games in the difference between the number of actions B and the number of actions T in the history; (ii) the correlation of the relative frequency of action L with the difference between the number of actions B and the number of actions T in the history is higher in the FR games than in the NR games for all $n > 2$.

To elicit how uninformed subjects react to the observed history of play, we regress the uninformed player’s decisions on the content of the history. To that end, we specify the observed decision to play L in a given stage, \hat{y} , as the observed counter-part of the latent variable y^* which measures the probability with which player 2 plays L . We measure the information available in the history as the difference between the number of observed Bottoms and the number of observed Tops in the current repeated game – variable denoted $Signal$. The first model we estimate thus relies on the latent equation: $y^* = \alpha + \beta Signal + \epsilon$. To account for some possible ordering effects, we also include the last stage decision: $Last = 1$ if B is the last element of the history (most recent decision). The second model we estimate is thus: $y^* = \alpha + \beta Signal + \gamma Last + \epsilon$. Instead of assuming a particular distribution for the error term ϵ , we estimate the coefficients through an OLS regression so that parameters measure the change in the probability to play L induced by a given change in the content of the history; we accordingly apply a robust estimation of standard errors. The results are presented in the Table below.

	NR Games					FR Games				
	All	$n = 2$	$n = 3$	$n = 4$	$n = 5$	All	$n = 2$	$n = 3$	$n = 4$	$n = 5$
N	1320	132	264	396	528	480	48	96	144	192
Model 1										
α	0.27***	0.06	0.28***	0.25***	0.27***	0.23***	0.12	0.32***	0.20***	0.13***
β	0.11***	0.34***	0.14***	0.10***	0.10***	0.14***	0.31***	0.21***	0.15***	0.13***
$R^2_{adj.}$	0.17	0.43	0.16	0.13	0.19	0.44	0.38	0.48	0.48	0.58
Model 2										
α	0.22***	0.06	0.26***	0.23***	0.21***	0.09***	0.12	0.25**	0.10	-0.00
β	0.06***	0.34***	0.09**	0.06***	0.06***	0.01	0.31***	0.02	0.00	0.01
γ	0.32***	—	0.18*	0.23**	0.34***	0.79***	—	0.68***	0.81***	0.88***
$R^2_{adj.}$	0.23	0.43	0.17	0.15	0.27	0.67	0.38	0.59	0.64	0.83

Legend. Significance levels: * 10%, ** 5%, *** 1%.

Note. OLS estimates of the parameters of $y^* = \alpha + \beta Signal(Model2 : +\beta Last) + \epsilon$. Each column presents the results of separate estimations for NR Games (left-hand side of the Table) and FR Games (right-hand side), on pooled data from all repeated games in the first column, and for each length n in the subsequent ones.

Part (i) is deduced from the sign and significance of the β parameter in all estimated equations: for both games and all lengths, an increase in the number of Bottoms (relative to the number of Tops) induces an increase in the probability that player 2 chooses Left.

In order to compare NR and FR games, three dimensions are of interest. First, as long as $n > 2$, the magnitude of the sensitivity of player 2 behavior to the content of the history is higher in FR than in NR ($\hat{\beta}^{FR} > \hat{\beta}^{NR}$). Second, an important difference between the two treatments is the sensitivity of the estimation results to the inclusion of the last stage decision. While in NR the last stage always improves the fit of the model ($R^2_{adj.}$ higher in model 2 than in model 1); in the FR games, as expected, the last observed decision absorbs all variations in the observed history – *i.e.*, the last decision is enough for the uninformed player to aggregate all the information available in the history. This makes sense, given the histories we observe at most contain one decision different from the others, and this always occurs at the beginning of the game. Because the number of Bottoms is more noisy than the last stage decision, and these variations have no effect on player 2 behavior, the fit of the model is drastically improved when the last stage is accounted for. Last, whatever the way the content of the history is measured (*i.e.* whether one focuses on model 1 or model 2) the predictive power on player 2 behavior is higher in FR than in NR in all games s.t. $n > 2$ ($R^2_{adj.}$ is higher in FR than in NR).

D Written instructions for the experiment

The instructions below are the translated instructions for the NR treatment (original instructions in French). The instructions for the other treatments are similar.

Welcome

You're about to participate in an experiment about strategic decisions. Your aim during the experiment will be to earn as much tokens as possible. The total amount of money you will get at the end of the experiment will depend on the number of tokens you have accumulated. A show-up fee of 5 Euros will be added to your earnings from the experiment.

All your answers will be anonymous and you will take decisions using the computer in front of you.

Before starting the experiment you will answer a short questionnaire designed to check that the instructions are well understood. We will answer any question (privately) before starting.

Description of the experiment

At the beginning of the experiment, **fixed groups of 2 participants** will be randomly formed. In each group, one participant plays as "Player 1", the other plays as "Player 2". You will know at the beginning of the experiment whether you are Player 1 or Player 2. Your role will remain the same during the whole experiment and you will play against the same opponent during the whole experiment. You are not able to identify who is your opponent among all the participants, he can be sited anywhere in the room.

The experiment consist in **20 rounds** of play. **Each round contains several decision stages**. Depending on the round, the number of stages is between 1 and 5. At the beginning of each round, you will be informed about the total number of stages to be played in the round. You will also know the stage number at each stage of the round.

In each round, two **payoff matrices** can be drawn, called **Table A** and **Table B**. One Table is randomly drawn at the beginning of each round. The odds of drawing either table are the same: **at the beginning of each round**, the probability that Table A is drawn is equal to 50%, and the probability that Table B is drawn is equal to 50%. The random draws from one round to the other are independent.

At the beginning of each round, Player 1 is informed about the table that has been drawn for this round. By contrast, Player 2 will not know the result of the random draw of the table before the end of the round. As soon as the table has been drawn, and Player 1 has been informed about the draw, stage 1 starts. The table remains the same in all stages until the end of the round.

Description of a round

In each stage of a round, each participant playing as **Player 1** has to choose between one of the rows, "Top" and "Bottom" , and each participant playing as **Player 2** has to choose one of the two columns, "Left" and "Right".

In the tables, the **first** number (in blue) indicates the number of tokens earned by **Player 1** depending on the decision of **Player 1** (in row), the decision of **Player 2** (in column), and the table drawn (Table A or Table B).

	Left	Right
Top	10, 0	0, 10
Bottom	0, 10	0, 10

Table A

	Left	Right
Top	0, 10	0, 10
Bottom	0, 10	10, 0

Table B

The **second** number (in red) indicates the number of tokens earned by **Player 2** depending on the decision of **Player 1** (in row), the decision of **Player 2** (in column), and the table drawn (Table A or Table B).

For instance, if **Player 1** chooses “Top”, **Player 2** chooses “Left”, and the payoff table randomly drawn for the round is Table A, then the payoff of **Player 1** is equal to 10 and the payoff of **Player 2** is equal to 0. If **Player 1** chooses “Bottom”, **Player 2** chooses “Left”, and the payoff table randomly drawn for the round is Table B, then the payoff of **Player 1** is equal to 0 and the payoff of **Player 2** is equal to 10.

The table drawn at the beginning of each round, and used to compute earnings, remains the same during the whole round, but is drawn again at random at the beginning of each new round.

At the end of each stage of a round, a screen appears once both players have made their choice. The screen displays the decisions of both players from all previous stages of the round. Once both players have confirmed they have seen the screen, the next stage in the round starts.

Value of tokens in euros

At the end of each round a screen appears to inform you about the table drawn (in case you did not know it already), the choices of both players at each stage of the round and the **average number of tokens** earned during the round. The average number of tokens earned during the round is computed as the sum of all tokens earned during the round, divided by the total number of stages. For instance, if the round consists in 5 stages and you earned 10 tokens during the first three stages and 0 tokens during the last two stages, your average number of tokens will be: $(10 + 10 + 10 + 0 + 0)/5 = 30/5 = 6$ tokens. One participant will be randomly chosen at the end of the experiment to draw randomly **3 rounds that will be paid**. Each round out the 20 has the same probability to be drawn at the end of the experiment. Your earnings will be computed as the sum of the average number of tokens earned in all three rounds, with the exchange rate **1 euro for 1 token**. For instance, if you earn 18 tokens during the three stages that are drawn to compute payoffs, your monetary earning will be: 18 euros plus the 5 euros show-up fee, *i.e.*; 23 euros.

At any time during the experiment, you will be able to look at feedback information regarding all previous rounds (your own decisions, the decisions of your opponent, the table drawn for the round and the tokens you have earned) by clicking on the button called “history”. All information about previous stages of the current round will appear directly on your decision screen. On Player 1’s screen (but not on the one of Player 2) the table that has been drawn at the beginning of the round will also appear.

If you wish to ask questions, please rise your hand and we will come to answer it privately. You are asked not to speak during the experiment.

Before starting the experiment, you will answer a short questionnaire designed so as to help you check that all instructions are clear enough.

You will then play a trial round, during which the tokens earned will not be recorded. This round lasts 4 stages. During this trial round, you will not play against your opponent. Rather, it is the computer that will take decisions: if you play as Player 1, then the computer will play as Player 2 and will follow the following arbitrary decision rule whatever your own choices: “Left” in stages 1 and 2, and “Right” in stages 3 and 4; if you play as Player 2, then the computer will play as Player 1 and will follow the following decision rule whatever your own choices: “Top” in stages 1 and 2, and “Bottom” in stages 3 and 4.

Good luck !

E Questionnaire for the experiment

The questionnaire below is the translated questionnaire for the NR treatment (original questionnaire in French). The questionnaires for the other treatments are similar.

1. You will face the same opponent throughout the 20 rounds of the experiment.
 True* False
2. Each round contains 5 periods.
 True False*
3. Within each round, the table remains the same from one period to the next, yet it has a 50% chance of changing at the beginning of every new round.
 True* False
4. Suppose you are player 1. In that case, you will have to choose a line in every period of each round while knowing which payoff table (A or B) was randomly selected. Your opponent (player 2), on the other hand, will not be informed about the randomly selected table.
 True* False
5. Suppose you are player 1, and table A was randomly drawn at the beginning of the round. If you choose the “Bottom” line during a given period, then your earnings for this period will be:
 10 points regardless of your opponent’s decision
 0 points regardless of your opponent’s decision*
 10 points if your opponent chooses “Left” and 0 if he/she chooses “Right”
 0 points if your opponent chooses “Left” and 10 if he/she chooses “Right”
6. Suppose you are player 1, and table B was randomly drawn at the beginning of the round. If you choose the “Bottom” line during a given period, then your earnings for this period will be:
 10 points regardless of your opponent’s decision
 0 points regardless of your opponent’s decision
 10 points if your opponent chooses “Left” and 0 if he/she chooses “Right”
 0 points if your opponent chooses “Left” and 10 if he/she chooses “Right”*

7. Suppose you are player 1 and table B was randomly drawn at the beginning of the round. Suppose the round lasts 3 periods. If you play “Top” in the first period, and “Bottom” in the 2nd and 3rd periods, and if your opponent plays “Right” in the first two periods and “Left” in the last period, then your average earnings for this round will be:
- $30/3 = 10$ points
 - $20/3 = 6.66$ points
 - $10/3 = 3.33$ points*
 - $0/3 = 0$ points
8. Suppose you are player 2. In that case, you will have to choose a column in every period of each round without knowing which payoff table (A or B) was randomly selected. Your opponent, on the other hand, will be informed about the randomly selected table.
- True* False
9. Suppose you are player 2 and table A was randomly drawn at the beginning of the round. Suppose the round lasts 5 periods. If you play “Left” in the first period and “Right” in periods 2 to 5, and if your opponent plays “Top” in every period, then your average earnings for this round will be:
- $50/5 = 10$ points
 - $40/5 = 8$ points*
 - $30/5 = 6$ points
 - $20/5 = 4$ points
 - $10/5 = 2$ points
 - $0/5 = 0$ points
10. Suppose that the three rounds randomly chosen to compute your final payment are rounds 5, 9, and 20. Suppose your average earnings were 8 points for rounds 5 and 9, and 3 points for round 20. In this case, your total gains in euros for this experiment will be:
- 24 euros*
 - 11.33 euros
 - 8 euros
 - 3 euros