

Naivete and Sophistication in Initial and Repeated Play in Games*

Bernardo García-Pola[†] Nagore Iriberry^{†‡}

July 7, 2022

Abstract

Compared to more sophisticated equilibrium theory, naive, non-equilibrium behavioral rules often better describe individuals' initial play in games. Additionally, in repeated play in games, when individuals have the opportunity to learn about their opponents' past behavior, learning models of different sophistication levels are successful in explaining how individuals modify their behavior in response to the provided information. How do subjects following different behavioral rules in initial play modify their behavior after learning about past behavior? This study links initial and repeated play in two different types of games (the 11-20 and 3×3 normal-form games) using a within-subject laboratory design. We classify individuals as following different behavioral rules in initial and repeated play and test whether and/or how strategic naivete and sophistication in initial play correlate with naivete and sophistication in repeated play. We find no evidence of a positive correlation between naivete and sophistication in initial and repeated play.

Keywords: Naivete, sophistication, strategic thinking, initial play, repeated play, level- k thinking, adaptive and sophisticated learning, mixture-of-types estimation

*This project was supported by a 2018 Leonardo Grant for Researchers and Cultural Creators, BBVA Foundation. The Foundation accepts no responsibility for the opinions, statements and contents included in the project and/or the results thereof, which are entirely the responsibility of the authors. We would also like to thank Vincent P. Crawford, Pedro Rey-Biel and seminar attendees of various presentations for their useful comments.

[†]AGORA Center for Market Design, Dept. of Economics, UNSW Business School (bernardo.garciapola@gmail.com).

[‡]University of the Basque Country, EHU-UPV, and IKERBASQUE, Basque Foundation for Research (nagore.iriberri@gmail.com)

1 Introduction

Nash equilibrium has been and still is the benchmark solution concept in game theory for predicting individual behavior in strategic environments. Since economics adopted the use of laboratory experiments, hundreds of experimental studies have tested whether individual behavior complies with the predictions of Nash equilibrium theory. These studies have shown that equilibrium theory has clear limitations in regard to its ability to describe how people behave in strategic environments. In response to ample experimental evidence, the important contributions of behavioral economics include models of bounded rationality that improve our understanding of how people actually behave in two different domains. First, when individuals make decisions for the first time with no previous experience or opportunity to learn, a scenario that is called initial play, naive, non-equilibrium, behavioral rules often outperform equilibrium theory in their ability to describe individual behavior (see for example, Goeree and Holt, 2001; Crawford et al., 2013).¹

Does behavior in initial responses relate in any way to behavior in repeated play in strategic environments? This is the central question of this paper.

In studies on initial play, models that explain how individuals begin playing games differ in their assumptions on the level of naivete or sophistication of individual thinking in strategic environments. We can order the behavioral rules in initial play from most naive to most sophisticated.² We propose that the most naive behavioral rules include processes that require no *strategic* thinking, meaning *no need* to predict an opponent's behavior, such that strategic settings are considered to be isomorphic to pure decision-making settings. For example, maxmax (optimistic) and maxmin (pessimistic)

¹For initial play, the crucial aspect is that individuals make decisions with no previous experience or opportunity to learn. This can include responses to one-shot games, as in Goeree and Holt (2001), or responses to multiple games that are similar but in which subjects are not provided with any feedback from game to game, as in, for example, Costa-Gomes et al. (2001). Second, given that people often do not start playing the Nash equilibrium strategy, bounded rationality models have been applied to repeated play to understand how people modify their behavior when provided with information on past behavior, that is, how individuals learn over time (see for example, Sobel, 2000).

²We use naivete and sophistication to refer to *strategic* naivete and sophistication, which can be different from *behavioral* naivete and sophistication. In other words, an individual showing behavior consistent with level-1 behavioral rule is naive regarding her revealed *strategic* sophistication but can be *behaviorally* the most sophisticated if all other opponents show random uniform behavior. This observation is related to work by Alaoui and Penta (2016) that tests whether individuals who show behavior consistent with a particular level- k thinking do so because of their own limitations or because of their beliefs about opponents' limited behavior.

behavioral rules fall into this category because maximizing over possible outcomes or maximizing over minimum possible outcomes does not require any ability to predict an opponent’s behavior. Level- k thinking models, which have been shown to be successful in explaining initial behavior in different settings (Stahl and Wilson, 1994, 1995; Nagel, 1995; Costa-Gomes et al., 2001; Camerer et al., 2004), illustrate different levels of strategic sophistication quite well. The level-1 behavioral type calculates the expected payoff associated with each of the available strategies, assuming that each of the opponent’s actions is equally likely, and takes the strategy with the highest expected payoff or, alternatively, sums her own payoffs across columns and takes the strategy that yields the maximum sum of payoffs. In the spirit of this latter interpretation, we also consider level-1 thinking to be a naive behavioral model.³ More sophisticated behavioral rules require individuals to best respond to some type of opponent behavior. Level-2 and level-3 represent assumptions of increasing sophistication about the opponent’s actions, as the level-2 type believes the opponent is behaving as a naive level-1 type and best responds to those beliefs, while level-3 type assumes that the opponent is behaving as a level-2 type and best responds to those beliefs. Finally, among the most sophisticated behavioral rules is the Nash equilibrium, which considers not only common knowledge of rationality but also rational expectations about beliefs.⁴

In studies focused on repeated play, models that explain how individuals modify their behavior in response to information on (their own and their opponent’s) past behavior also differ in strategic naivete or sophistication with respect to whether individuals use information on past behavior and, if they do how they use it. Learning models can also be ordered according to their sophistication level from most naive to most sophisticated in a hierarchical manner. We propose that the most naive learning model simply repeats the same strategy used in the past, having *no need* to use past strategies to predict the opponent’s strategy. We refer to this as the *No-Change* behavioral rule in repeated play. Adaptive learning models assume that individuals modify their behavior in response to information on past behavior, i.e., best responding to an opponent’s past behavior (illustrated best by fictitious play, as in Fudenberg and

³The cognitive hierarchy model (Camerer et al., 2004) assumes that level- k players best respond to combinations of existing lower levels. However, both level- k thinking and cognitive hierarchy models coincide in terms of their level-1 predictions.

⁴To be precise, according to Aumann and Brandenburger (1995), for two player games common knowledge of rationality is enough for equilibrium thinking.

Levine, 1998; Fudenberg et al., 1998). Note that adaptive learners assume that opponents indeed follow a *No-Change* type, as they assume that opponents will repeat the same strategy used in the past; therefore, adaptive learners will best respond to their opponents' past strategy. Finally, more sophisticated learning models assume that opponents indeed learn through an adaptive learning model and best respond accordingly (see, for example, Milgrom and Roberts, 1991; Selten, 1991; Conlisk, 1993a,b; Nagel, 1995; Camerer et al., 2002; Stahl, 2003).

Somewhat surprisingly, the studies on learning models (i.e., Cheung and Friedman, 1997; Erev and Roth, 1998; Fudenberg and Levine, 1998; Fudenberg et al., 1998; Camerer and Hua Ho, 1999) and on models to explain initial behavior (summarized in Crawford et al., 2013) have followed parallel paths.⁵ On the one hand, in studies on learning over time, initial play has been treated as a “black box”, an exogenous factor used only to initialize learning models, for example by estimating initial “attractions” associated with each of the particular strategies or, alternatively, simply assuming that initial “attractions” are the same across different strategies. On the other hand, models that aim to explain initial behavior have used mostly experimental designs that provide no feedback from game to game, precisely to suppress any opportunity to learn. Such models have been silent on learning over time.

However, it might seem natural that some type of relation exists between strategic behavior in initial and repeated play. As Costa-Gomes and Crawford (2006) note, modeling initial responses more precisely could yield insights into cognition that elucidate other forms of strategic behavior, such as learning and distinguishing between different levels of sophistication in rules and therefore influencing the implications for equilibrium selection and convergence. However, similar implications that seemed a priori intuitive have been empirically rejected (Costa-Gomes and Weizsäcker, 2008; Knoepfle et al., 2009). Gill and Prowse (2016) take a different approach by measuring cognitive ability exogenously with the Raven test; they then test whether more cognitively able subjects choose numbers closer to equilibrium, converge more frequently to equilibrium play and earn more. The question of whether the behavior in these two contexts is related is not only natural but also important. If such behavior is related, observing

⁵There are a few exceptions, as some models have been used to explain both initial behavior and learning behavior over time, such as the quantal response equilibrium model by McKelvey and Palfrey (1995), which simply estimates different noise levels or lambda-s for behavior in different stages.

the initial behavior of an individual would be informative about how her behavior will change and vice versa. Furthermore, this relation would allow a unified framework of behavior in games that incorporates both initial and repeated play (see, for example, Ho et al., forthcoming). If such behavior is not related, such that we observe very different levels of sophistication when the same individual faces a situation for the first time and in repeated play, characteristics that we sometimes measure as inherent to an individual, such as cognitive ability, may be more context dependent than previously believed.

We therefore address fundamental questions to propose a unified framework for studying initial and repeated play in games: How do strategic naivete and sophistication in initial play relate to naivete and sophistication in the use of information on past behavior in repeated play? Is a strategically naive player in initial responses, compared with a more sophisticated player, more likely to learn through a naive learning model in repeated play? We propose two laboratory experiments with mixture-of-types model econometric estimations to address these inherently empirical questions.⁶

We carry out two different experiments with similar designs with two different types of games, an 11-20 game, Arad and Rubinstein (2012), and 3×3 games. The experiments consist of two different parts. In the first part, the subjects played the games with no feedback, with our objective being to elicit their initial play (with no opportunity to learn or obtain experience). Subjects' behavior in the first part could not be affected by anything in the second part, as they did not know what they would do in the second part. Based on the subjects' decisions, we classified each subject as following one of multiple behavioral rules. For the 11-20 game, subjects could be easily classified in their sophistication with a single play of the game. However, this was not the case for the 3×3 games, so subjects in our experiment proceeded through 14 different 3×3 games (actually 7 asymmetric games, where the subjects played as both row and column players). Based on the subjects' profiles of 14 decisions, we classified each subject as following one of the multiple behavioral rules. This exercise is similar to those pioneered by Stahl and Wilson (1994, 1995) and later used by, for example,

⁶As the initial response can be equivalent to the response to one-shot games, our study can also be described as testing whether naivete/sophistication in behavior in one-shot games correlates in any way with naivete/sophistication in behavior in one-shot games when the individual is provided with information on past behavior. This is related to our design being able to capture the “initial model of learning”, as we refer to the objective of the repeated play part of the experiments that we carry out.

Costa-Gomes et al. (2001), Costa-Gomes and Crawford (2006), Rey-Biel (2009) and García-Pola et al. (2020).

In the second part of the experiment, subjects went over the same games again, but this time, in each of the games, they received information on both what they did and what their *current* opponent did in the first stage. For the 11-20 game, we implemented a third stage, in which subjects were provided with their own past strategy in the first part and asked to choose a strategy for each of the possible strategies that their current opponent could have chosen in the first part (a strategy method accounting for the information on opponent’s strategy in the first part). Using the subjects’ profiles of their decisions in repeated play and observed information on their own and current opponent’s past strategies in the first part, we classified each subject as following one of multiple behavioral rules in repeated play.

It is important to note that with this elicitation and identification of learning rules, we differ from studies that attempt to identify the ability of different learning rules to explain behavioral data (see, for example, Erev and Roth, 1998; Camerer and Hua Ho, 1999; Feltovich, 2000, and more recently, Kovářík et al., 2018). First, the learning models that we consider and identify vary in terms of which information individuals use (their own or their opponent’s) to modify past behavior and what individuals believe about how their opponents will use that same information on past behavior. Second, in our setting, for a particular game, subjects can learn about an opponent’s past actions just once, but we elicit how subjects learn from several different games or decisions based on their opponents’ past actions in those games. In other words, we elicit subjects’ learning rules using multiple different games or decisions in a way that does not allow the subjects themselves to evaluate the success of their learning model, which we refer to as the “initial model of learning”. These two important features considerably distinguish our approach to studying learning from existing work.

As this study is, as far as we know, the first empirical exercise to connect initial and repeated play, we use games that allow the highest separation among different behavioral rules in both initial and repeated play. The 11-20 game presents very good separation, as this game was specifically designed to separate different levels of strategic thinking, and we further designed the 3×3 games with this purpose in mind. The separation is the cornerstone for the use of a mixture-of-types model to identify and classify subjects into different behavioral rules in both initial and repeated

play. Finally, the within-subject design allows us to construct contingency tables to test whether naivete and sophistication in initial play are correlated with naivete and sophistication in repeated play. We find no evidence for a positive correlation between strategic naivete and sophistication in initial and repeated play.

Regarding initial behavior, consistent with previous findings, few Nash equilibrium players are found among the subject population. In the 11-20 game, the majority of subjects play 18 and 17 (corresponding with level-2 and level-3 sophisticated behavioral rules), with the two strategies accounting for 47% and 31% of the subject population, respectively. In the 3×3 games, we find that the majority of subjects, 60%, use a naive, non-strategic, behavioral rule. The second most frequent rule is a more sophisticated behavioral rule, level-2 type, which is used by 36% of the subjects. Furthermore, when we identify the behavioral rules that describe repeated play in both experiments, the majority of subjects, approximately 55%, show behavior consistent with adaptive learning, and an important number of subjects follow the most naive behavioral rule of ignoring their opponent's past action. Sophisticated learning models are more rarely used.

Most importantly and surprisingly, when we examine how naivete and sophistication compare between initial and repeated play, which is the central question of our study, we find little support for any positive correlation. In the 11-20 game, subjects are more likely to use adaptive learning models independently of the behavioral rule used in initial play. For the 3×3 games, subjects using a naive behavioral rule in initial play are, if anything, more likely to use a more sophisticated learning model than subjects using a more sophisticated model in initial play. In particular, 62% of individuals using a naive behavioral rule in initial play use an adaptive learning model, while only 45% of the subjects using a level-2 rule use an adaptive learning model.

The rest of this paper is organized as follows. Section 2 describes the theoretical framework defining the different behavioral rules and their classification into naive and sophisticated rules. Section 3 presents the experimental design for the empirical test. Section 4 contains the results, which are divided into the identification and classification of subjects according to their initial play, the identification and classification of subjects according to their repeated play, and the correlation between the naivete and sophistication displayed across the two settings. Finally, Section 5 concludes. Online Appendix A includes the econometric specification for mixture-of-types models, Online

Appendix B includes robustness checks on the experimental results, Online Appendix C includes additional tables, and finally Online Appendix D includes the experimental instructions.

2 Theoretical Framework: Naivete and Sophistication in Initial and Repeated Play

When analyzing initial play, we consider 8 behavioral types. We consider the leading behavioral models in the literature (Stahl and Wilson, 1994, 1995; Nagel, 1995; Costa-Gomes et al., 2001; Costa-Gomes and Crawford, 2006; García-Pola et al., 2020, among others).

The altruistic or social welfare maximizer type, A , simply sums her own and her opponent’s payoffs in each cell of the payoff matrix and applies the maxmax operator. The inequity averse type, IA , in a similar way, takes the absolute value of the difference between the her own and her opponent’s payoffs in each cell of the payoff matrix and applies the minmin operator. Although these two models resonate with interdependent preferences, which a priori are independent from models of strategic thinking, the actual naive implementation brings them close to a naive behavioral rule. The optimistic type ($MaxMax$) follows the strategy that results from applying the maxmax operator using only her own payoffs, while the pessimistic type ($MaxMin$) follows the strategy that results from applying the maxmin operator using only her own payoffs. The level-1 type ($L1$) sums her own payoffs across columns and takes the strategy that yields the maximum sum of her payoffs. The level-2 type ($L2$) expects her opponent to behave as a level-1 type and best responds to those beliefs. Level-3 and level-4 types ($L3$ and $L4$), similarly, expect their opponent to behave like level-2 and level-3 types, respectively, and best responds to those beliefs.⁷ Finally, NE players calculate the mutual best response required by equilibrium thinking.

How do we classify all these behavioral rules from the most naive to the most sophisticated? We take a simple approach and define the most naive behavioral rules as those that do *not need* to anticipate the opponent’s strategy, such that subjects following the most naive behavioral rule could treat strategic and pure decision-making

⁷Note that given that our games in Figure 1 do not have any dominated strategies, level- k rules and dominance- k rules, as defined in Costa-Gomes et al. (2001), coincide.

situations as isomorphic. Among the non-strategic behavioral rules, we include the altruistic or social welfare maximizer, inequity-averse, maxmax or optimistic, maxmin or pessimistic and level-1 types. Note that some of these behavioral rules can indeed be interpreted as individuals having beliefs and best responding to them (e.g., the level-1 or maxmin rule), but can also be interpreted as individuals simply following a naive behavioral rule as if they faced a pure decision-making setting (e.g., the level-1 type summing her own payoffs across columns or a maxmin type doing the maxmin operator over her own payoffs). As long as a behavioral rule does not need to anticipate the opponent’s strategy, we consider these behavioral rules to be non-strategic and naive. Once we define the most naive behavioral rule, we build on best response iterations to define higher levels of sophistication. On this basis, level-2 to level-4 rules are ordered immediately after the naive behavioral rules because a level-2 player anticipates that her opponent will behave like a level-1 type and best responds to those beliefs, while a level-4 anticipates that her opponent will behave like a level-3 type and best responds to those beliefs. Please see the robustness test in Section B.3.1 of the Online Appendix for a discussion of the similar hierarchical best response iterations for *A*, *IA* and *MaxMax* and *MaxMin*. Finally, the most sophisticated behavioral rule is the Nash equilibrium.

With respect to repeated play, we consider 4 main behavioral types. We again consider the leading behavioral models from the literature (Fudenberg and Levine, 1998; Fudenberg et al., 1998; Nagel, 1995; Camerer et al., 2002; Stahl, 2003).

The no-change type (*No-Change*) simply mimics the behavior undertaken in the first part of the experiment. Models based on adaptive learning theory (*Adaptive*) assume that individuals best respond and try to guess what their opponent will do (similarly to any belief-based learning model, as in Fudenberg and Levine (1998)). In our setting, as subjects are provided with their opponent’s past strategy, adaptive learning implies that the opponent will repeat her/his past strategy (that is, opponents are expected to follow a *No-Change* type and therefore, adaptive learners best respond to that behavior).⁸ When best responding to their opponents’ past behavior,

⁸Notice that in our repeated play setting, given that subjects are never provided with information on how successful their strategy in the first stage was, reinforcement learning (Erev and Roth, 1998) cannot be directly assessed. However, with a more flexible interpretation and assuming that subjects evaluate their own past strategy with their current opponent’s past strategy, reinforcement and adaptive learning models would predict the same strategy.

individuals may maximize their own payoffs (*Adaptive_S*), do maxmax over the sum of their own and their opponents' payoffs (*Adaptive_A*), or do minmin over the absolute difference between their own and their opponents' payoffs (*Adaptive_{IA}*). Sophisticated learning (*Sophisticated*) goes one step further and considers that her opponent follows adaptive learning behavior. Thus, the sophisticated learning rule uses her own past behavior, calculates the corresponding adaptive learning behavior (i.e., the opponent's best response to one's own past behavior), and then best responds to those beliefs regarding the opponent's expected behavior. Finally, we also consider one more degree of sophistication in repeated behavior (*Sophisticated 2*). The *Sophisticated 2* type assumes that her opponent will choose sophisticated learning behavior (i.e., choose the best response to her own behavior as an adaptive learner) and best responds to those beliefs. Note that all these behavioral types require not only a particular game to make predictions but also information on players' own and/or their opponents' past behavior, so the types are dependent on observed past behavior.

How should we classify all these behavioral rules from the most naive to the most sophisticated? We again take a simple approach and define the most naive behavioral rule in a repeated play setting as the one that does *not need* to use information on past strategies to have a model of how the opponent will behave. Among the learning rules that we consider, the *No-Change* rule is therefore the most naive, that is, the rule to simply repeat the strategy taken in the first stage. The rest of the behavioral rules build on this basis, increasing in sophistication as they take one additional step in the best response iteration on the use of information, such that the *Adaptive* rule is more sophisticated than the *No-Change* rule because adaptive learners are best responding to the *No-Change* rule and the *Sophisticated* rule is more sophisticated than the *Adaptive* rule because sophisticated learners are best responding to adaptive learners. Finally, the most sophisticated learning type, *Sophisticated 2*, assumes that her opponent is a sophisticated learner and best responds to those beliefs.

3 Experimental Procedures and Design

We carried out two independent experiments, one using the 11-20 game (199 subjects) and one using 3×3 normal-form games (198 subjects).

3.1 Procedures

Participants were recruited with the ORSEE system (Greiner, 2015). The sessions were conducted via computer with z-Tree software (Fischbacher, 2007). For the normal-form game experiment, two sessions with a total of 78 subjects were held in April and May 2019 in the Laboratory of Experimental Analysis (Bilbao LABEAN; <http://www.bilbaolabean.com>) at the University of the Basque Country, UPV/EHU. We conducted two additional sessions with the remaining 120 subjects in the Laboratory of Experimental Economics (LEE, <http://lee.uji.es>) at the University Jaume I of Castellón. For the 11-20 game experiment, five sessions were held in April 2022 in the Bilbao LABEAN.

The subjects were told that the experiment consisted of different parts and that payments would depend on both luck and their own and other subjects' decisions. Immediately before each part, subjects were given detailed instructions explaining the task involved, including examples of games, how they could make decisions, and how they would be matched and paid. Subjects were allowed to ask any questions that they might have during the presentation of the instructions. At the end of this presentation, the subjects were asked a few questions to guarantee that they had understood the instructions regarding each part. They could not start the experiment until they answered these questions correctly. A translated version of the instructions for the two experiments can be found in Online Appendix D.

For the 11-20 game, all subjects played the game three times with a different, randomly matched opponent each time. In the first part of the experiment, subjects played the 11-20 game with no feedback. In the second part, subjects were presented with the 11-20 game again, but this time, they were provided with information about their own past strategy and their current opponent's past strategy in the first part of the experiment. In the third part, subjects went over the 11-20 one last time, but this time they were provided with their own past strategy in the first part, and asked to choose a strategy for each of 10 possible strategies that their new opponent could have chosen in the first part (strategy method regarding opponent's strategy in the first part). We analyze individual behavior in parts 1 and 3 in Section 4.3 and relegate the analysis of the behavior in part 2 to the Section B in the Online Appendix.

For the 3×3 normal-form games, all subjects played the same seven games in the same order, first as the row player and then as the column player, playing a total of 14

games in each part.⁹ We did not inform the subjects that they were playing the same games in different roles: we showed all the games to all subjects from the perspective of row players. Subjects were randomly matched, such that within each part of the experiment, they were paired with a different opponent in each of the 14 games. In the first part of the experiment, subjects received no feedback from game to game so that we could elicit initial play in the 14 games. In the second part, subjects repeated the same 14 games in the same order but were provided with information about their own and their current opponent's past strategy in the first part of the experiment. The fact that subjects were provided with information on past actions in the second part was public knowledge, but they learned about the availability of this information *only* after they had finished the first part. In other words, the behavior in the first part of the experiment could not have been affected in any way by any of the experimental features in the second part. An example of how the games in both parts and the information provided in the second part were displayed in the experiment can be found in the instructions in the Online Appendix D.

When all subjects had submitted their choices in all parts, for each subject, the computer randomly chose one part from any of the three parts for payment in the 11-20 game, and two games from any of the two parts for payment in the normal-form game experiment. Thus, in the two experiments, each subject could be paid for different games. Before being paid, subjects completed a non-incentivized questionnaire on their demographic characteristics (gender, age, nationality, university entry grade and field of study), risk preferences following Eckel and Grossman (2002), and a cognitive reflection test. Descriptive statistics of all these variables can be found in Online Appendix Table A1. The subject pool showed characteristics typical of undergraduate students who are mostly studying for economics and business degrees, with a slightly higher presence of females, given that most were pursuing a degree in social sciences. We also requested free-format explanations of their choices and the expected choices of others in each of

⁹Any experiment using a within-subject design may raise concerns about potential experimenter demand effects. However, it is not clear to us how experimenter demand effects might have affected the results in our setting. On the one hand, individuals' natural taste for consistency in behavior (Eyster, 2003; Falk and Zimmermann, 2011) may have led subjects to choose the same behavior in both parts of the experiment. On the other hand, one might anticipate more pronounced reactions to the provided information if the subjects had identified our research question. The results clearly show that not all subjects repeated the same behavior and that not all subjects reacted to the provided information.

the parts of the experiment. We did not include these data in the analysis, but we did informally assessed the consistency between the subjects' explanations of what they did and the rule that we estimated using their elicited actions and frequently observed a clear alignment between the two. For work that attempts to relate subjects' free-format explanations of their actual actions and their actions, see Brañas Garza et al. (2011). Finally, we paid the subjects privately according to the two games selected plus a 3-euro show-up fee. The average payments were 21.11 euros and 15.76 euros, with standard deviations of 15.35 and 4.90, for the 11-20 and 3×3 normal-form game experiments, respectively. Each of the experiments lasted one hour and a half, including the presentation of instructions and payment.

3.2 Design of Games

In the 11-20 game, Arad and Rubinstein (2012), players choose numbers between 11 and 20. The chosen number is guaranteed as the payoff. Moreover, if a player's chosen number is exactly one number below the opponent's number, she earns the chosen number plus 80, and if the chosen number coincides exactly with the opponent's chosen number, then she earns the chosen number plus 10. This particular version is the one proposed by Alaoui and Penta (2016). The 11-20 game with such a large incentive to undercut has clear advantages for the study of naivete and sophistication in strategic thinking. First, it is straightforward to see that the lower the chosen number, the higher the iterative step in strategic thinking. Second, social-preference types of concerns are downplayed, as we confirm with the behavioral data.

In normal-form games, it is not as straightforward to match choices with strategic sophistication, so we designed seven 3×3 games for the experiment, as shown in Figure 1. We designed our own games instead of using games from other studies because we aimed to have high separation between different behavioral rules in both in initial and repeated play, which, as far as we know, was not the aim of any previous studies. The actual order in which the games were presented to the subjects was G1, G2... until G7 as row players, which we refer to as G11, G21, and so on until G71, and G1, G2... until G7 as column players, which we refer to as G12, G22, and so on until G72. As noted in the previous section, all subjects were shown the games as if they were row players, that is, we transposed the games when the subjects were playing as column players.

We chose this particular sequence, first as row players and then as column players, to prevent the subjects from realizing that they were making choices in the same games.

Figure 1: 3×3 Experimental Games

G1			G2		
4 , 20	20 , 12	18 , 2	6 , 18	22 , 4	4 , 16
6 , 8	8 , 14	22 , 16	20 , 6	2 , 24	16 , 4
18 , 14	14 , 6	2 , 18	12 , 12	2 , 6	18 , 22
G3			G4		
4 , 20	12 , 16	16 , 4	10 , 18	20 , 16	4 , 6
18 , 2	20 , 12	2 , 8	12 , 10	14 , 22	2 , 12
22 , 18	2 , 2	10 , 22	6 , 4	18 , 4	16 , 18
G5			G6		
8 , 16	16 , 14	20 , 12	14 , 16	2 , 20	12 , 22
16 , 8	18 , 12	4 , 4	6 , 18	20 , 4	10 , 6
14 , 6	16 , 4	2 , 20	22 , 4	14 , 18	4 , 10
G7					
4 , 20	22 , 14	18 , 4			
6 , 6	8 , 12	20 , 14			
18 , 16	14 , 8	4 , 18			

We chose 3×3 normal-form games because such games allow ample separation between the predictions of different behavioral rules. Note that with 14 3×3 games, there are 4,782,969 possible ways of playing the 14 games, while with 2×2 games, we would have *only* 16,384 possible combinations. Therefore, the use of 3×3 games substantially increases the a priori possibility of separation among the predictions of different behavioral rules. Additionally, we chose 3×3 games instead of, for example, 4×4 games to ensure that the number of strategies was relatively small such that it was easier to handle by subjects, which facilitated the explanation of the instructions.

3.3 Predictions in the 11-20 Game and 3×3 Normal-Form Games

In this section we review the predictions of all the behavioral rules that we consider in Section 2 for both initial responses and repeated play in the two types of games.

For the 11-20 games, the unique prediction for *NE* is to play 11. On the other hand, the prediction for *L1* is to play 19, for *L2* to play 18, for *L3* to play 17, and for *L4* to play 16. The higher the level-*k* is, the lower the number, and therefore, the closer

to *NE*. The predictions for the rest of the behavioral rules are concentrated on playing strategies 19 and 20. The prediction for the optimistic type is to play 19, while that for the pessimistic type is to play 20. The prediction for the altruistic type is indifferent between strategies 19 and 20, while the prediction for the inequity-averse type is that she would minimize differences in any of the strategy profiles in which both players coincide, although we assume that they would have a preference for coinciding at 20. See a summary of all these predictions in the first column of the top panel of Table 1.

For the 3×3 normal-form games, the last 14 columns in the top panel of Table 1 show the predictions for each of these behavioral rules in each of the games in Figure 1.

To predict strategies based on different learning models we need actual observed past behavior, so the bottom panel of Table 1, does not show the actual predicted strategies but, in general, the calculation that a behavioral rule requires in repeated play with the provided information, such that it is valid for both 11-20 and the 3×3 normal-form games.

Table 1: Predicted Strategies by Different Behavioral Rules: 11-20 Game and the 3×3 Normal-Form Games

	11-20	$G11$	$G12$	$G21$	$G22$	$G31$	$G32$	$G41$	$G42$	$G51$	$G52$	$G61$	$G62$	$G71$	$G72$
Initial Play															
A	19,20	2	3	3	3	3	1	1,2	2	1	3	1	3	1	2
IA	20	2	1	3	1	3	2	1,2,3	1,3	2	3	1	1	2	1
$MaxMax$	19	2	1	1	2	3	3	1	2	1	3	3	3	1	1
$MaxMin$	20	2	1	1	1	1	3	3	3	1	1	2	3	2	2
$L1$	19	1	1	2	3	2	1	3	2	1	3	3	2	1	1
$L2$	18	3	1	3	2	3	2	1	3	1	1	2	2	3	1
$L3$	17	3	3	1	3	2	3	3	1	2	1	2	1	3	3
NE	11	2	3	3	3	2	2	3	3	2	2	1	3	2	3
Repeated Play															
$No-Change$	“Same strategy as in the first part”														
$Adaptive_S$	“Best response to (opponent’s past strategy)”														
$Adaptive_A$	“A best response to (opponent’s past strategy)”														
$Adaptive_{IA}$	“IA best response to (opponent’s past strategy)”														
$Sophisticated$	“Best response to (opponent’s best response to (own past strategy))”														
$Sophisticated\ 2$	“Best response to (opponent’s best response to (best response to (opponent’s past strategy)))”														

Notes: The table reports the strategies predicted by the models for initial play (top panel) and repeated play (bottom panel). The numbers in the first column refer to the strategies in the 11-20 game. In the rest of the columns, 1, 2 and 3 refer to the first, second and third strategies in the 3×3 normal-form games, respectively. In a few instances, a behavioral rule is indifferent between multiple strategies, so we assume that the behavioral rule predicts any of these strategies with equal probability.

3.4 Assessment of the Design: Separation of Behavioral Rules across Games

On the one hand, as shown by the predictions described in Table 1, the 11-20 game is ideal for identifying and separating different level- k rules from the NE . On the other hand, the 3×3 normal-form games, shown in Figure 1, were carefully designed with the aim of yielding the largest separation between the predictions of different behavioral rules.

Table 2 shows the separation between the predictions corresponding to different behavioral rules for both initial play (panel A) and repeated play with information on

past actions (panel B) for the 3×3 games. Panel C shows the separation between the predictions corresponding to different behavioral rules for repeated play in the 11-20 game. The values in the table represent the proportion of decisions in which the predictions for two behavioral rules (the one in the row and the one in the column) are separated. The numbers can take any value between 0 (no separation at all, such that two behavioral rules predict exactly the same strategy in each of the decisions) and 1 (full separation, such that two behavioral rules predict a different strategy in all of the decisions).

Table 2: Separation of Different Behavioral Rules

Panel A: 3×3 Initial Play

	<i>A</i>	<i>IA</i>	<i>MaxMax</i>	<i>MaxMin</i>	<i>L1</i>	<i>L2</i>	<i>L3</i>
<i>A</i>	0.00						
<i>IA</i>	0.60	0.00					
<i>MaxMax</i>	0.46	0.62	0.00				
<i>MaxMin</i>	0.71	0.65	0.57	0.00			
<i>L1</i>	0.57	0.76	0.50	0.79	0.00		
<i>L2</i>	0.75	0.58	0.57	0.64	0.71	0.00	
<i>L3</i>	0.86	0.80	0.86	0.64	0.79	0.71	0.00
<i>NE</i>	0.57	0.51	0.86	0.64	0.79	0.79	0.57

Panel B: 3×3 Repeated Play

	<i>No Change</i>	<i>Adaptive_S</i>	<i>Adaptive_A</i>	<i>Adaptive_{IA}</i>	<i>Sophisticated</i>
<i>No Change</i>	0.00				
<i>Adaptive_S</i>	0.65	0.00			
<i>Adaptive_A</i>	0.60	0.52	0.00		
<i>Adaptive_{IA}</i>	0.62	0.81	0.64	0.00	
<i>Sophisticated</i>	0.71	0.60	0.71	0.62	0.00
<i>Sophisticated 2</i>	0.71	0.60	0.50	0.70	0.47

Panel C: 11-20 Repeated Play

	<i>No Change</i>	<i>Adaptive_S</i>	<i>Adaptive_A</i>	<i>Adaptive_{IA}</i>	<i>Sophisticated</i>
<i>No Change</i>	0.00				
<i>Adaptive_S</i>	0.90	0.00			
<i>Adaptive_A</i>	0.89	0.90	0.00		
<i>Adaptive_{IA}</i>	0.90	0.90	1.00	0.00	
<i>Sophisticated</i>	0.99	0.90	0.90	0.90	0.00
<i>Sophisticated 2</i>	0.95	0.80	1.00	0.90	0.99

Notes: The table reports the proportions of strategies for which the different behavioral models predict different strategies. The minimum possible separation value is 0, which occurs when the two models always prescribe the same strategy, and the maximum possible separation value is 1, which occurs when the two models always predict a different strategy.

In the 11-20 game, the separation values for initial play are equal to 1 for any level- k and *NE*. However, for all the naive behavioral models (*L1*, *A*, *IA*, *O* and

P) the separation is close to 0, as all of them are concentrated on strategies 19 and 20. To calculate the separation values for repeated play, we need to use past observed behavior. The separation values for repeated play are very high, as shown in panel C of Table 2.

In the 3×3 games the separation values for the initial play range between 0.46 and 0.86, which shows that each pair of behavioral rules is separated into at least 6 of 14 games and as many as 12 of 14 games. Regarding the separation values in repeated play, as for the 11-20 game, we could not calculate these values *ex ante*, as they depend on the particular observed past behavior of subjects.¹⁰ The values in panel B are therefore based on the actual observed behavior in the first part of the experiment. The values range between 0.47 and 0.81, which indicates that two behavioral rules for repeated play are separated into at least 6 of 14 games and as many as almost 12 of 14 games.

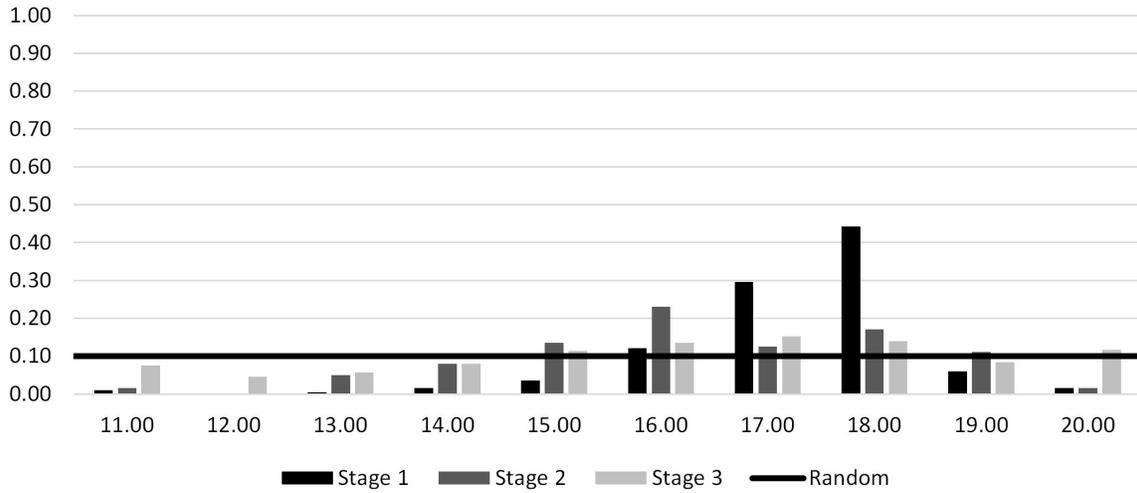
We therefore conclude that the goal of attaining large separation between the considered behavioral rules was achieved.

4 Results

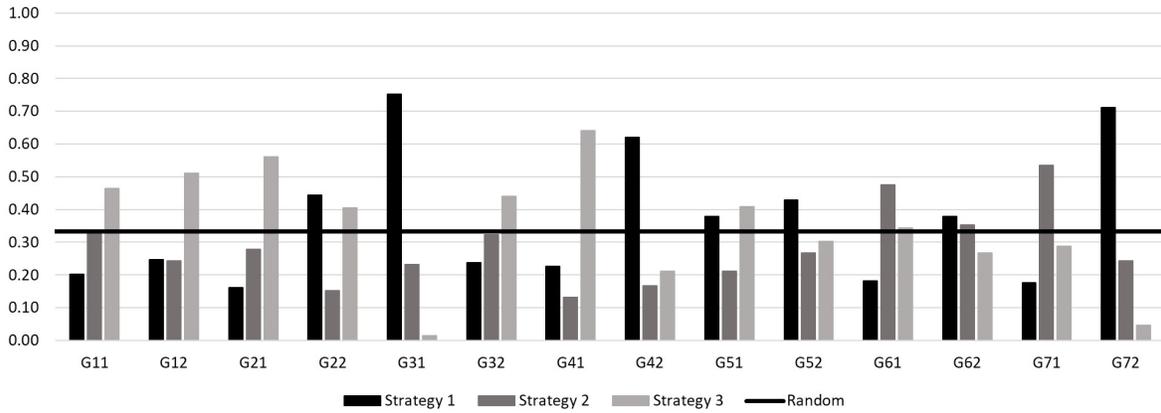
4.1 Descriptive Overview

We begin by considering the mean behavior in both initial and repeated play, which represents how individuals start playing in strategic environments with no feedback (first part) and how individuals react to both their own and their current opponent's past behavior (second part).

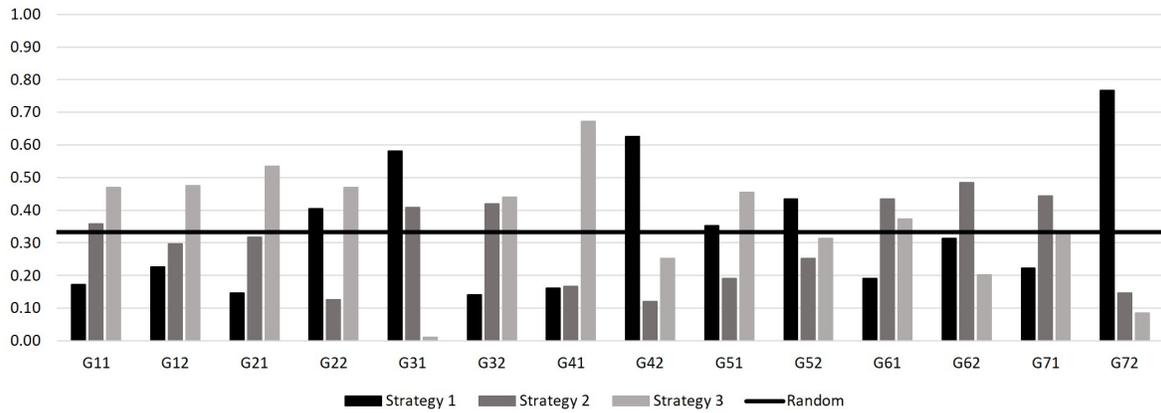
¹⁰However, we could use, as we actually did, the accumulated evidence from past studies (see Crawford et al., 2013, for example) that approximately half of subject populations show non-strategic behavior and a smaller proportion more sophisticated behavioral rules such as *L2* and *L3*, with a minority of subjects following the Nash equilibrium strategy. As the results in Section 4.2 show, we find a type distribution that is roughly consistent with existing findings in the literature.



(a) 11-20 Mean Behavior in Stages 1,2,3



(b) 3×3 Mean Behavior of Strategies 1,2,3 by Game: Initial Play



(c) 3×3 Mean Behavior of Strategies 1,2,3 by Game: I Repeated Play

Figure 2: Mean Behavior in Initial and Repeated Play

Figure 2 shows the results for the three stages of the 11-20 game (panel A), as well

as for the first (panel A) and second (panel B) parts of the 3×3 experiment.

Clearly, in the two experiments individual behavior is different from random play in both initial and repeated play; otherwise, we would observe that in each game, the 10 and 3 strategies are each played with equal probability (p -values less than 0.001 for both the first and second parts in the two experiments, based on a chi-square test against a uniform distribution).

In addition, for the 11-20 game we can reject that initial and repeated play are the same (p -values of 0.000 and 0.000, when comparing stages 1 and 2 and stages 1 and 3 are compared, respectively, from the two-sample chi-square test of the null that the two data samples come from the same distribution). We can also compare the repeated play in stages 2 and 3, where we can again reject that the behavior is the same (p -values of 0.002) However, for the 3×3 games, the mean behavior does not differ significantly between the first and second parts of the experiment, as we cannot reject that the behavior in both scenarios comes from the same distribution (p -value of 0.84 from the two-sample chi-square test that two data samples come from the same distribution), which may suggest that many subjects ignore the provided information on opponent's past behavior and follow the same strategy as in the first part. Note that, mean behavior can mask important differences with respect to individual heterogeneity. The key task in the next two subsections is to identify the relevant behavioral types that are able to reproduce the behavior in both parts of the experiment.

4.2 Naivete and Sophistication in Initial Play: Type Identification

Using the individual data on revealed choices from the first part of the experiment, we proceed to identify the behavioral type of each subject in initial play.

For the 11-20 game, we do not need any econometric model, as the behavioral types are readily inferred by their unique choices of numbers, as shown by panel A in Figure 2. Table 3, nevertheless, shows this type distribution. The most frequent type is $L2$, corresponding to 47% of the subject population, followed by $L3$, accounting for 31% of the subjects. The $L4$ behavioral type can be attributed to 13% of the subjects. Subjects adhering to the rest of the behavioral rules have a negligible presence.

For the 3×3 normal-form games, now necessarily using a mixture-of-types model

with uniform errors, we identify and classify each of the 198 subjects into a behavioral type.¹¹ The maximum likelihood function is estimated subject by subject. Please see Appendix A for a general description of the maximum likelihood function used to estimate behavioral types and for a particular derivation of the maximum likelihood function for estimating the behavioral types in initial play.

Table 3: Behavioral Type Identification for Initial Play: 11-20 Game and 3×3 Games

Model	11-20 Game	3×3 Games: Minimum Number of Perfect Guesses			
	(1)	No Constraints (2)	7 (3)	9 (4)	11 (6)
<i>Non-strategic</i>	0.08	0.56	0.57	0.51	0.62
<i>A</i>	–	0.14	0.14	0.10	0.14
<i>IA</i>	–	0.07	0.07	0.01	0.00
<i>MaxMax</i>	–	0.08	0.08	0.08	0.05
<i>MaxMin</i>	–	0.10	0.11	0.11	0.10
<i>L1</i>	–	0.17	0.18	0.21	0.33
<i>L2</i>	0.47	0.39	0.39	0.48	0.38
<i>L3</i>	0.31	0.02	0.02	0.00	0.00
<i>L4</i>	0.13	–	–	–	–
<i>NE</i>	0.01	0.02	0.02	0.01	0.00
No. of Subjects	188	174	169	91	21

Notes: The table displays the population frequencies estimated to be consistent with each of the behavioral rules listed in the Model column. In column 1, we simply report the type frequency for the 11-20 game. Note that we do not include subjects categorized in levels higher than $L4$, so we report the types for 188 subjects (out of 199) only. For the 3×3 games, we report different estimations for different numbers of perfect guesses: from all subjects (column 2) to subjects with 7, 9 and 11 (column 6) perfect guesses (alternatively, ε equal to 0.75, 0.53 and 0.32, respectively). We exclude subjects whose behavior is equally compatible with more than one type (ties, see footnote 12).

Table 3 shows the estimation results. We allow for different criteria on noise levels or alternatively perfect guesses, from 7 to 11 perfect guesses. We refer to a guess as perfect when a subject’s action coincides with a behavioral rule’s prediction. Note that, by chance, if individual play were random, any behavioral type that predicts a particular strategy combination across the 14 games would make 4.6 perfect guesses. Therefore, using this value as a benchmark, we consider both less and more stringent identification criteria: no constraints, at least 7, 9 and 11 perfect guesses (a 50%, 93% and 139% improvement over random, respectively). As expected, a trade-off exists between the number of perfect guesses required for identification and the number of subjects whom we can properly identify. Nevertheless, remarkably, when we impose

¹¹Alternatively, we could use the mixture-of-types model with logistic errors, as in for example Georganas et al. (2015). Estimation results are both quantitatively and qualitatively the same. These results are available upon request.

the criterion of 9 perfect guesses, which is a high threshold (a 93% improvement over random), we can identify 91 subjects.¹²

As observed in Table 3, when we focus on the overall population, in column 2, 56% of the subjects follow a non-strategic behavioral rule, followed by 39% who follow *L2*, while only a minority of subjects (4%) are identified as sophisticated (following *L3* and *NE*). Among the non-strategic behavioral types, the *L1* and *A* rules explain most of their behavior, followed by the pessimistic and optimistic behavioral rules. These results are roughly consistent with existing results, summarized in Crawford et al. (2013), although we find lower frequencies for *L1* and higher frequencies for *L2*. Furthermore, these conclusions do not change if we move across different columns (reflecting the criteria over the required perfect guesses). Only when we impose 11 correct guesses, for which we can only identify 21 subjects, do we find considerably more *L1* individuals to the detriment of the optimistic types. However, the overall conclusions remain unchanged: we still find that approximately 62% of the subject population is identified to follow a non-strategic behavioral rule, followed by 38% who follow *L2*. We cannot reject that the type distribution of the subjects does not depend on the constraints imposed regarding the number of perfect guesses (p -value of 0.22 for the chi-square test), so the estimation results are robust to different selections of criteria on the perfect guesses.¹³

In sum, for the 11-20 games, we observe mostly *L2* and *L3* behavioral types, while for the normal-form games, we observe mostly naive and *L2* behavioral types.

¹²Note that we do allow for the existence of the level-0 type in the estimation. When no model does better than random uniform, the estimated error would be equal to 1, which is interpreted as random uniform play describing best such subject's behavior. We find no subject who is best described as a level-0 type. Additionally, ties are possible between behavioral types, that is, when two behavioral types are equally good in describing a particular subject's action profile over the 14 games. We find 24, 17, 2, and 0 of those cases when we impose no constraints and when we impose the criteria of 7, 9 and 11 perfect guesses, respectively. When a tie occurred, we removed the subject from the analysis and therefore from Table 3 to avoid any potential bias.

¹³Despite subjects not receiving any feedback from game to game, it is still possible that they might learn to be more sophisticated as they play the 14 games. For robustness, we also estimated the type distribution using only the first half and only the second half of the 14 games. The estimated type frequency changes slightly, but we do not observe any increase in strategic sophistication from the first to the second half.

4.3 Naivete and Sophistication in Repeated Play with Information on Past Behavior: Type Identification

We now identify the learning model used by subjects by applying a mixture-of-types model with uniform errors to the individual data on revealed choices in repeated play for all subjects in each of the two experiments.

First, note that we consider the learning models, that is, the *No-Change*, *Adaptive*, *Sophisticated*, and *Sophisticated2* models, only for repeated play.¹⁴ Alternatively, we could re-estimate the same initial responses model as we did in Section 4.2, assuming that subjects ignore any provided information. See Section B.1 in the Online Appendix, where we show that on average learning models indeed do a better job in explaining the behavioral data on repeated play than the initial responses model. Second, in contrast to the case of initial play, for repeated play, note that for the two experiments we need to use the mixture-of-types model. The general description of the maximum likelihood function can be found in Online Appendix A, as can a particular derivation of the maximum likelihood function for estimating the behavioral types in repeated play. Additionally, for the 11-20 game, as we mentioned in Section 3.1 we use stage-3 data for repeated play. See the robustness test using stage-2 data in Section B.2 in the Online Appendix.

Table 4 shows the results for the 11-20 game (panel A) and for the 3×3 game (panel B) experiments. As in the first part of the experiment, we consider different criteria on noise levels, or alternatively perfect guesses, ranging from 5 to 7 (out of 10) perfect guesses for the 11-20 game and from 7 to 11 (out of 14) perfect guesses for the 3×3 games. Note that, by chance, if individual play were random, any behavioral type that predicts a particular strategy profile would make 1 and 4.6 perfect guesses in the 11-20 and 3×3 games. Therefore, using this value as a benchmark, we consider less stringent to more stringent criteria for identification of behavioral types. For the 11-20 game, we consider no constraints, and at least 5, 6 and 7 perfect guesses (corresponding to a 500%, 600% and 700% improvement over random, respectively). For the 3×3 games, we consider no constraints, and at least 7, 9 and 11 perfect guesses (corresponding to a 50%, 93% and 139% improvement over random, respectively). Again, a trade-off exists

¹⁴In addition, we have also considered an alternative naive learning model, which consists of choosing any of the available strategies after having discarded the one chosen in part 1. We find no empirical relevance for such naive learning model. Results are available upon request.

between the number of perfect guesses required for identification and the number of subjects that we can properly identify. However, the number of subjects whom we can cleanly identify is better than that in the first part. When we impose the criterion of 5, 6 or 7 perfect guesses in the 11-20 game, we now identify 121, 104 or 84 subjects, respectively. When we impose 7, 9 and 11 perfect guesses in the 3×3 games, we now identify 166, 117 and 46 subjects, respectively.¹⁵

For the 11-20 game, approximately 60% of the subjects follow the adaptive myopic best response, followed by the *No-Change* learning model. The presence of the *Sophisticated* learning model is residual. This identification of these types is also robust to the use of different criteria on perfect guesses (p -value of 0.18 for the chi-square test).

For the 3×3 games, the behavior of 34% of the subjects is best explained by the *No-Change* type, which reflects that an important number of subjects ignore opponent's past behavior and simply repeat their own past behavior. The most common behavior is adaptive behavior, which is displayed by 57% of the subjects, that is, those who choose the best response to their opponent's past behavior. Among the different adaptive learning types, the type that maximizes over the sum of her own and her opponent's payoffs seems to be the most frequent. Finally, very few subjects show sophisticated learning behavior. Consistent with previous findings, these conclusions do not change as we move across different columns. There is an exception when the highest threshold of 11 perfect guesses is imposed; the frequency of *No-Change* increases by 16 percentage points to the detriment of both the *Adaptive* and *Sophisticated* learning types. It is the result in this last column that, despite being marginally significant, makes us reject that the type distribution of the subjects varies based on the constraints imposed regarding the number of perfect guesses (p -value of 0.05 for the chi-square test).

In sum, in contrast to what we found for initial responses, for repeated play we find a more similar type distribution over the two games, with the majority of subjects

¹⁵As in the case when we identify behavioral models in initial play we allow for the existence of the level-0 type in the estimation. Additionally, we allow for ties between behavioral types, that is, two behavioral types that are equally good in describing a particular subject's action profile. For the 11-20 game, we find 29, 0, 0, and 0 of those cases when we impose no constraints, and when we impose the criteria of 5, 6 and 7 perfect guesses, respectively. For the 3×3 games, we find 22, 13, 3, and 0 of those cases when we impose no constraints, and when we impose the criteria of 7, 9 and 11 perfect guesses, respectively. When a tie occurred, we removed the subject from the analysis, and from Table 4 to avoid any potential bias.

Table 4: Behavioral Type Identification for Repeated Play

Panel A: 11-20 Game

Model	Minimum Number of Perfect Guesses			
	No Constraints (1)	5 (2)	6 (3)	7 (4)
<i>No-Change</i>	0.22	0.22	0.19	0.19
<i>Adaptive</i>	0.56	0.61	0.65	0.65
<i>Adaptive_S</i>	0.48	0.57	0.60	0.61
<i>Adaptive_A</i>	0.02	0.01	0.01	0.01
<i>Adaptive_{IA}</i>	0.07	0.03	0.04	0.02
<i>Sophisticated</i>	0.15	0.10	0.10	0.10
<i>Sophisticated 2</i>	0.07	0.07	0.06	0.06
No. of Subjects	165	116	101	88

Panel B: 3×3 Games

Model	Minimum Number of Perfect Guesses			
	No Constraints (1)	7 (2)	9 (3)	11 (4)
<i>No-Change</i>	0.34	0.35	0.41	0.50
<i>Adaptive</i>	0.57	0.56	0.54	0.45
<i>Adaptive_S</i>	0.16	0.17	0.21	0.28
<i>Adaptive_A</i>	0.28	0.28	0.28	0.15
<i>Adaptive_{IA}</i>	0.13	0.11	0.05	0.02
<i>Sophisticated</i>	0.06	0.06	0.02	0.04
<i>Sophisticated 2</i>	0.04	0.03	0.03	0.00
No. of Subjects	176	166	117	46

Notes: The table displays the population frequencies estimated to correspond to each of the behavioral rules listed in the Model column for 5, 6 and 7 perfect guesses (ϵ equal to 0.75, 0.53 and 0.32, respectively) for the 11-20 game and for 7, 9 and 11 perfect guesses (ϵ equal to 0.55, 0.44 and 0.33, respectively) for the 3×3 games. In both panels, we include only subjects who align with any behavioral model below *Sophisticated 2*, and we exclude subjects who are equally compatible with multiple behavioral models, see footnote 15.

following an adaptive learning model while most of the rest follows the very naive learning model of *No-Change*.

4.4 Correlation between Naivete and Sophistication in Initial and Repeated Play

We now study the central question of the paper, namely, whether there is a correlation between the type identification in initial and repeated play exploiting the fact that all subjects participated in the same two parts of the experiment. We use a contingency table in which the rows present the behavioral rules in initial play and the columns the behavioral rules in repeated play. Therefore, a particular cell in the contingency table shows the proportion of subjects identified as following the behavioral rule in that particular row in initial play who also follow the behavioral rule in that particular column in repeated play. The frequencies across the columns sum to 1 in each row. A positive correlation would show a higher frequency of naive, non-strategic behavioral types in initial play who use a *No-Change* or less sophisticated rule in repeated play than of level-2 or level-3 subjects, who in turn would show a higher frequency of adaptive or sophisticated learning. A no-correlation result would show independence in the distributions across different rows. A negative correlation would show that naive behavioral types in initial play use a more sophisticated learning model in repeated play than of more sophisticated type.

As observed in Tables 5a and 5b, for all subjects (panel A) or for a more precisely estimated sample of subjects (panel B), we see little evidence of a positive correlation between naivete and sophistication in initial and repeated play.

For the 11-20 game in Table 5a, the learning model followed by most subjects is the adaptive one, regardless of the model followed in initial play. *L2* and *L3* types show the most diverse classification in the repeated play models, being more likely to repeat their behavior but also to use more sophisticated rules. The few most sophisticated subjects in initial play (*L4* and *Nash* types) again show a clearer tendency to use adaptive learning models. Panel B shows the equivalent results for a reduced number of subjects when we impose the criterion of 6 out of 10 perfect guesses. In this case, the subjects show more consistency, and therefore, our identification of behavioral rules is improved despite the fact that we restrict the sample to 98 subjects. However, the

Table 5a: 11-20 Contingency Table

Panel A: No constraints

First Part Model	Second Part Model						No. of Subjects	
	<i>No-change</i>	<i>Adaptive</i>			<i>Soph</i>	<i>Soph 2</i>		
		<i>Adap_S</i>	<i>Adap_A</i>	<i>Adap_{IA}</i>				
<i>Non-strategic</i>	0.14	0.79	0.71	0.07	0.00	0.07	0.00	14
<i>L2</i>	0.25	0.54	0.45	0.01	0.07	0.15	0.06	71
<i>L3</i>	0.24	0.45	0.35	0.02	0.08	0.18	0.12	49
<i>L4</i>	0.16	0.58	0.58	0.00	0.00	0.21	0.05	19
<i>NE</i>	0.00	1.00	1.00	0.00	0.00	0.00	0.00	2
<i>No. of Subjects</i>	35	84	72	3	9	25	11	155

Panel B: Minimum of 6 correct guesses in each part

First Part Model	Second Part Model						No. of Subjects	
	<i>No-change</i>	<i>Adaptive</i>			<i>Soph</i>	<i>Soph 2</i>		
		<i>Adap_S</i>	<i>Adap_A</i>	<i>Adap_{IA}</i>				
<i>Non-strategic</i>	0.08	0.85	0.77	0.08	0.00	0.08	0.00	13
<i>L2</i>	0.22	0.64	0.60	0.00	0.04	0.11	0.02	45
<i>L3</i>	0.21	0.52	0.45	0.00	0.07	0.10	0.17	29
<i>L4</i>	0.20	0.70	0.70	0.00	0.00	0.10	0.00	10
<i>NE</i>	0.00	1.00	1.00	0.00	0.00	0.00	0.00	1
<i>No. of Subjects</i>	19	63	58	1	4	10	6	98

Notes: The table shows for each of the behavioral rules in initial play (by row) the proportion of subjects identified as following each of the behavioral rules in repeated play. For each row, the proportions across the four columns referring to the four main behavioral rules (*No-Change*, *Adaptive*, *Sophisticated* and *Sophisticated2*) should sum up to 1. Furthermore, for each row, the proportions across the three adaptive behavioral models (*Adaptive_S*, *Adaptive_A*, *Adaptive_S*) should sum to the value in the column for *Adaptive*.

Table 5b: 3×3 Contingency Table

Panel A: No constraints

First Part Model	Second Part Model						No. of Subjects	
	<i>No-change</i>	<i>Adaptive</i>				<i>Soph</i>		<i>Soph 2</i>
			<i>Adap_S</i>	<i>Adap_A</i>	<i>Adap_{IA}</i>			
<i>Non-strategic</i>	0.29	0.62	0.15	0.30	0.16	0.06	0.03	86
<i>A</i>	0.14	0.77	0.23	0.41	0.14	0.09	0.00	22
<i>IA</i>	0.27	0.55	0.18	0.27	0.09	0.00	0.18	11
<i>MaxMax</i>	0.15	0.77	0.15	0.31	0.31	0.00	0.08	13
<i>MaxMin</i>	0.57	0.43	0.14	0.07	0.21	0.00	0.00	14
<i>L1</i>	0.35	0.54	0.08	0.35	0.12	0.12	0.00	26
<i>L2</i>	0.45	0.45	0.18	0.23	0.05	0.06	0.03	62
<i>L3</i>	0.33	0.67	0.00	0.67	0.00	0.00	0.00	3
<i>NE</i>	0.33	0.67	0.00	0.67	0.00	0.00	0.00	2
<i>No. of Subjects</i>	54	85	24	42	19	9	5	153

Panel B: Minimum of 9 correct guesses in each part

First Part Model	Second Part Model						No. of Subjects	
	<i>No-Change</i>	<i>Adaptive</i>				<i>Soph</i>		<i>Soph 2</i>
			<i>Adap_S</i>	<i>Adap_A</i>	<i>Adap_{IA}</i>			
<i>Non-strategic</i>	0.38	0.63	0.17	0.42	0.04	0.00	0.00	24
<i>A</i>	0.17	0.83	0.17	0.67	0.00	0.00	0.00	6
<i>IA</i>	–	–	–	–	–	–	–	0
<i>MaxMax</i>	0.00	1.00	0.50	0.50	0.00	0.00	0.00	2
<i>MaxMin</i>	0.75	0.25	0.25	0.00	0.00	0.00	0.00	4
<i>L1</i>	0.42	0.58	0.08	0.42	0.08	0.00	0.00	12
<i>L2</i>	0.53	0.38	0.24	0.12	0.03	0.06	0.03	34
<i>L3</i>	–	–	–	–	–	–	–	0
<i>NE</i>	–	–	–	–	–	–	–	0
<i>No. of Subjects</i>	27	28	12	14	2	2	1	58

Notes: The table shows for each of the behavioral rules in initial play (by row) the proportion of subjects identified as following each of the behavioral rules in repeated play. For each row, the proportions across the four columns referring to the four main behavioral rules (*No-Change*, *Adaptive*, *Sophisticated* and *Sophisticated 2*) should sum up to 1. Furthermore, for each row, the proportions across the three adaptive behavioral models (*Adaptive_S*, *Adaptive_A*, *Adaptive_S*) should sum to the value in the column for *Adaptive*.

tendencies observed in panel B are similar to those in panel A. In either panel, we cannot reject that the distributions of the main types are independent across rows (p -values of 0.91 and 0.92, respectively, for the chi-square test).

Similar results are found for the 3×3 games in Table 5b. On the one hand, 62% of naive subjects in initial play follow an adaptive learning model, and 29% stick to their initial play. On the other hand, $L2$ subjects are equally likely to repeat their behavior and to follow an adaptive learning model. The few most sophisticated subjects in initial play ($L3$ and *Nash* types) do show a clear tendency to use adaptive learning models. If we focus on the two most frequent behavioral rules in initial play, naive and $L2$, these numbers show a clear absence of a positive correlation and, if anything, suggestive evidence of a negative correlation in naivete/sophistication between the initial and repeated play. Panel B shows the equivalent results for a reduced number of subjects when we impose the criterion of 9 out of 14 perfect guesses. In this case, the subjects show more consistency and thus our identification of behavioral rules is cleaner even though we restrict the sample to 58 subjects. However, the results regarding the correlation in panel B are very similar to those in panel A: 63% of naive subjects in initial play follow an adaptive learning model, and 38% of the $L2$ subjects use an adaptive learning model. If anything, the evidence of a negative correlation becomes even stronger. In neither panel can we reject that the distributions of the main types are independent across rows (p -values of 0.19 and 0.21, respectively, for the chi-square test).

We therefore conclude that there is no evidence of a positive correlation between naivete and sophistication in initial and repeated play.

In the Online Appendix, we carry out three additional robustness checks regarding the 3×3 experimental data. First, we include additional alternative behavioral rules in initial play; see section B.3.1. Second, we perform a specification test for omitted types in initial learning models; see section B.3.2. Third, and finally, we perform one additional specification test replacing the *No-Change* learning model with the initial responses model; see section B.3.3. All robustness checks lead us to the same conclusion: we find no evidence for a positive correlation between naivete and sophistication in initial and repeated play.

5 Discussion

In this paper, we have explored the relationship between the strategic sophistication and naivete of models in initial and repeated play. Is a strategically naive player in initial play more likely than a more sophisticated player to use a naive model in repeated play? We use an experimental design, based on two different types of games, 11-20 and 3×3 games, and a mixture-of-types model econometric estimation to answer this empirically motivated research question.

Consistent with previous findings, we find that the Nash equilibrium is not well suited to explaining the initial responses of individuals. The non-equilibrium rules that best explain individual behavior appear to be level-2 and level-3 rules in the 11-20 game, and level-2, level-1 and altruistic rules in normal-form 3×3 games. Additionally, consistent with previous findings, adaptive behavior appears to be the most common learning model in both the 11-20 and the 3×3 games, although a considerable number of individuals simply repeat their previously used strategy. Addressing the central question, and exploiting the within-subject design, we do not find any evidence that naivete and sophistication in repeated play are positively correlated with naivete and sophistication in initial play. The lack of positive correlation is robust to the alternative checks that we performed in the two experiments (contained in the Online Appendix).

The main result of our paper is reminiscent of the results of Costa-Gomes and Weizsäcker (2008) and Knoepfle et al. (2009). The former found an inconsistency between the behavior revealed by actions and elicited beliefs regarding opponents' expected behavior. The latter found that eye-tracking results favor much more sophisticated learning than do actual decision data, again indicating an inconsistency between the two. It could indeed be the case that, as in the case of actions and beliefs or actions and eye tracking, individuals treat initial and repeated play as different and/or independent tasks.

References

- Alaoui, L. and A. Penta (2016). Endogenous depth of reasoning. *The Review of Economic Studies* 83(4), 1297–1333. 2, 13
- Arad, A. and A. Rubinstein (2012). The 11-20 money request game: A level-k reasoning study. *American Economic Review* 102(7), 3561–73. 5, 13
- Aumann, R. and A. Brandenburger (1995). Epistemic conditions for nash equilibrium. *Econometrica: Journal of the Econometric Society*, 1161–1180.
- Brañas Garza, P., M. P. Espinosa, and P. Rey-Biel (2011). Travelers types. *Journal of Economic Behavior & Organization* 78(1-2), 25–36. 13
- Camerer, C. and T. Hua Ho (1999). Experience-weighted attraction learning in normal form games. *Econometrica* 67(4), 827–874. 4, 6
- Camerer, C. F., T.-H. Ho, and J.-K. Chong (2002). Sophisticated experience-weighted attraction learning and strategic teaching in repeated games. *Journal of Economic theory* 104(1), 137–188. 4, 9
- Camerer, C. F., T.-H. Ho, and J.-K. Chong (2004, aug). A cognitive hierarchy model of games. *The Quarterly Journal of Economics* 119(3), 861–898. 3
- Cheung, Y.-W. and D. Friedman (1997). Individual learning in normal form games: Some laboratory results. *Games and economic behavior* 19(1), 46–76. 4
- Conlisk, J. (1993a). Adaptation in games: Two solutions to the crawford puzzle. *Journal of Economic Behavior & Organization* 22(1), 25–50. 4
- Conlisk, J. (1993b). Adaptive tactics in games: Further solutions to the crawford puzzle. *Journal of Economic Behavior & Organization* 22(1), 51–68. 4
- Costa-Gomes, M., V. P. Crawford, and B. Broseta (2001, sep). Cognition and behavior in normal-form games: An experimental study. *Econometrica* 69(5), 1193–1235. 2, 3, 6, 8
- Costa-Gomes, M. A. and V. P. Crawford (2006, nov). Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review* 96(5), 1737–1768. 4, 6, 8, 41

- Costa-Gomes, M. A. and G. Weizsäcker (2008). Stated beliefs and play in normal-form games. *The Review of Economic Studies* 75(3), 729–762. 4, 31
- Crawford, V. P., M. A. Costa-Gomes, and N. Iriberri (2013, mar). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature* 51(1), 5–62. 2, 23
- Eckel, C. C. and P. J. Grossman (2002). Sex differences and statistical stereotyping in attitudes toward financial risk. *Evolution and human behavior* 23(4), 281–295. 12, 47
- Erev, I. and A. E. Roth (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American economic review*, 848–881. 4, 6, 9
- Eyster, E. (2003). Rationalising the past: A taste for consistency. 12
- Falk, A. and F. Zimmermann (2011). Preferences for consistency. 12
- Feltovich, N. (2000). Reinforcement-based vs. belief-based learning models in experimental asymmetric-information games. *Econometrica* 68(3), 605–641. 6
- Fischbacher, U. (2007, feb). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10(2), 171–178. 11
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic perspectives* 19(4), 25–42. 47
- Fudenberg, D., F. Drew, D. K. Levine, and D. K. Levine (1998). *The theory of learning in games*, Volume 2. MIT press. 4, 9
- Fudenberg, D. and D. Levine (1998). Learning in games. *European economic review* 42(3-5), 631–639. 3, 4, 9
- García-Pola, B., N. Iriberri, and J. Kovářík (2020). Non-equilibrium play in centipede games. *Games and Economic Behavior* 120, 391–433. 6, 8
- Georganas, S., P. J. Healy, and R. A. Weber (2015). On the persistence of strategic sophistication. *Journal of Economic Theory* 159, 369–400. 22
- Gill, D. and V. Prowse (2016, dec). Cognitive ability, character skills, and learning to play equilibrium: A level-kAnalysis. *Journal of Political Economy* 124(6), 1619–1676. 4

- Goeree, J. K. and C. A. Holt (2001, dec). Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review* 91(5), 1402–1422. 2
- Greiner, B. (2015, may). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association* 1(1), 114–125. 11
- Iriberry, N. and P. Rey-Biel (2013, nov). Elicited beliefs and social information in modified dictator games: What do dictators believe other dictators do? *Quantitative Economics* 4(3), 515–547. 36
- Knoepfle, D. T., J. T.-y. Wang, and C. F. Camerer (2009). Studying learning in games using eye-tracking. *Journal of the European Economic Association* 7(2-3), 388–398. 4, 31
- Kovářík, J., F. Mengel, and J. G. Romero (2018, mar). Learning in network games. *Quantitative Economics* 9(1), 85–139. 6
- McKelvey, R. D. and T. R. Palfrey (1995, jul). Quantal response equilibria for normal form games. *Games and Economic Behavior* 10(1), 6–38. 4
- Milgrom, P. and J. Roberts (1991). Adaptive and sophisticated learning in normal form games. *Games and Economic Behavior* 3(1), 82–100. 4
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review* 85(5), 1313–1326. 3, 4, 8, 9
- Rey-Biel, P. (2009). Equilibrium play and best response to (stated) beliefs in normal form games. *Games and Economic Behavior* 65(2), 572–585. 6
- Selten, R. (1991). Anticipatory learning in two-person games. In *Game equilibrium models I*, pp. 98–154. Springer. 4
- Sobel, J. (2000). Economists’ models of learning. *Journal of Economic theory* 94(2), 241–261. 2
- Stahl, D. O. (2003). Sophisticated learning and learning sophistication. *Available at SSRN 410921*. 4, 9
- Stahl, D. O. and P. W. Wilson (1994). Experimental evidence on players models of other players. *Journal of economic behavior & organization* 25(3), 309–327. 3, 5, 8

- Stahl, D. O. and P. W. Wilson (1995, jul). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior* 10(1), 218–254. 3, 5, 8
- Toplak, M. E., R. F. West, and K. E. Stanovich (2014). Assessing miserly information processing: An expansion of the cognitive reflection test. *Thinking & Reasoning* 20(2), 147–168. 47

A Mixture-of-Types Likelihood Function

We assume that a subject i employing rule k makes a type- k decision with probability $(1 - \varepsilon_i)$ but makes a mistake with probability $\varepsilon_i \in [0, 1]$. In such a case, she plays each of the three available strategies uniformly at random. As in most mixture-of-types model applications, we assume that the errors are identically and independently distributed across games and are subject-specific (as in, for example, Iriberri and Rey-Biel, 2013). The first assumption facilitates the statistical treatment of the data, while the second accounts for the fact that some subjects may be noisier and thus make more errors than others.

The likelihood of a particular individual corresponding to a particular type can be constructed as follows. Let $P_k^{g,j}$ be type k 's predicted choice probability for strategy j in game g . Some rules may predict more than one strategy in a particular game. This characteristic is reflected in the vector $P_k^g = (P_k^{g,1}, P_k^{g,2}, P_k^{g,3})$ with $\sum_j P_k^{g,j} = 1$.

For each individual in each game, we observe the choice and whether it is consistent with k . Let $x_i^{g,j} = 1$ if strategy j is chosen by subject i in game g in the experiment and $x_i^{g,j} = 0$ otherwise. The likelihood of observing a sample $x_i = (x_i^{g,j})_{g,j}$ given type k and subject i is then

$$L_i^k(\varepsilon_i | x_i) = \prod_g \prod_j \left[(1 - \varepsilon_i) P_k^{g,j} + \frac{\varepsilon_i}{3} \right]^{x_i^{g,j}}. \quad (1)$$

Finally, the likelihood function is given by the sum of all behavioral types that are considered.

$$L_i(\varepsilon_i | x_i) = \sum_k p_i L_i^k(\varepsilon_i | x_i) \quad (2)$$

p_i takes a value of 1 for the behavioral type k that best explains the individual behavior and 0 for the rest of the considered behavioral types.

To explain initial play, we consider $K = 8$ behavioral types or models: A , IA , $MaxMax$, $MaxMin$, $L1$, $L2$, $L3$ and NE , and use their revealed actions as input data. To explain repeated play with the information provided on past actions, we consider $K = 6$ different behavioral types: $No-Change$, $Adaptive_S$, $Adaptive_A$, $Adaptive_{IA}$, $Sophisticated$ and $Sophisticated 2$, and use their revealed actions and observed own and opponent's past action as input data.

B Robustness Checks

B.1 Do Learning Models Explain Repeated Play Better than the Initial Responses Model?

In this robustness check, we take the simplest approach to studying repeated play and re-estimate the models of initial responses without including the models of repeated play (*No-Change*, *Adaptive*, *Sophisticated* and *Sophisticated2*). We then compare the loglikelihood values under the two approaches.

For the 11-20 game, the loglikelihood values per subject are -18.51 for the models of initial responses models and -12.97 for the models of repeated play on average. In the same vein, the average perfect guesses per subject are 6.23 and 3.60, respectively.

For the 3×3 games, the loglikelihood values per subject are -13.11 for the models of initial responses models and -11.37 for the models of repeated play on average. The average perfect guesses per subject are 8.26 and 9.32, respectively.

We conclude that learning models that take into account how individuals use own and their opponent's past information are able to explain the repeated data better than the re-estimated models of initial responses.

B.2 Robustness for 11-20 Game

In the 11-20 game experiment, we included two repeated play stages, stages 2 and 3. Stage 2 is exactly equivalent to the repeated play in the 3×3 experiment. However, in contrast to the 3×3 experiment, we have only one unique decision to identify individuals' repeated play model. Having just one decision imposes important limitations on our ability to identify which learning model the subject is using. That is the reason why we included stage 3, where repeated play was elicited with the strategy method in the provision of information regarding the play in the first stage. In this robustness check we repeat the exercise shown in Section 4.4 but with stage 2 data. Given that we found little evidence for *Adaptive_A* and *Adaptive_{IA}* in this robustness check we include only *Adaptive_S*. We show that despite the estimated learning model type being different, the main finding of no correlation is maintained.

Table 7 shows the separation values when we use stage-2 data. Despite these separation values being relatively high, albeit lower than when we use stage-3 repeated play, ties are very frequent when we use stage-2 repeated play. The main reason for this is that stage-2 data use one unique decision to identify subjects' learning model, such

that it is enough that two models are fully confounded for not being able to identify the learning model the subject is using. Additionally, the removal of ties in the case of this analysis creates a clear bias: ties are more frequent with some learning models depending on the stage-1 data. For example, subjects who chose 20 in stage-1 and are *No-Change* in stage-2 will never be involve a tie. In contrast, if the same subject was *Adaptive_S* in stage 2, the subject may be removed from the analysis due to a tie if she is playing against a subject who chose 19 in stage 1 (if she ended up choosing 18, there would be a tie between *Adaptive* and *Sophisticated*). To solve this issue, we add an additional analysis, stage 2 unbiased, repeating the exercise but correcting stage-2 data with the average frequency of these ties.

Table 7: Separation of Different Behavioral Rules: 11-20 Game Stage 2

	<i>No Change</i>	<i>Adaptive_S</i>	<i>Adaptive_A</i>	<i>Adaptive_{IA}</i>	<i>Sophisticated</i>
<i>No Change</i>	0.00				
<i>Adaptive_S</i>	0.79	0.00			
<i>Adaptive_A</i>	0.80	0.98	0.00		
<i>Adaptive_{IA}</i>	0.67	0.99	1.00	0.00	
<i>Sophisticated</i>	0.99	0.80	0.97	0.91	0.00
<i>Sophisticated 2</i>	0.97	0.99	1.00	0.99	0.79

Notes: The table reports the proportions of strategies in which the different behavioral models predict different strategies. The minimum possible separation value is 0, which occurs when the two models always prescribe the same strategy, and the maximum possible separation value is 1, which occurs when the two models always predict a different strategy.

Table 8 shows the estimation results. We find that most subjects follow the *No – Change* rule repeating their past strategy and that the next largest group follows the *Sophisticated* rule. The presence of the adaptive myopic best response is much smaller than in the stage-3 analysis, being the least followed rule of the four.

Finally, Table 9 shows the correlation results with the stage-2 data. We find no evidence of positive correlation between naivete and sophistication in initial and repeated play. In neither panel can we reject that the distributions of the main types are independent across rows (p -values of 0.39 and 0.13, respectively, for the chi-square test).

Table 8: Behavioral Type Identification for Repeated Play: 11-20 Game Stage 2

Model	Stage 2 (1)	Stage 2 unbiased (2)
<i>No-Change</i>	0.40	0.32
<i>Adaptives</i>	0.14	0.17
<i>Sophisticated</i>	0.29	0.31
<i>Sophisticated 2</i>	0.17	0.20
No. of Subjects	93	93

Notes: The table displays the population frequencies estimated to be consistent with each of the behavioral rules listed in the Model column.. We exclude subjects whose behavior is equally compatible with more than one type.

Table 9: 11-20 Contingency Table: Stage 2

Panel A: Panel A: Stage 2 Raw

First Part Model	Second Part Model				No. of Subjects
	<i>No-change</i>	<i>Adaps</i>	<i>Soph</i>	<i>Soph 2</i>	
<i>Non-strategic</i>	0.50	0.30	0.10	0.10	10
<i>L2</i>	0.37	0.12	0.37	0.14	49
<i>L3</i>	0.29	0.12	0.29	0.29	17
<i>L4</i>	0.56	0.06	0.19	0.19	16
<i>NE</i>	–	–	–	–	0
<i>No. of Subjects</i>	37	12	27	16	92

Panel B: Panel B: Stage 2 Unbiased

First Part Model	Second Part Model				No. of Subjects
	<i>No-change</i>	<i>Adaps</i>	<i>Soph</i>	<i>Soph 2</i>	
<i>Non-strategic</i>	0.33	0.42	0.10	0.14	10
<i>L2</i>	0.28	0.15	0.41	0.17	49
<i>L3</i>	0.24	0.13	0.31	0.33	17
<i>L4</i>	0.55	0.06	0.19	0.20	16
<i>NE</i>	–	–	–	–	0
<i>No. of Subjects</i>	29.77	14.49	29.21	18.53	92

Notes: The table shows for each of the behavioral rules in initial play (by row) the proportion of subjects identified as following each of the behavioral rules in repeated play. For each row, the proportions across the four columns referring to the four main behavioral rules (*No-Change*, *Adaptive*, *Sophisticated* and *Sophisticated2*) should sum up to 1. We exclude subjects whose behavior is equally compatible with more than one type.

B.3 Robustness for 3×3 Games

One important concern in testing for a correlation between strategic sophistication and naivete in initial and repeated play is that the identification of behavioral types may be misspecified because some behavioral rules that are relevant to explaining subjects' behavior are not considered. With this concern in mind, we perform three robustness tests. First, we repeat the estimation with elicited behavior in the first part, including several alternative behavioral rules in addition to those that we already considered. Second, we perform an omitted-type specification test to alternatively confirm whether we obtain our result due to the omission of one or many relevant behavioral rules. Finally, as we find a high number of the *No-Change* type, accounting for almost 40% of the subjects, we perform an additional analysis replacing the *No-Change* type with all the behavioral rules we considered in the first part.

B.3.1 Addition of Alternative Behavioral Rules in Initial Play

We consider 4 alternative behavioral types for the initial play in addition to the 8 that we described in Section 2. All four types could be considered variations of *L1*, in which we alter the belief about the opponent's behavior. Given that we consider it to be plausible that subjects follow some simple non-strategic rules, it is also plausible that some subjects think in the same way. Consequently, we consider *L1* to refer to the best response to each of the other non-strategic rules that we initially included, that is, *L1_A*, *L1_{IA}*, *L1_{MaxMax}* and *L1_{MaxMin}*. Note that these alternative behavioral rules are clearly strategic and closer in spirit to *L2* in terms of strategic sophistication, as they predict a particular opponent's strategy and best respond to that strategy. Additionally, as shown in Table A2 in the Appendix, these additional behavioral types show good separation from the types that we initially considered.

As shown in Table A3 in the Appendix, the alternative models appear to show some relevance, although they do not alter the identified type distribution substantially. First, as expected, the new alternative behavioral rules steal frequency mostly from *L2* and the non-strategic types (mostly *A*). The other behavioral model that appears to be the most relevant is *L1_{MaxMax}*, which is followed by 11% of subjects. The contingency table displayed in Table 10 shows that subjects following these alternative models are best explained by *No-Change* and *Adaptive*, and only a minority are best explained by *Sophisticated* in repeated play. In summary, the consideration of additional alternative behavioral rules to explain initial play does not alter the main results: we find no evidence of a positive correlation between naivete and sophistication in initial and

repeated play.

Table 10: Contingency Table with Additional Alternative Behavioral Rules: All Subjects

First Part Model	Second Part Model						No. of Subjects	
	<i>No-Change</i>	<i>Adaptive</i>				<i>Soph</i>		<i>Soph 2</i>
			<i>Adap_S</i>	<i>Adap_A</i>	<i>Adap_{IA}</i>			
<i>Non-strategic</i>	0.31	0.61	0.14	0.30	0.17	0.04	0.03	70
<i>A</i>	0.21	0.71	0.29	0.29	0.14	0.07	0.00	14
<i>IA</i>	0.30	0.60	0.20	0.30	0.10	0.00	0.10	10
<i>MaxMax</i>	0.17	0.75	0.08	0.33	0.33	0.00	0.08	12
<i>MaxMin</i>	0.58	0.42	0.17	0.08	0.17	0.00	0.00	12
<i>L1</i>	0.32	0.59	0.05	0.41	0.14	0.09	0.00	22
<i>Alternative Models</i>	0.25	0.50	0.17	0.33	0.00	0.17	0.08	24
<i>L1_A</i>	0.20	0.60	0.00	0.60	0.00	0.00	0.20	5
<i>L1_{IA}</i>	0.20	0.60	0.20	0.40	0.00	0.20	0.00	5
<i>L1_{MaxMax}</i>	0.31	0.46	0.23	0.23	0.00	0.15	0.08	13
<i>L1_{MaxMin}</i>	0.00	0.00	0.00	0.00	0.00	1.00	0.00	1
<i>L2</i>	0.48	0.46	0.15	0.24	0.07	0.04	0.02	46
<i>L3</i>	–	–	–	–	–	–	–	0
<i>NE</i>	0.00	1.00	0.00	0.00	1.00	0.00	0.00	2
No. of Subjects	50	78	21	40	17	9	5	142

Notes: The table shows for each of the behavioral rules in initial play (by row) the proportion of subjects identified as following each of the behavioral rules in repeated play. For each row, the proportions across the four columns referring to the four main behavioral rules (*No-Change*, *Adaptive*, *Sophisticated* and *Sophisticated 2*) should sum up to 1. Furthermore, for each row, the proportions across the three adaptive behavioral models (*Adaptive_S*, *Adaptive_A*, *Adaptive_S*) should sum to the value in the column for *Adaptive*. We exclude 56 subjects whose behavior is equally compatible with more than one type in one part or the other.

B.3.2 Specification Test: Omitted Types

In a similar spirit to the previous robustness test, we also perform an omitted type specification test (as in Costa-Gomes and Crawford, 2006 to rule out the possibility that we did not consider relevant models.

In this test, instead of proposing alternative behavioral models, we let the actual subject behavior in our sample inform us of potential alternative rules. If we left out a rule that actually complies with the subjects' behavior, we would expect some of the subjects to show behaviors similar to those predicted by this rule. Therefore, we consider the observed behavior to identify potential new rules in the following manner. In addition to all 12 behavioral rules considered in the previous section, we add each subject's actual behavior as an additional behavioral rule, one subject at a time, and re-estimate the mixture-of-types model as many times as the number of subjects in our population, that is, 198 times. While conducting this exercise, we check whether the

added subject’s behavioral rule is able to explain other subjects’ behavior *better* than the existing 12 models and whether the rule can attract sufficient relevance, where we impose a threshold of 15% of the population frequency.

We find three such subjects (subject numbers 31, 85, and 86). What strategies do these subjects follow? First, we check for similarity (or, alternatively, separation) in these subjects’ behavior. These subjects appear to reflect the same type of behavior, as they show very little separation (0.21 between the behavior of subject 31 and subject 85, 0.14 between the behavior of subject 31 and subject 86, and 0.36 between the behavior of subject 85 and subject 86). Second, we check their separation from other existing behavioral rules, as shown in Table A4 in the Appendix. All three behavioral rules are well separated from all other considered rules, with the exception of *L2*, which shows a separation equal to or less than 0.43. Third, consistent with this finding, we also observe that when we consider these alternative models in the mixture-of-types model estimation, the behavioral rule that loses the most frequency is indeed *L2*, as shown by estimations in Table A5. Finally, we directly consider the actions of these subjects and find that their behavior is mostly consistent with *L2*, but in a few games mimics *L1*.¹⁶ In particular, the strategy profiles of subjects 85 and 31 diverge from *L2* or *L1* behavior in only two decisions and that of subject 86 diverges in only three decisions.

We conclude that these subjects show some variation from the existing *L2* behavioral type; however, none of them show a population frequency higher than that of *L2* when incorporated into the estimation together, as shown in Table A5, or one by one.

Does the result of the correlation between sophistication and naivete between initial and repeated play change when these new empirically motivated behavioral rules are considered? Table 11 shows that subjects following these alternative models are best explained by *No-Change* and by adaptive learners, with proportions similar to those in Table 5a. Therefore, we again conclude that we find no evidence for a positive correlation between naivete and sophistication in initial and repeated play.

B.3.3 Specification Test: Replacement of *No-Change* Type

We found that the large majority of subjects, almost 40% of them, followed the simplest *No-Change* behavioral type in repeated play, such that their behavior is best described as simply taking exactly the same strategy as the one that they took in the first part.

¹⁶In particular, the strategy profile of subject 31 is 3 1 3 3 3 1 3 3 1 2 1 2 3 1; the strategy profile of subject 85 is 3 1 3 3 3 1 1 3 1 2 3 3 3 1; and the strategy profile of subject 86 is 3 1 3 1 3 1 3 2 1 2 1 2 3 1.

Table 11: Contingency Table with the Addition of Three Subjects' Behavioral Rules:
All Subjects

First Part Model	Second Part Model						No. of Subjects	
	<i>No-Change</i>	<i>Adaptive</i>				<i>Soph</i>		<i>Soph 2</i>
			<i>Adap_S</i>	<i>Adap_A</i>	<i>Adap_{IA}</i>			
<i>Non-strategic</i>	0.35	0.57	0.11	0.28	0.17	0.04	0.04	46
<i>A</i>	0.17	0.67	0.17	0.50	0.00	0.17	0.00	6
<i>IA</i>	0.33	0.50	0.17	0.17	0.17	0.00	0.17	6
<i>MaxMax</i>	0.22	0.67	0.11	0.22	0.33	0.00	0.11	9
<i>MaxMin</i>	0.67	0.33	0.11	0.11	0.11	0.00	0.00	9
<i>L1</i>	0.31	0.63	0.06	0.38	0.19	0.06	0.00	16
<i>Alternative Models</i>	0.29	0.35	0.12	0.24	0.00	0.24	0.12	17
<i>L1_A</i>	0.00	0.50	0.00	0.50	0.00	0.00	0.50	2
<i>L1_{IA}</i>	0.33	0.33	0.00	0.33	0.00	0.33	0.00	3
<i>L1_{MaxMax}</i>	0.36	0.36	0.18	0.18	0.00	0.18	0.09	11
<i>L1_{MaxMin}</i>	0.00	0.00	0.00	0.00	0.00	1.00	0.00	1
<i>Subject 31</i>	0.57	0.14	0.00	0.00	0.14	0.14	0.14	7
<i>Subject 85</i>	0.25	0.75	0.25	0.38	0.13	0.00	0.00	16
<i>Subject 86</i>	0.38	0.57	0.19	0.29	0.10	0.05	0.00	21
<i>L2</i>	0.48	0.44	0.16	0.24	0.04	0.04	0.04	25
<i>L3</i>	-	-	-	-	-	-	-	0
<i>NE</i>	0.00	1.00	0.00	0.00	1.00	0.00	0.00	1
No. of Subjects	49	69	19	35	15	9	6	133

Notes: The table shows for each of the behavioral rules in initial play (by row) the proportion of subjects identified as following each of the behavioral rules in repeated play. For each row, the proportions across the four columns referring to the four main behavioral rules (*No-Change*, *Adaptive*, *Sophisticated* and *Sophisticated 2*) should sum up to 1. Furthermore, for each row, the proportions across the three adaptive behavioral models (*Adaptive_S*, *Adaptive_A*, *Adaptive_S*) should sum to the value in the column for *Adaptive*. We exclude 65 subjects whose behavior is equally compatible with more than one type in one part or the other.

We therefore question whether our results would significantly change if we replaced the *No-Change* type with all the behavioral rules that we considered in the first stage and added more sophisticated learning models such as *Adaptive*, *Sophisticated* and *Sophisticated 2*.

Table A6 in the Appendix shows these results. The behavioral rules from the first stage maintain relative frequencies similar to those in the original estimation see Table 3. The frequency of the models from the second stage decreases slightly as they have to compete with more alternative explanations, but more importantly they remain relevant and show the same frequency ordering. *Adaptive* is the most frequent non-naive model, followed by *Sophisticated* and *Sophisticated 2*, which show a much lower frequency, as before.

Table 12: Contingency Table with *No-Change* type replaced by Behavioral Models in Initial Responses: All Subjects

First Part Model	Models of initial play	Second Part Model			<i>Soph</i>	<i>Soph 2</i>	No. of Subjects	
		<i>Adaptive</i>	<i>Adap_S</i>	<i>Adap_A</i>				<i>Adap_{IA}</i>
<i>Non-strategic</i>	0.42	0.54	0.14	0.30	0.10	0.02	0.01	36
<i>A</i>	0.22	0.78	0.22	0.43	0.13	0.00	0.00	23
<i>IA</i>	0.27	0.73	0.18	0.45	0.09	0.00	0.00	11
<i>MaxMax</i>	0.46	0.46	0.08	0.23	0.15	0.00	0.08	13
<i>MaxMin</i>	0.43	0.57	0.21	0.21	0.14	0.00	0.00	14
<i>L1</i>	0.64	0.28	0.04	0.20	0.04	0.08	0.00	25
<i>L2</i>	0.36	0.53	0.24	0.24	0.05	0.07	0.05	59
<i>L3</i>	0.00	1.00	0.00	1.00	0.00	0.00	0.00	3
<i>NE</i>	0.50	0.50	0.00	0.00	0.50	0.00	0.00	2
<i>No. of Subjects</i>	58	82	26	43	13	6	4	150

Notes: The table shows for each of the behavioral rules in initial play (by row) the proportion of subjects identified as following each of the behavioral rules in repeated play. For each row, the proportions across the four columns referring to the four main behavioral rules (*No-Change*, *Adaptive*, *Sophisticated* and *Sophisticated 2*) should sum up to 1. Furthermore, for each row, the proportions across the three adaptive behavioral models (*Adaptive_S*, *Adaptive_A*, *Adaptive_S*) should sum to the value in the column for *Adaptive*. We exclude 48 subjects whose behavior is equally compatible with more than one type in one part or the other.

We can finally reproduce the contingency table replacing the *No-Change* type by all models that we considered in initial play, as shown in Table 12. The two most important models, *L2* and the non-strategic ones, show a very similar distribution over the models considered in repeated play. In other words, the non-strategic models and the *L2* behavioral types show frequencies very similar to the ones before of the models used in initial play and an adaptive learning model, showing once again that the naivete and sophistication in the first play show little correlation with naivete and sophistication in repeated play.

The full table, where we disaggregate the *No-Change* into the different behavioral types included in part 1, is shown in Table A7 in the Appendix. Sometimes, subjects whose behavior is best described by *No-Change* change type between the first and the second parts. However, if a subject is identified as *No-Change*, it is more likely that this subject is identified as using exactly the same behavioral type as in the first part. For example, subjects identified as following the *L2* behavioral rule are identified mostly as following one of two behavioral rules in the second part: either *L2* or *Adaptive* (either *Adaptive_S* or *Adaptive_A*).

More importantly, replacing *No-Change* does not change the estimated frequency of *Adaptive* or the correlation of sophistication between initial and repeated play. We therefore conclude that the results are robust to replacing the *No-Change* type with the models considered in initial play.

C Additional Tables

Table A1: Summary of Socio-Demographic Variables of the Subject Population

Variables	3 × 3 Games		11 – 20 Game	
	Mean Values	Stand. Dev.	Mean Values	Stand. Dev.
<i>Men</i>	0.41		0.39	
<i>Age</i>	21.73	2.99	20.97	2.76
<i>Spanish</i>	0.87		0.95	
<i>University Entry Grade</i> (out of 10)	6.85	1.16	7.58	1.46
Distribution over Field of Study:				
<i>Social Science</i>	0.77		0.91	
<i>Applied Science</i>	0.17		0.04	
<i>Natural Science</i>	0.04		0.03	
Distribution over risk choices:				
1.5€with 0.50 or 1.5€with 0.50	0.31		0.27	
1.3€with 0.50 or 1.8€with 0.50	0.11		0.13	
1.1€with 0.50 or 2.1€with 0.50	0.26		0.25	
0.9€with 0.50 or 2.4€with 0.50	0.07		0.05	
0.7€with 0.50 or 2.7€with 0.50	0.04		0.08	
0.6€with 0.50 or 2.8€with 0.50	0.04		0.04	
0.4€with 0.50 or 2.9€with 0.50	0.02		0.02	
0€with 0.50 or 3€with 0.50	0.16		0.17	
Cognitive reflection test:				
Percent correct in cognitive reflection test: Q1	0.28		0.31	
Percent correct in cognitive reflection test: Q2	0.17		0.39	
Percent correct in cognitive reflection test: Q3	0.41		0.56	

Notes: *Men* takes a value of 1 if the subject is male. *Age* reflects the age in years. *Spanish* takes a value of 1 if the subject is Spanish. *University Entry Grade* is normalized to a grade out of 10. *Social Science*, *Applied Science* and *Natural Science* take values of 1 if the subject is studying a social, applied or natural science. Risk Choice was elicited as in Eckel and Grossman (2002), where choices are ordered from safest to riskiest. Finally, the cognitive reflection test includes questions from Toplak et al. (2014) designed to avoid the possibility that the original test from Frederick (2005) is already known by the subjects. The questions are as follows: 1. If John can drink one barrel of water in 6 days, and Mary can drink one barrel of water in 12 days, how long would it take them to drink one barrel of water together? (correct answer 4 days; intuitive answer 9); 2. Jerry received both the 15th highest and the 15th lowest mark in the class. How many students are in the class? (correct answer 29 students; intuitive answer 30); 3. A man buys a pig for \$60, sells it for \$70, buys it back for \$80, and finally sells it for \$90. How much has he made? (correct answer \$20; intuitive answer \$10).

Table A2: Separation of Different Behavioral Rules with Additional Alternative Behavioral Models

	<i>A</i>	<i>IA</i>	<i>MaxMax</i>	<i>MaxMin</i>	<i>L1</i>	<i>L1_A</i>	<i>L1_{IA}</i>	<i>L1_{MaxMax}</i>	<i>L1_{MaxMin}</i>	<i>L2</i>	<i>L3</i>	<i>NE</i>
<i>A</i>	0.00											
<i>IA</i>	0.60	0.00										
<i>MaxMax</i>	0.46	0.62	0.00									
<i>MaxMin</i>	0.71	0.65	0.57	0.00								
<i>L1</i>	0.57	0.76	0.50	0.79	0.00							
<i>L1_A</i>	0.25	0.62	0.36	0.64	0.64	0.00						
<i>L1_{IA}</i>	0.64	0.98	0.64	0.71	0.50	0.57	0.00					
<i>L1_{MaxMax}</i>	0.75	0.73	0.64	0.57	0.79	0.50	0.64	0.00				
<i>L1_{MaxMin}</i>	0.71	0.65	0.93	0.64	0.71	0.71	0.64	0.57	0.00			
<i>L2</i>	0.75	0.58	0.57	0.64	0.71	0.57	0.79	0.5	0.79	0.00		
<i>L3</i>	0.86	0.80	0.86	0.64	0.79	0.71	0.43	0.50	0.50	0.71	0.00	
<i>NE</i>	0.57	0.51	0.86	0.64	0.79	0.57	0.50	0.86	0.57	0.79	0.57	0.00

Notes: The table reports the proportion of strategies across all 14 games in which the different behavioral models predict different strategies. The minimum possible separation value is 0, which occurs when the two models prescribe the same strategy in all 14 games, and the maximum possible separation value is 1, which occurs when the two models predict a different strategy in each of the 14 games.

Table A3: Behavioral Type Identification for Initial Play: Additional Behavioral Types

	Minimum Number of Perfect Guesses							
	No constraints		7		9		11	
	Main	Alt.	Main	Alt.	Main	Alt.	Main	Alt.
<i>A</i>	0.14	0.10	0.14	0.10	0.10	0.07	0.14	0.11
<i>IA</i>	0.07	0.06	0.07	0.05	0.01	0.01	0.00	0.00
<i>MaxMax</i>	0.08	0.10	0.08	0.10	0.08	0.06	0.05	0.04
<i>MaxMin</i>	0.10	0.12	0.11	0.11	0.11	0.09	0.10	0.07
<i>L1</i>	0.17	0.15	0.18	0.16	0.21	0.17	0.33	0.25
<i>L1_A</i>		0.06		0.06		0.04		0.00
<i>L1_{IA}</i>		0.04		0.04		0.03		0.04
<i>L1_{MaxMax}</i>		0.11		0.11		0.14		0.21
<i>L1_{MaxMin}</i>		0.01		0.01		0.01		0.00
<i>L2</i>	0.39	0.25	0.39	0.26	0.48	0.38	0.38	0.29
<i>L3</i>	0.02	0.00	0.02	0.00	0.00	0.00	0.00	0.00
<i>NE</i>	0.02	0.02	0.02	0.02	0.01	0.01	0.00	0.00
No. of Subjects	174	161	169	159	91	104	21	28

Notes: The table displays the population frequencies estimated for the main specification shown in Table 3 and when we add alternative models in initial play. We exclude 37, 30, 9 and 0 subjects whose behavior is equally compatible with more than one type when we impose no constraints and when we impose the criteria of 7, 9 and 11 perfect guesses, respectively.

Table A4: Separation of the Three Relevant Subjects' Behavior from other Behavioral Models

	<i>A</i>	<i>IA</i>	<i>MaxMax</i>	<i>MaxMin</i>	<i>L1</i>	<i>L1_A</i>	<i>L1_{IA}</i>	<i>L1_{MaxMax}</i>	<i>L1_{MaxMin}</i>	<i>L2</i>	<i>L3</i>	<i>NE</i>
<i>Subject 31</i>	0.57	0.58	0.71	0.71	0.50	0.57	0.57	0.57	0.64	0.36	0.71	0.57
<i>Subject 85</i>	0.54	0.65	0.50	0.71	0.57	0.50	0.50	0.64	0.79	0.36	0.79	0.64
<i>Subject 86</i>	0.57	0.55	0.64	0.71	0.50	0.57	0.57	0.50	0.64	0.43	0.79	0.71

Notes: The table reports the proportion of strategies across all 14 games in which the three subjects' behavioral models predict different strategies from the rest of the considered models. The minimum possible separation value is 0, which occurs when the two models prescribe the same strategy in all 14 games, and the maximum possible separation value is 1, which occurs when the two models predict a different strategy in each of the 14 games.

Table A5: Behavioral Type Identification for Initial Play: Additional Behavioral Types

	Minimum Number of Perfect Guesses							
	No constraints		7		9		11	
	Main	Alt.	Main	Alt.	Main	Alt.	Main	Alt.
<i>A</i>	0.10	0.05	0.10	0.05	0.07	0.06	0.11	0.07
<i>IA</i>	0.07	0.04	0.06	0.04	0.01	0.01	0.00	0.00
<i>MaxMax</i>	0.07	0.06	0.08	0.06	0.06	0.05	0.04	0.02
<i>MaxMin</i>	0.09	0.07	0.09	0.07	0.09	0.07	0.07	0.04
<i>L1</i>	0.16	0.13	0.16	0.13	0.17	0.13	0.25	0.16
<i>L1_A</i>	0.04	0.02	0.04	0.02	0.04	0.02	0.00	0.00
<i>L1_{IA}</i>	0.04	0.02	0.04	0.02	0.03	0.02	0.04	0.02
<i>L1_{MaxMax}</i>	0.11	0.10	0.11	0.10	0.14	0.11	0.21	0.13
<i>L1_{MaxMin}</i>	0.01	0.01	0.01	0.01	0.01	0.01	0.00	0.00
<i>L2</i>	0.30	0.17	0.30	0.17	0.38	0.20	0.29	0.18
<i>L3</i>	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
<i>NE</i>	0.02	0.01	0.02	0.01	0.01	0.01	0.00	0.00
<i>Subject 31</i>		0.05		0.05		0.06		0.07
<i>Subject 85</i>		0.13		0.13		0.12		0.16
<i>Subject 86</i>		0.14		0.14		0.14		0.16
No. of Subjects	161	150	159	150	104	121	28	45

Notes: The table displays the population frequencies estimated for the main specification shown in Table 3 and when we add alternative models in initial play. We exclude 48, 45, 21 and 0 subjects whose behavior is equally compatible with more than one type when we impose no constraints and when we impose the criteria of 7, 9 and 11 perfect guesses, respectively.

Table A6: Behavioral Type Identification for Repeated Play: *No-Change* replaced by Behavioral Rules from Initial Play

	Minimum Number of Perfect Guesses							
	No constraints		7		9		11	
	Main	Alt.	Main	Alt.	Main	Alt.	Main	Alt.
<i>No-Change</i>	0.34		0.35		0.41		0.50	
<i>A</i>		0.04		0.04		0.04		0.03
<i>IA</i>		0.02		0.02		0.01		0.00
<i>MaxMax</i>		0.04		0.04		0.04		0.03
<i>MaxMin</i>		0.05		0.05		0.05		0.06
<i>L1</i>		0.08		0.08		0.08		0.15
<i>L2</i>		0.13		0.13		0.13		0.03
<i>L3</i>		0.01		0.01		0.01		0.00
<i>NE</i>		0.01		0.01		0.01		0.00
<i>Adaptive</i>	0.57	0.55	0.56	0.55	0.54	0.56	0.46	0.65
<i>Adaptive_S</i>	0.16	0.17	0.17	0.17	0.21	0.23	0.28	0.38
<i>Adaptive_A</i>	0.28	0.29	0.28	0.29	0.28	0.30	0.15	0.24
<i>Adaptive_{IA}</i>	0.13	0.09	0.11	0.09	0.05	0.03	0.02	0.03
<i>Sophisticated</i>	0.06	0.04	0.06	0.04	0.02	0.02	0.04	0.06
<i>Sophisticated 2</i>	0.04	0.03	0.03	0.03	0.03	0.04	0.00	0.00
No. of Subjects	176	171	166	168	117	119	46	34

Notes: The table displays the population frequencies estimated for the main specification shown in Table 4 and when we replace the *No-Change* type with all the models included in initial play as in Table 3. We exclude 27, 23, 10 and 0 subjects whose behavior is equally compatible with more than one type when we impose no constraints and when we impose the criteria of 7, 9 and 11 perfect guesses, respectively.

Table A7: Full Contingency Table with No-Change type replaced by Behavioral Models in Initial Responses: All Subjects

First Part Model	Second Part Model											No. of Subjects			
	A	IA	MaxMax	MaxMin	L1	L2	L3	NE	Adapt	Adaps	Adapa		Adapia	Sophis	Sophis 2
<i>Non-strategic</i>	0.05	0.02	0.05	0.08	0.09	0.09	0.02	0.01	0.54	0.14	0.30	0.10	0.02	0.01	86
<i>A</i>	0.09	0.00	0.09	0.00	0.00	0.04	0.00	0.00	0.78	0.22	0.43	0.13	0.00	0.00	23
<i>IA</i>	0.00	0.09	0.00	0.09	0.00	0.00	0.09	0.00	0.73	0.18	0.45	0.09	0.00	0.00	11
<i>MaxMax</i>	0.00	0.00	0.08	0.23	0.08	0.08	0.00	0.00	0.46	0.08	0.23	0.15	0.00	0.00	13
<i>MaxMin</i>	0.00	0.07	0.00	0.21	0.07	0.07	0.00	0.00	0.57	0.21	0.21	0.14	0.00	0.00	14
<i>L1</i>	0.08	0.00	0.04	0.00	0.24	0.20	0.04	0.04	0.28	0.04	0.20	0.04	0.08	0.00	25
<i>L2</i>	0.05	0.02	0.03	0.03	0.07	0.15	0.00	0.00	0.53	0.24	0.24	0.05	0.07	0.05	59
<i>L3</i>	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	1.00	0.00	0.00	0.00	3
<i>NE</i>	0.00	0.00	0.00	0.00	0.50	0.00	0.00	0.00	0.50	0.00	0.00	0.50	0.00	0.00	2
No. of Subjects	7	3	6	9	13	17	2	1	82	26	43	13	6	4	150

Notes: The table shows for each of the behavioral rules in initial play (by row) the proportion of subjects identified as following each of the behavioral rules in repeated play. For each row, the proportions across columns referring to the 11 main behavioral rules (*A*, *IA*, *MaxMax*, *MaxMin*, *L1*, *L2*, *L3*, *NE*, *Adaptive*, *Sophisticated* and *Sophisticated 2*) should sum up to 1. Furthermore, for each row, the proportions across the three adaptive behavioral models (*Adaptive_S*, *Adaptive_A*, *Adaptive_S*) should sum to the value in the column for *Adaptive*. We exclude 48 subjects whose behavior is equally compatible with more than one type in one part or the other.

D Translation of Instructions

D.1 3×3 Games

The original instructions were in Spanish. Here we provide a translation of the instructions into English.

THANK YOU FOR PARTICIPATING IN OUR EXPERIMENT!

We will now start the experiment. From now on, you are not allowed to speak, look at what other participants do or walk around the room. Please turn off your phone. If you have any questions or need help, raise your hand and one of the researchers will talk with you. Please, do not write on these instructions. If you do not follow these rules, YOU WILL BE ASKED TO LEAVE THE EXPERIMENT, AND NO PAYMENT WILL BE GIVEN TO YOU. Thank you.

The university and the research projects have provided the funds for carrying of this experiment. You will receive 3 euros for having arrived on time. Additionally, if you follow the instructions correctly, you have the possibility to earn more money. This is a group experiment. The amount that you can earn depends on your decisions, the decisions of other participants, and chance. Different participants can earn different amounts.

No participant will be able to identify another by their decisions or by their profits in the experiment. The researchers will be able to observe the profits of each participant at the end of the experiment, but we will not associate the decisions that you have made with the identity of any participant.

EARNINGS:

During the experiment you can earn experimental points. At the end, each experimental point will be exchanged for euros, and 1 experimental point is worth exactly 0.5 euros. Everything that you win will be paid in cash in a strictly private way at the end of the experimental session.

Your final earnings will be the sum of the 3 euros that you receive for participating plus what you earn during the experiment.

Each experimental point equals 50 cents, so 2 experimental points equals 1 euro ($2 \times 0.5 = 1$ euro).

If, for example, you earn a total of 20 experimental points, you will receive a total of 13 euros (3 euros as payment for participation and 10 Euros from the conversion of the 20 experimental points to euros).

If, for example, you earn 4 experimental points, you will obtain 5 euros ($4 \times 0.5 = 2$ and $2 + 3 = 5$).

If, for example, you earn 44 experimental points, you will obtain 25 euros ($44 \times 0.5 = 22$ and $22 + 3 = 25$).

PARTS OF THE EXPERIMENT:

The experiment consists of two parts. You will participate by operating a computer. In the first part, there will be 14 rounds where you will make 14 decisions. In the second part, there will also be 14 rounds where you will make 14 decisions. At the end of the experiment, when you have completed the two parts of the experiment, the computer will randomly choose two of the 28 rounds, and you will be paid for the experimental points that you received in those two rounds chosen at random, plus the 3 euros for participating.

Before you begin each part of the experiment, we will explain in detail what kinds of decisions you can make and how you can obtain experimental points.

When we are all ready, we will start the first part of the experiment by explaining the instructions of this part in detail.

FIRST PART OF THE EXPERIMENT:

The first part of the experiment consists of 14 rounds. In each of the 14 rounds, you will be paired with a participant chosen at random from this session. The other participant will be different in each of the rounds, so you will never be paired with the same participant more than once. From now on, we will refer to you as "you" and the other participant as "the other participant".

In each round, you will have to make a decision by choosing among three possible options. Each decision will be presented in the form of a table similar to the one below (but with different values). You will see the corresponding table each time you have to choose an option. Each row of the table corresponds to an option that you can choose. The decision that you must make is to choose one option. The other participant will also have to choose, independently of you, from their options, which correspond to the columns of the table. That is, you choose from the rows, while the other participant chooses from the columns. However, to simplify things, the experiment is programmed in such a way that all the participants - including the person with whom you are matched - see their decision as shown in the example. That is, each of you will be presented with your possible actions in the rows of the table, and your experimental points will be shown in red. At the time of making your choice, you will not know

the option chosen by the other participant, and when the other participant is choosing their option, they will not know the option that you have chosen.

The number of experimental points that you earn in each of the rounds depends on the option that you have chosen and the option that the other participant has chosen.

The table of experimental points that you see below is an example of what you will see in each of the rounds.

Example:

The other player can choose:

	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	<input type="radio"/> 8 ;20	20;12	12; 4
You can choose:	<input type="radio"/> 2 ; 8	8 ;14	20;14
	<input type="radio"/> 18;14	12; 4	6 ;18

Click on one of the three white buttons to choose an option.
You can choose your decision clickin on another button as many times as you want. If you are sure about you choice click on "OK".

For example, if this round is chosen at random and you select the first option (row) and the other participant selects the second option (column), you will obtain 20 experimental points, and the other participant will receive 12 experimental points.

As another example, if this round is chosen at random and you select the third option (row) and the other participant selects the first option (column), you will obtain 18 experimental points and the other participant will receive 14 experimental points.

These are just two examples to better understand how decisions affect the experimental points that you can earn and do not aim to suggest what decisions you should make.

To make a selection, click on the white button next to the desired option. Then, the button will turn red to indicate which option you have selected. Once you have chosen an option, the choice is not final, and you can change your selection as many

times as you want by clicking on another button until you press the “OK” button that will appear in the lower right corner of each screen. Once you click “OK”, the selection will be final, and you will proceed to the next round. You will not be able to move to the next round until you have chosen an option and clicked “OK”. You will not have any time restrictions. Take as much time as you need in each round. When all of you have made your decisions in each of the 14 rounds, we will explain the second part of the experiment.

Summary:

- Your experimental points will be shown in red, and the experimental points of the other participant will be shown in blue.
- You will participate in 14 different rounds. In each of the rounds, the table of experimental points will be different, and you will be paired with a different participant chosen at random from this session.
- In each round, you can choose among three different options (rows), and the experimental points that you earn depend on the option that you select, the option that the other participant selects, and whether that round is chosen at random at the end of the experiment.

We will start the first part of the experiment in a few moments. Before starting the first part, you will see a new example, and you will have to answer several questions. If you have any questions or need help at any time during the experiment, please raise your hand, and one of the investigators will talk to you.

The other player can choose:

	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
	<input type="radio"/> 8 ;20	20;12	12; 4
You can choose:	<input type="radio"/> 2 ; 8	8 ;14	20;14
	<input type="radio"/> 18;14	12; 4	6 ;18

Click on one of the three white buttons to choose an option.
You can choose your decision clickin on another button as many times as you want. If you are sure about you choice click on "OK".

OK

Questions:

1. Please write the experimental points that you would earn in this round, if this round is randomly chosen for payment, and if you chose the second option and the other participant chose the third option.

2. Please write the experimental points that the other participant would earn in this round, if this round is randomly chosen for payment, and if you chose the third option and the other participant chose the second option.

3. Please state if the following statement is true or false: “Two rounds will be randomly selected for payment. The two rounds can be from part 1, from part 2, or 1 from part 1 and the other from part 2.”

SECOND PART OF THE EXPERIMENT:

The second part of the experiment also consists of 14 rounds and will work similarly to the first part. That is, the tables of experimental points that you will see in each of the 14 rounds in this second part will be the same as those you saw in the first part of the experiment. As in the first part, in each of the 14 rounds, you will be paired with a participant chosen at random from this session. However, in each of the rounds, the

other participant with whom you have been paired in this part does not have to be the same as the participant with whom you were paired in the first part. The pairing is performed again at random. In each of the rounds, the other participant, chosen at random, will be different, so you will never be paired with the same participant more than once.

As in the first part, both you and the other participant can choose among three possible options. The experimental points that you can earn in each of the rounds depend on the option that you select and the option that the other participant selects, as well as on whether that particular round is chosen at random at the end of the experiment.

Unlike in the first part, in this case, when you see the table of experimental points, you can also observe the option that you chose in the first part and the option that was chosen in the first part by the participant with whom you are paired in this part. The option that you both chose in the first part will be indicated by an arrow and will say “You chose” and “The other chose”. The information that you observe will be the same for the participants with whom you are paired.

The table of experimental points that you see is an example of what you will see in each of the rounds.

Example:

The screenshot shows an experimental interface with a table of points and choice options. The table is as follows:

		The other player can choose:		
			The other player chose	
		<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
You chose	<input checked="" type="radio"/>	8 ; 20	20 ; 12	12 ; 4
You can choose:	<input type="radio"/>	2 ; 8	8 ; 14	20 ; 14
	<input type="radio"/>	18 ; 14	12 ; 4	6 ; 18

Click on one of the three white buttons to choose an option.
 You can choose your decision clickin on another button as many times as you want. If you are sure about you choice click on "OK".

OK

As in the first part, if, for example, this round is chosen at random and you select the first option (row) and the other participant selects the second option (column), you will earn 20 experimental points, and the other participant will earn 12 experimental points.

As another example, if this round is chosen at random and you select the third option (row) and the other participant selects the first option (column), you will obtain 18 experimental points, and the other participant will receive 14 experimental points.

These are just two examples to help you understand how decisions affect the experimental points that you earn and are not intended to suggest what decisions you should make.

Unlike in the first part, in this part of the experiment, you can observe, as indicated in the example, which option you chose and which option the other participant chose in the first part. For example, in the example table, you chose the second option (row), and the other participant chose the second option (column). The other participant can also observe the option that you chose and the option that he/she chose; you both have the same information. Now you will have to make a choice again.

You can make your decision in the same way as in the first part by clicking on the button of the option that you want to choose and confirming by pressing “OK”. You will not have any time restrictions. Take as much time as you need in each of the rounds. When all of you have made your decisions in each of the 14 rounds, the experiment will end.

Summary:

- Your experimental points will be shown in red, and the experimental points of the other participant will be shown in blue.
- You will participate in 14 different rounds. In each round, the table of experimental points will be different, and you will be paired with a different participant chosen at random from this session.
- Unlike in the first part, you can now see which option you chose in the first part and which option the other participant chose in the first part. The other participant will also be able to observe the option that he or she chose and the option that you chose.
- In each round, you can choose among three different options (rows), and the experimental points depend on the option that you have chosen, the option chosen

by the other participant, and whether that round is chosen at random at the end of the experiment.

We will start the second part of the experiment in a few moments. If you have any questions or need help at any time during the experiment, please raise your hand, and one of the investigators will talk to you.

D.2 11-20 Game

The original instructions were in Spanish. Here we provide a translation of the instructions into English.

THANK YOU FOR PARTICIPATING IN OUR EXPERIMENT!

We will now start the experiment. From now on, you are not allowed to speak, look at what other participants do or walk around the room. Please turn off your phone. If you have any questions or need help, raise your hand and one of the researchers will talk with you. Please, do not write on these instructions. If you do not follow these rules, **YOU WILL BE ASKED TO LEAVE THE EXPERIMENT, AND NO PAYMENT WILL BE GIVEN TO YOU.** Thank you.

The university and the research projects have provided the funds for carrying out this experiment. You will receive 3 euros for having arrived on time. Additionally, if you follow the instructions correctly, you have the possibility to earn more money. This is a group experiment. The amount that you can earn depends on your decisions, the decisions of other participants, and chance. Different participants can earn different amounts.

No participant will be able to identify another by their decisions or by their profits in the experiment. The researchers will be able to observe the profits of each participant at the end of the experiment, but we will not associate the decisions that you have made with the identity of any participant.

EARNINGS:

During the experiment you can earn experimental points. At the end, each experimental point will be exchanged for euros, and 1 experimental point is worth exactly 0.4 euros. Everything that you win will be paid in cash in a strictly private way at the end of the experimental session.

Your final earnings will be the sum of the 3 euros you receive for participating plus what you earn during the experiment.

Each experimental point equals 40 cents.

If, for example, you earn a total of 11 experimental points, you will receive a total of 7.4 euros (3 euros as payment for participation and 4.4 euros from the conversion of the 20 experimental points to euros).

If, for example, you earn 30 experimental points, you will obtain 15 euros ($30 \times 0.4 = 12$ and $12 + 3 = 15$).

If, for example, you earn 99 experimental points, you will obtain 42.6 euros ($99 \times 0.4 = 39.6$ and $39.6 + 3 = 42.6$).

PARTS OF THE EXPERIMENT:

The experiment consists of three parts. You will participate by operating a computer. At the end of the experiment, when you have completed all three parts, the computer will choose one part at random and you will be paid for the money that you have received in that part, plus the 3 euros for participating.

Before we start each part of the experiment, we will explain in detail what kind of decisions you can make and how you can obtain experimental points.

Now we will go on to explain the instructions for part 1 of the experiment.

PART 1 OF THE EXPERIMENT:

In part 1, you will make a decision. You will be paired with a randomly chosen participant from this session. From now on, we will refer to you as “you” and the other participant as “the other participant” in these instructions.

You will be asked to choose a number. The experimental points that you can earn depend on the number you choose, the number that the other participant chooses and whether this part is randomly selected at the end of the experiment.

You will have to choose a number between 11 and 20. You will always receive the number of points equal to the number you choose. In addition:

- if you choose the same number as the other participant, you will receive 10 extra points.
- if you choose exactly one number less than the other participant, you will receive 80 extra points.

When choosing, you will not know the number chosen by the other participant, and when the other participant is choosing his or her number, he or she will not know the number that you have chosen either.

Example:

Part 1

The other participant
can choose:

11	12	13	14	15	16	17	18	19	20
----	----	----	----	----	----	----	----	----	----

You can choose:

11	12	13	14	15	16	17	18	19	20
----	----	----	----	----	----	----	----	----	----

Click on one of the white buttons to choose an option.
When you are sure of your choice click on the "OK" button. When you are sure of your choice click on the "OK" button.

OK

For example, if this round is chosen at random and...

- if you choose 17 and the other participant chooses 19, then you will receive 17 points and the other participant will receive 19 points.
- if you choose 12 and the other participant chooses 13, then you will receive 92 points and the other participant 13 points.
- if you choose 20 and the other participant chooses 19, then you will receive 20 points and the other participant 99 points.
- if you choose 16 and the other participant chooses 16, then you will receive 26 points and the other participant 26 points.

These are only examples to help you understand how your choices affect the experimental points that you can earn and are not intended to suggest what choices you should make.

To make your decision, click on the number that you want to choose. The number will then turn red to indicate which number you have selected. The choice is not final and you can change it as many times as you want by clicking on another number until you click on the “OK” button that will appear in the bottom right corner of each

screen. Once you have clicked “OK”, the selected number will be final. You will not have any time restrictions. Take as much time as you need. When you have all made your choices, we will move on to explain part 2 of the experiment.

Summary:

- You will have to choose a number between 11 and 20. You will always receive the number of points that you choose. Also,
- if you choose the same number as the other participant, you will receive 10 extra points.
- if you choose exactly one number less than the other participant, you will receive 80 extra points.

We will start part 1 in a few moments. Before we begin, you will see a new example and you will have to answer several questions. If you have any questions or need help at any time, please raise your hand and one of the researchers will come and talk to you.

Part 1

The other participant
can choose:

11	12	13	14	15	16	17	18	19	20
----	----	----	----	----	----	----	----	----	----

You can choose:

11	12	13	14	15	16	17	18	19	20
----	----	----	----	----	----	----	----	----	----

Click on one of the white buttons to choose an option.
When you are sure of your choice click on the "OK" button. When you are sure of your choice click on the "OK" button.

Questions:

1. Please write here your points earned in this part if you choose 17 and the other participant who has been chosen for this round chooses 14, if this part is chosen for your payment.

2. Please write here the points earned in this part by the other participant if you choose 19 and the other participant who has been chosen for this part chooses 18, if this part is chosen for your payment.

3. Please state whether the following statement is true or false: “One part will be randomly selected for payment.”

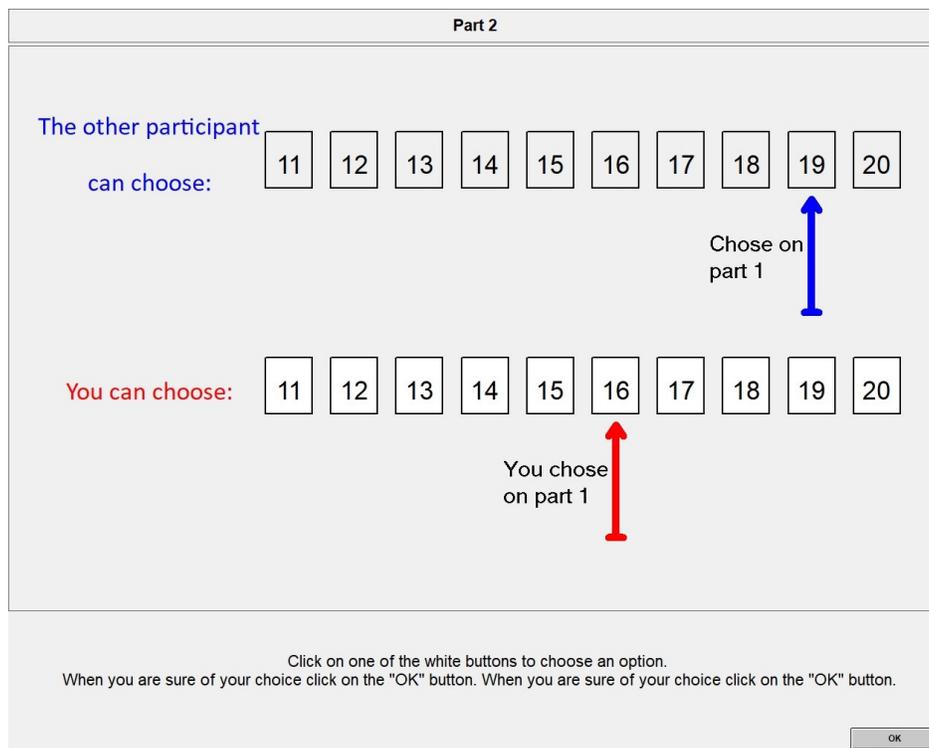
PART 2 OF THE EXPERIMENT:

Part 2 of the experiment will work the same as part 1. That is:

- How the experimental points are earned will be exactly the same.
- You will also be paired with another randomly chosen participant from this session. The other participant with whom you will be paired in this part will not be the same as the one in part 1.
- You will not know the number chosen by the other participant in this part 2 of the experiment, and when the other participant is choosing his or her number that he or she will not know the number you have chosen in part 2, either.

What is the difference? Now you will be able to see the number that the participant with whom you are paired in this part chose in part 1. In addition, you will see the number you chose in part 1. The information that you observe will be the same for the participants with whom you are paired.

Example:



In the example table, you chose 16 and the other participant chose 19 in part 1. The other participant can also observe the number that he or she chose and the number that you chose: you both have the same information and observe the same thing. Now you will have to decide which number you want to choose in this part.

These are only examples to help you understand how decisions affect the experimental points you can earn and are not intended to suggest what decisions you should make.

You can make your decision in the same way as in part 1 by clicking on the number that you want to choose and confirming by clicking “OK”. You will not have any time restrictions. Take as much time as you need. When you have all made your decision, we will move on to explain part 3.

Summary:

- As before, you will have to choose a number between 11 and 20. You will always receive the number of points that you choose. In addition,
 - if you choose the same number as the other participant in this part, you will receive 10 extra points.
 - if you choose exactly one number less than the other participant in this part, you will receive 80 extra points.

- The participant with whom you are paired in this part 2 is not the participant with whom you were paired in part 1.
- Unlike in part 1, you will now be able to see which number you chose in part 1 and the number that the other current participant chose in part 1. The other participant will also be able to see the number that he or she chose and the number that you chose.

We will begin part 2 in a few moments. If you have any questions or need help at any point in the experiment, please raise your hand and one of the researchers will come and talk to you.

PART 3 OF THE EXPERIMENT:

Part 3 of the experiment will work the same as part 2. That is,

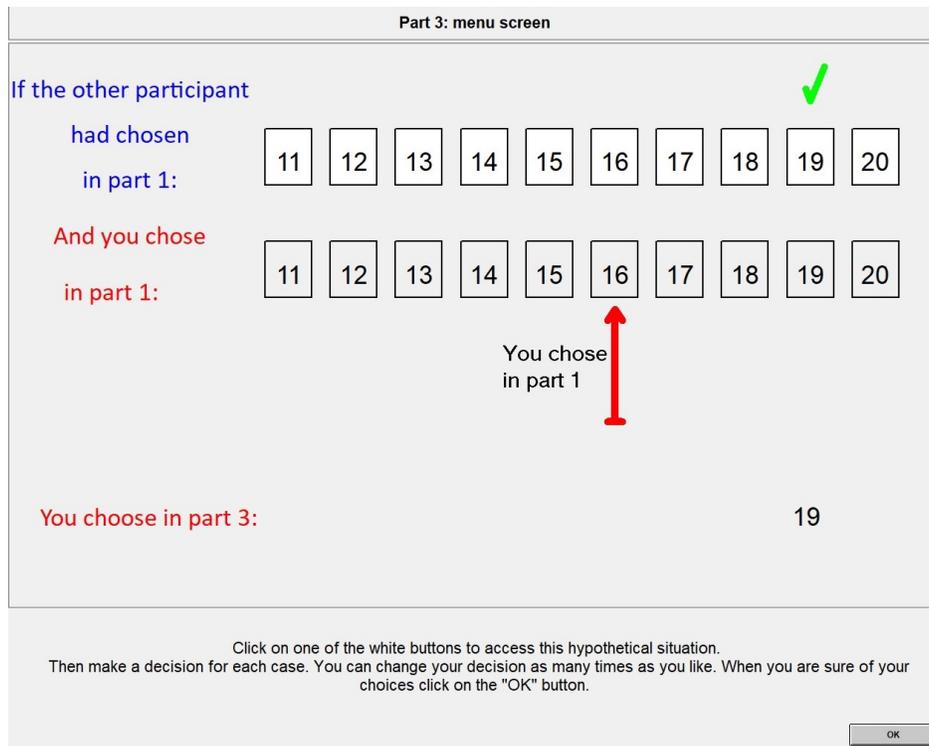
- How you obtain the experimental points will be exactly the same as usual.
- You will also be paired with another randomly chosen participant from this session. The other participant with whom you will be paired in this part will not be the same as the one in part 1 or part 2.
- As always, you will not know the number chosen by the other participant in this part 3 of the experiment, and when the other participant is choosing his or her number, he or she will not know the number you have chosen in part 3 either.
- You will see what number you chose in part 1.

What is the difference? Instead of making one decision, you will now make 10 decisions. Why 10? Because now you will not know what number the other participant chose in part 1. Since you do not know, you will make a decision for each hypothetical case of the number that he or she might have chosen in part 1. That is, you will make 10 decisions for each hypothetical case of your choice in part 1.

In short, it is like doing part 2 ten times, selecting which number you would choose in each hypothetical case, where only one will be the real one.

You will be able to switch between two types of screens: the menu screen and the decision screen.

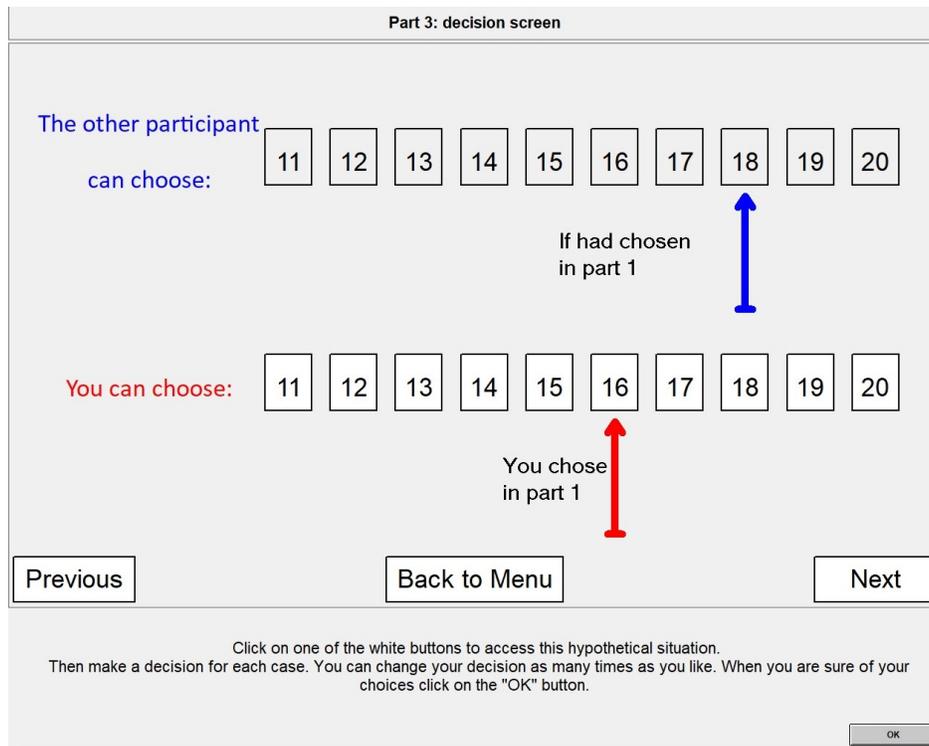
Example of the menu screen:



In the menu screen, in the first row you will see which possible numbers could have been chosen in part 1 by the other participant, from 11 to 20. The numbers that appear with a green tick are the cases for which you have already chosen a number. In the second row you will see what you chose in part 1 indicated with a red arrow. In the third and last row, you will see what you are choosing in this part 3 for each case. On the menu screen, if you click on each of the numbers in the first row of the other participant you will enter the corresponding decision screen.

In the example menu screen that we show you, you have made only one decision out of the possible 10, specifically, the case where the other participant chose 19 in part 1, because only that number has a green tick. You chose 16 in part 1, and if the other participant had chosen 19 in part 1, in part 3, you choose 19. To choose numbers for the other hypothetical cases, you must click on any of the other participant's numbers. This is just an example and is not meant to suggest how you should make your decisions.

Example of the decision screen:



In the decision screen, in the first row you will see the information corresponding to the hypothetical case that you have selected, and you must choose a number by clicking on a number in the “You can choose” row. In the example screen, you are in the decision screen for the case where the other participant chose 18, you chose 16, as indicated by the red arrow in your row. You can move between the hypothetical cases with the “Previous” and “Next” buttons, and if you click on the “Back to Menu” button, you will return to the menu screen, where you will see a summary of your decisions for this part and which decisions you have yet to make.

During the experiment, in the decision screen, you will see that you will have a choice already made, which is the one you made in part 2 for the hypothetical case that you came to observe. If you would like, you can also change that decision and any other decision in this part 3 as many times as you want. You can only finish the experiment when you give an answer for each of the 10 hypothetical cases.

Remember that you can choose any number you want from 11 to 20, and it can always be the same or different. These examples are simply to help you understand the screens that you will encounter and how to interpret them, they are not meant to suggest how you should make your decisions.

If this part of the experiment is chosen for payment, you will be paid only for the actual case of the 10 hypothetical decisions. The same will be true for the other

participant, who will be paid only for the actual case of the 10 hypothetical decisions.

You will not have any time restrictions. Take as much time as you need. When you have all made your decisions, the experiment will end.

Summary:

- As always, you will have to choose a number between 11 and 20. You will always receive the number of points you choose. Also,
 - if you choose the same number as the other participant chooses in this part, you will receive 10 extra points.
 - if you choose exactly one number less than the other participant in this part, you will receive 80 extra points. extra points.
- The participant with whom you are paired in this part 3 is different from the one in part 1 and the one in part 2.
- You will be able to see which number you chose in part 1, and you will be asked to choose a number for each hypothetical case of the possible number that the other participant could have chosen in part 1. The other participant will also be able to see the number he or she chose in part 1 and will also be asked to choose a number for each hypothetical case of the possible number that you could have chosen in part 1.
- If this part of the experiment becomes eligible for payoff, you will be paid for only one of the 10 decisions, i.e., for the decision in the real case.

We will start part 3 in a few moments. If you have any questions or need help at any point in the experiment, please raise your hand and one of the researchers will come and talk to you.