Learning to play in a stylized (Chinos) game: some preliminary results

Luis Pastor-Abia*, Emilio San-Fabián†, Enrique Louis** and Fernando Vega-Redondo‡

*Departamento de Química-Física
Universidad de Alicante, Alicante, Spain
†Departamento de Química-Física
and Unidad Asociada of the Consejo Superior de Investigaciones Científicas.
Universidad de Alicante, Alicante, Spain
**Departamento Física Aplicada
and Unidad Asociada of the Consejo Superior de Investigaciones Científicas.
Universidad de Alicante, Alicante, Spain
‡Departamento de Fundamentos del Análisis Económico
and Instituto Valenciano de Investigaciones Económicas
Universidad de Alicante, Alicante, Spain
and
Departament d'Economía i Empresa
Universitat Pompeu Fabra, Barcelona, Spain

Abstract.
In a previous paper, we studied the rich strategic and information-revelation considerations arising in simple guessing game which is very popular in some European countries: the so-called Chinos game. Here, we complement that analysis with a preliminary study of the learning dynamics induced when players no longer can assume that others play according to fixed (possibly equilibrium) rules but, instead, have to use past evidence to understand or interpret their opponent's behavior. We show that if all players are flexible learners, they will eventually succeed in playing according to the unique Nash equilibrium of the game. This outcome, however, is also found to be rather fragile: if only one player (say, a "noise agent") is unable or unwilling to learn, this wrecks havoc in the overall population dynamics. In that case, this dynamics never settles and a complex cyclical pattern typically ensues. The implications of this conclusion for the study of information-revelation mechanisms in the real world (e.g. asset markets) is quite apparent but yet to be tackled in full detail.

INTRODUCTION

The *Chinos* is a simple non-cooperative game in which every player hides a number of coins in his hand (which remains private information) and, subsequently and in turn, tries to guess the *total* number of coins held by the whole set of players involved in the game. Because of the sequential order of move, there is the potential for information revelation as the game unfolds. That is, each player's guess is presumably related to the coins he holds and thus may be used by other subsequent movers to shape their own guesses. The tension between the incentives to guessing right and the fear of not leaking too much information is the key strategic issue featured by this game.

In this paper, we propose a simplified version of this game, akin to that which was studied in a former paper[1]. In that paper, our main interest was to explore the dynamic implications of playing according to some simple and natural rules, and contrast this performance with the prediction of play embodied by the Nash equilibrium of the game. Here, our focus turns to a quite different issue, namely, that of learning. That is, we suppose that players use past evidence to learn how to play the game. This, of course, involves learning how to interpret adequately what others do. We show that learning is indeed possible in this context under rather weak informational assumptions, but that it is a phenomenon sharply disturbed when even small noise is present (e.g. a player that resists learning). A suitable understanding of the stylized Chinos game should prove very useful in understanding the issue of information revelation in economic and social contexts (e.g. in financial markets) – see ref. [1] for references and detailed discussion. In this sense, therefore, we believe that these insights may be valuably used in the the study and understanding of many socio-economic contexts.[1]

The rest of the paper is organized as follows. In the next section, we outline the model proposed to study the Chinos game. Next, we discuss some of the regularities we obtained in our former paper through numerical simulations. Finally, in the last section of the paper, we tackle what is our main present concern. We formulate a simple and natural model of learning and discuss the dynamic implication of full, as well as incomplete, learning.

THE MODEL

The *Chinos* game may be analyzed with the equilibrium tools of classical Game Theory. In particular, we could aim at characterizing its *Nash equilibria* (NE). Even though this can be readily done in simple cases (e.g. with up to one coin per agent or/and just two players – see [3] and [4]), the analysis for the general case is quite complex. Thus, to gain insight on the game, we focus on a simplified version of it where players *can repeat* the same guess as their predecessors. The analysis of this version will be used as a basis for our approach to the genuine Chinos game.

Players are indexed by $i = 1, 2, ...N$ according to their order of play. The number of coins ξ_i hid by player i is determined by a uniform random variable chosen on the set $\{0, 1, ..., m\}$. Let x_i denote the guess of player i. Then, given the realized vector of coins $\xi = (\xi_1, ..., \xi_N)$ that lie in the hands of the N players, the corresponding vector of guesses at the unique NE of the induced game is:

$$x_i = (x_{i-1} - \mu_0) + (\mu_1 + \xi_i) \qquad (i = 1, 2, ..., N) \tag{1}$$

where $\mu_0 \equiv Nm/2$ is the expected total number of coins by the whole population, $\mu_1 \equiv (N-1)m/2$ is the the expected total number of coins in the hands of $N-1$ players,

---

[1] A different approach to understanding this game is to undertake controlled experiments where real subjects are set to play under real (monetary) incentives. This is done in [2], where we find experimental confirmation of the strong tendency of players to learn how to play the game as prescribed by its unique Nash equilibrium.

and x_{i-1} is simply identified with μ_0 for $i = 0$. This strategy is formulated under the implicit assumption that m is even; for odd m, analogous strategies can be constructed.

The interpretation of Eq. (1) is appealing: 1) The first term accounts for the information passed to player i by player $i-1$. This term embodies *sufficient* information (at equilibrium!) on the extent to which the guesses of all preceding players deviate from the "average" uniformed guess. In that sense, this rule can be considered as *Markovian* and implies that, by observing x_{i-1}, player i is able to gain *precise* knowledge of the number of coins $\sum_{j<i} \xi_j$ in the hands of all predecessors. 2) The second term $\mu_1 + \xi_i$ reflects what player i can deduce from her exclusive information. For, if she only had that private information, player i would simply add her number of coins ξ_i to the expected number μ_1 held by the other $N-1$ players.

It follows from the above construction that, at her time of choice, player N is sure to guess correctly. In general, the probability P_i that player i guesses correctly must increase as i approaches N. For large $N-i$, the central limit theorem [5] gives $P_i = \sqrt{6/[m(m+2)\pi(N-i)]}$ which is in excellent agreement with numerical results [1].

At this point, it is worth emphasizing the obvious: the rule given by (1) is only an optimal course of action if and only if, each player i can be confident that players $j < i$ follow the rule as well. This feature, which is nothing but the stability condition, has particular implications in the present context. If any given player i harbors any doubt that others may react according to the equilibrium prescriptions, then she may be inclined to use a rule that generalizes (1) in the following two respects: i) all preceding observations are heeded (i.e. the rule is no longer Markovian) ii) past observations are weighed in some discretionary manner. A generalized rule of this kind is as follows:

$$x_i = \sum_{j<i} \lambda_{i,j}(x_j - \mu_0) + (\mu_1 + \xi_i) \tag{2}$$

where $\lambda_{i,j}$ are real numbers that reflect the sensitivity of player i to player j's guess.

We now return to the original *Chinos* game. This game is interesting because it captures the tension between strategic preemption and information leakage in ways that the simplified version cannot do – in the latter version, only the second consideration is present. However the game is so much more complex to analyze that it is difficult to obtain a closed analytical solution for its Nash equilibrium. Therefore, we approach it through numerical simulations, by building upon (2). We further reduce the number of parameters involved by making

$$\lambda_{i,j} = \frac{\tau_i}{\mid x_j - \mu_0 \mid}. \tag{3}$$

Thus, player i adds (or subtracts) τ_i if the guess of j is larger (or smaller) than average.

SIMULATIONS AND RESULTS

Numerical simulations for an arbitrary number of players with up to $m = 3$ coins and strategies defined by Eq. (3) have been carried out. We have also introduced what we

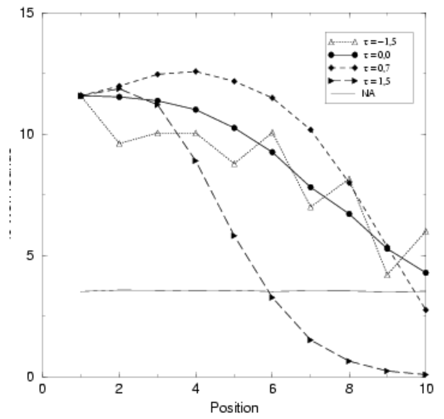

FIGURE 1. Percentage of rounds won by a given player vs. the position of the player in the round. The simulation is carried out as follows. First, a random number of coins with a uniform probability distribution between 0 and m is assigned to each player. Second, each player tries to guess the total number of coins in turn, according to she respective strategy, either Eq. (3), or the noise agent strategy. Third, if the guess chosen by the player is already taken, then she selects the nearest free guess. A lot of rounds (half million) are played in this way, and the average is computed and shown in the figures.

call a *noise agent* (NA), a kind of player who guesses random numbers with complete disregard of previous players guesses. The results are illustrated in figures 1 and 2.

Figure 1 shows the percentage of games won by each player from a set of 10 players. When $\tau = 0$, every player disregards the information from previous players. The percentage of games won by a given player smoothly decreases as the position of the player moves towards the end of the round. This is because, toward the end of the round, good guesses are already taken and the player is forced to guess far from the statistic average. When $\tau = 0.7$, the curve shows a maximum at intermediate positions. The reason is that intermediate players may use the information transmitted by previous players to make a better guess. However, this information again becomes useless towards the end of the round and therefore the percentage of games won drops steadily. For $\tau = 1.5$, the curve is similar, although now the maximum occurs at the second position, the drop after that being very fast and leading to a very low minimum. To understand why, we should remember that in our model, if a player guesses over (under) the average, it is assumed that she carries τ more (less) coins than the average of 1.5 per player. As soon as one player holds a number of coins above (below) the average, and guesses accordingly, she leads all the subsequent players astray, each one assuming that all the previous ones hold τ more coins than average. Not surprisingly, the odds of the players at the end of the round are very low, not because the good guesses are already taken, as was the case for $\tau = 0$ and $\tau = 0.7$, but because the sequence of guesses hopelessly drifts away from the statistic average (*herd behavior* [6]). Finally, we will deal with $\tau = -1.5$, which corresponds to the case of "distrustful" players. The players, accordingly, subtract (add) 1.5 whenever they detect overestimation (underestimation). This results in an oscillatory sequence of guesses as shown in the figure (see [1] for a full discussion).

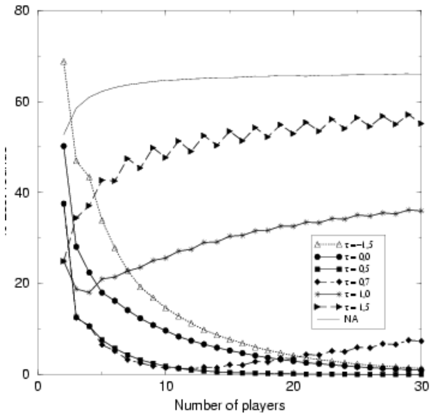

FIGURE 2. Percentage of lost rounds vs. total number of players N for sets of players with uniform $\tau_i = \tau$. The results obtained when all players are noise agents (see text) are also shown (continuous line); in this case the percentage of lost rounds can be calculated analytically, the result being $P_{\mathrm{lost}} = 100(N-1)/((m/(m-1))N-1)$, where m is the maximum number of coins hidden by each player.

Figure 2 shows the percentage of rounds with no correct guess, as a function of the total number of players (all with $\tau_i = \tau$). Aiming to provide a benchmark, the case where all players are NA is also shown. The curves in this figure can be understood after the previous discussion. Basically, they can be divided into two classes: either the percentage of lost games decreases as the number of players grows large or it increases. The first class ($\tau = 0$, $\tau = 0.5$, or $\tau = -1.5$) can be explained by the absence of herd behavior: players' guesses are reasonably close to the average, so that sooner or later some player ends up guessing the total number of coins, the greater the number of players the more likely. The shape of the second kind of curves ($\tau = 1.0$ or $\tau = 1.5$) is caused by herd behavior: as soon as a player gives a guess far enough from average, all subsequent players are led astray and the round will finish with no one guessing the right number of coins. Only the first players will have a fair shot at guessing correctly before herd behavior sets. The mildly oscillatory pattern observed in these cases has a trivial origin, as discussed in [1]. The curve with $\tau = 0.7$ shows a hybrid behavior, first decreasing to almost zero, then increasing. Combinations of trustful and distrustful players reduce herd behavior, and is more successful than a pure set of distrustful players [1].

We have investigated how the presence of a *single* NA distorts these results. Of course, the strongest effect is found when she plays in first place. The results indicate that trustful strategies are very sensitive to perturbations, while distrustful playing is more robust. Of course the less sensitive strategy is that which does not take into account the information from previous players ($\tau = 0$). Particularly remarkable is the early appearance of *herd behavior* when players follow trustful strategies ($\tau > 0$). In [1] a new kind of player, *the normalized agent*, was introduced for which $\tau_i = \tau/\sqrt{i-1}$. The proposal is based upon what the central limit theorem suggests. The results indicate that normalized agents not only do not show herd behavior but also give excellent overall performance.

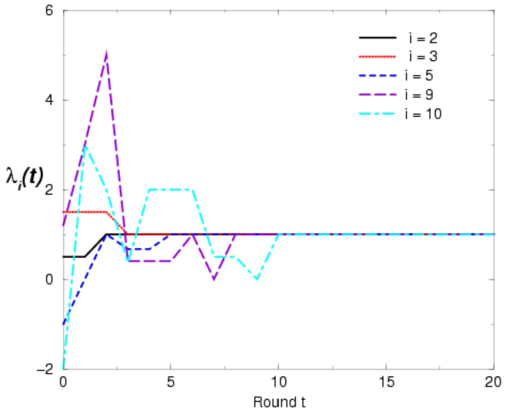

FIGURE 3. Evolution of parameter λ_i for various players in successive rounds of the simplified version of the Chinos game. Simulations were carried out with players starting the game (round 1) with randomly chossen λ_i and actualizing its value according to (4)

A FIRST APPROACH TO LEARNING

Learning is a very interesting aspect of game theory. Here we briefly discuss the results of a preliminary study of learning in the simplified version of the Chinos game that allows repetition of guesses. After a given round each player is forced to reveal the number of coins he hides. Then, before starting round $t+1$, all players are fully informed of the guesses and hid coins of the other players in the preceding round t, namely, they all know the set $[x_i(t), \xi_i(t)]$ for $i = 1,..,N$. Consider the simple case in which player i only takes note of the guess of the preceding player $i-1$, that is, the parameters $\lambda_{i,j}$ in (3) are taken equal to λ_i for $j = i-1$ and zero otherwise, and that all players start the game with λ_i chosen randomly (*i.e.*, away from NE). The question is whether there is a rule that all players can follow such that equilibrium is eventually achieved. A natural candidate for the dynamics of λ_i is inspired in the NE equilibrium strategies described in Eq. (1):

$$\lambda_i(t+1) = \frac{\sum_{j<i} \xi_j(t) - \mu_1 + \mu_i}{x_{i-1}(t) - \mu_0} \tag{4}$$

Numerical results for a case in which all players follow the dynamics prescribed by Eq. (4) are shown in Figure 3. It is clear that, after an oscillatory period, the above rule leads to Nash equilibrium play (i.e. $\lambda_i = 1$ for all i). This can be intuitively understood as follows. First, player always guesses according to NE. Then, the second player displays a “reacting factor” λ_2 that initially oscillates, but, as only the first player affects the way he plays, sooner or later he will play as in Nash equilibrium with $\lambda_2 = 1$. Once player 2 has arrived to that state of affairs, the third player will follow a similar path and eventually set on a value for λ_3 equal to one, and so on and so forth. Although the simulations of Figure 3 were carried out with all players in the same positions, we have checked that interchanging the positions in successive rounds only affects the time needed to restore equilibrium. This illustrates a certain robust aspect of Nash equilibrium.

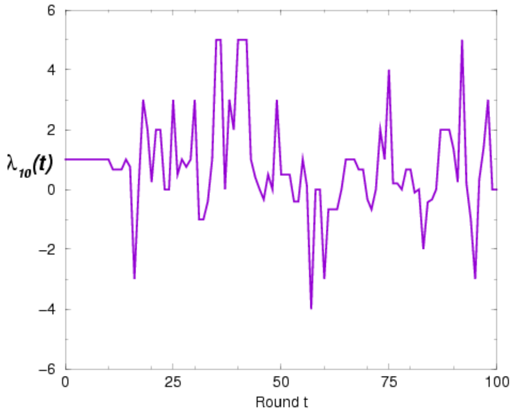

FIGURE 4. Parameter λ_i for player in position 10. In this case the second player had a constant $\lambda_2 = 3$, whereas the rest followed the dynamics of (4).

However, as shown in Figure 4, a single player who does not learn and always behaves according to a constant $\lambda \neq 1$ is enough to break the dynamics of convergence towards Nash equilibrium. Now the value of λ for the player in the tenth position oscillates around one and never stabilizes. The amplitude of the oscillations for player *i* increases with the distance to the "rigid" player. This indicates that a rule as simple as that of (1) is not enough for restoring (approximate) equilibrium in a general case where only few players fail to learn. It is likely that if one were to relax the Markovian rule implicit in (4) and allow players to take note of what *all* players (not only the predecessor) have done in the previous round, convergence towards equilibrium might be restored. Research on this and other related issues is currently in progress.

REFERENCES

1. Pastor-Abia, L., Pérez-Jordá, J., San-Fabián, E., Louis, E., and Vega-Redondo, F., *Advances in Complex Systems*, **4**, 177–190 (2001).
2. Feri, F., Meléndez-Jiménez, M., Ponti, G., and Vega-Redondo, F., *Universidad de Alicante* (2002).
3. Weibull, J., *unpublished*, (2000).
4. Guinea, F., and Martin-Delgado, M., *Quantum-physics/0201140*, (2001).
5. Sobol, I., *A primer for the Monte Carlo Method*, CRC Press, Boston, 1994.
6. Banerjee, A., *The Quarterly Journal of Economics*, **107**, 797–818 (1992).

ACKNOWLEDGMENTS

This work was supported in part by the Spanish CICYT (grants PB96-0085, BQU2001-0883 and BEC2001-0980) and by the European Commission through projects "TMR Network-Fractals c.n. FMRXCT980183" and "Fet Open project COSIN IST-2001-33555".