THREE ESSAYS ON MIXED STRATEGIES

By

Christian Diego Alcocer Argüello

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Economics - Doctor of Philosophy

2016

ABSTRACT

THREE ESSAYS ON MIXED STRATEGIES

By

Christian Diego Alcocer Argüello

I focus on the theoretical and empirical consequences of different types of bounded rationality assumptions and varying levels of economic rationality. The basic nature of mixed equilibria is intuitively plausible since it is often easy to recognize the advantages of being unpredictable: in sports, a pitcher tries to randomize his throws while the batter tries to outguess him; in crime and terrorism prevention, the time and place of attack are generally uncertain. And yet, we are sometimes at a loss of finding intuitive explanations for the specific probabilities in the prescribed equilibrium.

I start by postulating a new behavioral bias by exploring what happens when some players use simple rules of thumb to choose among pure strategies that they perceive to lead to the highest payoff. *Naïve players* were conceived as the simplest possible players that did not directly contradict the fundamental assumption of utility maximization. Then I find experimental evidence that some players are indeed affected by this proposed behavioral bias. Finally, I find evidence that there exist experimental conditions that increase this bias at the individual level. These findings contribute to our general understanding of the determinants of economic choices and the (actual) nature of economic rationality.

To illustrate the theory, one can imagine choosing between left or right when expected utility is constant. NE forecasts the choice depends on the payoffs of the other players. I relax typical assumptions on rationality and propose some players will flip a coin when indifferent. They can see payoffs, but they lack any strategic depth. I then prove that if their proportion is small enough, any NE of a game with no naïve players corresponds to an equilibrium of the generalized game with some naïve players where payoffs for all players are the same. The intuition is that the rest of the population compensates in the opposite direction of the distortion. Thus naïve mixers are not disadvantaged by using a behavioral rule of thumb. Nor are rational players in a position that allows them to increase their payoffs: there is no rent to be gained by this strategy restriction and no welfare loss to be fixed by a social planner.

In my second dissertation chapter, I experimentally confirm the existence of some players who consistently mix close to 50% in different settings. I first sort participants into naïve and non-naïve by letting them play variations of asymmetric matching pennies. Two weeks later, each group plays against varying proportions of automated computer players (bots) that follow changing off-equilibrium strategies, and I observe several measures of how they react to the distortions. Besides identification of naïve players, I show that the probability of being naïve can be modeled by a quantitative test: smart players today, play smart tomorrow, and smart players don't always mix uniformly. I also find evidence that the non-naïve population reacts better to off-equilibrium behavior, plays closer to NE and adjusts their behavior in the correct direction but not with the magnitude required to restore equilibria. I then employ this last result to design simple mechanisms to obtain above-equilibrium payoffs by taking advantage of naïve players.

In my final paper, I test the determinants that make an individual more (or less) naïve under different conditions related to game complexity and game stakes. I find evidence that players sometimes behave relatively close to coin flipping under a *distractor* which consists of adding weakly dominated strategies to matching pennies games. When they face a *focuser* which consists of monotonically increasing payoffs, players behave relatively close to the NE. This is further evidence of methods that enable players to attain above-equilibrium payoffs by taking advantage of this behavioral bias.

Some results from the first experimental paper were especially puzzling, and it was conjectured that altruistic components might be distorting the results. Using computer bots, I isolate philanthropic components of players' strategies. Adding a proportion of transparent bots that do not incentive any change in behavior but imply that surplus is wasted if they get any payoff, behavior gets closer to NE. When there is some chance players are matched against a bot, altruistic effects (utility gained by total surplus maximizing, even if it is another player who gets it) decrease and behavior moves towards what is predicted by utility maximization. This implies that, as a researcher, one way to verify if altruism is preventing players from reaching the NE is to compare behavior with and without bots that play the NE.

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to my dissertation chair, Thomas D. Jeitschko for believing in my research ideas since their early stages, and for his kind patience during our weekly meetings throughout the last four years of my life at Michigan State University. His influence, more than anybody else's, has determined the extent of my abilities as a researcher.

I also wish to thank my advisor Robert Shupp, without whose support this dissertation would not have existed. His technical skills which he was generous to share with me, his eye to discover the slightest mistake and his ability to and identify and solve big-picture problems reshaped my research countless times.

This dissertation also greatly benefitted from suggestions by my other committee members, Jeffrey M. Wooldridge and Jon Eguia, whose comments and lessons helped me complete this dissertation, as well as enhance its quality.

Mitchell Strahlman's help during the experimental sessions was invaluable. James Chen was an inspiration during the second half of my Ph. D. Thanks to professors Ting Liu, David Mayer-Foulkes, Jack Meyer and Arijit Mukherjee for their teachings, mentoring and advice.

I owe an enormous debt to Asja Asenova. Without her unconditional support, perhaps I would not have succeeded in this academic endeavor—this crazy adventure—that was my Ph. D.

Finally, none of this would have happened if I had not been lucky enough grow up with a loving family like mine. Thank you, dad, Cristy Alcocer, and Sandra Navarrete Argüello. I am sure you are proudly watching me too, mom.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1, Using Rules-of-Thumb:

A Note on Sophisticated vs. Simple Mixing in Two-Player Randomly Matched Games

Christian Diego Alcocer Argüello*

August 15, 2016

Abstract

We postulate a new behavioral bias in how people play mixed strategies by proposing the existence of *simple* players who lack strategic depth. We define them as those who, when indifferent between choices, follow a simple *rule-of-thumb* and assign a predetermined probability to each. We show that if they play 2×2 games, an equilibrium generally fails to exist.

However, under random matching within populations with some proportion of *simple* players, equilibrium is restored and is indistinguishable from Nash equilibria in games with unrestricted strategy choices, as long as the percentage of simple mixers is small enough. As such, players are unable to take advantage of the presence of *simple* mixers, and *simple* mixers do no worse than more sophisticated players.

Keywords:

JEL: C72, D03, D83

1 Introduction

Many games require that players mix among pure strategies in order to attain an equilibrium. Frequently such mixed strategy equilibrium configurations are intuitively plausible: children recognize the advantage of being unpredictable when playing games such rock-paper-scissors, and some descriptions of the 'matching

*Department of Economics, Michigan State University, East Lansing, MI 48824, alcocer@msu.edu.

pennies' game actually have players flip pennies (rather than choosing heads or tails) in order to determine payouts.

The logic of wanting to be unpredictable and wanting to keep one's rival(s) guessing about what to expect readily extends to more complex games. And yet, in games in which equilibrium requires players to mix, we are sometimes at a loss of finding intuitive explanations for the prescribed equilibrium; particularly since it may be difficult to determine the optimal mixing distribution. This task is complicated by the fact that, in equilibrium, the player is *ex definitione* indifferent between choosing any particular pure strategy that would be played with positive probability.

Our analysis is motivated by the apparent need for modeling of off-equilibrium behavior. In most of the experimental literature, players' behavior is significantly different than the strategies predicted by game theory; in real-life scenarios (Misirlisoy and Haggard, 2014) and particularly in normal-form games in experiments (Nyarko and Schotter, 2002; Stahl, 1995). Masiliunas et al. (2004) find that only 7% of choices are consistent with Nash equilibrium and cannot reject the hypothesis that in complex games, players' choices are uniformly distributed. Failure of NE holds even when the experimenters privately recommend optimal strategies (Cason and Sharma, 2007) or when participants are allowed to explicitly input mixed-strategy probabilities that are automatically implemented instead of having to go through the extra cognitive step of executing them manually (Bloomfield, 1994). Relatedly, Parkhurst et al. (2015) find that when individuals face information overload, they tend to use simplifying heuristics.

The characterization of behind totally-mixed Nash equilibria can be approached heuristically as a two-step decision process. First, a player maximizes his utility and determines, given his beliefs regarding his opponents' strategies, that he has no unique best response. He then proceeds to a more strategically sophisticated deliberation that determines what sort of mixing ensures the other players are indifferent as well and prevents them from having a unique best response.

In this paper, we explore what happens when some players, rather than ascertaining equilibrium mixing distributions, neglect the second step described above and shoot from the hip by using simple rules of thumb to choose among the pure strategies that lead to the highest payoff. We refer to such players as *simple mixers*. In contrast, a *sophisticated mixer* is a player who can freely determine and employ any mixed strategies.

We show that in 2×2 games, unless the rule of thumb that simple mixers are guided by happens to coincide with the equilibrium mixed strategy used by sophisticated players (so that their play is indistinguishable), an equilibrium fails to exist. However, under anonymous random matching within populations with *simple* and *sophisticated* mixers, when only the fraction of each population is common knowledge, then equilibrium is restored if the probability that a given player is a *simple* mixer is below an upper bound. This equilibrium is attained by sophisticated mixers adjusting their strategies to account for the distortions induced by the presence and behavior of simple mixers in the population. As a result, all players—sophisticated and simple alike—obtain the same payoffs that prevail in the equilibrium of the standard game in which there are no simple mixers. An implication of this is that sophisticated players are unable to take advantage of the presence of simple mixers who do no worse. In a sense, simple mixers free-ride off the sophistication of the other players.

2 Literature Review

The issue of complexity in the determination of equilibria has been extensively analyzed in the computational literature. Lipton et al. (2003) discuss how for two-player games, the known algorithms have polynomial or worse (exponential) running time in the number pure strategies. This keeps them in some class between P and NP since it is *easy* (i.e. in polynomial time) to verify if a given strategy profile is an equilibrium. Moreover, the complexity of solving a game can increase considerably, depending on the form in which it is represented (Koller et al., 1994) since the relation of the number of strategies to the size of the tree form is exponential. Furthermore, "the increase in size of transforming a game described by a set of rules into a matrix-form game is not even bounded in general, which often forces analysts to find solution techniques tailored to specific games" (Koller et al., 1994). Because of this, Lipton et al. (2003), for instance, consider *normal simple strategies*, strategies which are uniform on a small generic support set. They justify this way of thinking by considering pure strategies as resources. As such, an equilibrium is called "impractical" if a player has to randomize over an extensive set of resources.

Ever since Nash (1951), there have been attempts to extend the notion of Nash equilibrium (NE) to deal better with empirical data. Purification theory[1] (Harsanyi, 1973) and the concept of a quantal response equilibrium (QRE)[2] describe how mixed-strategy equilibria can be interpreted as the limit of pure-strategy

[1]Formally, in the most general terms, if player i's payoff given the (pure) strategy profile s is $u_i(s)$, then what he actually observes is $U_i(s;\theta_i) := u_i(s) + \varepsilon_i^s \theta_i^s$ where $\varepsilon_i^s \theta_i^s$ is the Gaussian i.i.d. perturbation scaled by his type and $E(\varepsilon_i^s) = 0$ (see Fudenberg and Tirole, 1991). The vector θ_i can be interpreted as related to levels of rationality: if all its components are zero, the player (perfectly) observes an unperturbed game. If $Var(\varepsilon_i^s) \to \infty$ or $\theta_i^s \to \infty$, he observes only noise and mixes uniformly across all–including dominated–strategies.

[2]In this case, if player i's payoff given the (pure) strategy profile s is $u_i(s)$, then what he actually observes is $U_i(s) := u_i(s) + \varepsilon_i^s$ where ε_i^s is the i.i.d. perturbation $E(\varepsilon_i^s) = 0$, (McKelvey and Palfrey, 1995; McKelvey and Palfrey, 1998; Goeree, Holt and Palfrey,

equilibria in games when perturbations to all payoffs are assumed. Ex-ante, all pure strategies receive a positive probability and probabilities are greater for those strategies that yield greater expected payoffs. A QRE can be approximated and estimated as a logit equilibrium (McFadden, 1973). [3]

Experimental evidence suggests QRE is sometimes a better predictor than NE and purification (Erlei and Schenk-Mathes, 2012). Under a given set of assumptions, a QRE can be approximated and estimated as a logit equilibrium (McFadden, 1973). Cognitive hierarchy theory or *level-k* thinking (Camerer et al., 2004; Van Damme, 1991) assumes that there are multiple levels of players with varying degrees of rationality. There are non-strategic *level*-0 players who always mix uniformly among all their pure strategies.[4] *Level-k* players best-respond to *level*-(*k*-1) populations.

As such, this paper attempts to enhance the literature on off-equilibrium play. The model below is conceptually different from purification and QRE where if the variances of the perturbations tend to zero, all strategies tend to the NE whereas a simple player will never mix with a non-degenerate probability different than $\breve{\sigma}_\rho$ nor will he play a strongly-dominated strategy. Also, in equilibrium the relative likelihood of outcomes under purification coincide with the relative likelihoods in the NE. Finally, a *level*-0 is distinct from a naïve player since the later only mixes uniformly amongst a subset of pure strategies.

3 Model

We establish the main findings by modifying a simple 2×2 game, but it will be readily apparent that the insights carry over to complex multi-player games straightforwardly. Consider a generic 2×2 payoff structure Γ such as that depicted in Figure 1. Superscripts denote players and subscripts denote pure strategies, with $u^i_{j,k}$ giving Player *i*'s payoff when Player 1 plays s^1_j and Player 2 plays s^2_k. We let $\sigma^i_1 := \sigma^i$ indicate the probability with which Player *i* chooses s^i_1, so the probability of picking s^i_2 is $\sigma^i_2 := \left(1-\sigma^i\right)$.

While maintaining the structure of the game regarding pure strategies $\left(s^i_m\right)$, we extend the analysis to include behavioral players–namely *simple mixers* that are restricted in their choice of mixing among their

2002).

[3] For a given $\lambda_i \geq 0$, the logistic quantal response is defined as $\sigma^j_i(s^j_i \mid \sigma_{-i}) = exp\left\{\lambda_i U^j_{i,-i}\right\} / \sum_{\kappa=1}^{K} exp\left\{\lambda_i U^\kappa_{i,-i}\right\}$ which is the probability that *i* assigns to action s^j_i (from some finite action set of cardinality K) given the other players' strategies, and $U^j_{i,-i}$ is the perturbed payoff he gets after s^j_i and σ_{-i}. The parameter λ_i can be interpreted as a measure of *i*'s rationality. As $\lambda \to \infty$, if there exists a unique NE (σ^*), as is the case throughout the games described below, then $(\sigma_1, ..., \sigma_N)$ converges to σ^*. As $\lambda \to 0$, players become blind and mix uniformly across all strategies such that the probabilities tend to $\frac{1}{K}$.

[4] This is one of the simplest possible behaviors in terms of computational complexity (Koller and Megiddo, 1996; Koller, Megiddo and Von Stengel, 1994; Lipton et al., 2003).

Player 2

		s_1^2	s_2^2	
Player 1	s_1^1	$u_{1,1}^1, u_{1,1}^2$	$u_{1,2}^1, u_{1,2}^2$	σ^1
	s_2^1	$u_{2,1}^1, u_{2,1}^2$	$u_{2,2}^1, u_{2,2}^2$	$1-\sigma^1$
		σ^2	$1-\sigma^2$	

Figure I.1: Generic 2 $\times$2 Game Γ

pure strategies. Instead of being allowed to assign any probability σ_m^i to any pure strategy s_m^i (as long as $\sigma_m^i \in [0,1]$ and $\sum\limits_{m \in M} \sigma_m^i \equiv 1$ where M is the set of available pure strategies), simple mixers can only utilize some predetermined mix $\left(\check{\sigma}_\rho\right)$ or pure strategies. We index such "rule-of-thumb" players by ρ and denote their behavior by $\sigma_\rho \in \left\{0, \check{\sigma}_\rho, 1\right\}$.

4 Analysis of Simple Play

Assume that it is common knowledge that either is a simple mixer and that Γ has a unique equilibrium.

Lemma 1 *Let* $\xi := \left(\sigma^{1*}, \sigma^{2*}\right)$ *denote an equilibrium in the generic* 2×2 *game with no simple players. Then,*

1. *the presence of a simple mixer i does not affect equilibrium behavior, if* $\sigma^{i*} \in \left\{0, \check{\sigma}_\rho, 1\right\}$;

2. *and if the equilibrium* ξ *is unique and* $\sigma^{i*} \notin \left\{0, \check{\sigma}_\rho, 1\right\}$*, then no equilibrium exists.*

Lemma 1 follows directly from the definitions of equilibria and simple mixing. First, if $\sigma^{i*} \in \left\{0, \check{\sigma}_\rho, 1\right\}$, the presence of a simple mixer does not affect equilibrium behavior since no player has any incentive to deviate from $\left(\sigma^{1*}, \sigma^{2*}\right)$ as expected payoffs are unchanged.

Second, to see why equilibria can be lost, suppose that Player i is a simple mixer where $i \in \{1,2\}$ and $\left(\tilde{\sigma}^1, \tilde{\sigma}^2\right)$ is an equilibrium such that $\tilde{\sigma}^i \notin \left\{0, \check{\sigma}_\rho, 1\right\}$. Since the equilibrium is unique and Player i is not allowed to play it given the definition of which mixed strategies are available to simple mixers, the game has no equilibrium.

As an illustration, consider the asymmetric matching pennies example, with $u_{1,1}^1 > 1$. The unique NE is $\sigma^{1*} = 1/2$ and $\sigma^{2*} = \frac{1}{1+u_{1,1}^1}$. If Player 2 is a simple mixer that follows the rule of thumb $\check{\sigma}_\rho = 1/2$,[5] then

[5]It will be readily apparent that, in more complex games, the main results hold with any predetermined mixing distribution, as long as it is public knowledge. Furthermore, behavioral, experimental and insufficient-reason arguments can be made to justify why uniform mixing ($\check{\sigma}_\rho = \frac{1}{2}$ in the 2×2 case) is a good starting point.

		Player 2		
		s_1^2	s_2^2	
Player 1	s_1^1	$u^1_{1,1}, 0$	$0, 1$	σ^1
	s_2^1	$0, 1$	$1, 0$	$1-\sigma^1$
		σ^2	$1-\sigma^2$	

Figure I.2: Asymmetric Matching Pennies

the game has no equilibrium, as depicted in Figure 3 that depicts the loss of convexity in i's strategy space.

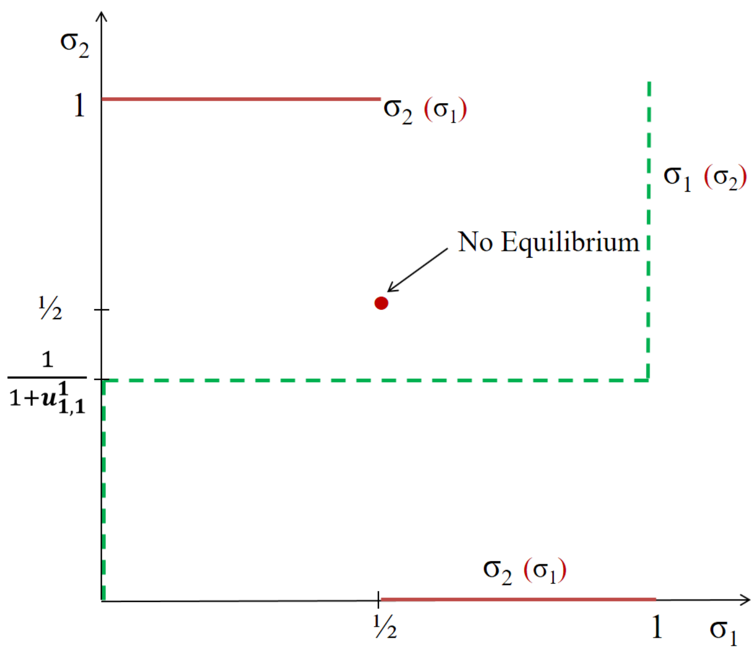

Figure I.3: Non-Existence of Equilibrium

5 Randomly Matched Play

Previously we assumed that all players were simple mixers. We now generalize the game by allowing proportions of both simple and sophisticated players. This *reconvexifies* strategy spaces. We will use inverted hats to distinguish the generalized game that includes simple and sophisticated mixers from the generic game that only includes sophisticated players (Γ) so $\check{u}^{i*}$ and $\check{\sigma}^{i*}$ will denote the equilibrium payoff and mixing of Player i in the generalized game.

Consider a large population of players who are randomly matched. For ease of exposition, we will focus on the case where the equilibrium is unique, Player 2 can be either simple or sophisticated, and

Player 1 is sophisticated with probability one. Below, we discuss how the multiple-equilibria analysis is not fundamentally different. Also, the case where Player 1 can be simple is symmetric to the proposition below. Let $\lambda \in [0,1]$ denote the fraction of simple mixers in the Player 2 population, and assume that this fraction is common knowledge. A simple player's mixing is restricted to the probabilities in $\{0, \breve{\sigma}_\rho, 1\}$.

Proposition 2 *Stability Theorem: For every equilibrium $\mathscr{E}$ of the generic 2×2 game Γ, there exists $\bar{\lambda}_\Gamma$ $\in (0,1]$ such that whenever $\lambda \in \left[0, \bar{\lambda}_\Gamma\right]$:*

1. There exists at least one equilibrium in the generalized game with simple and sophisticated players that is outcome-equivalent to the equilibrium ξ in that

   (a) Player 1 and Player 2, regardless of type, obtain the same payoffs as in ξ

   (b) The frequency in which pure strategies are chosen is the same.

2. Whenever $\sigma^{2*} \notin \{0, \breve{\sigma}_\rho, 1\}$, sophisticated Player 2's equilibrium strategy $\left(\breve{\sigma}^{2*}_{-\rho}\right)$ diverges from his strategy in $\mathscr{E}$ $\left(\sigma^{2*}\right)$ in the opposite direction of the distortion $\left(\breve{\sigma}_\rho\right)$:

$$\text{if } \breve{\sigma}_\rho < \sigma^{2*} \rightarrow \qquad \breve{\sigma}^{2*}_{-\rho} > \sigma^{2*} \text{ and}$$
$$\text{if } \breve{\sigma}_\rho > \sigma^{2*} \rightarrow \qquad \breve{\sigma}^{2*}_{-\rho} < \sigma^{2*}.$$

Proof. The proposition states that given any arbitrary 2×2 game,

$$\exists \lambda_\Gamma > 0 \mid \forall \lambda \in [0, \lambda_\Gamma]:$$
$$u^{1*} = \breve{u}^{1*}, u^{2*} = \breve{u}^{2*},$$
$$\sigma^{1*} = \breve{\sigma}^{1*}, \sigma^{2*} = \breve{\sigma}^{2*}.$$

We proceed by proposing an equilibrium $\left(\breve{\mathscr{E}}\right)$ that is in line with the restrictions imposed on simple players. Note that if the populations' expected strategies implied by $\breve{\mathscr{E}}$ coincide with those of $\mathscr{E}$, then payoff-equivalence follows directly. That is, for $i = 1,2: \sigma^{i*} = \breve{\sigma}^{i*}$ implies $u^{i*} = \breve{u}^{i*}$ as $\breve{u}^{i*}$ is just an expected value that depends on the probabilities $\breve{\sigma}^{i*}$ and the payoffs $u^i_{j,k}$.

First, suppose $\sigma^{2*} = \breve{\sigma}^{2*}$. If Player 1 is playing a pure strategy in $\mathscr{E}$ such that $\sigma^{1*} \in \{0,1\}$, then he will have no incentive to deviate. Similarly, if his equilibrium strategy is mixed, Player 1's equilibrium strategy $\breve{\sigma}^{1*}$ is determined by solving $\breve{\sigma}^{1*} u^2_{1,1} + \left(1 - \breve{\sigma}^{1*}\right) u^2_{2,1} = \breve{\sigma}^{1*} u^2_{1,2} + \left(1 - \breve{\sigma}^{1*}\right) u^2_{2,2}$, a function of $u^2_{j,k}$ which is the same condition for σ^{1*}. Thus, whether Player 1 was playing a pure or a mixed strategy, $\breve{\sigma}^{1*} = \sigma^{1*}$ holds as long as $\sigma^{2*} = \breve{\sigma}^{2*}$.

Now, suppose $\sigma^{1*} = \breve{\sigma}^{1*}$. If in the equilibrium ξ, Player 2 is playing a pure strategy then, as noted in the previous section, the presence of simple mixers does not affect his equilibrium behavior. Now suppose $\sigma^{2*} \in (0,1)$. The simple Player 2, being indifferent between his two actions (s_1^2, s_2^2), has $\breve{\sigma}_\rho^{2*} = \breve{\sigma}_\rho$ as a best response. Finally, the convolution that defines average (expected) equilibrium mixing in the restricted game of the sophisticated Player 2 ($\breve{\sigma}_{-\rho}^{2*}$) is determined by solving

$$\begin{aligned} \left[\breve{\sigma}_\rho \lambda + (1-\lambda)\,\breve{\sigma}_{-\rho}^{2*}\right] u_{1,1}^1 + \left[1 - \breve{\sigma}_\rho \lambda - (1-\lambda)\,\breve{\sigma}_{-\rho}^{2*}\right] u_{1,2}^1 &= \\ \left[\breve{\sigma}_\rho \lambda + (1-\lambda)\,\breve{\sigma}_{-\rho}^{2*}\right] u_{2,1}^1 + \left[1 - \breve{\sigma}_\rho \lambda - (1-\ \lambda)\,\breve{\sigma}_{-\rho}^{2*}\right] u_{2,2}^1 & \end{aligned}$$

with solution

$$\begin{aligned} \breve{\sigma}_{-\rho}^{2*} &= \frac{\breve{\sigma}_\rho \lambda u_{1,1}^1 + u_{1,2}^1 - \breve{\sigma}_\rho \lambda u_{1,2}^1 - \breve{\sigma}_\rho \lambda u_{2,1}^1 - u_{2,2}^1 + \breve{\sigma}_\rho \lambda u_{2,2}^1}{\lambda u_{1,1}^1 - u_{1,1}^1 + u_{1,2}^1 - \lambda u_{1,2}^1 + u_{2,1}^1 - \lambda u_{2,1}^1 - u_{2,2}^1 + \lambda u_{2,2}^1} \\ &= \frac{1}{1-\lambda}\left(\frac{u_{2,2}^1 - u_{1,2}^1}{u_{1,1}^1 - u_{1,2}^1 - u_{2,1}^1 + u_{2,2}^1} - \lambda \breve{\sigma}_\rho\right) \end{aligned}$$

which is decreasing in $\breve{\sigma}_\rho$. Note that $u_{1,1}^1 - u_{2,1}^1 \neq u_{1,2}^1 - u_{2,2}^1$ since $\sigma^{2*} \in (0,1)$ is unique by assumption.[6]

The outcome-equivalence stated by the Stability Theorem is verified by calculating the weighted average of the mixing of both populations, which coincides with Player 2's mixing in $\mathscr{E}$:

$$\begin{aligned} \breve{\sigma}^{2*} &= \\ \lambda \breve{\sigma}_\rho^{2*} + (1-\lambda)\breve{\sigma}_{-\rho}^{2*} &= \\ \frac{u_{2,2}^1 - u_{1,2}^1}{u_{1,1}^1 - u_{1,2}^1 - u_{2,1}^1 + u_{2,2}^1} &= \sigma^{2*}. \end{aligned}$$

Since $\breve{\sigma}^{1*} = \sigma^{1*}$ and $\breve{\sigma}^{2*} = \sigma^{2*}$, then $u^{1*} = \breve{u}^{1*}$ and $u^{2*} = \breve{u}^{2*}$ (for both the simple and sophisticated Player 2 populations).

[6]If $u_{1,1}^1 - u_{2,1}^1 = u_{1,2}^1 - u_{2,2}^1$, Player 1 would have a strictly dominant strategy. If, for instance, both sides of the equation are positive, then s_1^1 strictly dominates s_2^1. In turn, this would imply that if $\sigma^{2*} \in (0,1)$, then $u_{1,1}^2 = u_{1,2}^2$ so σ^{2*} is not unique. We assumed uniqueness, but it is clear that payoff invariance holds if uniqueness fails as in the case when $u_{1,1}^2 = u_{1,2}^2$. The frequency in which pure strategies are chosen is also the same since, as in the rest of this analysis, the sophisticated Player 2 has no incentive to deviate from $\breve{\sigma}_{-\rho}^{2*} = \frac{\sigma^{2*} - \lambda \breve{\sigma}_\rho^{2*}}{1-\lambda}$.

The value of the threshold λ_Γ will depend on whether $\sigma^{2*} \gtrless \check{\sigma}_\rho$. Since probabilities are bounded (in particular $\check{\sigma}^{2*}_{-\rho} \in [0,1]$), the compensation done by the sophisticated Player 2's when $\sigma^{2*} < \check{\sigma}_\rho$ can work up to the point where $\check{\sigma}^{2*}_{-\rho} = 0$ so in this case λ_Γ is the value of λ such that $\check{\sigma}^{2*}_{-\rho} = 0 : \lambda_\Gamma = \frac{1}{\check{\sigma}_\rho} \frac{u^1_{2,2} - u^1_{1,2}}{u^1_{1,1} - u^1_{1,2} - u^1_{2,1} + u^1_{2,2}}$. If $\sigma^{2*} > \check{\sigma}_\rho$, then the sophisticated Player 2 has to compensate in the opposite direction and λ_Γ is the solution for λ such that $\check{\sigma}^{2*}_{-\rho} = 1 : \lambda_\Gamma = \frac{1}{\check{\sigma}_\rho} \frac{u^1_{1,1} - u^1_{2,1}}{u^1_{1,1} - u^1_{1,2} - u^1_{2,1} + u^1_{2,2}}$. Finally, if $\sigma^{2*} = \check{\sigma}_\rho$, to have an equilibrium there is no need to have sophisticated players at all, and $\lambda_\Gamma = 1$. Therefore in the restricted game one has: $\lambda \leq \lambda_\Gamma = \min\left\{\frac{1-\sigma^{2*}}{1-\check{\sigma}_\rho}, \frac{\sigma^{2*}}{\check{\sigma}_\rho}\right\}$. ■

The intuition is that since $\sigma^{2*} = \lambda \check{\sigma}^{2*}_\rho + (1-\lambda)\check{\sigma}^{2*}_{-\rho}$, a sophisticated Player 2 has no incentive to deviate from

$$\check{\sigma}^{2*}_{-\rho} = \frac{\sigma^{2*} - \lambda \check{\sigma}^{2*}_\rho}{1-\lambda} = \frac{\sigma^{2*} - Pr(Player\ 2\ is\ Simple)\check{\sigma}^{2*}_\rho}{Pr(Player\ 2\ is\ Sophisticated)}$$

where the denominator is a measure of how much he compensates for the behavior of his simple counterpart. At the same time, the sophisticated Player 1, since he has no simple counterpart and expects no equilibrium deviation, has no incentive to deviate either.

Let $\Delta := \sigma^{2*} - \lambda \check{\sigma}^{2*}_\rho$, the numerator above. If $\Delta < 0$ or $\Delta > 1-\lambda$, then $\check{\sigma}^{2*}_{-\rho} \notin [0,1]$, implying there is no equilibrium. If $\Delta \in [0, 1-\lambda]$, then

$$\lambda_\Gamma = \begin{cases} \frac{\sigma^{2*}}{\check{\sigma}_\rho}, & \text{if } \sigma^{2*} < \check{\sigma}_\rho \quad \text{(the solution to } \check{\sigma}^{2*}_{-\rho} = 0) \\ 1, & \text{if } \sigma^{2*} = \check{\sigma}_\rho \quad \text{(no need to compensate)} \\ \frac{1-\sigma^{2*}}{1-\check{\sigma}_\rho}, & \text{if } \sigma^{2*} > \check{\sigma}_\rho \quad \text{(the solution to } \check{\sigma}^{2*}_{-\rho} = 1) \end{cases}$$

thus $\lambda_\Gamma \in \left\{1, \frac{1-\sigma^{2*}}{1-\check{\sigma}_\rho}, \frac{\sigma^{2*}}{\check{\sigma}_\rho}\right\}$. If the simple population is greater than this threshold (if $\lambda > \lambda_\Gamma$), then there is nothing the sophisticated population can do to restore any equilibrium and the theory is at a loss when trying to predict the outcome of this static game.

Note that the equilibrium

$$\begin{aligned} (\check{\sigma}^{1*}, \check{\sigma}^{2*}_{-\rho}, \check{\sigma}^{2*}_\rho) &= \\ (\sigma^{1*}, \frac{\sigma^{2*} - \lambda \check{\sigma}_\rho}{1-\lambda}, \check{\sigma}_\rho) &:= \check{\mathscr{E}} \end{aligned}$$

is not unique. The strategies in $\check{\mathscr{E}}$ are best-responses but, following our original motivation, the bounded rationality of a simple player could imply that he only plays pure strategies. (Alternatively, this could be seen as the special case where $\breve{\sigma}_\rho \in \{0,1\}$.) In this case, the analysis does not fundamentally change. If $\mathscr{E}$ is purely mixed and it is common knowledge that Player 2 will always play $\breve{\sigma}_\rho = 0$, then

$$\breve{\sigma}^{2*}_{-\rho} = \frac{\sigma^{2*}}{1-\lambda}$$

and $\lambda_\Gamma = 1 - \sigma^{2*}$. Likewise, if Player 2 will always play $\breve{\sigma}_\rho = 1$, then

$$\breve{\sigma}^{2*}_{-\rho} = \frac{\sigma^{2*} - \lambda}{1-\lambda}$$

and $\lambda_\Gamma = \sigma^{2*}$. These are the two cases where the distortion induced by simple players is at its maximum.

To illustrate, Figure 4 depicts how the Stability Theorem first states that for any parametrization $\left(u^i_{j,k}\right)$,the game is *stable* in the sense that all final payoffs are unchanged if the rule-of-thumb population changes, as long as its prevalence is small enough and common knowledge.

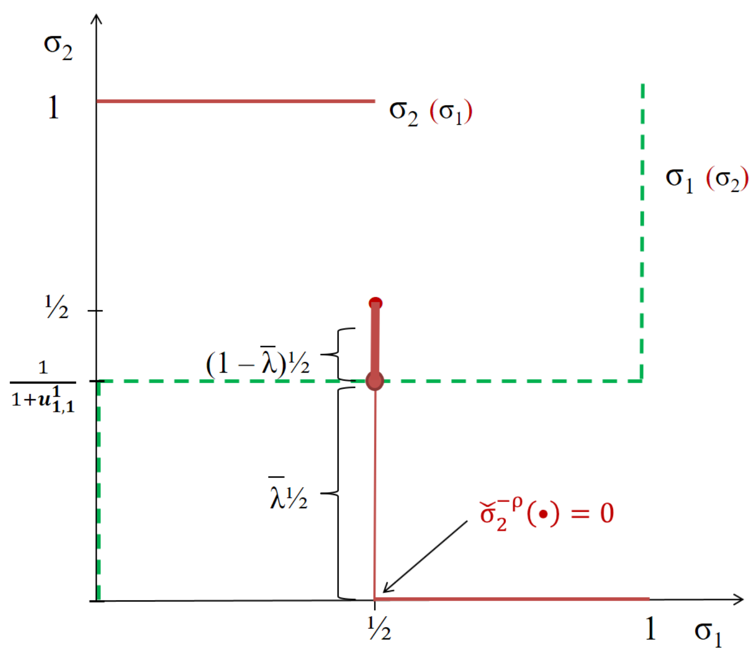

Figure I.4: Restoration of Equilibrium

Following the previous example, as shown in Figure 4, the theorem also predicts a compensating behavior by the sophisticated population. Since $\sigma^{2*} < 1/2$, then $\breve{\sigma}^{2*}_{-\rho} < \sigma^{2*}$. When $\lambda = \lambda_\Gamma$, the sophisticated Player 2 can compensate no more since we are at the point where $\breve{\sigma}^{2}_{-\rho} = 0$ and if the simple population was any larger $(\lambda > \lambda_\Gamma)$, there would exist no equilibrium.

6 Extensions

From an evolutionary perspective, in a symmetric 2×2 game (for $s,t \in \{1,2\} : u^t_{s,s} = u^{-t}_{s,s}$ and $u^t_{s,-s} = u^{-t}_{-s,s}$), pure strategy s^i_1 risk-dominates s^i_2 if, and only if, $\mathbb{E}(u^i \mid s^i_2, \sigma^{-i} = \frac{1}{2}) \leq \mathbb{E}(u^i \mid s^i_1, \sigma^{-i} = \frac{1}{2})$. It strictly-dominates it if the inequality is strict. In this case, this holds if $u^1_{1,1} + u^1_{1,2} \leq u^1_{2,1} + u^1_{2,2}$.**[?]**

It can be shown that the best response to uniform simple mixing coincides with playing a risk dominant strategy. As such, in the special case of $\check{\sigma}_\rho = \frac{1}{2}$, a risk-minimizing agent will have the same behavior as a simple player and no simple population will be able to invade a NE-mixing population successfully. Note that if both players are risk minimizers (instead of utility maximizers), even if they are simple mixers, then an equilibrium exists if, and only if λ is below some upper bound λ_Γ.

		Player 2		
		s^2_1	s^2_2	
Player 1	s^1_1	$u^1_{1,1}, u^1_{1,1}$	$u^1_{1,2}, u^1_{2,1}$	σ^1
	s^1_2	$u^1_{2,1}, u^1_{1,2}$	$u^1_{2,2}, u^1_{2,2}$	$1 - \sigma^1$
		σ^2	$1 - \sigma^2$	

Figure I.5: Generic 2 ×2 Symmetric Game

We can analyze learning in repeated games with simple players. Suppose two agents repeatedly play the asymmetric matching pennies game of Figure 2 as stable pairs with no rematching. Define λ as the (prior) probability that Player 2 is a simple mixer that follows $\check{\sigma}_\rho = 1/2$. Now let $\lambda^T := Pr(\text{Player 2 } \textit{is Simple} \mid s^2_T, s^2_{T-1}, \ldots, s^2_1)$ be the posterior probability where $s^2_T :=$ (Player 2's action at time T). Assume no discounting and perfect myopia (no strategic behavior across time periods).

If the game is repeated T times, there is a decreasing limit $\lambda^{T+1}_\Gamma < \lambda^T_\Gamma$ that converges to zero as $T \to \infty$ where , λ^T_Γ is an upper bound such that if $\lambda > \lambda^T_\Gamma$, the repeated game has no equilibrium. If $\lambda \leq \lambda^T_\Gamma$, the repeated game is stable in that all players get the same expected utility in the generalized game as in the generic game with no simple mixers: $u^{1*}_T = \check{u}^{1*}_T$ and $u^{2*}_T = \check{u}^{2*}_T$.

The inequality $\lambda^{T+1}_\Gamma < \lambda^T_\Gamma$ implies that as the game is repeated, agents update their priors based on observed outcomes. This requires a tighter limit on the simple mixer prevalence to preserve the equilibrium. Some initial proportion λ that allows for a stable equilibrium if the game is repeated T times might be too high if it is repeated $T+1$ times. Learning makes it harder to ensure the existence of a NE, but if it exists, it is payoff-equivalent to the equilibrium of the restricted generic game. Finally, since in this example $\sigma^{i*} < \check{\sigma}_\rho$

(the opposite case is symmetric) learning about the rationality of one's rival pushes posterior probabilities in the expected directions:

$$\lambda_{\Gamma}^{T+1} = \begin{cases} > \lambda_{\Gamma}^{T}, & \text{if } s_{T-1}^{2} = 1 \\ < \lambda_{\Gamma}^{T}, & \text{if } s_{T-1}^{2} = 0. \end{cases}$$

7 Conclusion

We postulate a new behavioral bias in how people play mixed strategies by proposing the existence of *simple* players who lack any strategic depth. We define them as those who, when indifferent between payoff-maximizing choices, follow a simple *rule-of-thumb* by assigning a predetermined probability to each.

We show that simple mixers are not at a disadvantage by using a behavioral rule of thumb; indeed, by simplifying their behavior, they free-ride off the computational complexity added to the problem of their sophisticated counterparts without any loss in rents. Nor are sophisticated players at an advantage that allows them to increase their payoffs: there is no rent to be gained by this strategy restriction and no welfare loss to be fixed by a social planner. Whenever they are indistinguishable from simple mixers, sophisticated players account for this and compensate for their behavior by including in their equilibrium calculation this added complexity.

Compensation holds because, under random matching, games are *stable* to the inclusion of *simple* players, up to a limit. We show that in 2×2 games, for any NE and any rule-of-thumb, if some proportion of the players is *simple* then the equilibrium is *stable* in the sense that all final payoffs are unchanged. The only assumptions are the percentage of *simple* players is below some endogenous threshold (otherwise equilibria do not generally exist) and common knowledge.

In real-world settings where individual behavior cannot be tracked, from any given sequence of player actions, the generalized game with simple and sophisticated players is, therefore, indistinguishable from the game where all players are sophisticated and follow the NE. In experimental settings where individual behavior can be observed under varying controls, simple and sophisticated mixers can perhaps be identified.

The empirical hypothesis is that in 2×2 games with a unique and totally mixed NE, if estimated behavior is statistically close to the NE, some players will be mixing statistically close to $\breve{\sigma}_{\rho} = \frac{1}{2}$ while the rest will compensate by playing in the opposite direction of the bias. Furthermore, this identification of simple and sophisticated players should be consistent throughout different experimental sessions, variations of the games and different experimental conditions.

REFERENCES

REFERENCES

Bloomfield, Robert. "Learning a Mixed Strategy Equilibrium in the Laboratory," *Journal of Economic Behavior & Organization* 25(3), 1994: 411--36.

Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong. "A Cognitive Hierarchy Model of Games," *The Quarterly Journal of Economics*, 2004: 861--98.

Cason, Timothy N., and Tridib Sharma. "Recommended Play and Correlated Equilibria: An Experimental Study," *Economic Theory* 33(1), 2007: 11--27.

Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberri. "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications," *Journal of Economic Literature* 51(1), 2013: 5--62.

Erlei, Mathias, and Heike Schenk-Mathes. "Bounded Rationality in Principal-Agent Relationships," *TUC Working Papers in Economics*, 2012.

Fudenberg, Drew, and Jean Tirole. *Game Theory*, Cambridge, Massachusetts, 393, 1991.

Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey. "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory* 104(1), 2002: 247--72.

Harsanyi, John C. "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points," *International Journal of Game Theory* 2(1), 1973: 1--23.

Koller, Daphne, and Nimrod Megiddo. "Finding Mixed Strategies with Small Supports in Extensive Form Games," *International Journal of Game Theory* 25(1), 1996: 73--92.

Koller, Daphne, Nimrod Megiddo, and Bernhard Von Stengel. "Fast Algorithms for Finding Randomized Strategies in Game Trees," *In Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*, 1994: 750--59.

Lipton, Richard J., Evangelos Markakis, and Aranyak Mehta. "Playing Large Games Using Simple Strategies," *Proceedings of the 4th ACM Conference on Electronic Commerce*, 2003: 36--41.

Masiliunas, Aidas, Friederike Mengel, and J. Philipp Reiss. "Behavioral Variation in Tullock Contests," *Working Paper Series in Economics*, Karlsruher Institut für Technologie (KIT), 2014.

McFadden, Daniel. *Conditional Logit Analysis of Qualitative Choice Behavior*. 1973.

McKelvey, Richard D., and Talbot Page. "Public and Private Information: An Experimental Study of Information Pooling," *Econometrica: Journal of the Econometric Society*, 1990: 1321--39.

McKelvey, Richard D., and Thomas R. Palfrey. "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics* 1(1), 1998: 9--41.

McKelvey, Richard D., and Thomas R. Palfrey. "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior* 10(1), 1995: 6--38.

Misirlisoy, Erman, and Patrick Haggard. "Asymmetric Predictability and Cognitive Competition in Football Penalty Shootouts," *Current Biology* 24(16), 2014: 1918--22.

Nash, John. "Non-Cooperative Games," *Annals of Mathematics*, 1951: 286--95.

Nyarko, Yaw, and Andrew Schotter. "An Experimental Study of Belief Learning Using Elicited Beliefs," *Econometrica: Journal of the Econometric Society* 70(3), 2002: 971--1005.

Rubinstein, Ariel. *Modeling Bounded Rationality*, Vol. 1. MIT Press, 1998.

Simon, Herbert A. "Models of Bounded Rationality: Behavioral Economics and Business Organization, Vol. 2," *The Massachusetts Institute of Technology*, 1982.

Stahl, Dale O., and Paul W. Wilson. "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior* 10(1), 1995: 218--54.

Van Damme, Eric. *Stability and Perfection of Nash Equilibria*. Vol. 339. Springer, 1991.

Chapter 2, Naïve versus Sophisticated Mixing: Experimental Evidence

Christian Diego Alcocer Argüello

Michigan State University[*‡]

2016-7-16

Abstract

We identify a new bias in how people behave when playing games with purely mixed equilibria. Previous research by Alcocer and Jeitschko (2014) define *naïve* players as those who, when indifferent between optimal choices, use a rule-of-thumb and assign an equal probability to each one of them. Based on their equilibrium results, we design an experiment to test for the existence of naïve players. In a first session, we sort participants into two groups: naïve and their *sophisticated* counterparts. In a second session, played two weeks later, each group plays against varying proportions of automated players (bots) that follow varying off-equilibrium mixed strategies. We find evidence of the existence of players that are relatively naïve and of the reaction by sophisticated players. This compensation is in the correct direction but smaller than what equilibrium restoration would require. This implies that when playing non-trivial games with a mixed equilibrium, there are predictable methods to attain above-equilibrium payoffs. Lastly, the analysis suggests that the probability of being naïve can be partially predicted by a simple quantitative test.

1 Introduction

Mixed strategy Nash equilibria (NE) in non-cooperative games can be approached heuristically as a two-step decision process. First, a player determines, given his beliefs regarding his opponents' strategies, that

[*]Keywords and terms: Experimental, Behavioral, Bounded Rationality, Compensated Equilibrium, Computer Bots, Heuristics, Mixed Equilibria, Modeling Cognitive Heterogeneity, Nonlinear Diff-in-Diff, Naïve and Sophisticated Players, Two-Stage Experiment. JEL: C72, C91, D03, D83.

[†]We wish to thank Jon X. Eguia, Thomas D. Jeitschko, Arijit Mukherjee, Jeffrey M. Wooldridge and Mitchell Strahlman for helpful comments.

he has no unique best response but rather has several optimal actions. Given that he is indifferent between these several optimal actions, he proceeds to the second step and determines what sort of mixing among these strategies prevents his opponents from seeking and attaining rents. In equilibrium, these beliefs about his opponents' strategies are confirmed. In other words, one can separate the decision process into two parts: the computation behind utility maximization and the more strategically sophisticated deliberation that ensures the opponent is indifferent as well. In this paper, we find evidence that some individuals appear to deliberate imperfectly during the first step, and others completely neglect the second one and instead *shoot from the hip*. This leaves the door open for other players to compensate for these missteps and potentially gain rents. We find evidence of the existence of this type of sophisticated players – that is, ones who appear to compensate for that type of off-equilibrium behavior.

The results in this paper confirm the assumptions and results of Alcocer and Jeitschko (2014) (*AJ*) who explore the implications of relaxing the typical rationality assumptions on players' decision processes and propose the existence of *naïve* players (or *rule-of-thumb mixers*). In their work, these naïve players correctly determine their own expected payoffs and will, if there exists a unique best option, select it, as the standard theory predicts. However, naïve players have a rationality bound such that – when they are indifferent between two or more maximum payoffs – mix according to some (off-equilibrium) distribution instead of going through the levels of deliberation required to find a NE. The principle of indifference (e.g., Keynes (1921); it is also called *principle of insufficient reason* in probability theory) suggests that such players will assign the same weight to each best response. As such, in the case of 2×2 games, indifferent players would flip a coin, even when this mixing is not part of a NE.

Alcocer and Jeitschko's main result is that, in finite games, for any totally-mixed NE and any *rule-of-thumb*, if some proportion of the players is *naïve* then, as long as this proportion is small enough and common knowledge, the equilibrium is *rule-of-thumb-stable*. This means the generalized game is an isomorphic setting to the NE since the convolution resulting from the linear combination of the mixing of naïve and sophisticated players coincides with the mixing distribution of the standard perfect-rationality case. This implies that the expected final payoffs for all players are the same as in the standard perfect rationality case. The only necessary assumptions are that players cannot recognize the rationality level of their rivals (such that all they know about their opponent is the probability that he is naïve) and that the proportion of naïve players is publicly known; otherwise equilibria do not generally exist.

One important implication of *AJ*'s equilibrium results is that sophisticated players and external agents, such as policy makers, are unable to take advantage of the presence of simple mixers, who, in a sense, free-ride off the cognitive complexity of their sophisticated counterparts. As such, if it is publicly known that a small proportion of the players of a population are rationality-bound in the fashion described above, then there is no justification for the intervention of a regulator (or some other third agent) to enter the market and restore equilibrium, or coordinate a welfare improvement. This holds true even if it is possible to influence the naïve population and determine their off-equilibrium strategy. Additionally, *AJ*'s results regarding naïve mixing may contribute to the theoretical explanation of why, on average, populations may play a NE while few, if any, individuals actually do (Harsanyi, 1973; McKelvey and Palfrey, 1995).

On the other hand, in line with the stylized fact that populations generally do not play according to NE in experimental settings, our understanding and forecasting power of behavior can be enhanced by combining the concepts of imperfect naïve mixing and compensation rather than relying on NE alone. The empirical questions of whether equilibrium is restored and whether some players correctly compensate for the off-equilibrium behavior of other players emerge and are the basis of this experimental investigation. We also inquire if this compensation is robust to changes in the game's payoffs, the fashion in which the naïve population (simulated by computer bots) mixes and the prevalence of the naïve population. Relatedly, we test if there exist any theoretically and empirically sound methods for identifying all players of a population as either naïve or sophisticated based on their behavior while playing simple games, and whether this cognitive heterogeneity can be modeled conditional on observable demographic and academic characteristics.

Specifically, we utilized an experiment with two sessions. In the first session, players participate in multiple rounds of mixed strategy games and are identified as either naïve or sophisticated according to their behavior. The second session incorporates computer-simulated players (bots) that play various known off-equilibrium strategies designed to isolate and measure the relative abilities of sophisticated and naive players to react to the presence of different proportions of these naïve-like automated players. We find evidence of the existence of a relatively naïve population that tends to mix uniformly and that the relatively more sophisticated players react better, but still imperfectly, to off-equilibrium behavior. This confirms that populations generally do not play the NE and, as will be discussed in the *Main Results* subsection, this also reveals a general potential method to take advantage of the off-equilibrium mixing consistently observed in typical experimental settings with non-trivial, random-matching games. Using the results from a pre-play demographic survey and cognitive test, we show that the probability of being sophisticated can be predicted based on observable characteristics.

2 Literature Review

2.1 Purification

Ever since Nash (1951), there have been attempts to extend the concept of NE to deal better with mixed strategy games. In this section, we briefly describe some of the best known theoretical extensions and review several results regarding mixed strategy games from the experimental economics literature.

Purification theory (Harsanyi, 1973) describes how mixed-strategy equilibria can be interpreted as the limit of pure-strategy equilibria in games when perturbations to all payoffs are assumed.[1] Each player chooses their best response (unique with probability one) given the probability distribution over the other players' actions. Ex-ante, all pure strategies receive a positive probability which is increasing in expected payoffs.

This is conceptually different from *AJ*'s *naïve* mixing. Under purification if the variances of the perturbations tend to zero, all strategies tend to the NE while a naïve player will never play an equilibrium with a non-uniform mixing strategy. In addition, consider the game in Figure 2.1: using the purification method, the strongly-dominated strategy **C** is played with some strictly positive probability whereas this does not occur among naïve (as defined by *AJ*) players. Finally, in equilibrium the relative likelihood of outcomes **A** and **B** under purification coincide with the relative likelihoods in the NE $(1:3)$, while the relative likelihood by a naïve player (at the individual level) will be 1 : 1.

	A	**B**	**C**	(NE)
A'	$3,0$	$0,3$	$-1,-1$	$(\frac{1}{4})$
B'	$0,1$	$1,0$	$-1,-1$	$(\frac{3}{4})$
C'	$-1,-1$	$-1,-1$	$-2,-2$	(0)

Figure II.2.1: Bordered Asymmetric Matching Pennies

2.2 Quantal Responses and Logit Equilibrium

The concept of a quantal response equilibrium (*QRE*), like *purification*, is introduced by adding perturbations to a game (McKelvey and Palfrey, 1995; McKelvey and Palfrey, 1998; Goeree, Holt and Palfrey,

[1] Formally, in the most general terms, if player i's payoff given the (pure) strategy profile s is $u_i(s)$, then what he actually observes is $U_i(s;\theta_i) := u_i(s) + \varepsilon_i^s \theta_i^s$ where $\varepsilon_i^s \theta_i^s$ is the gaussian i.i.d. perturbation scaled by his type and $E(\varepsilon_i^s) = 0$ (see Fudenberg and Tirole, 1991). The vector θ_i can be interpreted as related to levels of rationality: if all its components are zero, the player (perfectly) observes an unperturbed game. If $Var\left(\varepsilon_i^s\right) \to \infty$ or $\theta_i^s \to \infty$, he observes only noise and mixes uniformly across all–including dominated–strategies.

2002).[2] Experimental evidence suggests QRE is sometimes a better predictor than NE and purification (Erlei and Schenk-Mathes, 2012).

Under a given set of assumptions,[3] a *QRE* can be approximated and estimated as a *logit equilibrium* (McFadden, 1973). In a logit equilibrium, for a given $\lambda_i \geq 0$, the *logistic quantal response* is defined as $\sigma_i^j(s_i^j \mid \sigma_{-i}) = exp\left\{\lambda_i U_{i,-i}^j\right\} / \sum_{\kappa=1}^{K} exp\left\{\lambda_i U_{i,-i}^\kappa\right\}$ which is the probability that i assigns to action s_i^j (from some finite action set of cardinality K) given the other players' strategies, and $U_{i,-i}^j$ is the perturbed payoff he gets after s_i^j and σ_{-i}.

The parameter λ_i can be interpreted as a measure of i's rationality. As $\lambda \to \infty$, if there exists a unique NE (σ^*), as is the case throughout the games described below, then $(\sigma_1, ..., \sigma_N)$ converges to σ^*. As $\lambda \to 0$, players become *completely blind* and mix uniformly across all strategies such that the probabilities tend to $\frac{1}{K}$. Again, note that this differs from *AJ*'s naïve mixing which only yields uniform mixing across pure strategies if they all yield the same maximum utility.

2.3 Level-K Thinking

An alternative family of theories often used to model off-equilibrium behavior is *level-k* thinking or *cognitive hierarchy* theory (Camerer, Ho and Chong, 2004; Van Damme, 1991). This class of theories generally assumes that there are multiple levels of players with varying levels of sophistication. Most begin by assuming that there are non-strategic *level*-0 players who always mix uniformly among all their pure strategies.[4] The next level (*level*-1) players best-respond to this *anchor*. Iteratively, *level-k* players best-respond to *level*-(*k*-1) populations.[5]

Using the game in Figure 1 as an example, a *level*-0 player would play $Pr(A) = Pr(B) = Pr(C) = \frac{1}{3}$, which, while similar, is distinct from what a naïve player would do: $Pr(C) = 0$ and $Pr(A), Pr(B) \in \left\{0, \frac{1}{2}, 1\right\}$. Note that no distribution of levels exists to replicate the strategy of the naïve mixer.

[2] In this case, if player i's payoff given the (pure) strategy profile s is $u_i(s)$, then what he actually observes is $U_i(s) := u_i(s) + \varepsilon_i^s$ where ε_i^s is the i.i.d. perturbation $E(\varepsilon_i^s) = 0$.

[3] Independence of irrelevant alternatives, irrelevance of alternative set effects and $with\ little\ loss\ of\ generality$ positivity (agents always assign a strictly positive probability to all strategies).

[4] Using the terms of strategy purification, we can think of them as players who face perturbations with infinite variance or players of type *infinity*. This is one of the simplest possible behaviors in terms of computational complexity (Koller and Megiddo, 1996; Koller, Megiddo and Von Stengel, 1994; Lipton, Markakis and Mehta, 2003).

[5] The empirical challenge is usually to determine the number of levels that best fits the data. Camerer (2004), for instance finds that *an average of 1.5 steps fits data from many games*. See also Stahl and Wilson (1994) who assume players' types include level-0, level-1, level-2, *NE*, *worldly* (who best-responds to the previous ones) and *rational*, who best-responds to all. They use numeric methods to estimate best responses in 3×3 games and then estimate the prevalence of each population to show their models yield better log-likelihoods than an estimation that imposes all players are *rational*.

2.4 Mixed Nash Equilibria in Experimental Economics

Many experiments have explored play in mixed strategy games. One conclusion of this literature is the stylized fact that when facing games with a unique, totally-mixed NE, players' behavior is significantly different than the strategies predicted by game theory. This is especially so when dealing with complex games but also holds for games with a unique pure-strategy NE. In fact, this general finding has been one of the primary motivations for the theoretical literature on extensions to the NE (like those above) that allow for all pure strategies to have positive probabilities.

There is some experimental evidence of learning in that behavior gets closer to the NE in later periods but, generally, it shows no convergence to the NE (Erlei and Schenk-Mathes, 2012; Misirlisoy and Haggard, 2014; Ochs, 1995). Moving away from normal form games, there is also empirical evidence that confirms departures from NE in principal-agent and moral hazard (Erlei and Schenk-Mathes, 2012) games, even those with a (unique) subgame perfect NE (Haptonstahl, 2009).

Masiliunas et al. (2014) show that reduced game complexity in the context of experimental contests and auctions (induced with bots[6] that disclose their future actions) correlates with *less behavioral variation* and closeness to the NE. Relatedly, Parkhurst, Koford and Grijalva (2015) find that when individuals face information overload, they tend to use simplifying heuristics and in some cases, like Masiliunas et al. (2014) cannot reject the hypothesis that in complex games, players' choices are uniformly distributed.[7]

Models and experiments involving belief estimation and elicitation (Nyarko and Schotter, 2002; McKelvey and Page, 1990; Offerman, Sonnemans and Schram, 1996) have also attempted to address the lack of NE play in mixed strategy games. Although belief estimation can help explain some of the variation and belief elicitation tends to drive behavior closer to equilibria, these techniques leave the question of what motivates strategy choice by individuals and populations unanswered. Finally, distinct types of rationality models have been tested to account for these deviations from NE, including Stahl (1995) and others.[8] This present investigation tests whether the bounded rationality assumptions underlying *AJ*'s *naïve mixing*, might be able to shed some light on these issues while having explanatory power.

[6]See Cason and Sharma (2006), and Houser and Houser and Kruzban (2002) for further examples of computer bots used to generate experimental evidence in games where the *NE* is shown to not be the best forecasting tool.

[7]Their observed percentage of choices that coincide with the *NE* is close to the percentage that a uniform mixer would get.

[8]Including, but not restricted to research on, fairness, framing effects, hyperbolic discounting, inequity aversion, reference points and welfare (social) preferences. See Rubinstein (1998).

3 Experimental Design and Procedures

The experiment presented here is designed to test if relatively naïve and sophisticated players can be consistently identified through a sorting of their observed behavior and whether they behave differently when playing against computer simulated off-equilibrium strategy players. As noted earlier, we implement an experiment consisting of two sessions where we use the first session to sort players into either a naïve or a sophisticated category. There are then two follow-up sessions; one for the naïve category players and another for the sophisticated. Participants were unaware of the function of the sorting session or that they had been categorized for the follow-up session. It should be noted that our paper focuses on a relatively small population – undergraduate students. As such, the effects of differences in educational experiences on the determination of cognitive ability are diminished.

All sessions involved the same basic setting. Specifically, after arriving and logging on to the computer, half the players were randomly assigned as P_1 ("row" players) and half as P_2 ("column" players). While we differentiate row and column players for discussion in the paper, all players viewed the game as a row player on the computer. Players maintained their type throughout a session and played multiple rounds of the following matching pennies game where either $a = 1$ (the symmetric case, denoted as $\tilde{\Gamma}$) or $a = 3$, denoted as Γ. Players were anonymous and were randomly re-matched after each round. Each session included a tutorial and practice round.

		P_2		
		Left	**Right**	
P_1	**Up**	a, 0	0, a	$(1-\beta)$
	Down	0, 1	1, 0	(β)
		$(1-\gamma)$	(γ)	

Figure II.3: Base Game

$a \in \{1,3\}$, NE: $\beta^* = \gamma^* = \frac{a}{a+1}$

3.1 Sorting Session

The primary purpose of the sorting session was to provide a mechanism to sort players into naïve and sophisticated categories. Each participant played the following games: 10 rounds of $\tilde{\Gamma}$ and 50 rounds of Γ in a random order. In the symmetric version of the *matching pennies* game $(\tilde{\Gamma})$, the (unique, totally mixed) NE is $\beta_1^* = \gamma_1^* = \frac{1}{2}$. Although Γ is not symmetric, it was chosen such that the unique NE was the same for both types of players: $\beta_3^* = \gamma_3^* = \frac{3}{4}$. After the completion of all 60 rounds of play, the participants completed a short questionnaire that included standard socioeconomic variables, then answered a short multiple-choice quantitative test.

3.2 Bots Session

After participating in the sorting session, players were sorted and invited to one of the two follow-up bot sessions which were otherwise identical. As in the sorting session, participants were randomly assigned to be either P_1 (row player) or P_2 (column player): Participants played the same games (Γ and $\tilde{\Gamma}$) as described above, but bots were included as players in varying percentages and were played in the order shown in Table 3.2. The bots in these sessions played with some probability $\overline{\gamma}^N$or $\overline{\beta}^N$ as described in column 5.[9] All percentages and probabilities were common knowledge.

Stage	Rounds	Game	Bot Prevalence (λ)	Bot Randomization	
				If P_1	If P_2
1) $\tilde{\Gamma}$_25_50	10	$\tilde{\Gamma}$	25%	$\overline{\beta}^N = \frac{1}{2}$	$\overline{\gamma}^N = \frac{1}{2}$
2) Γ_25_50	30	Γ	25%	$\overline{\beta}^N = \frac{1}{2}$	$\overline{\gamma}^N = \frac{1}{2}$
3) Γ_50_50	30	Γ	50%	$\overline{\beta}^N = \frac{1}{2}$	$\overline{\gamma}^N = \frac{1}{2}$
4) Γ_50_0	10	Γ	50%	$\overline{\beta}^N = 0$	$\overline{\gamma}^N = 0$

Table II.3.2: Stages during the Bots Session

In stage one (denoted as $\tilde{\Gamma}$_25_50) participants played 10 rounds of $\tilde{\Gamma}$ while knowing that 25% of the total population consisted of bots that played either *Up* or *Down* (*Left* or *Right*) with a $50/50$ chance. This was the only stage where the bots were programmed to play the NE strategy. This stage was designed to test if the inclusion of NE-playing bots in a simple game had any effect on players' behavior.

During the last three stages (Γ_25_50, Γ_50_50 and Γ_50_0) participants played Γ with the knowledge that they would face a computer bot playing some known strategy with some probability. During stages two and three, the bots were coin-flipping ($\overline{\beta}^N = \overline{\gamma}^N = \frac{1}{2}$) and their prevalence varied: in the second stage the bots made up 25% of the population, whereas in the third they made up 50% of the population. Similarly, in stages three and four, bots made up 50% of the population; and their behavior was either coin-flipping ($\overline{\beta}^N = \overline{\gamma}^N =$

[9] β and γ are the strategies played by P_1 and P_2, and $\overline{\gamma}^N$ or $\overline{\beta}^N$ are their arbitrarily-selected strategies.

$\frac{1}{2}$) or always playing *down/right* ($\overline{\beta}^N = \overline{\gamma}^N = 0$) Note that from the participants' perspective, these stages are, in principle, not qualitatively distinct but just a series of games that include stronger distortions from the NE of Γ. As such, these stages are where we expect to see compensating behavior to the bots' play (as described in the next section) and predict that the *sophisticated* group compensates more accurately or completely than the *naïve* group.

3.2.1 Compensated Equilibrium Predictions

Following *AJ*, Table 3.2.1 shows the 'Compensated Nash Equilibrium' (CNE) for all stages. The CNE is, in essence, the Nash equilibrium after taking into account the presence of players that are following publicly known off-equilibrium strategies. During most of the treatments of the bots session, it consists of unique and interior (purely mixed) strategies, implying that any deviation collapses the opponent's best-response relation into a degenerate distribution.

Stage	**Player 1** (β^*)	**Player 2** (γ^*)
1) $\tilde{\Gamma}$_25_50	0.50	0.50
2) Γ_25_50	0.83	0.83
3) Γ_50_50	0.00*	1.00
4) Γ_50_0	0.00	1.00

* Denotes Trembling-Hand Perfect Equilibrium

Table II.3.2.1: Compensated Nash Equilibrium (CNE)

As mentioned above, for stage 1 in which the participants are playing game $\tilde{\Gamma}$, the presence of the bots should have no impact on the CNE since the bots are just playing the NE. As such, the CNE is the same as the NE in Stage 1. For the rest of the stages (2-4), the participants are playing game Γ where the CNE has player type P_1 playing *up* with probability $\frac{1}{4}$ (and, thus, *down* with probability $\frac{3}{4}$) while player type P_2 plays *left* with probability $\frac{1}{4}$ (*right* with probability $\frac{3}{4}$): $\beta^* = \gamma^* = \frac{3}{4}$. If the population of coin-flipping bots (i.e., they play $\overline{\beta}^N = \overline{\gamma}^N = \frac{1}{2}$) is 25% as in stage 2, the unique CNE is to play $\beta^{S*} = \gamma^{S*} = 0.83$. The CNE has the characteristic that the predicted equilibrium behavior compensates by moving in the opposite direction. That is, since $\overline{\beta}^N < \beta^*$, then $\beta^{S*} > \beta^*$, where β^{S*} represents the CNE strategy and the same holds for P_2.

When we increase the prevalence of bots to 50% as in stage 3, then both player types are indifferent (and, thus, reach an equilibrium) if they play $\beta^{S*} = \gamma^{S*} = 1$. This is not, however, the only equilibrium since if $\gamma = 1$, then any $\beta \in [0,1]$ is a best response and $\gamma = 1$ is always a best response. Additionally, any expected perturbation by P_2 would make P_1 have a unique best response: if P_1 believes $\gamma < 1$, then his profit-maximizing action is $\beta = 0$. This is important since, typically, experimental populations do not play the NE and thus it would often be optimal to play $\beta = 0$ if P_1 or $\gamma = 1$ if P_2.

In stage 4, the bot behavior moves even further away from the NE. In this stage, players' compensation is expected to be greater than in the previous two stages. Specifically, the bots play $\overline{\beta}^N = \overline{\gamma}^N = 0$ while their prevalence relative to the previous stage is unchanged (50%). In a sense, this is the easiest stage to play. With or without perturbations, the game is no longer strategic in that now each player has a strongly dominant strategy. In this case, players are no longer expected to mix, and their best responses are $\beta = 0$ and $\gamma = 1$.

4 Results

A total of forty-four undergraduate students from Michigan State University enrolled in the experiment.[10] All participants attended one of the two sorting sessions.[11] Participants were then divided into two groups (naïve and sophisticated) consisting of 20 and 24 participants, respectively. Despite efforts to ensure that participants could and would come to the follow-up or bots sessions (including making sure all participants indicated they could attend either of the follow-up sessions, and delaying payment of the first session show-up fee until the second session), only 16 of the naïve participants and 20 of the sophisticated participants attended the follow-up bots sessions.[12]

[10]Participants earned on average \$24.17 in the sorting sessions and \$28.89 in the bot sessions. In an effort to maximize the cognitive effort of the participants and to remove potential wealth effects, 3 (4) rounds were randomly selected at the end of the session for payment in the sorting (bots) sessions. Participants earned tokens (0, 1 or 3) worth \$5 so the maximum they could earn in a round was \$15.

[11]It was verified that players who self-selected into both sessions (by choosing the date they wished to participate) were not systematically different. Simple statistical tests found no evidence for apparent differences on behavior (means), cognitive levels (test scores) or observable characteristics (questionnaire) between both groups.

[12]By comparing earnings, test scores, and estimated behavior, it was verified that attrition was not systematic.

4.1 Sorting Session

For the $\tilde{\Gamma}$ games in the sorting sessions we fail to reject the null that participants on average played the NE (the NE is $\frac{1}{2}$, the average mixing was 0.533 and the *p-value* related to the test for equality is 0.851).[13] As such, during these rounds, it would not have been profitable for players to deviate.[14] Both populations seemed to understand the standard symmetric matching pennies game and play it, on average, as the theory in its simplest form would forecast. On the other hand and consistent with the experimental literature, both player types (P_1 and P_2) do not appear to have understood the more complex game Γ as clearly. In fact, every P_2 would have benefited by deviating to a pure strategy in response to P_1's play. Specifically, both populations played a mixed strategy estimated to be less than the NE (P_1 played $\hat{\beta} = 0.601$,[15] P_2 played $\hat{\gamma} = 0.707$, NE $= 0.75$), but only the P_1 population's play was statistically different (*p-value* $= 0.003$) than the NE.[16]

4.1.1 Sorting Process

The P_1 and P_2 populations were divided in half; those *relatively naïve* and those *relatively sophisticated*. Within each population, each player was classified as *sophisticated* if their behavior while playing Γ was relatively close to the NE, or as *naïve* if it was relatively close to uniform mixing. Since Γ is not symmetric, it was necessary to perform the sorting separately within each player type population. The approach taken was to generate a dummy (s_i) equal to one if each individual i's mixing was above the median amongst its type, so that half of each population was defined as naïve ($s_i = 0$) and the other half as sophisticated ($s_i = 1$). The median mixing for P_1 was playing *Down* with probability 0.588. The median mixing for P_2 was playing right with a probability of 0.725.

[13]Each choice is a Bernoulli trial and we focus on its parameter p. As such, our analysis is essentially non-parametric and our results coincide numerically with those from the Mann-Whitney U test. This is because the relevant estimators (i.e. differences between averages) are exactly the same in both cases. Furthermore, when doing hypothesis testing, their assumed distribution (normal) is also the same since both tests invoke the central limit theorem.

[14]More specifically, when using OLS to estimate the model $\hat{M}_{1,i} = \delta_0 + \delta_1 Type_i + \delta_2 Mixing_{1,i} + \delta_3 Type_i * Mixing_{1,i} + u_i$ where $M_{1,i}$ is individual i's average experimental payoffs in all $\tilde{\Gamma}$ games of the sorting session, $Type_i$ is a dummy equal to one if i is type P_1 and $Mixing_{1,i}$ is his average behavior throughout all the Γ_1 games, no evidence was found to reject the null hypothesis $H_0 : \delta_2 = \delta_3 = 0$. For details on this model, estimates and *p-values* of the Wald tests, contact the authors.

[15]This observed mixing was thus in the interval between the *NE* and the *level*-0 or *purely naïve* mixing that coincides with the simplest uniform distribution that assigns a probability of 0.50 to each available action or pure strategy. For P_1 and P_2 it was also verified that the average mixing was statistically greater than 0.50 (the *p-values* are, respectively, 0.033 and 0.000).

[16]Throughout all the experiments and treatments, one of the robustness checks that was done was to ignore the first five (also, where applicable, the first ten) game rounds to verify if taking *learning* into consideration fundamentally changed the results. It never did: if only the behavior during the last iterations of the games is measured, it is still the case that players' strategies coincided with the Nash equilibrium in Γ_1 but was statistically lower than the Nash equilibrium in Γ_3.

We chose to use the relatively simple method described above because it is intuitive and quite robust when compared to many alternative methods.[17] Under these alternative sorting methods, the division of participants did not change considerably, and the correlation between s_i and the dummies generated by alternative sorting methods was never statistically negative. Likewise, where applicable, in the regression modeling the probability of being identified as sophisticated (described below), the estimated partial effect of test scores was never negative. Also, the differences in behavior between sophisticated and naïve players during the bots session were qualitatively very similar when distinguishing them by the alternative sorting methods. Overall, we cannot reject the hypothesis that the labeling was consistent: players sorted as either sophisticated or naïve seemed to behave as such throughout.

4.2 Bots Session, Main Results

As described above, the experiment finds evidence of the existence of stable sets of relatively naïve and sophisticated players.[18] Naïve players' mixing is often closer to uniform mixing than the sophisticated players' mixing; not only during the sorting session (true by construction) but, as hypothesized, also in the bots sessions under different bot prevalence levels and off-equilibrium bot behaviors. Sophisticated players' mixing is also often closer to the CNE.[19] Furthermore, sophisticated players frequently react "better" to varying off-equilibrium distortions induced by the bots.

Table 4.2 presents the main results. It includes the mean strategy (average choices are β for P_1 and γ for P_2) for each of the four stages, for each of the four player types: NP_1, SP_1, NP_2 and SP_2. In Stage $\tilde{\Gamma}$_25_50 behavior is not statistically different than the NE of $1/2$.[20] Introducing NE-playing bots does not seem to, in itself, cause players to change their behavior.

During the last three stages, mixing is always statistically different than the CNE for all player types[21]

[17] The full list of alternative sorting mechanisms investigated included sorting based on: questionnaire results, payoffs earned during the screening session, payoffs won during the second half of the screening session, behavior during the screening session (ignoring the first 10 or 20 iterations of Γ_3 to account for learning), behavior during the screening session (pooling P_1 and P_2 players together), behavior compared to their best response (as opposed to the *NE*) given their opponent's behavior during either the whole or the last 40 iterations of the screening session, and dropping from the dataset all but the *most sophisticated* and the *most naïve* (according to the original screening criterium) players.

[18] That is, evidence of a continum of *sophistication* (or *naiveté*) levels. We found no *purely naïve* nor *purely sophisticated* players. For practical purposes, though, we will refer to the whole population as consisting of only two distinct sets with no intersection even though subtle distinctions were observed on the data (and this could be the basis of further research). The results were robust to relaxing the implied assumptions.

[19] This distinction compares behavior to equilibria, not to the best responses that correspond to actual play. Best responses are described in Appendix 4.

[20] *P-values*: $0.266, 0.551, 0.374, 0.426$.

[21] The *p-values* are 0.000 in all cases except NP_2 in stage Γ_50_0, where *p-value* $= 0.013$.

	Player 1 (β)			Player 2 (γ)		
Stage	THP CNE	Naïve (NP_1)	Sophisticated (SP_1)	THP CNE	Naïve (NP_2)	Sophisticated (SP_2)
1) $\tilde{\Gamma}$_25_50	0.50	$0.56^{\dagger}$	$0.53^{\dagger}$	0.50	$0.55^{\dagger}$	$0.54^{\dagger}$
2) Γ_25_50	0.83	0.57	0.62	0.83	0.67^{*}	0.76^{*}
3) Γ_50_50	0.00	0.43	0.43	1.00	0.75^{*}	0.82^{*}
4) Γ_50_0	0.00	0.23^{*}	0.11^{*}	1.00	0.93^{*}	$1.00^{*\dagger}$

* Denotes Naïve and Sophisticated Mixing are Different at the 5% Significance Level
† Denotes Mixing Equal to CNE at the 5% Significance Level
# Naïve Players: 16; # Sophisticated Players: 20
Table II.4.2: Main Results, Estimated Mixed Strategies

with the exception of the SP_2 players who play exactly the CNE in stage Γ_50_0. Every one of the SP_2 players played right in all 10 rounds of this last stage. From a behavioral perspective, it could be argued that it is harder for P_1 to compensate to $\beta = 1$ than for P_2 to compensate to $\gamma = 1$: Γ_1's asymmetry implies it is P_1 who has the responsibility of determining *the size of the cake* (so surplus maximization goes in the opposite way of compensation) whereas P_2 only determines *who gets the cake.* Pursuing this question is beyond the scope of this analysis, but the data suggests the strength of the effects of being naïve do not completely overpower other effects like altruistic preferences, surplus maximization or loss aversion.

The data, therefore, suggests that in games that are complex enough such as stages 2 and 3 (i.e., those games where the NE is different that $1/2$ and without a strongly dominant strategy), the sophisticated players do not compensate enough to restore the equilibrium given the off-equilibrium behavior of bots. This is, in some ways, not inconsistent with the stylized fact that participant populations do not play the NE even without bots. This further implies the existence of a unique best response and it is remarkable that the evidence suggests this best response might be predictable. If what we find here is confirmed, then when complex-enough games as described above, the optimal strategy is to assume simply that you will be matched against a perfectly naïve player. That is, in every variations of Γ above, any player could obtain above-equilibrium payoffs by playing the mixed strategy derived from assuming the opponent population will be mixing uniformly. In a sense, due to the naïve component of their mixing, experimental populations *can* be taken advantage of.

With the exception of the P_1 players during stage 3 **(**Γ_50_50) where NP_1 and SP_1 effectively behaved the same,[22] the mixing of the sophisticated players in the more complex stages (2 and 3) is statistically

[22] $\hat{\beta}_{NP1} = 0.4267, \hat{\beta}_{SP1} = 0.433$ and the *p-value* of $H_0 : \hat{\beta}_{NP1} = \hat{\beta}_{NP1}$ is 0.877.

closer to the CNE (and further from 50/50 mixing) than the respective play of their naïve counterparts.[23][24] Figures 1 and 2 below illustrate this for both player types (P_1 in Figure 1 and P_2 in Figure 2) across stages 2 and 3 (bot prevalence 0.25 and 0.5) and include the data from the sorting sessions (bot prevalence of 0) as a reference. Note that, overall, sophisticated players play closer to the CNE and naïve players always mix closer to $1/2$ (although not statistically so for player 1 types . The consistency with which the naïve players in the more complex games in our experiment mix closer to $1/2$ – both across the different stages and the two-week separated sorting and bots sessions – supports the hypothesis of the existence of a stable set of relatively naïve players in the population. Furthermore, as will be discussed in the next subsection, this set of players may be able to be identified through lower scores on quantitative tests. The effects of interacting with bots and the distinction between *naïve* and *sophisticated* players is illustrated in Figures 1-4.

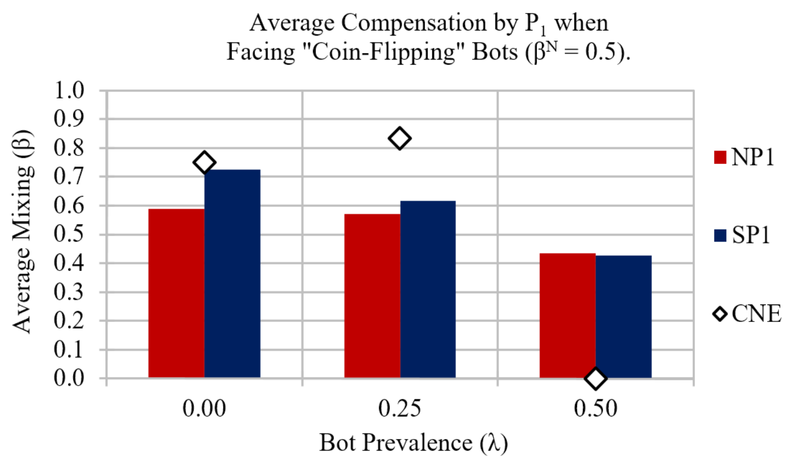

Figure II.4.1: *CNE* and P_1's Behavior (Naïve and Sophisticated)

as the Prevalence of the Bot Population Increases

[23] *P-values*: 0.028, 0.021 and 0.036.

[24] To test for learning, these two stages were separated into halves and analyzed separately. Although there were certainly strong dynamic interactions (see *serial correlation*, below), the evidence for learning is inconclusive: there appeared no consistent or systematic adjustment of the populations' behavior when compared the first to the second half of the stages. The adjustments, if any, were sometimes done in opposite directions.

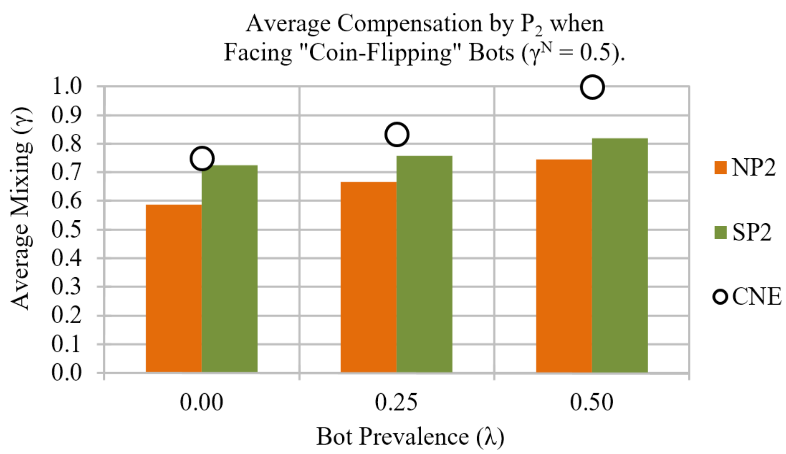

Figure II.4.2: *CNE* and P_2's Behavior (Naïve and Sophisticated) as the Prevalence of the Bot Population Increases

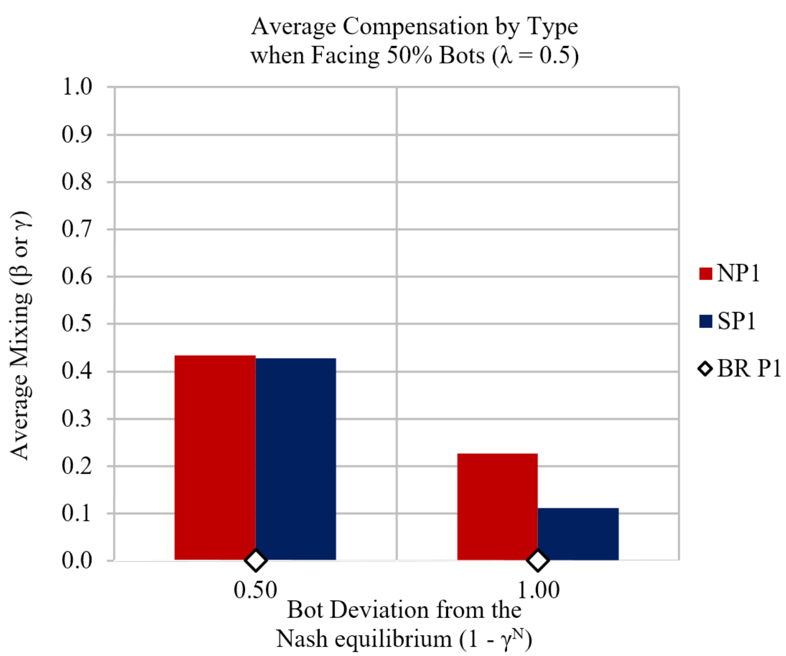

Figure II.4.3: Best Responses and P_1's Behavior (Naïve and Sophisticated) as the Play of the Bot Population Deviates from the Equilibrium

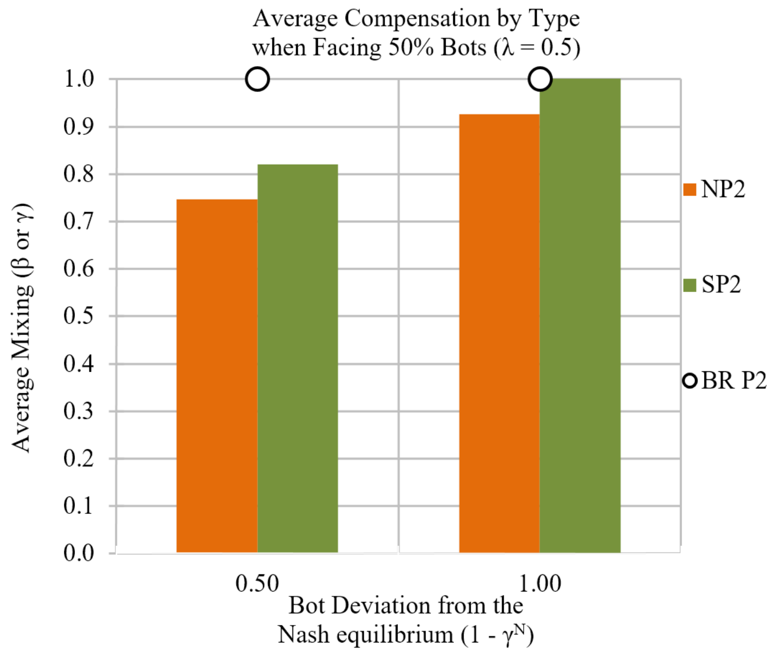

Figure II.4.4: Best Responses and P_2's Behavior (Naïve and Sophisticated) as the Play of the Bot Population Deviates from the Equilibrium

We now formally compare the impact of varying bot distortions on the naïve versus the impact on the sophisticated players. First, focusing on the Γ games with coin-flipping bots, we can contrast the effect of an increase of their population from 25% to 50% (we call this treatment *T1*) by comparing naïve P_1's behavior in stage 2 ($\hat{\beta} = 0.57$) with stage 3 ($\hat{\beta} = 0.43$). The difference is 0.14 and the same calculation with the sophisticated population yields a difference of 0.19. The difference-in-differences (*DiD*) estimate for the treatment effect is thus -0.05, implying the sophisticated population reacted more to treatment *T1*.

As seen in Table 4.2b, doing the same derivations for P_2 yields an estimated effect with the same sign. Yet, the interpretation is the opposite since the expected compensation went up: the CNE is 0.83 in stage 2 and 1.00 in stage 3. The evidence suggests that the P_2 naïve population reacted more to *T1*. However, all *DiD* point estimates are insignificant at 10% or higher levels. Only P_1 after treatment *T2* (comparing stages 3 and 4: games with 50% bots that go from coin-flipping to playing $\beta = \gamma = 0$), with an estimated partial effect of -0.108, has a relatively low *p-value* of 0.127. The sign of this effect is negative which implies the sophisticated population reacted more to *T2*.

	P_1		P_2	
Treatment	DiD	Nonlinear DiD	DiD	Nonlinear DiD
$T1$	-0.052 (0.386)	-0.053 (0.384)	-0.016 (0.766)	-0.003 (0.948)
$T2$	-0.108 (0.127)	-0.111 (0.088)	-0.001 (0.986)	0 *NA*

Table II.4.2b: Treatment Effects

There is a caveat to this analysis and its intuition is analogous to the case of dependent binary variables: a desirable feature of probit and logit models versus linear regressions is diminishing marginal magnitudes of the partial effects. To illustrate, consider P_2's reaction to *T1*: the average adjustment of NP_2 (0.08) is of larger size than the change by SP_2, 0.06. However, when adjusting away from uniform mixing it is probably behaviorally easier if one's mixing starts close to uniformity and harder when mixing is already close to either bound: zero or one. This is another way of saying that any variable that has an effect on players' choices will likely have diminishing marginal effects.

As such, a positive adjustment from $\beta = 0.67$ (NP_2's behavior in stage 2) would be easier than the same modification from $\gamma = 0.76$ (SP_2). This explains why the nonlinear *DiD* estimation for the treatment effect of *T1* on P_2 shown in Table 4.2b is much closer to zero than the linear *DiD*.[25] Under a different specification instead of Φ, with locally faster diminishing marginal effects, this coefficient would turn positive, in line with previous results that seem to show sophisticated players understand the games better than the naïve. The coefficients for P_1 after treatments *T1* and *T2* are also slightly enhanced by the nonlinear estimation, and the effect of *T2* on P_2 cannot be estimated with this method due to colinearity. The estimation of the rate at which these marginal effects decrease is, however, beyond the scope of this analysis.

4.3 Modeling Cognitive Heterogeneity

Next, we investigate if there is a link between general analytic abilities and being a *naïve* player. Our measure of intellectual skills is a ten-question test that participants answered immediately after first session

[25] With a linear model, we would have

$$choice = \alpha + \beta s + \gamma T_k + \delta s \cdot T_k + u$$

where $s_i := 0[i \text{ sophisticated}] + 1[i \text{ is naïve}]$ and $T_k := 0[\text{Stage } (k+1)] + 1[\text{Stage } (k+2)]$. (Note the estimation of δ is numerically equivalent whether we define s or T_k as the treatment variable.) We now specify

$$Pr(choice = 1 \mid s, T_k) = \Phi(\alpha + \beta s + \gamma T_k + \delta s \cdot T_k).$$

These are all dummy variables and Puhani (2012) shows the treatment effect is $\Phi(\alpha+\beta+\gamma+\delta) - \Phi(\alpha+\beta+\gamma)$ and standard errors can be consistently estimated.

(two weeks before the bots session; see Appendix 2). Our results find that high scores relate to playing relatively close to the CNE. We interpret this as further evidence that the proposed concepts of *naiveté* and *sophistication* are empirically grounded: relatively naïve players do exist. This also suggests, first, that the cognitive levels determined in the sorting phase can be partially captured by the proposed analytical test. Second, by conditioning on one's test score, there is no statistically significant impact on the probability of being naïve of other characteristics such as income, demographics or academics.

The hypothesis is that the probability of being sophisticated depends non-linearly but positively on every individual i's cognitive ability (c_i). That is, $\forall i : \Pr(s_i = 1 \mid c_i, x_i) = f(\delta_c c_i + x_i \delta_x)$, where x_i is a vector of observables, and $f(\bullet)$ is an increasing function. The variables in x_i were gathered in a survey filled after the test. We discuss the test and the questionnaire in Appendix 2.

We report these results in Tables 4.2.3a, 4.2.3b and 4.2.3c. Let c_i be the total score on the test. The simplest probit regression, with $f(z) = \Phi(z)$ (the standard normal c.d.f.) and with no other covariates, estimates $\hat{\delta}_c = 0.25$ (*p-value* of 0.042). The related marginal effect is 0.090 suggesting that a one point increase on the ten-question test is associated with a 9 percentage point increase in the probability that the player is sophisticated. These results are robust to several other specifications: logit and linear regressions, defining c_i as some subset of the questionnaire, or adding a different set of controls in x_i.

Model	Probit	Logit	Linear
Response Variable	s	s	s
Test Score	**0.253**	**0.419**	**0.092**
[Marginal Effect]	[**0.090**]	[**0.091**]	[**0.092**]
(p-value)	(0.042)	(0.051)	(0.040)
(Robust)	(0.053)	(0.060)	(0.029)

Table II.4.2.3a: Regressions of s on *Test Score*

Model	Probit	Probit	Probit	Probit	Probit	Probit	Probit
Response Variable	s	s	s	s	s	s	s
Test Score	$\mathbf{0.253}$	$\mathbf{0.347}$	$\mathbf{0.296}$	$\mathbf{0.263}$	$\mathbf{0.227}$	$\mathbf{0.254}$	$\mathbf{0.296}$
(p-value)	(0.042)	(0.018)	(0.025)	(0.037)	(0.095)	(0.049)	(0.027)
Gender	–	$\mathbf{1.064}$	–	–	–	–	–
		(0.034)					
Race	–	–	$-\mathbf{0.146}$	–	–	–	–
			(0.300)				
Age	–	–	–	$-\mathbf{0.161}$	–	–	–
				(0.524)			
GPA	–	–	–	–	$\mathbf{0.259}$	–	–
					(0.618)		
Reported Income	–	–	–	–	–	$\mathbf{0.271}$	–
						(0.288)	
Weekly Expenditures	–	–	–	–	–	–	$\mathbf{0.370}$
							(0.213)

Table II.4.2.3b: Regressions of s including different sets of regressors

Model	Probit	Probit	Probit	**Model**	Probit
Response Variable	s	s_All	s_L_40	**Response Variable**	s
Test Score	$\mathbf{0.253}$	$\mathbf{0.232}$	0.172	**GRE Questions**	0.222
(p-value)	(0.042)	(0.056)	(0.150)	**(p-value)**	(0.214)

Table II.4.2.3c: Regressions of s including different sets of regressors (2)

We interpret this as evidence that high cognitive abilities (and none of the other tested observables) are indeed positively correlated with being *sophisticated*. A simple way to visualize these correlations without the aid of modeling is to compare the histograms of the test scores for the $s_i = 0$ (*naïve*) and $s_i = 1$ (*sophisticated*) populations; these graphs provide evidence for the positive relationship between (s) and test scores, as shown in the figure below.

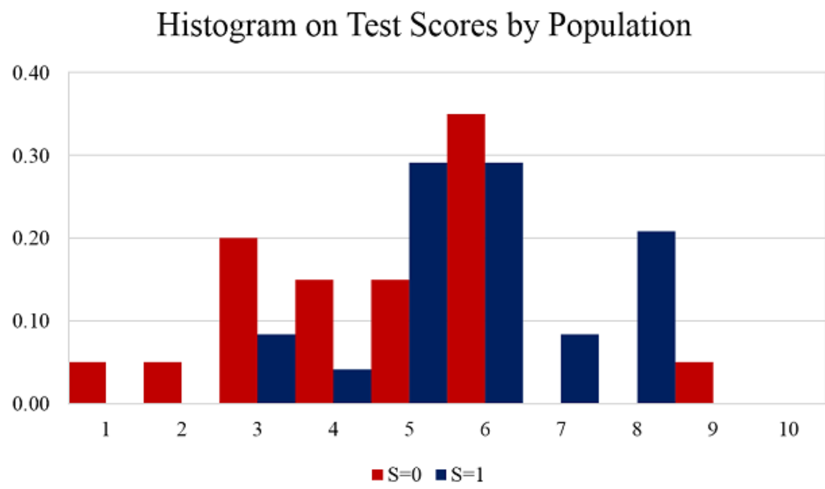

Figure 5: Naïve and Sophisticated Players Quantitative Test Results

4.4 Serial Correlation

All players' choices exhibit positive serial correlation from one round to the next, measured across the 80 rounds they played in the second session. Table 4.2.4 shows the serial correlation estimations for each player and stage (*p-values* in parenthesis). Even though the game is asymmetric, as discussed before, no obvious distinction can be made between the P_1 and P_2 populations in this regard; none exhibited systematically different measures of serial correlation. Interestingly, *naïve* players $(s_i = 1)$ played with a serial correlation closer to zero. This is perhaps more clearly seen by looking at the regression results below, and could be interpreted as an indication that *sophisticated* players' understood the game better and played closer to the corner solutions that maximized their payoffs. Also, since a negative serial correlation can be interpreted as an effect of trying to prevent being predictable, a positive serial correlation is probably related to playing against a population with bots, when being unpredictable is not considered as important.

Stage	P_1	P_2
1) $\tilde{\Gamma}$_25_50	0.204 (0.006)	0.283 (0.000)
2) Γ_25_50	0.378 (0.000)	0.360 (0.000)
3) Γ_50_50	0.392 (0.000)	0.476 (0.000)
4) Γ_50_0	0.309 (0.000)	0.310 (0.000)

Table II.4.2.4: Serial Correlation

$$\widehat{Choice_i} = \underset{(0.000)}{0.409} + \underset{(0.000)}{0.302} \; L.Choice_i - \underset{(0.000)}{0.098} \; s_i + \underset{(0.000)}{0.194} \; L.Choice_i \cdot s_i$$

5 Conclusions

In this article, we describe the design and results of a two-session experiment that included computer-controlled bots. We test and confirm the existence of players whose strategy mixing is persistently closer to uniform mixing than to pure or mixed NE (i.e. those relatively *naïve*), and those relatively *sophisticated*. Consistent methods to identify both types of players were developed and tested. Moreover, the probability of being naïve can be partially predicted by the score on a quantitative test, taken two weeks before the main games.

We also found evidence that the sophisticated population reacts better to off-equilibrium behavior in the theoretically-predicted direction but not with the magnitude required to restore equilibria. This confirms the stylized fact that NE are not played in complex games and, benefiting from the naïve component of populations' mixing, the direction of the deviations can thus be predicted. We employed this to design simple mechanisms to obtain above-equilibrium payoffs under these experimental conditions.

These results open up several research possibilities. From a theoretical perspective, we can foresee applied models where naïve players can be taken advantage of by players with a more moderate rationality bound. This way, a policy maker can potentially generate a Pareto-improvement in games with our postulated assumptions on rationality. Empirically, we can test the determinants that make a player be naïve relative to him/herself (as opposed to relative to the rest of the population) under different experimental conditions including varying game complexity and time constraints, belief elucidation, access to randomization devices, white noise, and information overload. Lastly, we conjecture that the methods described above to obtain above-equilibrium payoffs can be generalized to other normal and extensive form games.

APPENDICES

APPENDIX A, Instructions

Hello, and thank you for participating!

By coming here, you already earned $10 (ten US dollars). In a few moments, you will be given the opportunity to add to these earnings. The money you win will be paid to you at the end of the experiment. Every token you earn will earn you $5.00 (five US dollars). You will be playing simple games with a randomly chosen player from this room. This session will last an estimated 60 minutes. It will consist of 60 iterations of simple games and THREE of these will be randomly chosen to determine your total earnings.

Example: suppose you earn 1, 3 and 0 tokens in the three randomly chosen games. In this case, your total earnings will be $30 ($10 plus four times $5). Irrespective of your results, you will be invited to participate in a similar session in two weeks.

This will be the screen you and the rest of the participants will be seeing. You have two options: choose A or B. Likewise, the player you were paired with can pick either C or D. Your decision will be made before you know what the other player did. Suppose you press B and then it is revealed that the other player chose D. In this case you will earn 1 token and the other player will receive 0 tokens for this game. Note your payoff is represented in red and is the first number in each cell, and the other player's payoff is the second number, in green.

Further instructions will appear here. Please keep your attention on your own computer screen and stay silent throughout this experiment. If you have any questions, please raise your hand and ask the experiment administrators. PRESS OK TO START A SAMPLE ROUND.

APPENDIX B, Questionnaire and Test

Below is the questionnaire and the test that the participants answered after the sorting session. There was an emphasis on confidentiality. Players were identified through login names they created. The demographic variables (1-9) follow standard literature. One example is Fryer and Levitt (2005). The socioeconomic variables (10-13) were meant to be a proxy for income.

Questions $(1,2)$ from the test are the Linda Paradox and the Wason Selection Task, standard in the literature, used for instance in Charness and Sutter (2013).[26] Questions $(3,4)$ are the CNE mixing (last equation of the solution) in Γ and $\tilde{\Gamma}$ respectively. Questions $(5,10)$ are basically the same, and the answer is trivial; some of the information they provide is to be disregarded, given how the setup is worded. The purpose was to measure if there is a contradictory behavior in these two answers (answering correctly when playing against "bots" in videogames but incorrectly when facing a real opponent with strategic thinking) that reflected similar behavior when playing against computer bots. (No evidence for this kind of contradictory behavior was found.)

Finally, questions $(6-9)$ were taken from the Practice Book for the Paper-based GRE Revised General Test, Second Edition, by ETS (2012), available online. From the two Quantitative Reasoning sample exams, the questions with the highest percentage of examinees who answered correctly, not counting questions involving graphs, were chosen. High-percentage (their range was in 82% – 88%) questions were preferred since the population tested to get these percentages was mostly students trying to get into graduate school who have practiced specifically for the exam, whereas the participants of these experiments were students who had not yet finished their undergraduate education. Moreover, the aim was to choose GRE-type questions that would be correctly answered about 50% of the time (to maximize the score variance). Graph questions were disregarded because they are probably less related than other types of questions to the abilities we wanted to measure and, more importantly, these sections typically involve five questions on the same graph(s).

Questionnaire:

Thank you for participating!

[26] Other similar options include the *Cognitive Reflection Test* (CRT) by Frederick (2005).

Please fill out this survey and answer the test at the back of this page. As it has already been explained, the cash you have already earned depended only on your game results. Also, irrespective of these, you will be invited to a follow-up session one or two weeks from today with a higher expected payoff. You will have 20 minutes.

Please remember that all the information you provide will remain confidential. You do not need to write your name. These questions follow standard labor, education, development and health economics literature and are meant to identify your socioeconomic and demographic characteristics.

1) Gender: ☐ Male ☐ Female

2) Race (select one or more):

☐ American Indian or Alaskan Native ☐ Asian

☐ Black or African American ☐ Hispanic or Latino

☐ Native Hawaiian ☐ White

☐ Other

3) Age: _______ 4) Education (GPA): _______

5) Education (Major): _________ 6) Education (Current Semester): _______

7) Education Level (Father): _____ 8) Education Level (Mother): _________

9) Citizenship: ☐ U.S. Citizen ☐ Other (Specify): ________________

10) How would you classify your parents regarding income (lower, lower-middle, middle, upper-middle, or upper class)?

☐ Lower ☐ Lower Middle ☐ Middle ☐ Upper Middle ☐ Upper

11) Do you own a car? ☐ Yes ☐ No

12) Do you live alone or share? ☐ Alone ☐ Share

13) Have you traveled out of Michigan in the last six months?

☐ Yes ☐ No

14) How much money do you usually spend every week? US$_______

Write any name, word, number or combination of characters that will allow us to anonymously identify you in the next session (please don't forget it!). Examples: "Gandalf" , "789" , "GreenSpartan" , etc. _____________________________________

Test

You have 20 minutes to finish this test. Please select or write the best answer.

1) Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations. Which is more probable?

a. Linda is a bank teller.

b. Linda is a bank teller and is active in the feminist movement.

2) You are shown a set of four cards placed on a table, each of which has a number on one side and a colored patch on the other side. The visible faces of the cards show 3, 8, red and green. Which card(s) must you turn over to test the truth of the following proposition? *If a card shows an even number on one face, then its opposite face is red.*

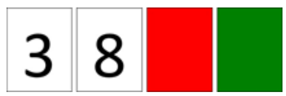

a. To be certain, you only need to turn over the 3 card.

b. To be certain, you only need to turn over the 8 card.

c. To be certain, you only need to turn over the red card.

d. To be certain, you only need to turn over the green card.

e. To be certain, you only need to turn over the 8 and 3 cards.

f. To be certain, you only need to turn over the 8 and red cards.

g. To be certain, you only need to turn over the 8 and green cards.

h. To be certain, you need to turn over all cards.

3) The solution for P in the equation 3P = 1 – P is:

P = __0.25__

4) If x = Q, and x = 1 – Q, then:

Q = __0.5__

5) You are going to shoot a penalty kick, and if the goalie does not guess the direction, you are going to score for sure. You have three options, shooting to the left, center or right of the goalie. He is a lefty, and you are sure he will jump to his left with a 35% probability, stay at the center with a 35% probability and jump to the right with a 30%

probability. You cannot shoot with your left foot so it is well known you shoot penalty kicks to the left of goalkeepers only 5% of the time, 5% to the center and 90% to their right. In what direction should you shoot?

a. To the left of the goalie

b. To the center

c. To the right of the goalie

6) At Company Y, the ratio of the number of female employees to the number of male employees is 3 to 2. If there are 150 female employees at the company, how many male employees are there at the company?

_______100______ male employees.

7) The floor space in a certain market is rented for $15 per 30 square feet for one day. In the market, Alice rented a rectangular floor space that measured 8 feet by 15 feet, and Betty rented a rectangular floor space that measured 15 feet by 20 feet. If each woman rented her floor space for one day, how much more did Betty pay than Alice?

a. $27

b. $36

c. $54

d. $90

e. $180

8) A business owner obtained a $6,000 loan at a simple annual interest rate of r percent to purchase a computer. After one year, the owner made a single payment of $6,840 to repay the loan, including the interest. What is the value of r?

a. 7.0

b. 8.4

c. 12.3

d. 14.0

e. 16.8

9) Working at their respective constant rates, machine I makes 240 copies in 8 minutes and machine II makes 240 copies in 5 minutes. At these rates, how many more copies does machine II make in 4 minutes than machine I makes in 6 minutes?

a. 10

b. 12

c. 15

d. 20

e. 24

10) In the final battle of a videogame you really want to beat, when Frankenstein attacks, you can parry, block or dodge. Damage is only prevented by picking the right defense, but he is so fast you have to choose without that information. The order of his attacks is random, but you know that parrying is correct slightly more often than blocking and that blocking is correct a little more often than dodging. So far you have been parrying more often than blocking and dodging, but only because you liked how it looked. What should you do next to try to beat him?

a. Parry.

b. Block.

c. Dodge.

APPENDIX C, Interface

Care was taken to diminish framing and other external effects. Apart from the inclusion of bots that publicly announced their prevalence and behavior, and game-order randomization as described above, the games were coded in zTree (Fischbacher, 2007) and had the following characteristics.

- Before the first (*sorting*) session, the players were presented with a brief tutorial describing the games. After the tutorial, a practice round was played. The second (*bots*) session included a new tutorial featuring examples which introduced the computer bots and another practice round. At the end of each session, all players were shown a table with their results.

- Regardless of whether they were P_1 or P_2, players always saw the games as a row player playing against a column player: P_2 type players saw a transposed version of the games. During each round, the players pressed one of two buttons labeled **A** and **B** to select their action. An **OK** confirmation button press was required. After each round, the results were showed using colors to highlight the players' and their opponents' actions, including a verbal description of the payoffs. The figure below shows a screen capture of the interface developed for this paper.

- During the invitations, tutorials and rounds, care was taken to use neutral and simple language like *other player* or *playing with* instead of *opponent* or *playing against.* Care was also taken to avoid technical terms.

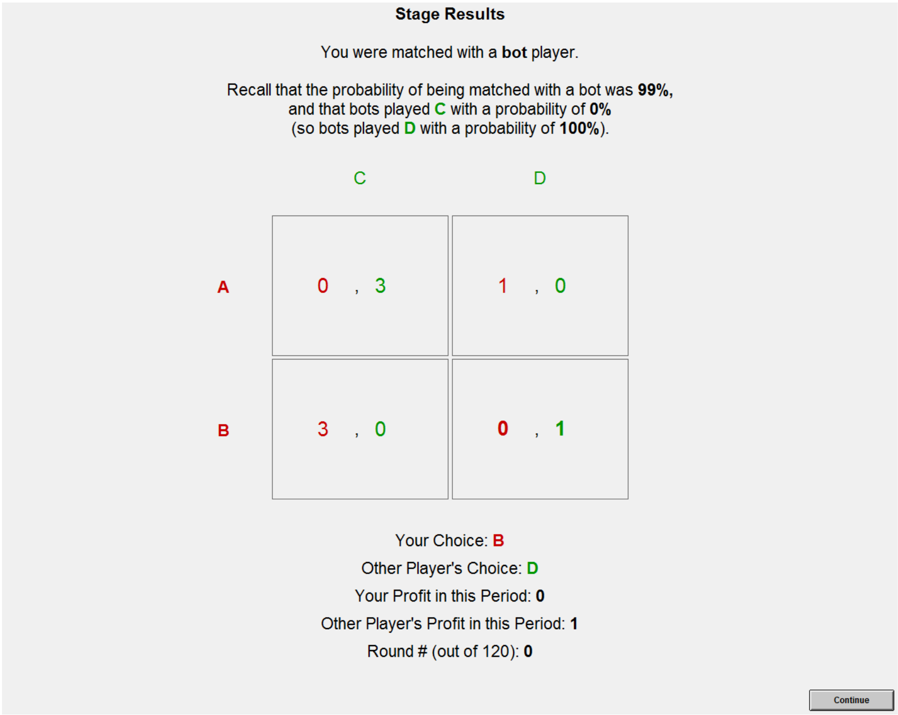

Figure II.A: Interface with Bots

APPENDIX D, Best Responses

Table A6 shows the best response of each population to the strategy of the respective opponent population (the naïve players from the P_1 population, NP_1 only played against NP_2, for instance). It refers to the best response of this population to the reported behavior of its relevant opponents. We consider it a corner solution in $\{0,1\}$ only if the opponent population played a strategy that was statistically different than the compensated NE (if the opponent's mixing was outside the 99% confidence interval).

	Player 1 (β)		**Player 2** (γ)	
Stage	**Naïve** (NP_1)	**Sophisticated** (SP_1)	**Naïve** (NP_2)	**Sophisticated** (SP_2)
1) $\tilde{\Gamma}$_25_50	Any	Any	Any	Any
2) Γ_25_50	0	0	1	1
3) Γ_50_50	0	0	1	1
4) Γ_50_0	0	0	1	1

Table II.A6: Best Responses

REFERENCES

REFERENCES

Alcocer, Christian and Thomas D. Jeitschko. "Using Rules-of-Thumb: Sophisticated vs. Simple Mixing in Two-Player Randomly Matched Games," *Working paper*, 2014.

Agnew, Julie, and Lisa R. Szykman. "Asset Allocation and Information Overload: The Influence of Information Display, Asset Choice, and Investor Experience," *The Journal of Behavioral Finance* 6(2): 2005: 57--70.

Aumann, Robert. "Correlated equilibrium as an expression of Bayesian rationality," *Econometrica*, 55(1):1-18, 1987.

Aumann, Robert, and Adam Brandenburger. "Epistemic conditions for Nash equilibrium," *Econometrica*, 63(5):1161-1180, 1995.

Benjamin, Daniel J., Sebastian A. Brown, and Jesse M. Shapiro. "Who Is 'Behavioral'? Cognitive Ability and Anom

alous Preferences," *Cognitive Ability and Anomalous Preferences*, 2006.

Bloomfield, Robert. "Learning a Mixed Strategy Equilibrium in the Laboratory," *Journal of Economic Behavior & Organization* 25(3), 1994: 411--36.

Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong. "A Cognitive Hierarchy Model of Games," *The Quarterly Journal of Economics*, 2004: 861--98.

Camerer, Colin, and Dan Lovallo. "Overconfidence and excess entry: An experimental approach." *The American Economic Review* 89.1 (1999): 306-318.

Cason, Timothy N., and Tridib Sharma. "Recommended Play and Correlated Equilibria: An Experimental Study," *Economic Theory* 33(1), 2007: 11--27.

Charness, Gary, and Matthias Sutter. "Groups Make Better Self-Interested Decisions," *The Journal of Economic Perspectives*, 2012: 157--76.

Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberri. "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications," *Journal of Economic Literature* 51(1), 2013: 5--62.

Dasgupta, Partha, and Eric Maskin. "The Existence of Equilibrium in Discontinuous Economic Games, I: Theory," *The Review of Economic Studies*, 1986: 1--26.

Duffy, Sean, and John Smith. "Cognitive load in the multi-player prisoner's dilemma game: Are there brains in games?." *Journal of Behavioral and Experimental Economics* 51 (2014): 47-56.

Epper, Thomas, Helga Fehr-Duda, and Adrian Bruhin. "Viewing the Future through a Warped Lens: Why Uncertainty Generates Hyperbolic Discounting," *Journal of Risk and Uncertainty* 43(3), 2011: 169--203.

Erlei, Mathias, and Heike Schenk-Mathes. "Bounded Rationality in Principal-Agent Relationships," *TUC Working Papers in Economics*, 2012.

ETS. *Practice Book for the Paper-based GRE Revised General Test*, Second Edition, 2012.

Feldman, Julian, "An analysis of predictive behavior in a two-choice situation," Unpublished Ph.D. dissertation, Carnegie Institute of Technology, 1959.

Fischbacher, Urs. "Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments," *Experimental Economics* 10(2), 2007: 171--78.

Frederick, Shane. "Cognitive Reflection and Decision Making," *Journal of Economic Perspectives*, 2005: 25--42.

Fryer, Roland G., and Steven D. Levitt. "The Black-White Test Score Gap through Third Grade. NBER Working Paper No. 11049," *National Bureau of Economic Research*, 2005.

Fudenberg, Drew, and Jean Tirole. *Game Theory*, Cambridge, Massachusetts, 393, 1991.

Gilad, Dalia, and Doron Kliger. "Priming the Risk Attitudes of Professionals in Financial Decision Making," *Review of Finance* 12(3), 2008: 567--86.

Gill, David, and Victoria L. Prowse."Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis." *Journal of Political Economy*, forthcoming (2014).

Grether, David M., and Louis Wilde. "Consumer Choice and Information," *Information Economics and Policy, Vol. 1*, 1983: 115---144.

Goeree, Jacob K., and Holt Charles A. "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions." *The American Economic Review* 91.5 (2001): 1402-422. Web.

Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey. "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory* 104(1), 2002: 247--72.

Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey. "Risk averse behavior in generalized matching pennies games." *Games and Economic Behavior* 45.1 (2003): 97-113.

Healy, Paul J. "Epistemic foundations for the failure of Nash equilibrium," Unpublished manuscript (2011).

Haptonstahl, Stephen R. "Competing Solutions to the Principal-Agent Model," Saint Louis: Washington University, 2009.

Harsanyi, John C. "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points," *International Journal of Game Theory* 2(1), 1973: 1--23.

Harstad, Ronald M., and Reinhard Selten. "Bounded-Rationality Models: Tasks to Become Intellectually Competitive," *Journal of Economic Literature* 51(2), 2013: 496--511.

Houser, Daniel, and Robert Kurzban. "Revisiting Kindness and Confusion in Public Goods Experiments," *American Economic Review*, 2002: 1062--69.

Kahneman, Daniel, and Amos Tversky. "Prospect Theory: An Analysis of Decision under Risk," *Econometrica: Journal of the Econometric Society*, 1979: 263--91.

Keynes, John Maynard. Chapter IV. The Principle of Indifference. *A Treatise on Probability*, 4. Macmillan and Co. 1921.

Koller, Daphne, and Nimrod Megiddo. "Finding Mixed Strategies with Small Supports in Extensive Form Games," *International Journal of Game Theory* 25(1), 1996: 73--92.

Koller, Daphne, Nimrod Megiddo, and Bernhard Von Stengel. "Fast Algorithms for Finding Randomized Strategies in Game Trees," *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*, 1994: 750--59.

Lattimore, Pamela K., Joanna R. Baker, and Ann D. Witte. "The Influence of Probability on Risky Choice: A Parametric Examination," *Journal of Economic Behavior & Organization* 17(3), 1992: 377--400.

Lipton, Richard J., Evangelos Markakis, and Aranyak Mehta. "Playing Large Games Using Simple Strategies," *Proceedings of the 4th ACM Conference on Electronic Commerce*, 2003: 36--41.

Mas-Colell, Andreu, Michael Dennis Whinston, Jerry R. Green. *Microeconomic Theory*. Vol. 1, Oxford university press New York, 1995.

Masiliunas, Aidas, Friederike Mengel, and J. Philipp Reiss. "Behavioral Variation in Tullock Contests," *Working Paper Series in Economics*, Karlsruher Institut für Technologie (KIT), 2014.

McFadden, Daniel. "Conditional Logit Analysis of Qualitative Choice Behavior," 1973.

McKelvey, Richard D., and Talbot Page. "Public and Private Information: An Experimental Study of Information Pooling," *Econometrica: Journal of the Econometric Society*, 1990: 1321--39.

McKelvey, Richard D., and Thomas R. Palfrey. "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics* 1(1), 1998: 9--41.

McKelvey, Richard D., and Thomas R. Palfrey. "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior* 10(1), 1995: 6--38.

Mann, Henry B.; Whitney, Donald R. " On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other," Annals of Mathematical Statistics 18(1): 50-60, 1947.

Misirlisoy, Erman, and Patrick Haggard. "Asymmetric Predictability and Cognitive Competition in Football Penalty Shootouts," *Current Biology* 24(16), 2014: 1918--22.

Mookherjee, Dilip, and Barry Sopher. "Learning Behavior in an Experimental Matching Pennies Game," *Games and Economic Behavior* 7(1), 1994: 62--91.

Nash, John. "Non-Cooperative Games," *Annals of Mathematics*, 1951: 286--95.

Nguyen, Quang, and PingSun Leung. "Do Fishermen Have Different Attitudes toward Risk? An Application of Prospect Theory to the Study of Vietnamese Fishermen," *Journal of Agricultural and Resource Economics*, 2009: 518--38.

Nyarko, Yaw, and Andrew Schotter. "An Experimental Study of Belief Learning Using Elicited Beliefs," Econometrica: *Journal of the Econometric Society* 70(3), 2002: 971--1005.

Ochs, Jack. "Games with Unique, Mixed Strategy Equilibria: An Experimental Study," *Games and Economic Behavior* 10(1), 1995: 202--17.

Ockenfels, Axel, and Reinhard Selten. "Impulse Balance Equilibrium and Feedback in First Price Auctions," *Games and Economic Behavior* 51(1), 2005: 155--70.

Offerman, Theo, Joep Sonnemans, and Arthur Schram. "Value Orientations, Expectations and Voluntary Contributions in Public Goods," *The Economic Journal*, 1996: 817--45.

Parkhurst, Gregory, Brandon Koford, and Therese Grijalva. "The influence of positive and negative framed information load: An experimental investigation," Working Paper, 2015.

Puhani, Patrick A. "The treatment effect, the cross difference, and the interaction term in nonlinear 'difference-in-differences' models," *Economics Letters*, 115(1), 2012: 85-87.

Rubinstein, Ariel. *Modeling Bounded Rationality*, Vol. 1. MIT Press, 1998.

Schram, Arthur, and Joep Sonnemans. "How individuals choose health insurance: An experimental analysis," *European Economic Review*, 55(6): 799-819, 2011.

Selten, Reinhard, and Chmura Thorsten. " Stationary Concepts for Experimental 2x2 Games." *The American Economic Review* 98.3 (2008): 938-66.

Shachat, Jason, and J. Todd Swarthout. "Do we detect and exploit mixed strategy play by opponents?" *Mathematical Methods of Operations Research* 59.3 (2004): 359-373.

Simon, Herbert A. "*Models of Bounded Rationality: Behavioral Economics and Business Organization, Vol. 2*," The Massachusetts Institute of Technology, 1982.

Stahl, Dale O., and Paul W. Wilson. "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior* 10(1), 1995: 218--54.

Van Damme, Eric. *Stability and Perfection of Nash Equilibria*. Vol. 339. Springer, 1991.

Chapter 3, Determinants of Naïve versus Sophisticated Mixing

Christian Diego Alcocer Argüello

Michigan State University[*,†]

2016-7-16

Abstract

Alcocer and Jeitschko (2014) postulate a behavioral bias in how individuals play mixed strategies. They define *naïve* players as those who, when indifferent, assign an off-equilibrium predetermined probability to each action. Alcocer and Shupp (2016) find evidence of the existence of these players: those whose strategy mixing is consistently closer to uniform mixing than the observed behavior of the rest of the population.

We now focus on individual responses to changing experimental conditions. We find evidence that there exist *distractors* (and *focusers*) that push players' towards (away) *naïve* mixing in matching pennies games. This allows for methods to take advantage of this bias and attain above-equilibrium payoffs.

Using computer bots, we also isolate altruistic components of players' strategies. We look at games where we previously found evidence of surplus-maximizing behavior that is different from equilibrium mixing. Adding a proportion of *transparent* bots that (ex-ante) do not incentivize any change in behavior but imply that surplus is wasted if they get any payoff, behavior gets closer to Nash equilibria.

1 Introduction

We investigate experimental conditions–*distractors* and *focusers*–that push players' behavior towards and away what Alcocer and Jeitschko (*AJ*, 2016) identify as either relatively *naïve* or *sophisticated* mixing. Our main analysis focuses on observing measures of how players' strategies respond to these controls and move away or closer to Nash equilibria (*NE*). This allows us to forecast the direction of this deviation in

[*]Keywords: Experimental, Behavioral, Bounded Rationality, Compensated Equilibrium, Computer Bots, Heuristics, Mixed Equilibria, Naïve and Sophisticated Players. JEL: C72, C91, D03, D83.

[†]We wish to thank Jon X. Eguia, Thomas D. Jeitschko, Arijit Mukherjee and Jeffrey M. Wooldridge for helpful comments.

games with totally-mixed equilibria. Doing so can allow one to design and test simple mechanisms that take advantage of this off-equilibrium mixing behavioral bias, and earn above-equilibrium payoffs. The evidence suggests that best-responding to naïve mixing yields above-equilibrium rents and that these rents increase (decrease) if the populations are distracted (focused).

Relatedly, we also isolate altruistic components of players' strategies by having populations randomly play against varying prevalences of computer players (bots) with varying publicly-announced strategies. We show that when bots' behavior does not directly incentive any particular response, it still pushes players towards equilibrium in cases where we have evidence of behavior that has altruistic components in that it maximizes total-surplus. We argue that like the identification of distractors and focusers, identification of altruism biases allows for the existence of mechanisms to take advantage of some predicted off-equilibrium mixing. Also, we discuss how these findings help explain some puzzling results from Alcocer and Shupp (AS, 2016).

Our results are consistent with the stylized fact that behavior is increasingly different than NE when dealing with increasingly complex games, in particular in those with unique, totally-mixed NE. This stylized fact has been one of the primary motivations for the theoretical literature on equilibrium extensions that allow non-NE strategies to have positive probabilities. However, the question of which mixed or pure strategies individuals actually follow remains open.

AJ postulate a behavioral bias or rationality bound in how people play mixed strategies. As an illustration, one can imagine an agent that can choose between two actions, left or right, knowing that given other the players' strategies, his expected payoff is the same in either case. According to neoclassical NE theory, the probability with which they will choose either of the two options depends on what they know about the other players' payoffs. Their paper relaxes the typical assumptions on rationality and proposes the theoretical existence of naïve players who lack any strategic depth and thus, following the principle of indifference (e.g., Keynes (1921); it is also called principle of insufficient reason in probability theory), will always flip a coin when indifferent.

The main result of AJ's paper is that, in finite games, if some proportion of the players is naïve (and follow some off-equilibrium mixing when indifferent), then expected final payoffs for all players are the same as in the standard perfect rationality case. The only assumption is that the percentage of naïve players is small enough otherwise equilibria do not generally exist. The sophisticated players essentially play off-equilibrium in response to the naïve players and the combined result is to bring us back to payoffs equivalent to the NE payoffs on average. This compensation is in the opposite direction of the distortion induced by the naïve mixing bias. This means the generalized game is an isomorphic setting to the NE since the convolution resulting from the linear combination of the mixing of naïve and the responding sophisticated players coincides with the mixing distribution of the standard perfect-rationality case.

AS use a set of laboratory experiments to test (and confirm) the existence of some naïve like players who consistently flip a coin, or mix close to 50%, in different settings. They develop consistent models to identify them within a population and also show that the probability of being naïve can be partially predicted by a simple quantitative test. They first sort participants (without their knowledge) into two groups based on behavior in a set of simple normal form 2×2 games with a unique and totally mixed NE (*asymmetric matching pennies* games): naïve and their sophisticated counterparts. They then have each group separately play against changing proportions of automated players (bots) that follow varying off-equilibrium mixed strategies. Overall, they find evidence of the existence of players that are relatively naïve and that the sophisticated players react in the direction predicted by AJ. Also, their analysis suggests that the probability of being naïve can be partially predicted by a simple quantitative test.

This paper builds on AS' initial work and uses similar experiments to determine; 1) if given experimental conditions can make individuals more or less naïve relative to themselves (analogous to within-estimation, we investigate what makes players more or less naïve relative to themselves; as opposed to analysis between populations, as in AS, who investigate how to identify players who are naïve relative to a population), and 2) if there is a social component in utilities that can be identified with the aid of bots. The hypothesis behind (1) is that players sometimes behave relatively close to coin flipping under an experimental control we label *distractor* and consists of adding weakly dominated strategies to matching pennies games, than when they face a different control we denominate *focuser* which consists of monotonically increasing payoffs such that equilibria are not modified. The hypothesis behind (2) is that when there is some chance players are matched against a bot, altruistic effects (utility gained by total surplus maximizing, even if it is another player who gets it) decrease and behavior moves towards what is predicted by utility maximization.

2 Literature Review

Mixed strategies are prevalent in everyday life. An illustrative example from sports occurs when in baseball a batter tries to outguess what will the pitcher throw next. They are also relevant in crime prevention situations where the location of an attack cannot be perfectly predicted because agents are mixing their strategies. In an epistemic game theory approach about NE with mixed strategies, players have unobservable prior beliefs about their own strategies, about players' beliefs, and so on (*hierarchies of beliefs*).[1] In two seminal papers, Aumann (1987) and Aumann and Brandenburger (1995), show that mutual belief in rationality and common knowledge of the game's mixed strategies and payoff functions are sufficient conditions for NE. Mixed strategies are not conscious randomizations, but self-fulfilling *conjectures* as to what other players will do. Consequently, if a NE fails to be observed, one of these assumptions failed too.

Healy (2011) links these epistemic foundations with the experimental evidence that populations tend not to play NE. By eliciting subjects' beliefs, he identifies the sources of failure to play the NE in five classic 2×2 games: asymmetric matching pennies, dominance solvable game (where the NE is not Pareto-dominated), prisoners' dilemma, symmetric coordination (battle of the sexes) and asymmetric coordination. He finds that the Aumann and Brandenburger (1995) assumption that generally fails is that players have imperfect beliefs about others' payoffs, even when the game clearly specifies them.

This result is close to the fundamental assumptions of *purification* theory (Harsanyi, 1973), *quantal response equilibrium* (McKelvey and Palfrey, 1995; McKelvey and Palfrey, 1998; Goeree, Holt and Palfrey, 2002) and *logit equilibrium* (McFadden, 1973) where mixed-strategy equilibria can be interpreted as the limit of pure-strategy equilibria in games when perturbations are added to payoffs.[2] Goeree, Holt and Palfrey (2003) and Selten and Chmura (2008) show that *quantal response equilibrium* is a better predictor than NE in several variants of 2×2 games with unique, mixed NE that they call asymmetric matching pennies.

In most of the experimental literature, players' behavior is significantly different than the strategies predicted by game theory (Selten and Chmura, 2008). Experimental evidence suggests that information increases are correlated to decision variance (Schram and Sonnemans, 2011). The impact of excessive information and cognitive loads varies across individuals (Swanson et al., 2011) and, for asset choices, is greater for those with less background in finances (Agnew and Szykman, 2005). Relatedly, Camerer and Lovallo (1999) run experiments of a game of simultaneous firm entry with congestion where the NE

[1] The problem behind belief formation of other players' strategies in mixed equilibria (let alone other players' beliefs) is known to be non-trivial. Feldman (1959) ran a series of experiments where individuals were shown sequences of zeros and ones. Even though the sequences were random, agents would try to come up with theories and heuristics to predict them; actually seeing patterns where they do not exist.

[2] An alternate family of theories often used to model off-equilibrium behavior is *level-k* thinking or *cognitive hierarchy* theory (Camerer, Ho and Chong, 2004; Van Damme, 1991).

is to enter with a non degenerate probability (i. e., it is a mixed-strategy NE). They find evidence on overconfidence: excess entry compared to the NE.

Asymmetric or generalized matching pennies, or, more generally, 2×2 games with unique, mixed NE are extensively analyzed in the literature. A well-known example that is very relevant to the discussion in the next sections is Goeree and Holt (2001). They confirm that populations do not play the NE but, interestingly and opposed to the theory (but perhaps not surprisingly), they find that changes in one player's payoffs in one outcome increase the probability that the related strategy is played. (Whereas in a mixed NE a change in one's payoffs only affects the other player's equilibrium strategy.) Shachat and Swarthout (2004) have individuals play asymmetric matching pennies games against computer bots and while it is confirmed that individuals generally detect deviations from Nash equilibrium and have an intuition of how to exploit them, consistently with the rest of the literature, players do not perfectly follow the resulting best responses.

Gill and Prowse (2014) also find a correlation between cognitive ability (measured by a Raven test), and economic rationality (measured in *level-k* terms). Parkhurst et al. (2015) find that when individuals face information overload, they tend to use simplifying heuristics (Grether and Wilde, 1983). When players have two options and are overloaded with information, they tend to mix uniformly mixing in an experimental setting. Similarly, Duffya and Smith (2014) find that when playing a prisoner's dilemma, players who are heavily distracted by having to memorize a 7 digit number play worse than those facing the smaller cognitive load of remembering a 2 digit number.

3 Experimental Design and Procedures

We implement an experiment consisting of three treatments designed to test if players' behavior varies predictably when playing in games that include what we denominate *distractors* (Treatment 1) and *focusers* (Treatment 2), and against computer-simulated players (bots) with varying behavior and prevalence (Treatment 3). At the end of each session, participants answer a simple quantitative test identical to that used in AS.[3] Each treatment involved two sessions (see Table 3 for details).

[3] The test and its discussion is included in an appendix.

Treatment	**Name**	**Rounds**	**Show Up**	**$ Per Token**	**Paying Rounds**
1	Focuser: Monotonic Increase	50	10	5	2
2	Distractor: Bordered Game	50	10	3	2
3	Bots: Altruism Measure	75	10	5	3

Table III.3: Experimental Treatments

Across the three treatments, all sessions involve the same basic setting. After arriving and logging on to the computer, half the players are randomly assigned as P_1 ("row" players) and half were designated as P_2 ("column" players). While we differentiate row and column players in our discussion and although the game is designed such that the NE is numerically the same for both, the asymmetric distinction between both is fundamental for the analysis below. That said, all players viewed the game as a row player on the computer whether they were a row player or not.

Players maintained their type throughout a session and played multiple rounds of the following matching pennies game along with the 25 to 50 rounds of the game variants described below. Note that, following AS (2016), in the context of this game naïve mixing is defined as occurring when players assign a probability statistically close to 50% to each available action. Similarly, sophisticated or non-naïve mixing is defined as occurring if play is closer to the NE: $\beta^* = \gamma^* = 3/4$. Players were anonymous and were randomly re-matched after each round. Each session included a tutorial and practice round.

	P_2			
		Left	**Right**	
P_1	**Up**	3, 0	0, 3	$(1-\beta)$
	Down	0, 1	1, 0	(β)
		$(1-\gamma)$	(γ)	

NE: $\beta^* = \gamma^* = 3/4$

Figure III.3: Base Game

3.1 Treatment Descriptions

3.1.1 Focuser: Monotonic Increase

The Focuser Treatment sessions include 25 rounds of the base game and 25 rounds of the monotonically increased game (see below) in random order. Note that under this control, payoffs are multiplied by three such that the *NE* remains the same $\left(\beta^* = \gamma^* = \frac{3}{4}\right)$.

The hypothesis is that players will consistently play closer to the NE (and away from simple 50-50 mixing) when they face a *monotonically expanded* game. The idea is straightforward: when there is more at stake, players concentrate more on their choices and their behavior gets closer to the one predicted by neoclassical theory (which assumes high rationality levels) than to naïve mixing or random noise.[4]

	Left	**Right**	
Up	9, 0	0, 3	$(1-\beta)$
Down	0, 9	3, 0	(β)
	$(1-\gamma)$	(γ)	

NE: $\beta^* = \gamma^* = \frac{3}{4}$

Figure III.3.1.1: Monotonically Increased Game

3.1.2 Distractor: Bordered Game

As in the Focuser Treatment, the Distractor Treatment sessions include 25 rounds of the base game and 25 rounds of a modified bordered game in random order. The bordered game adds a weakly-dominated pure

[4]Note that in a Bernoulli trial, variance is maximized when $p = 1/2$.

strategy to the base game such that the NE of this new game coincides with the NE of the base game.

As in previous research this strategy addition, while dominated, is expected to add a layer of complexity to the choice and thus potentially this cognitive load will push behavior away from the NE. The hypothesis is that when playing a *bordered* version of the base game, behavior will move towards naïve (coin-flipping) mixing. To be clear, our hypotheses do not only imply that distractors push mixing away from the NE. Throughout all experiments, we also cannot statistically reject the hypothesis that the direction of this distortion is predictable by AJ's results.

	Left	**Right**	**z'**	
Up	$3,0$	$0,3$	$0,0$	$(1-\alpha-\beta)$
Down	$0,1$	$1,0$	$0,0$	(β)
z	$0,0$	$0,0$	$0,0$	(α)
	$(1-\gamma-\delta)$	(γ)	(δ)	

NE: $\beta^* = \gamma^* = 3/4, \alpha^* = \delta^* = 0$

Figure III.3.1.2: Bordered Game

3.2 Bots Treatment

During the two sessions of the Bots Treatment, participants play the base game during 75 rounds. Twenty five of these rounds include no bots while 50 rounds do. In the 50 rounds with bots, the bots played the NE during 25 rounds and were coin-flippers (simulating *perfectly naïve*) for the other 25 rounds. Note that the order in which the 75 rounds were played was random and that the bot strategies and their prevalence (each player had a 25% probability of being matched against a bot during the bot rounds) were announced before each round and were public information.

Game Type	Number of Rounds	Bot Prevalence	Bot Randomization		NE or CNE	
			If P_1	If P_2	If P_1	If P_2
Base	25	0%	NA	NA	$\beta^* = 0.75$	$\gamma^* = 0.75$
2	25	25%	$\beta^B = 0.75$	$\gamma^B = 0.75$	$\beta^* = 0.75$	$\gamma^* = 0.75$
3	25	25%	$\beta^B = 0.50$	$\gamma^B = 0.50$	$\beta^* = 0.83$	$\gamma^* = 0.83$

Table III.3.2: Bots Treatment

The off-equilibrium strategy the bots used during the 25 coin-flipping rounds imply a distortion allowing us to observe several measures of how different players react to it. Specifically, these bots help us to isolate a potential social component of utilities that may impact play. The intuition behind this is that the base game is asymmetric so, from a behavioral perspective, it is harder for P_1 to compensate and react to the distortions created than for P_2. This is because it is P_1 who has the responsibility of determining the *size of the cak*e, whereas P_2 only determines *who gets the cake.* In this case, for P_1, surplus maximization and loss aversion (playing Up often) go in the opposite way of compensation (playing Down often).

3.2.1 Compensated Nash Equilibrium Predictions

Following AJ, Table 3.2 also shows the 'Compensated Nash Equilibrium' (*CNE*) for the rounds with off-equilibrium bots. The *CNE* is, in essence, the Nash equilibrium after taking into account the presence of players that are following publicly known off-equilibrium strategies. During all stages, it consists of unique and interior (purely mixed) strategies, implying that any deviation can be taken advantage of and collapses the opponent's best-response relation into a degenerate distribution.

For type 2 games in which the bots are playing the NE, their presence should have no impact on the *CNE*. As such, the *CNE* for type 2 games coincides with the NE of the base game. For the type 3 game, the bot population is *coin-flipping* (i.e., they play $\tilde{\beta} = \tilde{\gamma} = 1/2$). Since their prevalence is 25%, the unique *CNE* is to play $\beta^{S*} = \gamma^{S*} = 0.83$. The *CNE* has the characteristic that it predicts behavior that compensates by moving in the opposite direction of the distortion. That is, since human players know some automated players will mix with a probability that is less than the equilibrium probability ($\tilde{\beta} < \beta^{NE}$), then the equilibrium prediction is that human players compensate ($\beta^{NE} < \beta^{CNE}$).

4 Results

A total of one hundred and eight undergraduate students from Michigan State University participated in the experiment.[5] In an effort to maximize the cognitive effort of the participants and to remove potential wealth effects, two rounds (three in the longer bots sessions) were randomly selected at the end of the session for payment. Participants earned tokens (0 or 3) worth \$5 in the bots and distractor sessions, or (0, 3 or 9) worth

[5]The analysis is on a relatively narrow population. As such, the effects of differences in educational experiences on the determination of cognitive ability are diminished.

\$3 in the focuser sessions. All participants attended exactly one session.[6]

Treatment	Name	Participants	Avg. Earn	Avg. Tokens
1	Focuser	38	\$23.16	2.19
2	Distractor	34	\$18.40	0.84
3	Bots	36	\$21.38	0.76
	Total	**108**	**\$20.98**	**1.26**

Table III.4: Three Experimental Treatments, Main Statistics

4.1 Focuser: Monotonic Increase Game Results

As noted above, the Focuser Treatment tests if multiplying payoffs by three (a monotonic expansion that maintains the NE configuration of the base game: $\beta^* = \gamma^* = 3/4$), compels players to pay more attention to choices and thus play closer to the NE. The intuition is that when wins (or loses, if players think of earning zero as an opportunity cost) are greater, the incentives to take advantage of opponents' *mistakes* and to try and reduce their own increase too (we can consider off-equilibrium behavior as a mistake in that it allows opponents to extract rents). If this is true, then both player types (P_1 and P_2) should get closer to the NE. Moreover, the forecast is that this approach will be from below when comparing the base to the expanded game.[7] That is $\forall \Omega \in \{\beta, \gamma\} : \Omega^M < \Omega^E$, where Ω^M and Ω^E are each population's strategies in the base and expanded games.

Table 4.1 and Graph 4.1 summarize the results. First, we confirm one of AS' primary results: observed behavior is consistently between the NE and uniform (or naïve) mixing. While the results listed in Table 4.1 seem to confirm our hypotheses that $\Omega^M < \Omega^E$, we should proceed with caution because of the large standard errors which are probably related to the small sample size, relative to the weak effect we are measuring. As such, hypothesis testing results, are weak: the *p-values* of the one-side t-tests of the null hypotheses $\beta^M < \beta^E$ and $\gamma^M < \gamma^E$ are respectively 0.121 and 0.120.[8] And we cannot yet reject the null

[6] It was verified that players who self-selected into each session (by choosing the date they wished to participate) were not systematically different. Simple statistical tests found no evidence for obvious differences on behavior (means) or cognitive levels (test scores) between the three groups.

[7] Alcocer and Shupp (2016) confirm the stylized fact that players do not play the NE and show there is evidence of a bias such that players behavior in this kind of games is on average between 50% and the NE. Also, as will be discussed in the subsection on the results from the bots sessions, the observation that in all stages $\underline{\beta} < \overline{\gamma}$ is consistent with previous experimental results as well as being forecasted by our conjectures.

[8] Each choice is a Bernoulli trial and we focus on its parameter p. As such, our analysis is essentially non-parametric and our results coincide numerically with those from the Mann-Whitney U test. This is because the relevant estimators (i.e. differences between averages) are exactly the same in both cases. Furthermore, when doing hypothesis testing, their assumed distribution (normal) is also the same since both tests invoke the central limit theorem.

hypotheses that $\Omega^M = \Omega^E$. However, we can do better since the NE is the same for both player types. By pooling, we can almost reject at the 5% significance level (*p-value* = 0.051) the null that behavior is the same in both stages.

Treatment*	$\mathbf{P}_1$ Mixing ($\overline{\beta}$)	$\mathbf{P}_2$ Mixing ($\overline{\gamma}$)	Pooled $(\overline{\beta}+\overline{\gamma})/\mathbf{2}$
Base	0.497	0.655	0.576
(Std. Err.)	(0.023)	(0.022)	(0.016)
Expanded	0.535	0.690	0.613
(Std. Err.)	(0.023)	(0.021)	(0.016)
Difference	0.038	0.035	**0.037**
(Std. Err.)	(0.016)	(0.015)	(0.011)

*NE is 0.750

Table III.4.1: Focuser Treatment Results

The weakness of these results may indicate that other factors such as *endowment effects*, *loss aversion*, *tendency to avoid loses*, and other kinds of behavior might be involved that are not taken into account in simple utility maximization calculations. This, however, would not necessarily contradict our initial interpretation that different stakes are correlated with varying naiveté levels and that when players *focus* more on their actions, their behavior corresponds better to NE predictions. More generally, this is experimental evidence that there are simple experimental controls, that push players' behavior away from NE in the direction of naïve mixing.

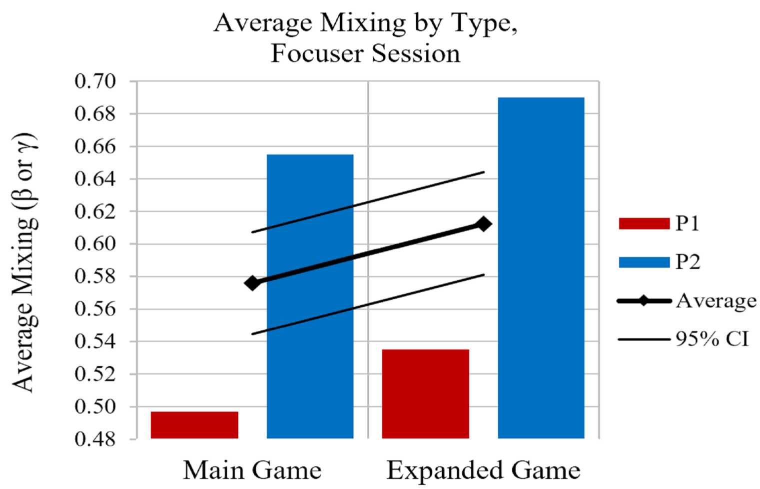

Figure III.4.1: Focuser Stage Results

4.2 Distractor: Bordered Game Results

While the focuser treatment is designed to investigate experimental conditions that may make players play closer to the NE, the distractor treatment aims to investigate experimental conditions, *distractors*, that can push players' behavior closer to uniform (or *naïve*) mixing. Specifically we hypothesize that if players are overloaded with unnecessary information, then they will tend to default to the simpler naïve mixing strategy. To test this, we randomly add a dominated action to the base game (what we call the *bordered* game) that acts as a distractor. While this addition leaves the NE effectively unchanged, we hypothesize that $\forall \Omega \in \{\beta, \gamma\} : \Omega^M > \Omega^E$, where Ω^M and Ω^B are each population's strategies in the base and bordered games.

Table 4.2 shows the results. These are not as clear-cut as those from the Focuser Treatment. Focusing first on columns 2 and 3, a straightforward t-test yields no statistical difference in behavior between bordered and not-bordered games. For both P_1 and P_2, estimated strategies are essentially the same when comparing the base and the bordered games. The respective *p-values* are 0.502 and 0.431.[9]

[9]Pooling the data from both the row and the column players did not help: $p\text{-}value = 0.306$.

Stage	P_1 Mixing ($\bar{\beta}$)	P_2 Mixing ($\bar{\gamma}$)	$\tilde{P}_1$ Mixing ($\tilde{\beta}$)	$\tilde{P}_2$ Mixing ($\tilde{\gamma}$)
Base	0.595	0.704	0.602	0.741
(Std. Err.)	(0.024)	(0.022)	(0.024)	(0.023)
Bordered	0.619	0.729	0.607	0.693
(Std. Err.)	(0.026)	(0.024)	(0.025)	(0.025)
Difference	0.024	0.026	0.005	**-0.048**
(Std. Err.)	(0.017)	(0.016)	(0.016)	(0.015)

Table III.4.2: Distractor Game Results

However, looking at the data more closely, we find that there are two outlier players (out of 34) that either did not understand the game or were not trying to maximize their payoffs. The behavior of these two outliers did not show any apparent pattern. One of them played the dominated action six times (out of 25) whereas the other played it five times. One of them was type P_1 and the other P_2. If we remove the two outliers, behavior for P_1 does not change ($p\text{-}value = 0.8875$) but for P_2 we have some evidence that it does, and in the right direction ($p\text{-}value = 0.0789$).

This provides some evidence that players get distracted when weakly dominated strategies are added to a game and that this distraction pushes their behavior towards uniform mixing. As we do with the results of the Focuser Treatment, we interpret this as confirmation (albeit weak) that under a distractor control, behavior can be pushed away from the NE. This is to be expected and coincides with the experimental literature. Our contribution is evidence that without these controls, behavior is distorted in the direction of naïve mixing (uniform mixing amongst all pure strategies included in the NE).

This treatment also tests whether naïve mixing can better forecast behavior than the equilibrium concepts of *purification*, *quantal response equilibria*, and *level-k thinking* which all place positive probability even on dominated strategies. As such, these concepts predict that the pure strategies **z** and **z'** will be played with some strictly positive probability whereas a naïve player is assumed to be able to (always) avoid a (conspicuously) dominated action and play either *z* or *z*' with zero probability.

The evidence is inconclusive. We want to verify if weakly dominated strategies are played with any strictly positive probability as this never happens under naïve mixing. After the removal of the two outliers, the estimated probability that the dominated action is played is 1.33%.[10] This is evidence against the behavioral concept of naïve mixing at the population level.

Still, it must be noted that it was 7 (out of 34) players that played (z) or (z') once (five of them) or twice (two players). This implies 79.4% of the players have an estimated behavior of exactly $\bar{\alpha} = 0$ (if P_1) or $\bar{\delta} = 0$ (if P_2). So, after looking at them individually, we cannot reject that these are partially naïve players even though the other 20.6% of the population are not.

4.3 Bots Treatment Results

When facing the base game, AS find that, even though the NE mixing is mathematically the same for both player types ($\beta^* = \gamma^* = 3/4$), P_2 type players' actual estimated behavior ($\bar{\gamma}$) is close to the NE whereas P_1 type players mixed closer to naïve mixing $(1/2)$. They also find that this behavioral distinction holds when playing games that include off-equilibrium bots. Why do the two player types' behavior differ?

We conjecture that when playing the base game, P_1 faces a problem that is behaviorally harder since, in a sense, he has the conflicting responsibility of keeping his opponent indifferent (heuristically, this is the condition to calculate totally-mixed equilibria) and maximizing total surplus. If he chooses *Up*, total surplus is three times greater than if he plays *Down* so, irrespective of what his opponent does, he might have an extra behavioral incentive to play *Up* or $\beta = 0$. On the other hand, P_2 types action only determines who gets this surplus.

If for any given round there is a publicly-known probability that one is matched with a bot–who gains no utility (i.e., will not earn any money)–then any altruistic and surplus-maximizing incentives are reduced in favor of straightforward utility maximization (i.e., keeping his opponent indifferent). That is, the expected altruistic utility diminishes if the probability of being matched with a human is reduced. Moreover, if these bots are playing the NE, they are *transparent* in that their presence does not imply a distortion that needs to be compensated if an equilibrium is to be restored. On the other hand, AJ show that *coin-flipping* bots like those included in Treatment 3 are expected to induce a compensation and the unique *CNE* is to play $\beta^{S*} = \gamma^{S*} = 0.83$.

If altruistic/total surplus maximizing incentives play a role in the difference between P_1 and P_2 behavior, then one should expect when comparing treatments 1 (no-bots) and 2 (transparent-bots) to find that P_1's behavior moves closer to the NE and our results support this. Denoting $\bar{\beta}^k$ as P_1's estimated behavior when

[10]No statistical test is needed since the null is that the mixing is *exactly* zero.

playing Stage k as described in Table 2, we get: $\overline{\beta}^1 = 0.57, \overline{\beta}^2 = 0.64$ and $\overline{\beta}^3 = 0.52$.[11] The crucial relation is that $\overline{\beta}^1 < \overline{\beta}^2$ and that they are statistically different. The *p-value* of the one-sided t-test that compares behavior with no bots and with NE bots is : 0.030 allowing the rejection of the null hypothesis $\beta^1 = \beta^2$.

If we restrict ourselves to game-theoretical tools like equilibrium concepts or utility-maximization criteria, the magnitude of the difference $(\overline{\beta}^2 - \overline{\beta}^1 > 0)$ is not as easy to interpret as its sign. This is because, in these games, players' mixing is only expected to adjust to their opponents' payoffs, not to their own payoffs since in an internal equilibrium they are indifferent amongst their pure strategies. We still interpret the finding of $\beta^1 < \beta^2$ as evidence that playing against NE-bots makes P_1 care less about the surplus impact and thus move closer to the NE. When he faces bots that play the NE, this altruistic/surplus-maximizing behavior diminishes and his strategy gets closer to the NE. Heuristically, P_1's incentive to play *Up* is reduced, and thus he simply plays it less. As such, $(\overline{\beta}^2 - \overline{\beta}^1)$ is a positive measure of his surplus-maximizing, social or altruistic preferences.

Consistent with AS and in general with the experimental literature on mixed strategies, no population's mixing is statistically equal to the equilibrium in any treatment reported here. Additionally, in all but one case $(\gamma^1 = 0.79)$[12] this bias is in the direction of naïve mixing $(1/2)$. Since P_1 was overplaying *Up*, every P_2 would have benefited by best-responding to playing *Left* with a probability of one $(\gamma = 0)$. Likewise, in all but Treatment 1, P_1's best-response was $\beta = 0$. This result, that it is optimal to assume one is matched with a naïve player in these games with unique and totally-mixed NE, is consistent throughout all the experiments.

5 Conclusions

Analyzing the behavioral bias of *naïve* mixing at the individual level and using computer players (bots) during experiments allows us to generate two sets of conclusions:

1. As a player, when facing games with a mixed equilibrium, one can follow two rules to maximize earnings. First, *distract* other players (or, equivalently, prevent them from *focusing*). Second, best-respond to naïve equilibrium mixing (i.e. to uniform mixing amongst all pure strategies included in

[11] About $\overline{\beta}^3$, it is worth noting that it is very similar to Alcocer and Shupp's (2016) findings. In their case, the equivalent calculation yielded $\overline{\beta}^3 = 0.59$. In general, our results were comparable to Alcocer and Shupp's (2016) and, when not related to our analysis, are not reported.

[12] The *p-value* related to the test for $\gamma^1 = \gamma^{1*}$ is 0.027.

the NE).[13] In this paper, we confirm that, in some kinds of games, populations' behavior is generally between the NE and naïve mixing, and we provide evidence that this distortion can be exacerbated by distractors and diminished by focusers.

2. As a researcher, one way to verify if altruism is preventing players from reaching the NE is to compare behavior with and without bots that play the NE. We find evidence that since bots that play the NE reduce surplus-maximization incentives (if there is some probability that one is matched with a bot and the bot gets all the utility in a round, then with some probability this utility is lost), these *transparent* bots function as focusers.

[13] When playing the base game from this paper, the best responses are the corner solutions $\beta = \gamma = 0$. Laboratory distractors could include switching the lights off or poking other players, but we cannot generally recommend that.

APPENDICES

APPENDIX A, Game Instructions

Hello, and thank you for participating!

By coming here, you already earned $10 (ten US dollars). In a few moments, you will be given the opportunity to add to these earnings. The money you win will be paid to you at the end of the experiment. Every token you earn will earn you $5.00 (five US dollars). You will be playing simple games with a randomly chosen player from this room. This session will last an estimated 60 minutes. It will consist of 50 rounds of simple games and THREE of these will be randomly chosen to determine your total earnings.

Example: suppose you earn 1, 3 and 0 tokens in the three randomly chosen games. In this case, your total earnings will be $30 ($10 plus four times $5).[14]

This will be the screen you and the rest of the participants will be seeing. You have two options: choose A or B. Likewise, the player you were paired with can pick either C or D. Your decision will be made before you know what the other player did. Suppose you press B and then it is revealed that the other player chose D. In this case you will earn 1 token and the other player will receive 0 tokens for this game. Note your payoff is represented in red and is the first number in each cell, and the other player's payoff is the second number, in green.

Further instructions will appear here. Please keep your attention on your own computer screen and stay silent throughout this experiment. If you have any questions, please raise your hand and ask the experiment administrators. PRESS OK TO START A SAMPLE ROUND.

[14]The number of rounds, dollars per token and number of paying rounds vary from treatment to treatment, as detailed in Section 3. This example was edited accordingly.

APPENDIX B, Questionnaire and Test

Below is the questionnaire and the test that the participants answered after each session. There was emphasis on confidentiality.

Questions $(1,2)$ are the Linda Paradox and the Wason Selection Task, standard in the literature, used for instance in Charness and Sutter (2013).[15] Questions $(3,4)$ are the *CNE* mixing (last equation of the solution) in the base game. Questions $(5,10)$ are basically the same and the answer is trivial; some of the information they provide is to be disregarded, given how the setup is worded. The purpose was to measure if there is a contradictory behavior in these two answers (answering correctly when playing against "bots" in videogames but incorrectly when facing a real opponent with strategic thinking) that reflected analogous behavior when playing against computer bots. (No evidence for this kind of contradictory behavior was found.)

Finally, questions $(6-9)$ were taken from the Practice Book for the Paper-based GRE Revised General Test, Second Edition, by ETS (2012), available online. From the two Quantitative Reasoning sample exams, the questions with the highest percentage of examinees who answered correctly, not counting questions involving graphs, were chosen. High-percentage (their range was in 82% – 88%) questions were preferred since the population tested to get these percentages was mostly students trying to get into graduate school who have practiced specifically for the exam, whereas the participants of these experiments were students who had not yet finished their undergraduate education. Moreover, the aim was to choose GRE-type questions that would be correctly answered about 50% of the time (to maximize the score variance). Graph questions were disregarded because they are probably less related than other types of questions to the abilities we wanted to measure and, more importantly, these sections typically involve five related questions.

Questionnaire:

Thank you for participating!

You have 20 minutes to finish this test. Please select or write the best answer.

1) Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was

[15] Other similar options include the *Cognitive Reflection Test* (CRT) by Frederick (2005).

deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations. Which is more probable?

a. Linda is a bank teller.

b. Linda is a bank teller and is active in the feminist movement.

2) You are shown a set of four cards placed on a table, each of which has a number on one side and a colored patch on the other side. The visible faces of the cards show 3, 8, red and green. Which card(s) must you turn over to test the truth of the following proposition? *If a card shows an even number on one face, then its opposite face is red*.

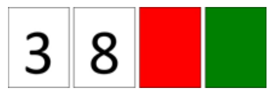

a. To be certain, you only need to turn over the 3 card.

b. To be certain, you only need to turn over the 8 card.

c. To be certain, you only need to turn over the red card.

d. To be certain, you only need to turn over the green card.

e. To be certain, you only need to turn over the 8 and 3 cards.

f. To be certain, you only need to turn over the 8 and red cards.

g. To be certain, you only need to turn over the 8 and green cards.

h. To be certain, you need to turn over all cards.

3) The solution for P in the equation 3P = 1 – P is:

P = __0.25__

4) If x = Q, and x = 1 – Q, then:

Q = __0.5__

5) You are going to shoot a penalty kick, and if the goalie does not guess the direction, you are going to score for sure. You have three options, shooting to the left, center or right of the goalie. He is a lefty, and you are sure he will jump to his left with a 35% probability, stay at the center with a 35% probability and jump to the right with a 30% probability. You cannot shoot with your left foot so it is well known you shoot penalty kicks to the left of goalkeepers only 5% of the time, 5% to the center and 90% to their right. In what direction should you shoot?

a. To the left of the goalie

b. To the center

c. To the right of the goalie

6) At Company Y, the ratio of the number of female employees to the number of male employees is 3 to 2. If there are 150 female employees at the company, how many male employees are there at the company?

_______100______ male employees.

7) The floor space in a certain market is rented for $15 per 30 square feet for one day. In the market, Alice rented a rectangular floor space that measured 8 feet by 15 feet, and Betty rented a rectangular floor space that measured 15 feet by 20 feet. If each woman rented her floor space for one day, how much more did Betty pay than Alice?

a. $27

b. $36

c. $54

d. $90

e. $180

8) A business owner obtained a $6,000 loan at a simple annual interest rate of r percent to purchase a computer. After one year, the owner made a single payment of $6,840 to repay the loan, including the interest. What is the value of r?

a. 7.0

b. 8.4

c. 12.3

d. 14.0

e. 16.8

9) Working at their respective constant rates, machine I makes 240 copies in 8 minutes and machine II makes 240 copies in 5 minutes. At these rates, how many more copies does machine II make in 4 minutes than machine I makes in 6 minutes?

a. 10

b. 12

c. 15

d. 20

e. 24

10) In the final battle of a videogame you really want to beat, when Frankenstein attacks, you can parry, block or dodge. Damage is only prevented by picking the right defense, but he is so fast you have to choose without that information. The order of his attacks is random, but you know that parrying is correct slightly more often than blocking and that blocking is correct a little more often than dodging. So far you have been parrying more often than blocking and dodging, but only because you liked how it looked. What should you do next to try to beat him?

a. Parry.

b. Block.

c. Dodge.

APPENDIX C, Interface

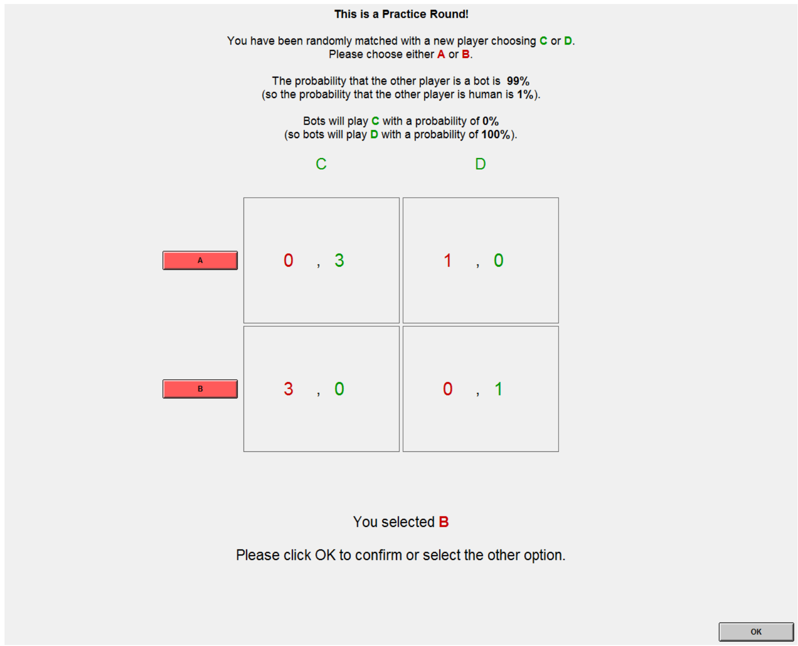

Figure III.A: Interface with Bots

Care was taken to diminish framing and other external effects. Apart from the inclusion of bots that publicly announced their prevalence and behavior, and game-order randomization as described above, the games were coded in zTree (Fischbacher, 2007) and had the following characteristics.

- Players were presented with a brief tutorial describing the games. The bots sessions included another tutorial featuring examples which introduced the computer bots. After the tutorial, a practice round was played. At the end of each session, all players were shown a table with their results.

- Regardless of whether they were P_1 or P_2, players always saw the games as a row player playing against a column player: P_2 type players saw a transposed version of the games. During each round, the players pressed one of two buttons labeled **A** and **B** to select their action. An **OK** confirmation button press was required. After each round, the results were showed using colors to highlight the players' and their opponents' actions, including a verbal description of the payoffs. The figure below shows a screen capture of the interface developed for this paper.

- During the invitations, tutorials and rounds, care was taken to use neutral and simple language like *other player* or *playing with* instead of *opponent* or *playing against.* Care was also taken to avoid technical terms.

REFERENCES

REFERENCES

Alcocer, Christian and Thomas D. Jeitschko. "Using Rules-of-Thumb: Sophisticated vs. Simple Mixing in Two-Player Randomly Matched Games," *Working paper*, 2014.

Alcocer, Christian and Robert Shupp. "Naïve versus Sophisticated Mixing: Experimental Evidence," *Working Paper*, 2014.

Agnew, Julie, Lisa R. Szykman. "Asset Allocation and Information Overload: The Influence of Information Display, Asset Choice, and Investor Experience," *The Journal of Behavioral Finance* 6(2): 2005: 57--70.

Aumann, Robert. "Correlated equilibrium as an expression of Bayesian rationality," *Econometrica*, 55(1):1-18, 1987.

Aumann, Robert, and Adam Brandenburger. "Epistemic conditions for Nash equilibrium," *Econometrica*, 63(5):1161-1180, 1995.

Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong. "A Cognitive Hierarchy Model of Games," *The Quarterly Journal of Economics*, 2004: 861--98.

Camerer, Colin, and Dan Lovallo. "Overconfidence and excess entry: An experimental approach." *The American Economic Review* 89.1 (1999): 306-318.

Duffy, Sean, and John Smith. "Cognitive load in the multi-player prisoner's dilemma game: Are there brains in games?." *Journal of Behavioral and Experimental Economics* 51 (2014): 47-56.

Feldman, Julian, "An analysis of predictive behavior in a two-choice situation," Unpublished Ph.D. dissertation, Carnegie Institute of Technology, 1959.

Gill, David, and Victoria L. Prowse."Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis." *Journal of Political Economy*, forthcoming (2014).

Grether, David M., and Louis Wilde. "Consumer Choice and Information," *Information Economics and Policy, Vol. 1*, 1983: 115---144.

Harsanyi, John C. "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points," *International Journal of Game Theory* 2(1), 1973: 1--23.

Healy, Paul J. "Epistemic foundations for the failure of Nash equilibrium," Unpublished manuscript (2011).

Goeree, Jacob K., and Holt Charles A. "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions." *The American Economic Review* 91.5 (2001): 1402-422. Web.

Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey. "Quantal Response Equilibrium and Overbidding in Private-Value Auctions," *Journal of Economic Theory* 104(1), 2002: 247--72.

Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey. "Risk averse behavior in generalized matching pennies games." *Games and Economic Behavior* 45.1 (2003): 97-113.

Mann, Henry B.; Whitney, Donald R. " On a Test of Whether One of Two Random Variables is Stochastically Larger than the Other," *Annals of Mathematical Statistics* 18(1): 50-60, 1947.

McFadden, Daniel. "Conditional Logit Analysis of Qualitative Choice Behavior," 1973.

McKelvey, Richard D., and Thomas R. Palfrey. "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics* 1(1), 1998: 9--41.

McKelvey, Richard D., and Thomas R. Palfrey. "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior* 10(1), 1995: 6--38.

Parkhurst, Gregory, Brandon Koford, and Therese Grijalva. "The influence of positive and negative framed information load: An experimental investigation," Working Paper, 2015.

Schram, Arthur, and Joep Sonnemans. "How individuals choose health insurance: An experimental analysis," *European Economic Review*, 55(6): 799-819, 2011.

Selten, Reinhard, and Chmura Thorsten. " Stationary Concepts for Experimental 2x2 Games." *The American Economic Review* 98.3 (2008): 938-66.

Shachat, Jason, and J. Todd Swarthout. "Do we detect and exploit mixed strategy play by opponents?" *Mathematical Methods of Operations Research* 59.3 (2004): 359-373.

Van Damme, Eric. *Stability and Perfection of Nash Equilibria*. Vol. 339. Springer, 1991.