جامعة نيويورك أبوظبي
NYU | ABU DHABI

مركز التصميم السلوكي المؤسساتي
CBID
CENTER FOR BEHAVIORAL INSTITUTIONAL DESIGN

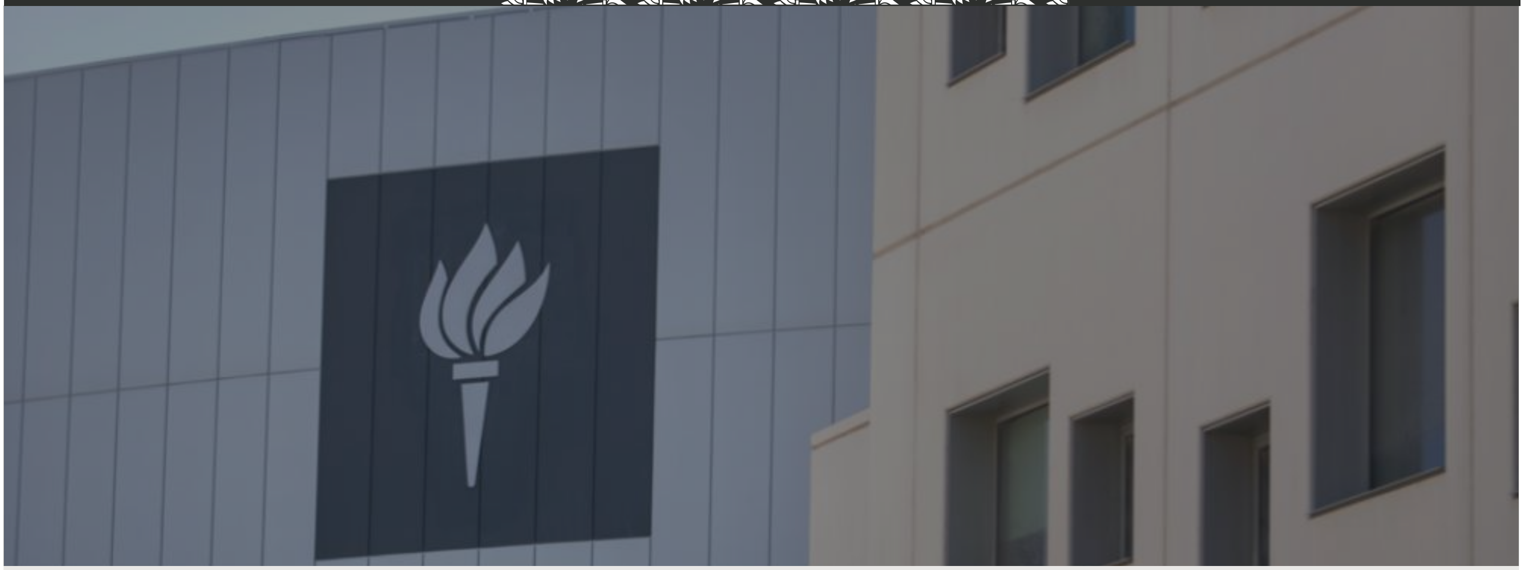

Limited Strategic Thinking and the Cursed Match

Olivier Bochet and Jacopo Magnani

Working Paper # 71
September 2021

Division of Social Science Working Paper Series
New York University Abu Dhabi, Saadiyat Island P.O Box 129188, Abu Dhabi, UAE
http://nyuad.nyu.edu/en/academics/academic-divisions/social-science.html

Limited Strategic Thinking and the Cursed Match*

Olivier Bochet[†] Jacopo Magnani[‡]

August 13, 2021

Abstract

In vertically differentiated matching markets with private information, agents face an acceptance curse: being accepted as a partner conveys bad news. We experimentally investigate whether individuals anticipate the acceptance curse in such an environment. We test the effect of an exogenous change in reservation values which, by making some types more selective, induces significant changes in the posterior distribution of match qualities. Consistent with limited strategic sophistication, subjects do not respond to this manipulation. Through additional investigation and structural estimations, we suggest a mechanism explaining the lack of subjects' response: out-of-equilibrium beliefs are quantitatively more important than limited conditional thinking.

Keywords: Strategic Thinking, Matching, Cursed Equilibrium, Experiments.

JEL codes: C9, D03 C78

*The authors warmly thank Nikos Nikiforakis, Ernesto Reuben and John Wooders for stimulating discussions. The authors also thank participants at various seminars for their helpful comments and suggestions. Olivier Bochet acknowledges financial support from Tamkeen under the NYU Abu Dhabi Research Institute Award for project CG005.

[†]Division of Social Science, New York University Abu Dhabi, and the Center for Behavioral Institutional Design (C-BID). Email: *olivier.bochet@nyu.edu*

[‡]Emlyon Business School, Lyon, France. Email: *jacopo.magnani@gmail.com*

1 Introduction

Motivation: In matching markets such as marriage and labor markets, agents often only have imperfect information about the true attributes of potential partners. However, in these markets, useful information can be inferred from others' willingness to match. In vertically differentiated markets, being accepted as a partner by another agent conveys bad news about their quality, a phenomenon dubbed as the *acceptance curse* (Chade, 2006). In extreme cases, this form of adverse selection may be so powerful that a rational agent should choose to remain unmatched, an unraveling strikingly illustrated by the famous quote of Groucho Marx: "I refuse to join any club that would have me for a member." More generally, to evaluate potential partners one needs to appropriately adjust the information obtained through observable characteristics by conditioning on acceptance. The failure to account for adverse selection when choosing potential partners can result in mismatches, which in turn can lead to match dissolution when information is revealed at a later stage. While mismatches and breakups can happen even under full rationality (Becker et al., 1977), limited strategic thinking can exacerbate the inefficiencies caused by imperfect information. It is thus important to evaluate the extent to which individuals are able to understand adverse selection in matching markets.

Methodology and Main Findings: To test whether individuals anticipate the acceptance curse, we implement in the laboratory a one-shot, two-player version of the market model of Chade (2006). This game replicates a market where matching with higher quality types yields higher payoffs and higher quality types have higher reservation values. Qualities are private information, but before deciding whether to propose to form a match, each player receives a noisy signal about the quality of their potential partner. In the equilibrium of this game, medium types face an acceptance curse: conditioning on acceptance lowers the probability that the quality of their potential partner is high. Moreover, given an exogenous increase in the reservation values of high types, the acceptance curse faced by medium types becomes more severe, because high types are more selective. Thus, when the reservation value of high types is sufficiently high, medium types should cease to propose to form matches. We test this comparative statics in the $BASE$ treatment of our lab experiment and find that subjects propose 66% of the time when they play the role of medium types, irrespective of the reservation value of high types. Thus, not only they propose to form matches when equilibrium predicts they should not, but they do not adjust their strategy in response to an exogenous change in reservation values of other players.

Based on the experimental literature on games with private information (discussed below), we posit two possible explanations of why people do not anticipate the acceptance curse: mistaken beliefs or limited conditional thinking. Subjects have mistaken beliefs if they do not correctly anticipate how the proposal decision of others depends on others' reservation values. Subjects suffer from limited conditional thinking if they cannot properly condition expectations on the hypothetical scenario in which a potential partner chooses to propose, no matter whether their beliefs about the strategies of others are correct or not. To disentangle these two potential explanations, we use two treatments, BEL (for "beliefs") and $COND$ (for "conditioning"), which sequentially eliminate the conditions that allow these explanations to account for behavior. In both treatments, each subject plays the role of player 1 and faces an individual choice problem obtained by assigning exogenous strategies to player 2. In our BEL treatment, we exogenously vary the likelihood that a high-quality player 2 proposes to form a match across rounds and we provide full information about this to each player 1 subject. In this way, the BEL treatment removes the effects of limits in the subjects' ability to form consistent beliefs, while allowing scope for limited conditional thinking. In the $COND$ treatment, we program player 2 to always propose, thus removing the need for the subject to condition on hypothetical events, and we vary the signal likelihoods to match the posterior distribution of qualities in each round of the BEL treatment. Apart from the need for subjects to condition on acceptance in BEL, BEL and $COND$

are comparable environments. Thus, we can interpret any difference between BEL and $COND$ as the result of limits in conditional thinking.

We find that proposal rates are highly responsive to changes in the posterior distribution of qualities in both BEL and $COND$. This is in contrast to the absence of responsiveness in $BASE$. We find small differences between BEL and $COND$, suggesting limited conditional thinking plays a smaller role than inconsistent beliefs. To further interpret our findings, we estimate a structural model based on cursed equilibrium. We find a players behaves as if fully cursed in $BASE$ but cursed to a much smaller degree in BEL. Thus, we argue that although both out-of-equilibrium beliefs and limited conditional thinking matter, beliefs seem quantitatively more important than conditioning.

Contributions and Related Literature: Our paper makes two contributions. First, we test whether individuals understand the acceptance curse in two-sided matching markets. Previous experiments show individuals often fail to account for the information of other players in settings like auctions (e.g. Kagel and Levin, 1986) and elections (e.g. Esponda and Vespa, 2014), but there is no previous evidence on two-sided matching.[1] One structural difference between our game and the environments studied in the previous literature is that in our framework a player's type consists of two different pieces of information: each player has private information about his own quality as well as a signal about the quality of potential partners.[2] The essence of the game is that a player needs to use information about his own quality to interpret signals about potential partners. How individuals perform this task, which is central to two-sided matching markets, has not been studied in previous papers. Second, we contribute to the literature on the role of beliefs and cognitive limits in games with private information. This

[1]Araujo et al. (2018) study an environment that resembles a one-sided matching market. Their paper focuses on dynamic adverse selection and does not consider the specific form of adverse selection arising in two-sided markets.

[2]In the compromise game of Carrillo and Palfrey (2009) each of the two players has private information about his attribute (strength, in their context), but players do not observe signals about the attribute of their opponents.

literature aims at understanding the mechanisms driving the widely observed deviations from Bayes-Nash equilibrium (BNE). The literature has focused on two types of mechanisms: limits in forming consistent beliefs and limits in conditional thinking.[3,4] A typical approach to test for belief-based explanations is the "robot protocol" which we adopt in our experiment: a subject's opponent is replaced by an automated player who follows a known strategy. A number of papers have shown that even in a robot protocol subjects deviates substantially from optimality (Charness and Levin, 2009; Ivanov et al., 2010; Esponda and Vespa, 2014) and conclude this casts doubt on the validity of belief-based explanations (but for a counterargument see Camerer et al. (2016) and see Ali et al. (2019) for a recent example where the robot protocol is effective). An explanation for why people deviate from optimal behavior even under the robot protocol is limited conditional thinking, that is the notion that individuals are limited in their ability to condition on hypothetical events (such as being pivotal in a simultaneous election or winning a sealed-bid auction). Several experiments have shown that subjects deviate from equilibrium less in sequential move games, where contingent thinking is not required, than in simultaneous move games (see for example Esponda and Vespa (2014), Ngangoué and Weizsäcker (2021) and Levin et al. (2016)).[5] Our $COND$ treatment is similar to the sequential-moves treatments of these experiments, but stronger in the sense that we fully remove the need for subjects to condition on others' decisions.

Our paper aims at evaluating to what extent deviations from equilibrium are driven

[3]Several theoretical models have been used to explain empirical deviations from BNE, such as cursed equilibrium (Eyster and Rabin, 2005), level-k (Nagel, 1995; Stahl and Wilson, 1995; Crawford and Iriberri, 2007; Brocas et al., 2014), analogy-based expectation equilibrium (Jehiel, 2005; Jehiel and Koessler, 2008) and behavioral equilibrium (Esponda, 2008). These models are usually interpreted as relaxations of the requirement of belief consistency imposed by BNE. Some of these models clearly fit this interpretation, as in the case of level-k models, but others can be thought of reduced forms for a variety of different mechanisms. For instance, although the cursed equilibrium model can be given a literal interpretation in terms of beliefs, one can also interpret the cursedness parameter of the model as a measure of limited strategic thinking more broadly defined. This is the interpretation we use in our paper.

[4]There other possible explanations. For example limits in Bayesian updating (Levin et al., 2016) and heterogeneity in processing the available information (Charness et al., 2019).

[5]Another experiment documenting failures in contingent reasoning is Martínez-Marquina et al. (2019).

by inconsistent beliefs and limited conditional thinking, respectively. This approach is similar to Koch and Penczynski (2018), who decompose the winner's curse in a common value auction into the two channels. Our methodology differs from theirs in the way we remove scope for limited conditional thinking: while they rely on a transformation of the rules of the game (including state space and payoff function), we achieve this goal by adjusting the information structure. The other methodological innovation we make in this paper is measuring how different treatments affect the players' responsiveness to *changes* in adverse selection. We manipulate the degree of adverse selection faced by subjects by changing other players' payoffs in $BASE$, by changing the decision rules of automated players in BEL and by changing the the information structure in $COND$. Observing the comparative statics of subjects' behavior across treatments allows us to make inferences about inconsistent beliefs and limited conditional thinking.[6]

The rest of the paper is organized as follows. In section 2 we introduce the matching game that we implemented in the baseline version of our experiment and derive comparative statics predictions. In section 3 we describe the lab implementation and we explain how our design aims to identify the mechanism driving possible deviations from equilibrium. In section 4 we present the results from our three treatments and estimate the structural model. Section 5 concludes.

2 Theoretical Framework

In section 2.1, we describe the matching game that we implemented in the baseline version of our experiment. In section 2.2 we derive equilibrium predictions under the standard assumption of fully sophisticated players. In section 2.3 we derive predictions when players do not fully anticipate the acceptance curse.

[6]Our comparative statics exercise is related to the experimental methodology used to identify individual levels of rationality in games, see for example Kneeland (2015). In the language of Kneeland (2015), subjects in our $BASE$ treatment seem to satisfy rationality but not second-order rationality, as they respond to changes in zeroth-order payoffs (their own) but not to changes in first-order payoffs (the payoffs of their counterparts).

2.1 A Matching Game

Consider the following Bayesian game. There are two players: $i = 1, 2$. Each player has a quality $q_i \in \{H, M, L\}$ (we will denote by q, q', q'' realizations of q_i). Each type has the same probability (1/3). Each agent i in a pair receives a noisy signal s_i about the partner's quality q_j, $s_i \in \{h, m, l\}$ (we will denote by s, s', s'' realizations of s_i). We call the pair (q_i, s_i) the type of player i. We denote the signal likelihoods by $\delta(q, s) \equiv Pr[s_i = s | q_j = q]$ and the posterior probability of q_j conditional on signal s_i by: $\pi(q, s) \equiv Pr[q_j = q | s_i = s]$. In our experiment, we use the parameterization illustrated in Table 1. We chose this information structure because it simplifies the problem of the players: an h signal reveals that the true quality of player j is H and similarly an l signal reveals player j's quality is L. Thus the only uncertainty faced by a player is how to interpret an m signal.

Signal Likelihoods $\delta(q_j, s_i)$

		s_i: h	m	l
q_j	H	1/2	1/2	0
	M	0	1	0
	L	0	1/2	1/2

Posterior Probabilities $\pi(q_j, s_i)$

		s_i: h	m	l
q_j	H	1	1/4	0
	M	0	1/2	0
	L	0	1/4	1

Table 1: Signal Likelihoods and Posterior Probabilities

After observing the signal about the partner's quality, each player i chooses whether to propose to form a match ($a_i = P$) or not ($a_i = N$). These decisions are simultaneous. If each agent in a pair chooses to propose then a match occurs and agents earn match payoffs. In this case, each agent's payoff depends only on his partner's quality.[7] We denote player i's match payoff by $\mu(q_j)$. If at least a player chooses not to propose,

[7]Our assumption that match payoffs depend only on the partner's quality follows Chade (2006) and other seminal search and matching models such as Burdett and Coles (1997). This assumption provides the simplest model of a market where preferences over partners are homogenous, that is a vertically differentiated market. We believe that our findings are also applicable to situations where match payoffs depend on both q_i and q_j (e.g. displaying complementarities in qualities) as long as the ranking of potential partners is the same for all players.

then each agent earns a reservation value that depends only on his own quality: the reservation payoff of player i is denoted $\rho(q_i)$. The game is illustrated in Figure 1.

1 \ 2	P	N
P	$\mu(q_2), \mu(q_1)$	$\rho(q_1), \rho(q_2)$
N	$\rho(q_1), \rho(q_2)$	$\rho(q_1), \rho(q_2)$

Figure 1: Actions and Payoffs

Motivated by the search and matching literature (and in particular Chade (2006)), we look at situations where higher quality players have higher reservation values, that is:[8]

$$\mu(H) > \mu(M) > \mu(L) \text{ and } \rho(H) > \rho(M) > \rho(L)$$

The actual values used in the experiment are reproduced in Table 2. We consider two versions of this game, called A and B. The only difference between game A and game B is the reservation value of H-quality players, which is higher in game A than in game B. We chose these values in order to obtain different equilibrium predictions in the two versions of the game, as explained in the next section.

	q		
	H	M	L
$\mu(q)$	160	80	40
$\rho(q)$, game A	100	75	25
$\rho(q)$, game B	80	75	25

Table 2: Payoff Values

[8] In search and matching models such as Chade (2006), reservations values are endogenous and determined by future search outcomes. Because we use a one-shot game, we impose this payoff structure. We return to this point in the conclusions of the paper.

2.2 Equilibrium Analysis

We denote by $\sigma_i(q, s) \equiv Pr[a_i = P|q_i = q, s_i = s]$ the probability that player i chooses to propose given his quality and observed signal. Because the game is symmetric, we focus on symmetric equilibria and drop the i index, denoting a generic player's strategy by $\sigma(q, s)$. At this point it is also useful to define player i's "acceptance" probability, defined as the probability of player j proposing conditional on player i's type:

$$\alpha(q, s) \equiv Pr[a_j = P|q_i = q, s_i = s]$$

In deciding whether to propose or not, a player needs to compare the expected match payoff with his reservation value. Denote by $v(q, s)$ player i's expected match payoff conditional on the observed signal $s_i = s$, on own quality $q_i = q$ and on player j proposing. This is computed as:

$$v(q, s) = \sum_{q'} \beta(q', q, s)\mu(q')$$

where $\beta(q', q, s)$ is the posterior belief that player j's quality is q', conditional on the observed signal $s_i = s$, on own quality $q_i = q$ and on player j proposing, that is:

$$\beta(q', q, s) \equiv Pr[q_j = q'|s_i = s, a_j = P, q_i = q]$$

In Appendix A we provide expressions for the posterior beliefs and the acceptance probability.

Finally, in deciding whether to propose or not a player needs to compute the expected gain from proposing relative to not proposing:

$$\Delta(q, s) \equiv \underbrace{\alpha(q, s)v(q, s) + [1 - \alpha(q, s)]\rho(q)}_{\text{expected payoff from P}} - \rho(q) = \alpha(q, s)[v(q, s) - \rho(q)]$$

Choosing P is a best response for type (q, s) whenever $\Delta(q, s) \geq 0$.

Now, we can solve for the equilibrium of the games. To describe pure strategy equilibria of the game, it is useful to define the set of signals for which a player with quality q chooses P:

$$\mathcal{P}_q \equiv \{s|\sigma(q, s) = 1\}$$

In a pure-strategy equilibrium, whenever $s \notin \mathcal{P}_q$ then $\sigma(q, s) = 0$. Using the definitions above, we can obtain the following propositions (proofs are in Appendix B and C).[9]

Proposition 1. *In game A, in any Bayes-Nash equilibrium:* $\sigma(H, m) = 0$ *and* $\sigma(M, m) = 0$*. The pure-strategy BNE where most types propose is such that:* $\mathcal{P}_H = \{h\}, \mathcal{P}_M = \{h\}, \mathcal{P}_L = \{h, m, l\}$.

Proposition 2. *In game B, the pure-strategy BNE where most types propose is such that:* $\mathcal{P}_H = \{h, m\}, \mathcal{P}_M = \{h, m\}, \mathcal{P}_L = \{h, m, l\}$.

Propositions 1 and 2 state that while it can be an equilibrium for H- and M-quality players to propose after observing an m signal in game B, H- and M-quality players should never propose after observing an m signal in game A. In the equilibrium of game A provided in proposition 1, M-quality players face an acceptance curse: conditioning on acceptance lowers the likelihood that a potential partner's quality is H after observing an m signal:

$$\underbrace{Pr[q_j = H|s_i = m, a_j = P, q_i = M]}_{\beta(H,M,m)} = 0 < 0.25 = \underbrace{Pr[q_j = H|s_i = m]}_{\pi(H,m)}$$

This is not the case in the equilibrium of game B provided in proposition 2, where these two probabilities are equal. Anticipating the acceptance curse, M-quality players optimally respond by being more selective in game A than in game B.

[9]While in propositions 1 and 2 we focus on the BNE where most types propose, there are other BNE in both games (for example, where no player ever proposes). However, we focus on these equilibria because: 1) always proposing is a weakly dominant strategy for L-quality players and 2) proposing conditional on an h signal is a weakly dominant strategy for any quality. As discussed below, our data show subjects propose with frequency above 80% whenever proposing is a weakly dominant strategy.

2.3 Cursed Equilibrium

So far we have implicitly assumed that players are sophisticated. The model provides distinctive predictions under the assumption that players fail to fully anticipate the acceptance curse. To show this, we use the cursed equilibrium model of Eyster and Rabin (2005).[10] In a cursed equilibrium, a player with type (q, s) correctly anticipates the acceptance probability $\alpha(q, s)$ but fails to fully condition his posterior belief about player j's quality on acceptance. Formally, the posterior belief of a χ-cursed player is:

$$\beta^c(q', q, s; \chi) \equiv \chi\pi(q', s) + (1 - \chi)\beta(q', q, s) \tag{1}$$

In other words, posterior beliefs are a weighted average of the "naive" posterior belief $\pi(q', s)$ (which only conditions on signals) and the "sophisticated" posterior belief $\beta(q', q, s)$ (which also takes acceptance into account). The parameter $\chi \in [0, 1]$ is a measure of the player's failure to condition on acceptance. Under full rationality, $\chi = 0$, the cursed equilibrium coincides with BNE. Equation (1) can be derived from the assumption that player i with $q_i = q, s_i = s$ believes that the probability player j with $q_j = q', s_j = q'$ chooses $a_j = P$ is $\alpha(q, s)$ with probability χ and $\sigma(q', s')$ with probability $1 - \chi$. However, another interpretation of equation (1) is that players fail to fully condition on the hypothetical scenario in which others propose due to cognitive limits (while correctly anticipating the acceptance probability). Independent of the interpretation one can give to equation (1), it is possible to derive theoretical predictions about proposals in the cursed equilibrium of our games. To do this, first we define the

[10] Other models can be used to derive predictions about this case, such as level-k models (Nagel, 1995; Stahl and Wilson, 1995), analogy-based expectation equilibrium (Jehiel, 2005; Jehiel and Koessler, 2008) and behavioral equilibrium (Esponda, 2008). We focus on cursed equilibrium because it provides a one-parameter reduced-form model that is highly tractable for the purpose of the structural estimation approach we will adopt to analyze the experimental data. The cursed equilibrium model is also formally equivalent to a special case of analogy-based expectation equilibrium (Miettinen, 2009).

expected match payoff of a cursed player as:

$$v^c(q,s;\chi) = \sum_{q'} \beta^c(q',q,s;\chi)\mu(q')$$

and then the expected gain from proposing for a cursed player as:

$$\Delta^c(q,s;\chi) \equiv \alpha(q,s)[v^c(q,s;\chi) - \rho(q)]$$

In a cursed equilibrium, each type (q, s) chooses their proposal strategies to maximize $\Delta^c(q, s; \chi)$. In Appendix D we prove the following:

Proposition 3. *In game A, for $\chi \geq \frac{5}{14}$, there is a cursed equilibrium such that: $\mathcal{P}_H = \{h\}, \mathcal{P}_M = \{h, m\}, \mathcal{P}_L = \{h, m, l\}$.*

When players are sufficiently cursed, they simply fail to anticipate the acceptance curse. This leads M-quality players to overestimate the likelihood of player j being an H type after observing an m signal. As a result, M-quality players choose to propose conditional on an m signal. Note however that even in the cursed equilibrium H-quality players do not propose conditional on m signals. It is also possible to show that in game B, the BNE identified in proposition 2 is a cursed equilibrium for any χ. Thus, failures to anticipate the acceptance curse can results in M-quality players proposing conditional on m signals in both games.

While the previous analysis provides useful qualitative insights about the comparative statics of the model, the predictions are unlikely to be borne out in the lab due to noise in subjects' decisions. Subjects in the lab are likely to play actions P and N with interior probabilities even when BNE or cursed equilibrium are in pure strategies. To obtain more realistic predictions, we use a model that allows for both stochastic decisions and failures in strategic thinking.

To model stochastic decision we use the logit formulation of the quantal-response equilibrium model (see for example Goeree et al., 2016) and combine this with the cursed equilibrium model. A similar approach has been used in Camerer et al. (2016) and Carrillo and Palfrey (2009). In this framework, the probability of choosing action P depends on the perceived gain from proposing according to the following equation:

$$\sigma(q,s) = \left[1 + e^{-\lambda\Delta^c(q,s;\chi)}\right]^{-1}$$

where the parameter $\lambda \geq 0$ measures the responsiveness of decisions to perceived payoffs. When $\lambda = 0$ all actions are played with probability 50% and as $\lambda \to \infty$ the model converges to a cursed equilibrium. To illustrate the predictions, we solve the model numerically for different values of λ and χ and plot the results in Figures 2 and 3. The plots show the predicted proposal rates $\sigma(\cdot)$ for types (H, m) and (M, m) in game A and game B. These are the key variables in our empirical analysis. For clarity we do not show the proposal strategies of other types: all these types always have a weakly dominated strategy across both games (such as N for L-quality types) and therefore the comparative statics are less interesting.

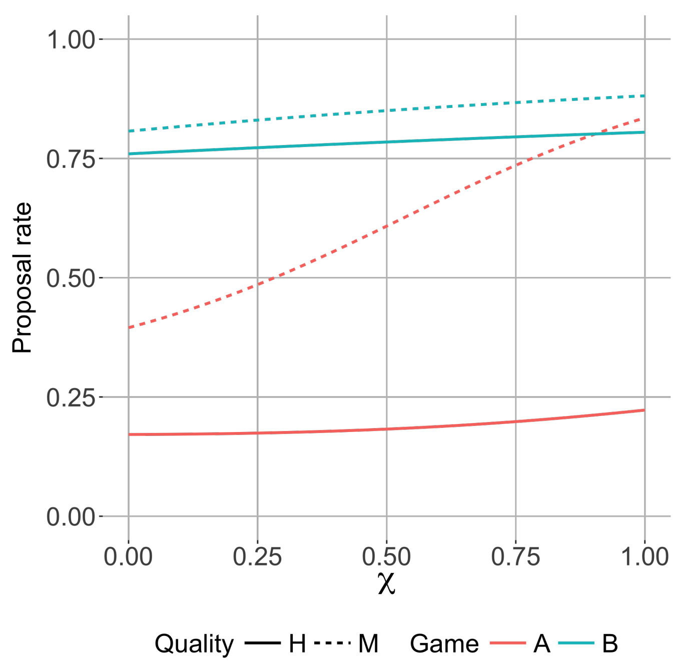

Figure 2: Predicted proposal rates for m signals ($\lambda = 0.15$).

Figure 2 plots the predicted proposal rates against χ, for a medium value of λ. Solid lines represent H-quality types and dashed lines represent M-quality types. The proposal probability $\sigma(H, m)$ is above 50% in game B (solid blue line) and below 50% in game A (solid red line). The intuition for this is consistent with our previous equilibrium analysis: H-quality types are more selective in game A than in game B. Moreover, the behavior of H-quality players is not significantly affected by the χ parameter. The proposal probability $\sigma(M, m)$ in game B (dashed blue line) is above 50% and again does not vary much with χ. This follows from the fact that in this game H-players are not selective and thus ignoring the limited degree of adverse selection present in this environment has only minor effects. The proposal probability $\sigma(M, m)$ is instead highly sensitive to χ in game A (dashed red line). When $\chi = 0$, fully sophisticated players correctly perceive that not proposing is the best response for type (M, m), given that H-players are very selective. As a result, (M, m) types propose less than 50% of the time. As χ increases, the perceived gain from proposing increases and thus the (M, m) proposal rate rises. When $\chi = 1$, the proposal rate of (M, m) types in game A is just a few percentage points below game B.

The comparative statics of this model are qualitatively similar for different values of λ. To illustrate this point, Figure 3 plots proposal rates for a lower λ (= 0.05) and a higher λ (= 0.25). By comparing these plots, it is possible to observe several regularities. First, $\sigma(H, m)$ and $\sigma(M, m)$ are lower in game A than in game B. Second, the difference in $\sigma(H, m)$ between games is roughly constant in χ. Third, the difference in $\sigma(M, m)$ across games is decreasing in χ. Figure 3 also helps us to illustrate the main effect of the parameter λ. When λ is lower, the proposal rates become closer to 50% and less responsive to both changes in reservation values and changes in χ.

It is important to note that the parameters λ and χ have different effects on predicted behavior. Our comparative statics exercises suggest a distinctive prediction of cursedness: large values of χ can result in striking differences between the behavior of H- and M-quality players. To make this statement more precise we solve the model

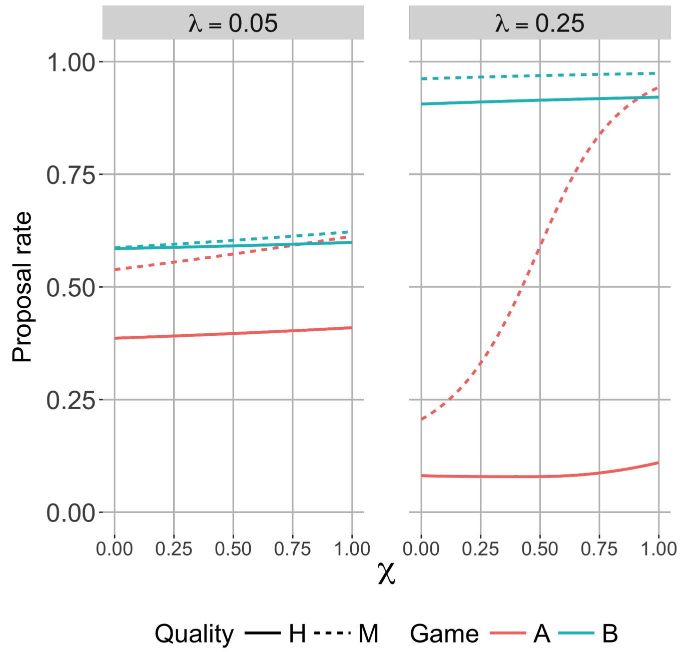

Figure 3: Predicted proposal rates for m signals (low and high λ).

numerically and obtain the following:

Proposition 4. *Consider a profile of proposal probabilities that simultaneously satisfies the following conditions:*

1. *differences in* (H, m) *proposal rates between game A and game B are large* ($>$ 25%)*;*
2. *differences in* (M, m) *proposal rates between game A and game B are small* ($<$ 5%)*.*

In order to rationalize such a profile of proposal probabilities with a cursed QRE model, the value of χ *must be large* (> 0.5)*.*

3 Experimental Design

We implemented the model of section 2 in the lab to test whether subjects anticipate the acceptance curse. Our experiment consisted of three treatments, summarized in Table 3. In section 3.1 we describe the $BASE$ treatment. In section 3.2, we describe two additional treatments, BEL and $COND$, that are aimed at identifying the mechanism driving possible deviations from equilibrium in the baseline environment. In section 3.3 we provide more details about the experiments.

Treatment	Number of subjects	Number of rounds
$BASE$	48	60
BEL	48	60
$COND$	48	60

Table 3: Experimental Design

3.1 The $BASE$ Treatment

Our $BASE$ treatment directly implements the games described in the previous section. Each subject participated in 60 repetitions of the game, called rounds. The 60 rounds are divided into 30 A games and 30 B games, but the actual sequence is random (drawn in advance and constant across all the $BASE$ sessions). In each round, subjects are randomly matched in pairs within a group of 8 participants. Although each subject interacts with the same partner more than once during a session, all the interactions are anonymous.

In each round, the quality of a player (called "type" in the experiment) is randomly determined. To keep a neutral language, in the experiment qualities H, M and L are relabelled as X, Y and Z respectively, but here we keep our original labels for clarity. As in the theoretical analysis, the quality of a player is equally likely to be H, M or L in a given round. In a round, each subject is informed of his quality. Although subjects are not informed of the quality of their partners, they each receive a clue about it, as in the game.

Subjects are told that in order to determine the clue, the computer will digitally draw a random ball from an urn containing 24 balls of different colors. The randomly drawn ball can be either red, yellow or blue. These clues correspond to signals h, m and l in our theoretical framework, respectively. The number of blue, yellow and red balls in the urn is determined by the quality of the subject's partner. The urns used in the experiment are illustrated in Figure 4. If the partner's quality is H, the urn

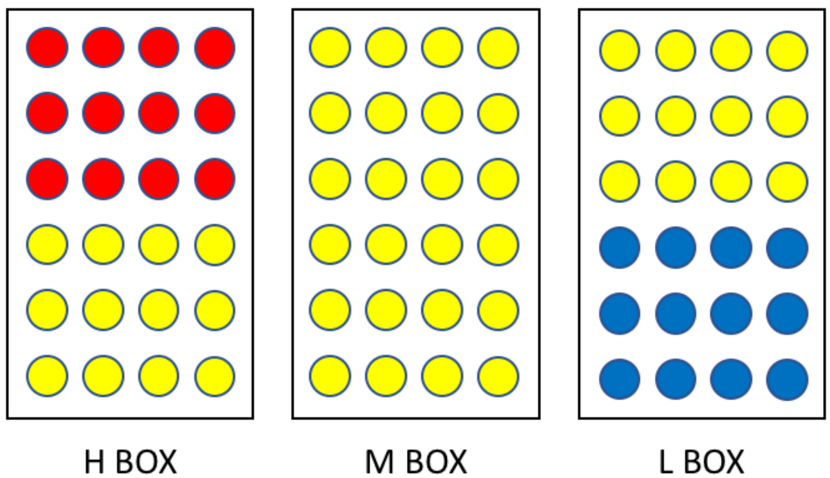

Figure 4: Urns Used to Explain the Information Structure in the Experiment.

contains no blue balls, 12 yellow balls and 12 red balls. If the partner's quality is M, the urn contains 24 yellow balls but no blue or red balls. If the partner's quality is L, the urn contains 12 blue balls, 12 yellow balls and no red balls. Subjects are informed that the computer will first determine which urn to use given their partner's quality and then it will digitally draw a random ball: each single ball in the urn has the same probability of being selected, equal to 1/24. Note that this process exactly replicates the information structure of our model, presented in Table 1.

After receiving the clue, each subject chooses whether they want to form a match (called "partnership" in the experiment) or not. Only if both players in a pair choose to form a match, a match is formed. Subjects are informed of the exact amount of payoffs they can earn by matching with different types of partners and of the payoffs players with different qualities will receive if no match occurs. Participants are informed of the payoff structure for both game A and game B and are aware of which game they are playing in the current round. The payoff values used in the experiment are exactly those reproduced in Table 2 (see below for more details on subject payments).

3.2 The *BEL* and *COND* Treatments

When designing our $BASE$ treatment, we anticipated the possibility of deviations from equilibrium. We thus designed additional treatments aimed at testing two potential mechanisms proposed by the literature: limits in forming consistent beliefs and limits in conditional thinking. Subjects may deviate from equilibrium behavior if they do not anticipate how the proposal decisions of others depend on others' qualities (and their reservation payoffs). Subjects may also fail to properly condition expectations about the match quality on the hypothetical scenario in which a potential partner chooses to propose (no matter whether their beliefs about the strategies of others are correct or not). We aim at separating these two potential mechanisms through two treatments, BEL and $COND$, which sequentially eliminate the conditions that allow these explanations to account for behavior. Our BEL treatment aims at eliminating the effects of limits in the subjects' ability to form consistent beliefs, while our $COND$ treatment eliminates scope for both inconsistent beliefs and limited conditional thinking.

In the two treatments, each subject plays the role of player 1 and faces an individual choice problem obtained by assigning exogenous strategies to player 2. Because our goal is to explain non-equilibrium behavior of M-quality players, we fix player 1's quality to be M. In the Beliefs treatment, after observing an m signal, the automated player 2 always proposes if his quality is either M or L, but proposes with probability p if his quality is H:[11]

$$\sigma_2(H,m) = p, \sigma_2(M,m) = \sigma_2(L,m) = 1$$

By varying p we change the degree of adverse selection faced by player 1: a lower value of p corresponds to a higher degree of adverse selection. We consider five situations: $p \in \{0, 0.25, 0.5, 0.75, 1\}$. We call each of this five settings a BEL-task. In the experiment, subjects are aware of the the strategies followed by automated players. Thus, our $COND$ treatment is similar to the robot protocol used in the literature (Charness and Levin, 2009; Ivanov et al., 2010; Esponda and Vespa, 2014; Koch and Penczynski, 2018).

In the $COND$ treatment, we design five decision problems, called $COND$-tasks, that are formally equivalent to the five BEL-tasks except that we remove the need for subjects to condition on acceptance. We do this by making two changes to the task. First, the strategy of player 2 is set to always propose after an m signal, independently of his own type:

$$\sigma_2(H,m) = \sigma_2(M,m) = \sigma_2(L,m) = 1$$

Second, the information structure is adjusted so that the posterior distribution of qualities conditional on signal $s_1 = m$ is equal to the posterior distribution obtained conditioning on $s_1 = m$ *and acceptance* in the equivalent BEL-task. Before we explain the exact way in which we adjust the information structure, it may be useful to consider the following simple example. In the BEL-task with $p = 0$, the posterior probability that player 2's quality is H conditional on drawing a yellow ball and on acceptance is zero. The reason is that player 1's quality is M and in this setting an H player never proposes to an M player. In the equivalent $COND$-task, we replace all the yellow balls in the H urn with red balls, so that the likelihood of obtaining a yellow ball from this urn is zero. As a result, even though player 2 always proposes, the posterior probability that player 2's quality is H conditional on drawing a yellow ball is zero, as in the original

[11] In both BEL and $COND$ treatments, player 2's strategy is to always propose after receiving an h signal and to propose after receiving an l signal if and only if its type is L. Note however that since the quality of player 1 is always equal to M, the only contingency that realizes is $s_2 = m$.

BEL-task.

More generally, for each value of p, we calculate the $COND$-task likelihoods in the following way. First, we compute the posterior probability in the BEL-task, denoted $\phi(p)$:

$$\phi(p) \equiv Pr[q_2 = H | s_1 = m, q_1 = M, a_2 = P; BEL\text{-task}]$$

Then we require the posterior belief in the $COND$-task to equal the posterior belief in the BEL-task:

$$Pr[q_2 = H | s_1 = m; COND\text{-task}] = \phi(p)$$

To achieve this we only adjust the likelihoods of observing h and m signals when player 2's quality is H, while leaving all other likelihoods unchanged. In terms of the urns displayed to subjects, we only change the number of red and yellow balls in the H urn. We derive expressions for the $COND$-task likelihoods as a function of p in Appendix E.

Table 4 summarizes our choice of parameter values for the five tasks in the BEL and $COND$ treatments. The BEL column shows the five value of p, that is the probability player 2 proposes conditional on m signal and his type being H. The four middle columns show information for the $COND$ treatment, namely the likelihoods of an m signal when $q_j = H$, the likelihoods of an h signal when $q_j = H$, the number of yellow balls in the H urn and the number of red balls in the H urn. The "Posterior" column shows the resulting value of $\phi(p)$, that is the posterior probability that player 2's quality is H, for both treatments. When $p = 1$, an H-quality player 2 always proposes in the BEL treatment, just like an H-quality player 2 in the $COND$ treatment. Thus, no adjustment in the $COND$ treatment likelihood is required. Lower values of p correspond to a lower likelihood of observing an m signal when $q_2 = H$ in the $COND$ treatment. When $p = 0$, as in our earlier example, the probability of observing an m signal when $q_2 = H$ in the $COND$ treatment falls to zero. In both the BEL and $COND$ treatments, subjects repeat each of the five tasks 12 times, for a total of 60 rounds (as in $BASE$). The actual sequence of tasks is random but constant across sessions of the same treatment.

BEL	*COND*				Posterior
p	$Pr[m\|H]$	$Pr[h\|H]$	# yellow	# red	$\phi(p)$
1	0.5	0.5	12	12	0.25
0.75	0.375	0.625	9	15	0.20
0.5	0.25	0.75	6	18	0.14
0.25	0.125	0.875	3	21	0.08
0	0	1	0	24	0

Table 4: Parameters for the *BEL* and *COND* treatments

In both the BEL and $COND$ treatments, we use the payoff structure of game A. One potential concern with adopting the standard implementation of the "robot protocol" in our design is that the BEL and $COND$ treatments would remove scope not only for strategic thinking but also for social preferences: while player 1's decision affects the payoffs of player 2 in $BASE$ and thus may be affected by distributional preferences, this is not possible when player 1's opponent is a robot. To make sure any difference between $BASE$ and $BEL/COND$ is not driven by distributional preferences, in both the BEL and $COND$ treatments we assign the realized payoff of player 2 to another human subject. We do this by randomly matching subjects in pairs in each round, within a fixed group of 8 participants, as in the $BASE$ treatment. In each round, one of the two players in a pair is randomly assigned the role of player 1 (we call this an "active" player in the experiment) and the other is assigned the role of player 2 (a "passive" player role). Participants know that the computer will make a decision on behalf of player 2 following the given rules, but both players in a pair earn payoffs.

3.3 Details

We ran this experiment using oTree (Chen et al., 2016) in December 2018 at the LINEEX lab of the University of Valencia. We recruited 48 subjects in each treatment, divided into 6 groups of 8 participants. Subjects were seated at visually isolated terminals and read instructions on their screens (reproduced in Appendix F). We then ran subjects through 2 practice rounds (without pay), before they participated in the full 60-round sequence of the experiment. In all treatments, after each round subjects received feedback about the decision of their partner, the outcome of the game (match or no match), the quality of their partner and their payoffs. At the end of the experiment, one random round was selected for payment. Before finishing the experiment, subjects participated in an incentivized multiple-price list for eliciting risk-aversion (Holt and Laury, 2002), an incentivized three item cognitive reflection test (Frederick, 2005) and a brief survey (about their sex and undergraduate major). See the Appendix for details. Total payoffs were converted into Euros at an exchange rate of 1/6 Euros per point and summed to a 5 Euros show-up fee. The average final payment was 18 Euros and the experiment lasted on average 75 minutes.

4 Experimental Results

In section 4.1 we present results from the *BASE* treatment. In section 4.2 we present results from the *BEL* and *COND* treatments. Finally, in section 4.3 we explain our structural estimation approach and provide estimates of the model using data from all treatments. In the following sections, we use only observations from round 21 to 60 in each session. We do this to allow for some learning, but the results are similar if we use the whole sample.

4.1 Base Treatment

We begin presenting the experimental data by analyzing subjects' decisions in the *BASE* treatment. Figure 5 and Figure 6 plot the aggregate proposal rates in game A and game B respectively, pooling across subjects and rounds. We first argue that the aggregate proposal rates show subjects understood the basic rules of the game since they often played a dominant strategy whenever one was available to them. For example,

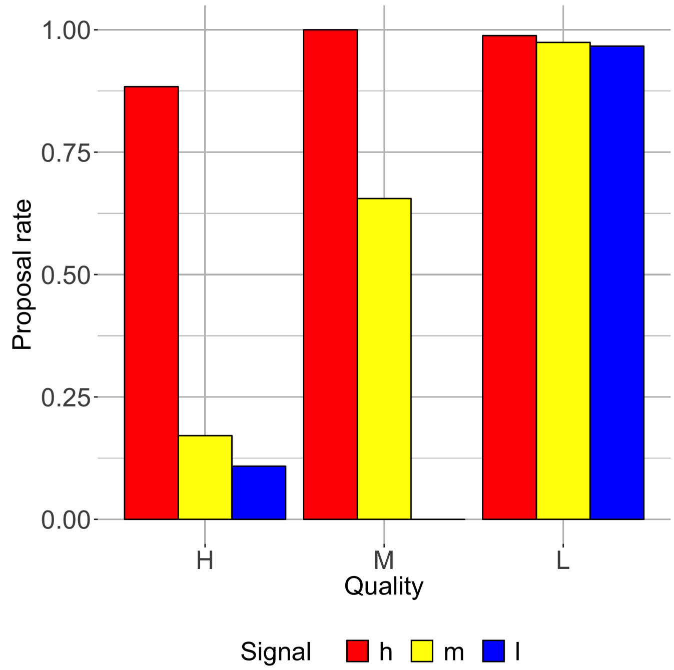

Figure 5: Treatment *BASE*, game A : proposal rates

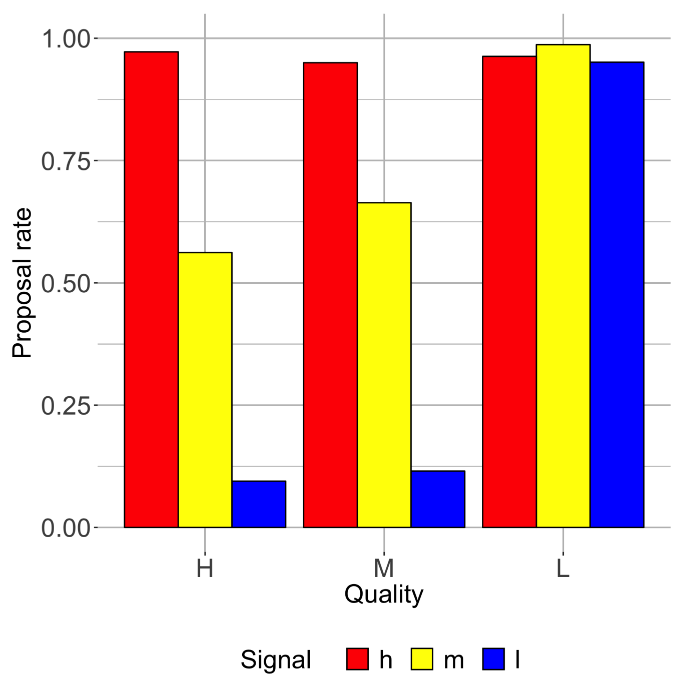

Figure 6: Treatment *BASE*, game B : proposal rates.

both Figure 5 and Figure 6 show that L-quality players propose almost 100% of the time independently of the signal they receive. This is reassuring because proposing is indeed a weakly dominant strategy for these types. Similarly, H- or M-quality players proposed almost all the times they observed h signals, but only few of the times they observed an l signal.

Result 1 (BASE)**.** *The proposal rates of types for whom proposing is a weakly dominant strategy are above 80%. The proposal rates of types for whom proposing is a weakly dominated strategy are below 20%.*

Another way to assess subjects' decisions is to determine whether their strategies were empirical best responses, given the actual distribution of strategies in our $BASE$ sessions. To do this, we compute the expected gain from proposing, $\Delta(q, s)$, for different types using the empirical distribution of strategies. We find most of the strategies that are played with a frequency above 50% are indeed empirical best responses, with one important exception. We have already pointed out that L-quality players almost always chose their weakly dominant strategies and this immediately implies they are almost always best-responding. Similarly, H- and M-quality types are best responding at very high rates conditional on h and l signals. More interestingly, we find that H- and M-quality players propose most of the time (with frequency 56% and 66%, respectively) conditional on an m signal in game B, where this strategy is indeed an empirical best response. At the same time, H-quality players choose to propose only 17% of the time in game A, where proposing is not an empirical best response. The only exception to this pattern of frequent best responses is the behavior of M-quality players in game A: while in this game proposing conditional on an m signal is not an empirical best response for M-quality players, they propose 66% of the time.

Result 2 (BASE)**.** *The proposal rates of types for whom proposing is (is not) an empirical best response are above (below) 50%, with only one exception: proposal rates for* (M, m) *types in game A. In this game, proposing is not an empirical best response for type* (M, m) *but these types of players propose 66% of the time.*

We note that the behavior of (M, m) types is consistent with limited strategic thinking. The high frequency with which M-quality propose after observing m signals could be also due to risk-seeking attitudes. However, this explanation seems inconsistent with the fact that our multiple-price list task reveals subjects to be risk-averse on average. More importantly, risk preferences cannot explain our next finding, namely the fact the proposal rates of (M, m) types are essentially constant between game A and B. Indeed, (M, m) types propose 66% of the time in both games. Unsurprisingly, we cannot reject the null hypothesis of no difference in (M, m) proposal rates between game B and game A using a Wilcoxon test (p-value: 0.8438).[12]

To provide further evidence about the comparative statics of our experiment, we use regression analysis. Restricting attention to observations consisting of an H- or M-quality player receiving an m signal, we estimate the following equation:

$$\textit{propose} = b_0 + b_1 H + b_2 A + b_3 H \times A + \epsilon \tag{2}$$

In this equation, the variable *propose* equals one if the subject chose to propose conditional on an m signal. The variable H equals one if the subject's quality was H (the excluded category is M quality). The variable A equals one if the subject is playing in game A (the excluded category is game B). We estimate equation (2) using either a linear probability model or a logit regression and always clustering errors at both the individual and group levels. Table 5 reports the results from regression (2). The H coefficient estimates show that H-quality players are significantly less likely to propose (conditional on an m signal) than M-quality players in game B. More interestingly, our estimate of the coefficient on the dummy variable for game A is small and insignificant. Thus, M-quality players do not adjust their proposal rate between game A and game B on average. On the other hand, the estimate on the interaction term $H \times A$ is large in magnitude and highly significant, showing that H-types propose at significantly lower

[12] For this test, each observation is the group-level average (M, m) proposal rate. We have 6 groups in the $BASE$ treatment, resulting in 12 paired A-B observations.

	Dependent variable: *propose*	
	Linear	Logit
H	-0.105^{***}	-0.447^{***}
	(0.002)	(0.012)
A	-0.008	-0.034
	(0.029)	(0.129)
$H \times A$	-0.383^{***}	-1.787^{***}
	(0.018)	(0.190)
Constant	0.668^{***}	0.701^{***}
	(0.005)	(0.022)
Clustering	Yes	Yes
Observations	839	839
Log Likelihood	-530	-504
Akaike Inf. Crit.	1,068	1,016

Note: This table reports estimates from regression (2). The regressions are estimated using only observations of (H, m) and (M, m) types. Standard errors are clustered at individual and group level in parentheses. *p<0.1; **p<0.05; ***p<0.01

Table 5: Proposal decisions in *BASE*

rates in game A relative to game B. To sum up, we obtain the following:

Result 3 (BASE). *While proposal rates of (H, m) types differ significantly between game A and game B, proposal rates of (M, m) types do not. This is qualitatively consistent with highly cursed behavior.*

Our finding that aggregate (M, m) proposal rates are similar between game A and game B is not an artifact of averaging different behaviors across subjects. To prove this statement, we show how proposal rates are affected by observable characteristics such as gender or experience. We estimate an augmented version of regression (2):

$$propose = \sum_{k=0}^{K} b_0^k X^k + \sum_{k=0}^{K} b_1^k H X^k + \sum_{k=0}^{K} b_2^k A X^k + \sum_{k=0}^{K} b_3^k H A X^k + \epsilon \qquad (3)$$

where X^k denotes one of $K+1$ variables. This model includes $K = 5$ observables, namely $CRT, RiskAversion, Female, STEM, Round$ and a constant ($X^0 = 1$). The variable CRT is the subject's standardized score in the Cognitive Reflection Test (Frederick, 2005). To compute the variable $RiskAversion$ we use the subject's switching point in the multiple price list (Holt and Laury, 2002) and then we standardize it.[13] Thus, a subject with a higher $RiskAversion$ measure is more risk averse. The variable $Female$ equals one for women and zero for men. The variable $STEM$ is a dummy equal to one if the subject's undergraduate major is science, technology, engineering or mathematics and 0 otherwise (excluded categories from our survey: social sciences, humanities). The variable *Round* is computed by subtracting 30 to the actual round: thus, when $Round = 0$ the subject has experienced half of the 60 rounds in the experiment.

As before, we estimate regression (3) using both a linear probability model and a logit model, while clustering errors at the individual and group levels. The estimates are reproduced in Table 6. From these results it is possible to draw several conclusions. We are particularly interested in how the behavior of M-quality types varies between game B and game A conditional on different observables X^k, as measured by the interaction terms $A \times X^k$. First, note that the coefficient on A is not statistically significant. This means that there are no differences in (M, m) proposal rates of non-STEM male students with average CRT score, average risk-aversion and with 30 rounds of experience between game A and game B. Moreover, this pattern is not affected by gender nor by CRT score (the coefficients on $A \times Female$ and $A \times CRT$ are not statistically significant).[14] Thus, Result 3 appears fairly robust as it holds not only in the aggregate but also for different types of individuals.

Being a STEM student, having more experience in the game or being more risk-averse increase the responsiveness of (M, m) proposal rates to the game in the direction consistent with equilibrium: the coefficients on $A \times STEM$, $A \times Round$ and $A \times$

[13] To standardize individual i's raw measure x_i, we compute $\frac{x_i - mean}{sd}$, where $mean$ is the mean score in our population and sd is the standard deviation of the score in our population. We do this for both CRT and $RiskAversion$.

[14] Interestingly, although the CRT score does not affect the responsiveness of M-types, it does affect the behavior of H-types. The coefficient on $H \times A \times CRT$ is negative and highly significant. Thus, subjects with higher CRT scores propose less often conditional on an m signal when their reservation value is higher, relative to subjects with lower CRT scores. This suggests that CRT score is correlated with a subjects' ability to understand the rules of the game or pay attention to his own payoff, but not with the subject's ability to think strategically about others.

Dependent variable: *propose*		
	Linear	Logit
H	0.117 (0.169)	0.573 (0.761)
A	0.103 (0.093)	0.469 (0.433)
CRT	0.042 (0.046)	0.204 (0.255)
Round	−0.0005 (0.006)	−0.002 (0.028)
Risk Aversion	0.003 (0.025)	0.007 (0.140)
Female	0.123 (0.181)	0.569 (0.878)
STEM	0.080 (0.121)	0.365 (0.577)
$H \times A$	-0.622^{***} (0.154)	-3.145^{***} (0.812)
$H\times$ CRT	0.092^{***} (0.015)	0.409^{***} (0.053)
$A\times$ CRT	0.002 (0.021)	0.002 (0.121)
$H\times$ Round	−0.006 (0.007)	−0.025 (0.033)
$A\times$ Round	-0.003^{*} (0.002)	-0.013^{**} (0.007)
$H\times$ Risk Aversion	0.032 (0.104)	0.149 (0.413)
$A\times$ Risk Aversion	-0.031^{***} (0.001)	-0.133^{***} (0.012)
$H\times$ Female	−0.133 (0.129)	−0.611 (0.590)
$A\times$ Female	−0.052 (0.084)	−0.243 (0.410)
$H\times$ STEM	−0.158 (0.274)	−0.709 (1.291)
$A\times$ STEM	-0.084^{**} (0.042)	-0.390^{*} (0.234)
$H \times A\times$ CRT	-0.218^{***} (0.054)	-1.353^{***} (0.130)
$H \times A\times$ Round	0.007^{**} (0.003)	0.023^{**} (0.012)
$H \times A\times$ Risk Aversion	−0.021 (0.047)	−0.103 (0.089)
$H \times A\times$ Female	0.014 (0.094)	−0.163 (0.656)
$H \times A\times$ STEM	0.305 (0.260)	1.892 (1.247)
Constant	0.559^{***} (0.118)	0.217 (0.513)
Clustering	Yes	Yes
Observations	839	839
Log Likelihood	−508	−481
Akaike Inf. Crit.	1,064	1,009

Note: This table reports estimates from regression (3). The regressions are estimated using only observations of (H, m) and (M, m) types. Standard errors are clustered at individual and group level in parentheses. *p<0.1; **p<0.05; ***p<0.01

Table 6: Proposal decisions and individual characteristics in *BASE*

$RiskAverse$ are negative and significant. A possible interpretation of these estimates is that subjects with $STEM$-related cognitive skills, more experience or a stronger dislike for risk are more likely to be aware of the acceptance curse. However, the magnitudes of these effects are not large: for example, the difference in proposal rates between A and B for STEM students is only 8.4% larger (in absolute value) than non-STEM students according to our linear model estimates (see the $A \times STEM$ interaction coefficient in the "Linear" column). This is a small magnitude compared to the difference in proposal rates between game A and game B for H-quality players: the coefficient on the $A \times H$ interaction in our linear model suggests subjects are 62% less likely to propose (conditional on an m signal) in game A than in game B when they play the role of H-types. We summarize this discussion in the following:

Result 4 (BASE)**.** *Even after controlling for individual characteristics and experience in the game, we find that the magnitude of the difference in* (M, m) *proposal rates between game A and game B is much smaller than the difference in* (H, m) *proposal rates.*

4.2 *BEL* and *COND* Treatments

Results from the $BASE$ treatment show subjects behave in a way consistent with limited strategic thinking. The rate at which M-quality players propose conditional on an m signal does not respond significantly to adverse selection. The BEL and $COND$ treatments allow us to observe how behavior changes when we remove scope for limits in forming consistent beliefs and for limits in conditional thinking respectively. We focus on the behavior of M-quality players who observe an m signal and analyze how their proposal rate responds to variation in the degree of adverse selection across tasks. Recall that the degree of adverse selection in BEL and $COND$ tasks is determined by the parameter p: when p is lower, the posterior probability that player 2's quality is H is lower. Thus, we use $1 - p$ as our measure of adverse selection. In both BEL and

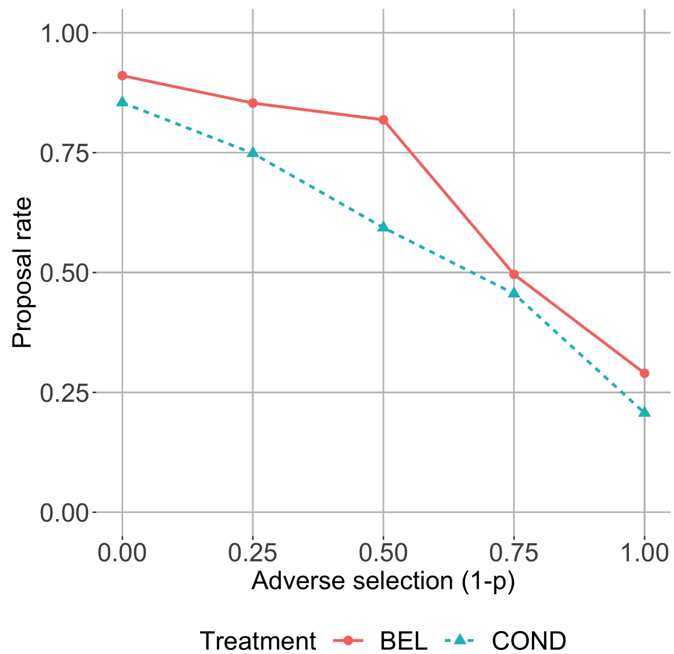

Figure 7: Proposal rates conditional on signal m in BEL and $COND$

$COND$, subjects experience five tasks corresponding to $(1-p) = 0, 0.25, 0.5, 0.75, 1$ (each repeated multiple times).

In Figure 7 we plot the aggregate proposal rates of M-quality players conditional on m signals in BEL and $COND$. As this figure shows, proposal rates appear responsive to the degree of adverse selection $(1-p)$. To provide further evidence of this we estimate regressions of the probability of proposing in each treatment:

$$propose = \sum_{k=0}^{K} a_0^k X^k + a_1 AdverseSelection + \epsilon \tag{4}$$

where the X^k variables include a constant and five observables (CRT, $RiskAversion$, $Female$, $STEM$, $Round$) and $AdverseSelection = (1-p)$. We estimate equation (4) separately for BEL and $COND$. For each treatment we run a linear model and a logit model, always clustering errors at the individual and group levels. The results are reproduced in Table 7. The coefficient on $AdverseSelection$ is negative and highly significant, thus leading to the following:

Result 5. *Proposal rates of (M, m) types respond significantly to changes in the degree of adverse selection $(1 - p)$ in both BEL and $COND$.*

	Dependent variable: *propose*			
	BEL		COND	
	Linear	Logit	Linear	Logit
AdverseSelection	−0.600***	−3.501***	−0.617***	−3.023***
	(0.024)	(0.131)	(0.014)	(0.387)
Constant	1.085***	3.503***	1.321***	5.054***
	(0.151)	(0.919)	(0.234)	(1.510)
Individual Characteristics	Yes	Yes	Yes	Yes
Clustering	Yes	Yes	Yes	Yes
Observations	624	624	579	579
Log Likelihood	−316	−300	−338	−321
Akaike Inf. Crit.	648	617	693	658

Note: This table reports estimates from regression (4). The regressions are estimated using only observations of subjects in the role of player 1 with (M, m) type. The variable *AdverseSelection* is given by the round-specific parameter $1 - p$. All regressions control for individual characteristics. Standard errors are clustered at individual and group level in parentheses. *p<0.1; **p<0.05; ***p<0.01

Table 7: Proposal decisions in *BEL* and *COND*

This result suggests that limited strategic thinking is an important driver of behavior in $BASE$. In $BASE$ behavior is unresponsive to changes in adverse selection induced by manipulation of reservation values. When we provide information on strategies, subjects are highly responsive. Back on the envelope calculations based on our estimates of regression (4) imply that (M, m) proposal rates in $BASE$ should change by more than 20% points from game B to game A, when in fact we cannot find any significant effect in our data.

The second observation about Figure 7 is that the differences between BEL and $COND$ appear to be small. This is also suggested by the fact that estimated coefficient in regression (4) is very similar between BEL and $COND$. To formally test this hypothesis, we run non parametric tests of differences in (M, m) proposal rates between BEL and $COND$. For each value of the $AdverseSelection$ parameter (i.e. $1 - p$), we use a Wilcoxon test to evaluate the null hypothesis that there are no differences in group-level proposal rates between BEL and $COND$. The results are reported in Table 8. Only for two values of our $AdverseSelection$ parameter (0.25 and 0.5) we can reject the null hypothesis at the 10% significance level. We thus have the following result:

Result 6. *The differences in proposal rates of* (M, m) *types between* BEL *and* $COND$ *are small.*

Providing information about the strategies of potential partners makes subjects more responsive to adverse selection. Further removing the need for subjects to engage in conditional thinking has small cumulative effects on behavior. This result suggests that conditional thinking limits play a minor role in explaining deviations from equilibrium relative to inconsistent beliefs.

$AdverseSelection$ parameter	0	0.25	0.5	0.75	1
Wilcoxon test p-value	0.12	0.06	0.02	0.47	0.21

Note: This table reports p-values of a one-sided Wilcoxon test of differences in proposal rates between BEL and $COND$. For each value of the $AdverseSelection$ parameter, the null hypothesis is that there are no differences in group-level proposal rates between BEL and $COND$. The alternative hypothesis is that proposal rates are higher in BEL than $COND$. For each of these tests, an observation is the average proposal rate in a group. Thus, for each value of $AdverseSelection$, we have 6 observations in BEL and 6 observations in $COND$.

Table 8: Differences in proposal rates between BEL and $COND$

4.3 Structural Estimation

The results presented in the previous sections suggest that limits in forming consistent beliefs are an important driver of deviations from BNE in our game, while limited conditional thinking plays a secondary role. To provide further evidence about the relative importance of these two mechanisms, we adopt a structural estimation approach based on the Cursed Quantal Response Equilibrium model. This framework allows us to quantify the subjects' limits in anticipating adverse selection, as measured by the cursedness parameter χ, while allowing for stochastic choice. By estimating different χ parameters for different treatments, we can separately estimate the effect of inconsistent beliefs and limited conditional thinking. In this approach, we interpret χ as a reduced-form measure of limits in anticipating adverse selection that can potentially arise from different mechanisms. This approach is also suggested by Eyster and Rabin (2005) who note that, although the cursed equilibrium model can be given a literal interpretation in terms of beliefs, the "primary motivation for defining cursed equilibrium is not based on learning or any other foundational justification, but rather on its pragmatic advantages as a powerful empirical tool to parsimoniously explain data in a variety of contexts." To further motivate this approach, recall that the key assumption of the cursed equilibrium model is that posterior beliefs are a weighted average of the "naive" posterior belief (which only conditions on signals) and the "sophisticated" posterior belief (which also takes acceptance into account):

$$\text{belief} = \chi \Pr[q_j = q'|s_i] + (1 - \chi) \Pr[q_j = q'|s_i, q_i, a_j = P]$$

We argue that this assumption can be used to model both a player who wrongly believes others do not fully use information and a player who fails to fully condition on the hypothetical scenario in which others propose (as well as a player who suffers from both mistakes).

In order to identify the effect of each mechanism, we exploit our treatment design. Our estimate of cursedness in the $BASE$ treatment, χ^{BASE}, reflects both limits in forming consistent beliefs and limits in conditional thinking. On the other hand, our estimate of cursedness in the BEL treatment, χ^{BEL}, only reflects limited conditional thinking since it is derived from an environment where there is no scope for inconsistent beliefs (subjects know the strategies of their potential partners). Thus, we can interpret the difference $\chi^{BASE} - \chi^{BEL}$ as a measure of limited ability to form consistent beliefs. Note that our interpretation does not rely on estimates of χ in $COND$; indeed this parameter is unidentified because in $COND$ the naive and sophisticated posterior beliefs are identical by design (since player 2 always proposes).

The other structural parameter is λ, measuring the responsiveness of proposal probabilities to the perceived gains from proposing. Lower values of λ results in proposal rates closer to 50%. We assume the same value of λ across treatments $BASE$, BEL and $COND$. Thus, even though data from $COND$ does not directly identify the χ parameters, it is used to estimate the model. We denote the three parameters to be estimated by $\theta \equiv [\lambda, \chi^{BASE}, \chi^{BEL}]$.

Each observation, indexed by $n \in \{1, ..., N\}$, is at the round-subject level and consists of variables y_n, q_n, s_n, ω_n. The variable $y_n \in \{0, 1\}$ is the proposal decision of the subject in that round (in BEL and $COND$ we use only active subjects who play the role of player 1) while $q_n \in \{H, M, L\}$ and $s_n \in \{h, m, l\}$ denote the quality and signal of the subject in that round.[15] The variable ω_n is an indicator for the treatment and game/task (recall that there are two games in $BASE$ and five tasks in $BEL/COND$, indexed by p):

$$\omega_n \in \{BASE_A, BASE_B, BEL_{p=0}, ..., BEL_{p=1}, COND_{p=0}, ..., COND_{p=1}\}$$

We use ω_n to link to treatment- and round-specific parameters of the game, such as reservation values, match values, signal likelihoods and decision rules of automated players. We estimate the parameters θ by maximum likelihood. With the notation introduced above, the log-likelihood of our model can be written as:

$$\ell = \sum_n \{y_n \log \sigma(q_n, s_n, \omega_n; \theta) + (1 - y_n) \log[1 - \sigma(q_n, s_n, \omega_n; \theta)]\}$$

where $\sigma()$ is the predicted probability of proposing for observation y_n, q_n, s_n, ω_n, given

[15] While in the regressions presented in previous sections we have often restricted attention to rounds in which M- or H-quality players observed m signals, we estimate the structural model using data for all types. The only restriction is that we use data for rounds 21 to 60, as before.

parameters θ. In $BASE$, we compute $\sigma()$ as a fixed-point solution of the QRE model. In BEL and $COND$, player 2's strategy is given and thus $\sigma()$ can be computed directly. We optimize ℓ numerically.

Parameter	Estimate and s.e.
λ	0.074***
	(0.002)
χ^{BASE}	0.990***
	(0.121)
χ^{BEL}	0.265***
	(0.082)
-2 log Lik.	5257
Observations	5760

Note: This table reports maximum likelihood estimates of the structural model parameters. Asymptotic standard errors are reported in parentheses. *p<0.1; **p<0.05; ***p<0.01

Table 9: Maximum likelihood estimation results

Structural estimation results are reported in Table 9. The estimated λ parameter is relatively low, implying stochastic decisions are an important feature of our data. It is however statistically different from zero: thus, subjects respond to perceived payoffs. More interestingly, we find that subjects behave as if they are nearly fully cursed in $BASE$ since χ^{BASE} is close to 1 (and highly significant). This is consistent with the fact that the difference in (M, m) proposal rates between game A and game B is minimal. Our estimate of cursedness in BEL is also statistically significant, but much lower in magnitude: $\chi^{BEL} = 0.265$. Thus, our estimates suggest that mistaken beliefs have a large effect on cursedness, which can be measured as $\chi^{BASE} - \chi^{BEL} = 0.725$. This effect is not only large but also statistically significant: we run a (one-sided) likelihood-ratio test of the null hypothesis that $\chi^{BASE} = \chi^{BEL}$ and reject with a p-value of 10^{-12}. We summarize our findings from the structural estimation in the following:

Result 7. *Structural estimation shows that: 1) subjects behave as if they are fully cursed in $BASE$, 2) they behave as if they are cursed to a much lower degree in BEL, implying a large role of mistaken beliefs.*

5 Conclusion

Do people understand adverse selection in matching markets? If they fail, why? To answer these questions, we have implemented a two-player, one-shot version of a vertically-differentiated market where qualities are private information and agents receive noisy signals about others' qualities. We have found that subjects propose to form matches when Bayes-Nash equilibrium predicts they should not and do not respond to changes in reservation values of other types. These findings suggest most subjects fail to anticipate the acceptance curse in matching markets. To understand the driving mechanism, we have designed two additional treatments. In one treatment, we remove scope for mistaken beliefs by providing information about the strategies of automated players. In this treatment, subjects are highly responsive to changes in adverse selection. Further removing scope for limited conditional thinking has small additional effects on subjects' behavior. Overall, our data show that mistaken beliefs are the key driver of failures to anticipate adverse selection in two-sided matching, while limited conditional thinking plays a secondary role.

In our experiment, providing information on others' strategies seems to be very effective in fixing the subjects' failure to anticipate adverse selection. By contrast, several papers have shown that deviations from equilibrium persist under similar "robot protocols" (Charness and Levin, 2009; Ivanov et al., 2010; Esponda and Vespa, 2014). Why is information about strategies so effective in our experiment compared to other studies? One possible explanation is that using information about others' strategies to form conditional expectations is easier in two-sided matching markets than in previously studied environments, such as auctions and elections. We conjecture that once subjects are informed about how selective are different types, it is easy for them to get a sense of the posterior distribution of types conditional on matching. On the contrary, deriving implications about adverse selection from information about others' strategies is a complex cognitive task in auctions and elections, since one needs to combine this information with market clearing rules or voting rules. Even bilateral trading settings like the acquiring-a-company game (Charness and Levin, 2009) are more complicated than our setting because the buyer's offer affects payoffs both directly (as the price he pays) and indirectly through selection. Matching with non-transferable utilities and exogenous signals is a much easier problem in this respect.

In order to provide the clearest possible test of strategic thinking in matching markets, we have used a one-shot game. Although implementing a dynamic search-and-matching market in the lab poses several challenges, it could shed light on issues that we cannot explore within our static framework. In our static game, limited strategic thinking can only result in players proposing when they should not. In a search and matching market, limited strategic thinking can also result in players being excessively selective. Cursed players may turn down profitable matches if they overestimate their option-value of searching. For instance, a low-quality player who fails to anticipate adverse selection and thus overestimates the expected payoff from matching conditional on future m signals may choose to keep searching rather than proposing after observing an l signal.[16] Testing the implications of limited strategic thinking for search behavior is clearly an interesting venue for future research.[17]

[16] A similar mechanism is formally analyzed in Antler and Bachi (2020).

[17] An interesting study on this topic is Araujo et al. (2018), whose findings show subjects fail to understand the dynamics of adverse selection in a non-stationary environment.

References

Ali, S Nageeb, Maximilian Mihm, Lucas Siga, and Chloe Tergiman, "Adverse and Advantageous Selection in the Laboratory," 2019. Forthcoming at *American Economic Review*.

Antler, Yair and Beniamin Bachi, "Searching Forever After," 2020. Forthcoming at *American Economic Journal: Microeconomics*.

Araujo, Felipe A, Stephanie W Wang, and Alistair J Wilson, "The times they are a-changing: dynamic adverse selection in the laboratory," 2018. Forthcoming at *American Economic Journal: Microeconomics*.

Becker, Gary S, Elisabeth M Landes, and Robert T Michael, "An economic analysis of marital instability," *Journal of political Economy*, 1977, *85* (6), 1141–1187.

Brocas, Isabelle, Juan D Carrillo, Stephanie W Wang, and Colin F Camerer, "Imperfect choice or imperfect attention? Understanding strategic thinking in private information games," *The Review of Economic Studies*, 2014, p. rdu001.

Burdett, Ken and Melvyn G Coles, "Marriage and class," *The Quarterly Journal of Economics*, 1997, *112* (1), 141–168.

Camerer, Colin, Salvatore Nunnari, and Thomas R. Palfrey, "Quantal response and nonequilibrium beliefs explain overbidding in maximum-value auctions," *Games and Economic Behavior*, 2016, *98*, 243 – 263.

Carrillo, Juan D and Thomas R Palfrey, "The compromise game: two-sided adverse selection in the laboratory," *American Economic Journal: Microeconomics*, 2009, *1* (1), 151–181.

Chade, Hector, "Matching with noise and the acceptance curse," *Journal of Economic Theory*, 2006, *129* (1), 81–113.

Charness, Gary and Dan Levin, "The origin of the winner's curse: a laboratory study," *American Economic Journal: Microeconomics*, 2009, *1* (1), 207–36.

_ , _ , **and David Schmeidler**, "An experimental study of estimation and bidding in common-value auctions with public information," *Journal of Economic Theory*, 2019, *179*, 73–98.

Chen, Daniel L, Martin Schonger, and Chris Wickens, "oTree?An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 2016, *9*, 88–97.

Crawford, Vincent P and Nagore Iriberri, "Level-k auctions: can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions?," *Econometrica*, 2007, *75* (6), 1721–1770.

Esponda, Ignacio, "Behavioral equilibrium in economies with adverse selection," *American Economic Review*, 2008, *98* (4), 1269–91.

_ **and Emanuel Vespa**, "Hypothetical thinking and information extraction in the laboratory," *American Economic Journal: Microeconomics*, 2014, *6* (4), 180–202.

Eyster, Erik and Matthew Rabin, "Cursed equilibrium," *Econometrica*, 2005, *73* (5), 1623–1672.

Frederick, Shane, "Cognitive reflection and decision making," *Journal of Economic perspectives*, 2005, *19* (4), 25–42.

Goeree, Jacob K, Charles A Holt, and Thomas R Palfrey, *Quantal Response Equilibrium: A Stochastic Theory of Games*, Princeton University Press, 2016.

Holt, Charles A and Susan K Laury, "Risk aversion and incentive effects," *American economic review*, 2002, *92* (5), 1644–1655.

Ivanov, Asen, Dan Levin, and Muriel Niederle, "Can relaxation of beliefs rationalize the winner's curse?: an experimental study," *Econometrica*, 2010, *78* (4), 1435–1452.

Jehiel, Philippe, "Analogy-based expectation equilibrium," *Journal of Economic theory*, 2005, *123* (2), 81–104.

_ **and Frédéric Koessler**, "Revisiting games of incomplete information with analogy-based expectations," *Games and Economic Behavior*, 2008, *62* (2), 533–557.

Kagel, John H and Dan Levin, "The winner's curse and public information in common value auctions," *The American economic review*, 1986, pp. 894–920.

Kneeland, Terri, "Identifying Higher-Order Rationality," *Econometrica*, 2015, *83* (5), 2065–2079.

Koch, Christian and Stefan P Penczynski, "The winner's curse: Conditional reasoning and belief formation," *Journal of Economic Theory*, 2018, *174*, 57–102.

Levin, Dan, James Peck, and Asen Ivanov, "Separating Bayesian updating from non-probabilistic reasoning: An experimental investigation," *American Economic Journal: Microeconomics*, 2016, *8* (2), 39–60.

Martínez-Marquina, Alejandro, Muriel Niederle, and Emanuel Vespa, "Failures in Contingent Reasoning: The Role of Uncertainty," *American Economic Review*, 2019, *109* (10), 3437–74.

Miettinen, Topi, "The partially cursed and the analogy-based expectation equilibrium," *Economics Letters*, 2009, *105* (2), 162–164.

Nagel, Rosemarie, "Unraveling in guessing games: An experimental study," *The American Economic Review*, 1995, *85* (5), 1313–1326.

Ngangoué, M. Kathleen and Georg Weizsäcker, "Learning from Unrealized versus Realized Prices," *American Economic Journal: Microeconomics*, May 2021, *13* (2), 174–201.

Stahl, Dale O and Paul W Wilson, "On players' models of other players: Theory and experimental evidence," *Games and Economic Behavior*, 1995, *10* (1), 218–254.

Appendices

THE FOLLOWING PAGES CONTAIN APPENDICES A TO D. ALL APPENDICES ARE FOR ONLINE PUBLICATION ONLY.

A Useful Expressions

In this Appendix we provide expressions for some of the key variables of the model, namely $\pi(q', s)$, $\alpha(q, s)$ and $\beta(q', q, s)$.

First, the posterior probability of $q_j = q'$conditional on signal $s_i = s$ is given by:

$$\pi(q', s) \equiv Pr[q_j = q'|s_i = s] = \frac{\delta(q', s)}{\sum_{q''}[\delta(q'', s)]}$$

where the second equality follows by Bayes' rule.

Second, the acceptance probability, defined as the probability of player j proposing conditional on player i's type, is given by:

$$\alpha(q, s) \equiv Pr[a_j = P|q_i = q, s_i = s] = \sum_{q'}\sum_{s'}\pi(q', s)\delta(q, s')\sigma(q', s')$$

Finally, the posterior belief that player j's quality is q', conditional on the observed signal $s_i = s$, on own quality $q_i = q$ and on player j proposing, can be computed as:

$$\beta(q', q, s) = \frac{\sum_{s'}\delta(q, s')\sigma(q', s')}{\alpha(q, s)}\pi(q', s)$$

To derive this expression start with the definition of the posterior belief:

$$\beta(q', q, s) \equiv Pr[q_j = q'|s_i = s, a_j = P, q_i = q]$$

By Bayes' Theorem we can rewrite the posterior probability as:

$$\beta(q', q, s) = \frac{Pr[a_j = P|q_j = q', q_i = q, s_i = s]}{Pr[a_j = P|q_i = q, s_i = s]}Pr[q_j = q'|q_i = q, s_i = s]$$

Getting rid of redundant conditioning variables, we can rewrite the posterior belief as:

$$\beta(q', q, s) = \frac{Pr[a_j = P | q_j = q', q_i = q]}{Pr[a_j = P | q_i = q, s_i = s]} Pr[q_j = q' | s_i = s] \tag{A.1}$$

The probability of player j proposing conditional on the actual qualities of the players is given by:

$$Pr[a_j = P | q_i = q, q_j = q'] = \sum_{z'} Pr[s_j = s' | q_i = q] Pr[a_j = P | q_j = q', s_j = s']$$

Using this we can rewrite equation (A.1) as:

$$\beta(q', q, s) = \frac{\sum_{s'} Pr[s_j = s' | q_i = q] Pr[a_j = P | q_j = q', s_j = s']}{Pr[a_j = P | q_i = q, s_i = s]} Pr[q_j = q' | s_i = s] \tag{A.2}$$

Using our notation we rewrite equation (A.2) as:

$$\beta(q', q, s) = \frac{\sum_{s'} \delta(q, s') \sigma(q', s')}{\alpha(q, s)} \pi(q', s)$$

B Proof of Proposition 1

Proposition 1 has two parts. First, we prove that in game A, in any Bayes-Nash equilibrium: $\sigma(H, m) = 0$ and $\sigma(M, m) = 0$.

By contradiction, assume there is a BNE where $\sigma(H, m) > 0$. In this case, proposing conditional on an h signal is a best response for all types: $\sigma(q, h) = 1 \ \forall q$. Similarly, it is a best response for L-quality players to propose upon receiving an m signal: $\sigma(L, m) = 1$.

For $\sigma(H, m) > 0$ to be part of a BNE, it must be $\Delta(H, m) \geq 0$. This condition implies:

$$\alpha(H,m)[v(H,m)-\rho(H)] \geq 0;$$
$$v(H,m) \geq \rho(H);$$
$$\sum_{q'} \beta(q',H,m)\mu(q') \geq \rho(H)$$

Substituting the actual matching and reservation values we get:

$$160\beta(H,H,m) + 80\beta(M,H,m) + 40\beta(L,H,m) \geq 100$$

Because posterior beliefs sum to one, this can be rewritten as:

$$160[1-\beta(M,H,m)-\beta(L,H,m)] + 80\beta(M,H,m) + 40\beta(L,H,m) \geq 100$$

which yields

$$4\beta(M,H,m) + 6\beta(L,H,m) \leq 3 \tag{B.1}$$

We then compute the two posterior beliefs in condition (B.1) using the expressions provided in Appendix A, the given likelihoods and using the fact that $\sigma(q,h)=1\ \forall q$ and $\sigma(L,m)=1$. This yields:

$$\beta(M,H,m) = \frac{1}{4}\frac{1+\sigma(M,m)}{\alpha(H,m)}$$

and

$$\beta(L,H,m) = \frac{1}{4}\frac{1}{\alpha(H,m)}$$

where

$$\alpha(H,m) = \frac{5}{8} + \frac{1}{8}\sigma(H,m) + \frac{1}{4}\sigma(M,m)$$

Using the last three expressions and condition (B.1) gives:

$$\frac{5}{8} + \frac{1}{4}\sigma(M,m) \leq \frac{3}{8}\sigma(H,m)$$

This inequality cannot be satisfied if $\sigma(M,m)$ and $\sigma(H,m)$ are between 0 and 1. Thus, we have reached a contradiction and we have proved that in any BNE of game A it must be: $\sigma(H,m) = 0$.

The next step involves showing that in any BNE: $\sigma(M,m) = 0$. By contradiction assume that $\sigma(M,m) > 0$. Then it must be $\Delta(M,m) \geq 0$. This condition implies:

$$\begin{aligned} \alpha(M,m)[v(M,m) - \rho(M)] \geq 0; \\ v(M,m) \geq \rho(M); \\ \sum_{q'} \beta(q',M,m)\mu(q') \geq \rho(M) \end{aligned}$$

Substituting the actual matching and reservation values we get:

$$160\beta(H,M,m) + 80\beta(M,M,m) + 40\beta(L,M,m) \geq 75 \tag{B.2}$$

To compute the posterior beliefs in condition (B.2) we use the expressions provided in Appendix A, the given likelihoods and the fact that $\sigma(q,h) = 1\ \forall q$, $\sigma(L,m) = 1$ and $\sigma(H,m) = 0$. This yields:

$$\beta(H,M,m) = 0$$

$$\beta(M,M,m) = \frac{\sigma(M,m)}{\sigma(M,m) + 1/2}$$

$$\beta(L,M,m) = \frac{1/2}{\sigma(M,m) + 1/2}$$

Substituting these expressions in condition (B.2) and solving for $\sigma(M,m)$ we get:

$$\sigma(M,m) \geq \frac{7}{2}$$

which is not possible. Having reached a contradiction, it must be $\sigma(M, m) = 0$.

The second part of Proposition 1 states that in game A the pure-strategy BNE where most types propose is such that: $\mathcal{P}_H = \{h\}, \mathcal{P}_M = \{h\}, \mathcal{P}_L = \{h, m, l\}$. We have already shown that in any pure-strategy BNE $m \notin \mathcal{P}_H$ and $m \notin \mathcal{P}_M$. Moreover, $l \notin \mathcal{P}_H$ and $l \notin \mathcal{P}_M$ because proposing after observing l is not a best response for H- and M-types. We can show that the other strategies are part of a BNE. For type (H, h), P is the best response because: $v(H, h) = \mu(H) = 160 > \rho(H) = 100$. For type (M, h), P or N are both best responses (because H types do not accept m signals). For type (L, s_i), P is the best response for any signal s_i, because $\mu(q) > \rho(L) \forall q$.

C Proof of Proposition 2

In game B, the pure-strategy BNE where most types propose is such that: $\mathcal{P}_H = \{h, m\}, \mathcal{P}_M = \{h, m\}, \mathcal{P}_L = \{h, m, l\}$. To show this first note that in any pure-strategy BNE $l \notin \mathcal{P}_H$ and $l \notin \mathcal{P}_M$ because proposing after observing l is not a best response for H- and M-types. Then we show that the other strategies in this profile are best responses. For type (H, h), P is the best response because: $v(H, h) = \mu(H) = 160 > \rho(H) = 100$. For type (H, m), P is now the best response because: $v(H, m) = 0.25 \times \mu(H) + 0.5 \times \mu(M) + 0.25 \times \mu(L) = 90 > \rho(H) = 80$. For type (M, h), P or N are both best responses (because H types do not accept m signals). For type (M, m), P is the best response because: $v(M, m) = 0.25 \times \mu(H) + 0.5 \times \mu(M) + 0.25 \times \mu(L) = 90 > \rho(M) = 75$. For type (L, s), P is the best response for any signal s.

D Proof of Proposition 3

In game A, for $\chi \geq \frac{5}{14}$, there is a cursed equilibrium such that: $\mathcal{P}_H = \{h\}, \mathcal{P}_M = \{h, m\}, \mathcal{P}_L = \{h, m, l\}$. To prove this, we check whether this strategy profile satisfies the conditions for a cursed equilibrium. First, proposing is a best response for type (L, s) for all s and for any $\chi \in [0,1]$.

Next, consider an H-quality type. As before, P is the best response for type (H, h) and N is the best response for type (H, l), for any $\chi \in [0,1]$. To show that N is the best response of type (H, m), we first compute the cursed posterior beliefs. These are given by:

$$\beta^c(H, H, m) = \frac{1}{7} + \frac{3}{28}\chi$$
$$\beta^c(M, H, m) = \frac{4}{7} - \frac{1}{14}\chi$$
$$\beta^c(L, H, m) = \frac{2}{7} - \frac{1}{28}\chi$$

Then it is possible to show that $\Delta^c(H, m) \geq 0$ only if $\chi \geq 2$. Thus, N is the best response of type (H, m) for any $\chi \in [0,1]$.

Finally, consider an M-quality type. As before, P is a best response for type (M, h) and N is the best response for type (M, l), for any $\chi \in [0,1]$. We then check when P is the best response of type (M, m). The cursed posterior beliefs are given by:

$$\beta^c(H, M, m) = \frac{1}{4}\chi$$
$$\beta^c(M, M, m) = \frac{2}{3} - \frac{1}{6}\chi$$
$$\beta^c(L, M, m) = \frac{1}{3} - \frac{1}{12}\chi$$

Then it is possible to show that $\Delta^c(M, m) \geq 0$ only if $\chi \geq \frac{5}{14}$. Thus, P is the best response of type (H, m) for $\chi \geq \frac{5}{14}$.

E Derivation of Likelihoods in the *COND* Treatment

In this Appendix we provide expressions for the modified signal likelihoods used in the design of the $COND$ treatment. For each value of $p \in \{0, 0.25, 0.5, 0.75, 1\}$, we calculate the $COND$-task likelihoods in the following way. First, we compute the posterior probability in the BEL-task, denoted $\phi(p)$:

$$\phi(p) \equiv Pr[q_2 = H|s_1 = m, q_1 = M, a_2 = P; BEL\text{-task}]$$

Then we require the posterior belief in the $COND$-task to equal the posterior belief in the BEL-task:

$$Pr[q_2 = H|s_1 = m; COND\text{-task}] = \phi(p) \tag{E.1}$$

To achieve this we only adjust the likelihoods of observing h and m signals when player 2's quality is H, while leaving all other likelihoods unchanged.

Applying Bayes' rule on equation (E.1) we obtain:

$$\frac{Pr[s_1 = m|q_2 = H; COND\text{-task}] \times 1/3}{Pr[s_1 = m|q_2 = H; COND\text{-task}] \times 1/3 + 1 \times 1/3 + 0.5 \times 11/3} = \phi(p)$$

It follows that:

$$Pr[s_1 = m|q_2 = H; COND\text{-task}] = \frac{\frac{3}{2}\phi(p)}{1 - \phi(p)}$$

Since $Pr[s_1 = l|q_2 = H; COND\text{-task}] = 0$, then we also know:

$$Pr[s_1 = h|q_2 = H; COND\text{-task}] = 1 - Pr[s_1 = m|q_2 = H; COND\text{-task}] = \frac{1 - \frac{5}{2}\phi(p)}{1 - \phi(p)}$$

F Experiment Instructions

F.1 *BASE*

This appendix reproduces the instructions for the *BASE* treatment of the experiment.

INSTRUCTIONS

You are about to participate in an experiment in the economics of decision-making. If you follow these instructions carefully, you can earn an amount of money which will be paid to you in cash at the end of the experiment.

Your computer screen will display useful information. Remember that the information on your computer screen is private. To insure best results for yourself and accurate data for the experimenters, please do not communicate with the other participants at any point during the experiment. If you have any questions, or need assistance of any kind, raise your hand and one of the experimenters will come.

PARTS and PAYMENTS

This experiment will consist of four parts. At the end of the experiment, you will be paid 5€, plus earnings based on the points you have earned during the experiment. Your points will be converted to Euros at an exchange rate of 1/6 Euro per point. To sum up, your final payment in Euros is given by the following formula:

$$Your\ points/6 + 5 \quad €$$

1

PART 1

THE BASIC IDEA

In this part, you have to decide whether to form a partnership with other players. The benefit of forming a partnership with another player will depend on the other player's type, which is randomly assigned by the computer. During the experiment, you will not know the type of your potential partners. However, you will receive clues about their types before you decide to form a partnership. If you form a partnership, you receive an amount of points that depends on your partner's type. If you do not form a partnership, you receive an amount of points that depends only on your own type.

ROUNDS

The experiment will be divided into **60 rounds**. In each round, you are randomly paired with another player. Decisions and points you make in one round do not affect other rounds.

PLAYERS AND TYPES

In each round, each participant is randomly assigned a type. A participant's type can be X, Y, or Z. In each round your type is **equally likely** to be X, Y or Z. In each round, you will have an opportunity to form a partnership with another player. The computer will randomly pair you and another player. The way in which pairs are formed is random and does not depend on the players' types. This means that you are **equally likely** to be paired with an X, Y, or Z-type player, independently of your own type.

CLUES

You will not know the type of the participant you are paired with in a round. However, you will receive a clue about your potential partner's type. Clues are determined in the following way. The **computer will digitally draw a random ball from a box containing 24 balls of different colors**. Each ball can be either blue, yellow or red. The number of blue, yellow and red balls in **the box depends on your partner's type**. The boxes used in the experiment are illustrated in the figure below.

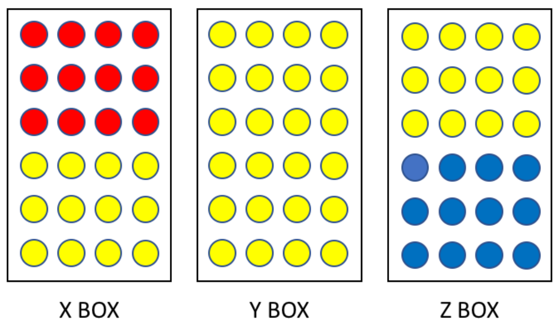

If your partner's type is X, the box contains no blue balls, 12 yellow balls and 12 red balls. If your partner's type is Y, the box contains 24 yellow balls but no blue or red balls. If your partner's type is Z, the box contains 12 blue balls, 12 yellow balls and no red balls. To give you a clue about the type of your potential partner, the computer will first determine which box to use given your partner's type. Then it will digitally draw a random ball: each single ball in the box has the same probability of being selected, equal to 1/24. The clue you receive is the color of this randomly drawn ball.

FORMING A PARTNERSHIP

After you have received a clue about your potential partner's type, you can decide whether you want to form a partnership or not. **Only if you and the other player agree to form a partnership, a partnership is formed.** For example, if you want to form a partnership but the other player does not, the partnership is not formed.

POINTS

Every time you form a partnership, you earn an amount of points that depends only on the type of the other player. Whenever you do not form a partnership, you earn an amount of point that depends only on your own type. Consider the example illustrated in the following table. In this example, if you form a partnership with an X type you earn 160 points, if you form a partnership with a Y type you earn 80 points and if you form a partnership with a Z type you earn 40 points. If you do not form a partnership and your type is X, you earn 100 points. If you do not form a partnership and your type is Y, you earn 75 points. If you do not form a partnership and your type is Z, you earn 25 points.

Points if you **form** a partnership:			
Partner's type	X	Y	Z
Your points	160	80	40
Points if you do **not form** a partnership:			
Your type	X	Y	Z
Your points	100	75	25

The exact amounts of points you can earn will depend on the game you are playing. There are two versions of this game, called A and B. There will be 30 rounds for each game, but the exact sequence of games will be random. The following table reports the actual amounts of points you can earn in each game. Note that Game A is the example discussed above. The only difference between game A and B is the payoff a type X player receives if he does not form a partnership.

GAME A

Points if you **form** a partnership:			
Partner's type	X	Y	Z
Your points	160	80	40
Points if you do **not form** a partnership:			
Your type	X	Y	Z
Your points	**100**	75	25

GAME B

Points if you **form** a partnership:			
Partner's type	X	Y	Z
Your points	160	80	40
Points if you do **not form** a partnership:			
Your type	X	Y	Z
Your points	**80**	75	25

PAYMENT

At the end of the experiment, your payment will depend on the points you have earned in this part. In particular, the computer will randomly select one round out of the 60 rounds in part 1. At the end of the experiment, the points you have earned in the selected round will be converted to Euros at an exchange rate of 1/6 Euros per point.

PART 1 APP PAGES

The experiment app will show you several pages, described below. Between a page and the next you may have to wait for other participants to make their choices. At the beginning of Part 1 you will see a page like this.

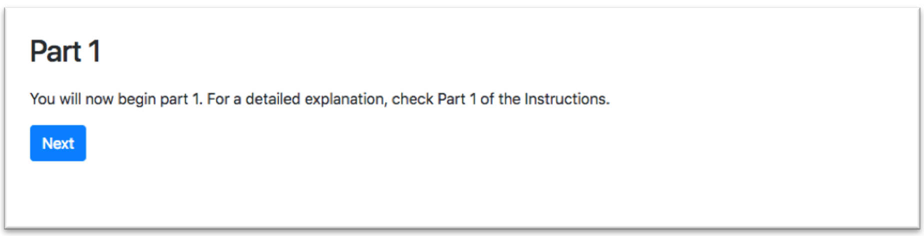

At the beginning of a new round, you will see a page informing you of whether you are playing game A or game B.

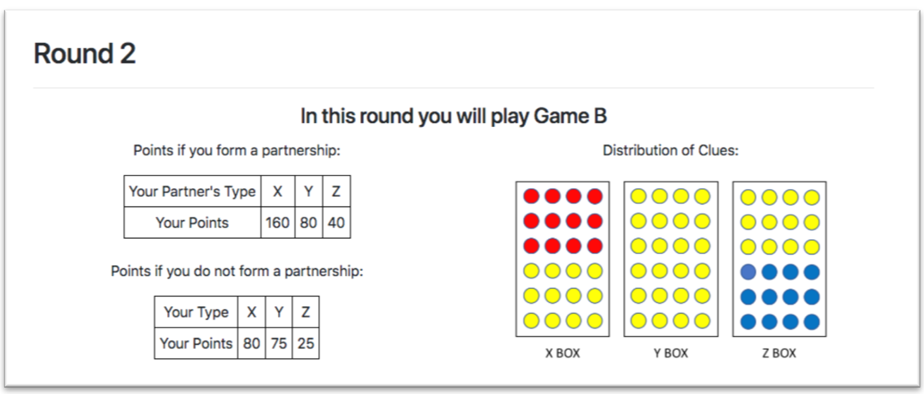

Your Partner's Type	X	Y	Z
Your Points	160	80	40

Your Type	X	Y	Z
Your Points	80	75	25

In this page, you receive information about your own type. You are also given a clue about your potential partner's type and you decide whether you want to form a partnership or not.

Make a choice

You are a Y type (if you do not form a partnership, you earn 75 points).

You have received a yellow clue.

Do you want to form a partnership?

○ Yes ○ No

Next

At the bottom of this page you can see a table with a history of all the previous rounds.

History Table

Round	Game	Your Type	Your Clue	Your Choice	Partner's Choice	Partner's Type	Your points
1	A	Z	Red	Yes	No	X	25
2	B	Y	Yellow	-	-	-	-

If you do not click the "next" button on this page, after one minute the app will move you to the next page. In the next page you are told your points in this round and the actual type of your partner.

Round 2

Your formed a partnership. Your partner's type is: Z Your points are: 40.

Next

When you have played the last round of Part 1, you will see a page informing you that this part is over. In this page you will find out the Part 1 paying round. At the bottom of this page you can see the history table summarizing all the rounds.

Part 1 Is Over

Part 1 of the experiment is over. The paying round was round 1. Your payoff in that round is: 25 points.

Next

History Table

Round	Game	Your Type	Your Clue	Your Choice	Partner's Choice	Partner's Type	Your points
1	A	Z	Red	Yes	No	X	25
2	B	Y	Yellow	Yes	Yes	Z	40

9

PART 2

In this part, you will face 10 decisions listed on your screen. In each decision you have to choose between "Option A" and "Option B". If you choose Option A, you will earn either 6 or 5 points. If you choose Option B, you will earn either 10 points or 1 point. After you choose one option, whether you earn the higher payoff or the lower payoff is randomly determined by the computer. Before making a choice, you will know the exact probability of earning the higher payoff rather than the lower payoff in each option. For example, in one decision Option A will give you 6 points with a probability of 30% and 5 points otherwise, while Option B will give you 10 points with a probability of 30% and 1 point otherwise.

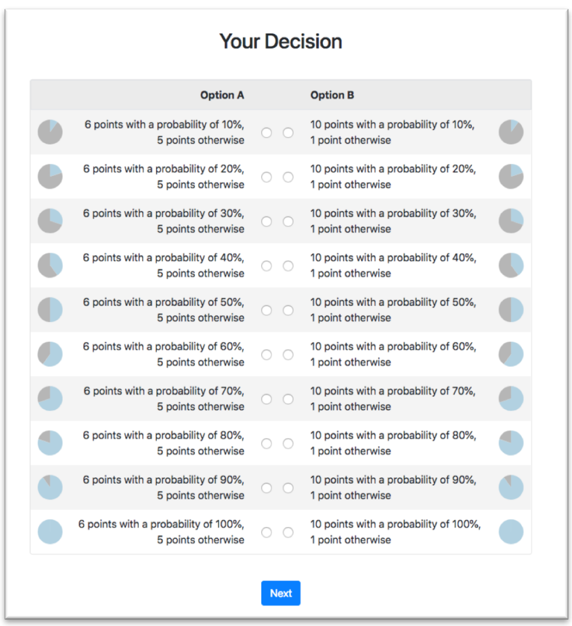

Your Decision

Option A	Option B
6 points with a probability of 10%, 5 points otherwise	10 points with a probability of 10%, 1 point otherwise
6 points with a probability of 20%, 5 points otherwise	10 points with a probability of 20%, 1 point otherwise
6 points with a probability of 30%, 5 points otherwise	10 points with a probability of 30%, 1 point otherwise
6 points with a probability of 40%, 5 points otherwise	10 points with a probability of 40%, 1 point otherwise
6 points with a probability of 50%, 5 points otherwise	10 points with a probability of 50%, 1 point otherwise
6 points with a probability of 60%, 5 points otherwise	10 points with a probability of 60%, 1 point otherwise
6 points with a probability of 70%, 5 points otherwise	10 points with a probability of 70%, 1 point otherwise
6 points with a probability of 80%, 5 points otherwise	10 points with a probability of 80%, 1 point otherwise
6 points with a probability of 90%, 5 points otherwise	10 points with a probability of 90%, 1 point otherwise
6 points with a probability of 100%, 5 points otherwise	10 points with a probability of 100%, 1 point otherwise

Next

As in this example, in any one of the 10 decisions, the probability you will earn the higher payoff (6 if Option A is chosen or 10 if Option B is chosen) is the same between option A and option B. In the first decision, at the top of the list, the probability you will earn the higher payoff is 10%. As you move down the table, the chances of the higher payoff for each option increase. In fact, for decision 10 in the bottom row, each option pays the highest payoff for sure. So, your choice in decision 10 is simply between 6 points (Option A) or 10 points (Option B).

For each of the ten decisions, you will be asked to choose Option A or Option B by clicking on the appropriate button. The computer will ensure that you switch at most once from Option A to Option B. If you choose Option A in one decision, the computer will automatically select Option A for all the previous decisions. If you choose Option B in one decision, the computer will automatically select Option B for all the following decisions. Once you have made a choice in all decisions, you can click on the Next button to submit your choices.

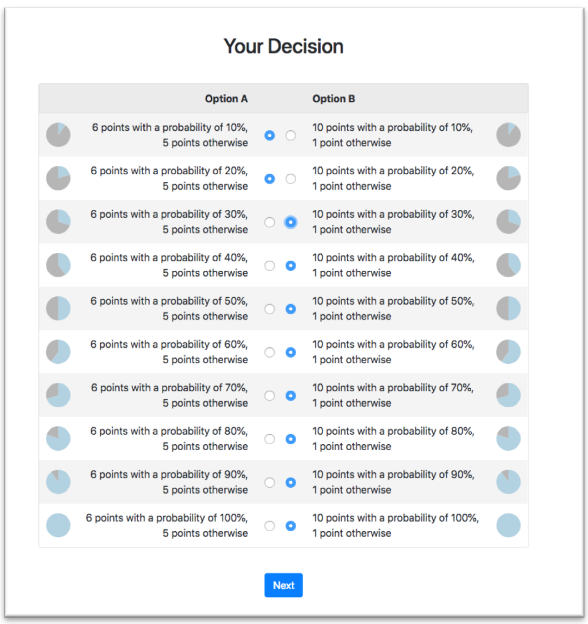

After you have submitted your choices, one of the 10 decisions will be randomly chosen for your payment. For the option you chose, A or B, in this decision, the computer will randomly determine whether you earn the higher or lower payoff. To determine the outcome of your choice, the computer will digitally draw a random number between 0 and 100. If the random number is below the probability of earning the higher payoff, then you receive the higher payoff. If the random number is above the probability, then you receive the lower payoff. For example, assume you chose Option A in the first decision and this decision is selected for payment. If the computer randomly draws a 60, you will earn 5 points.

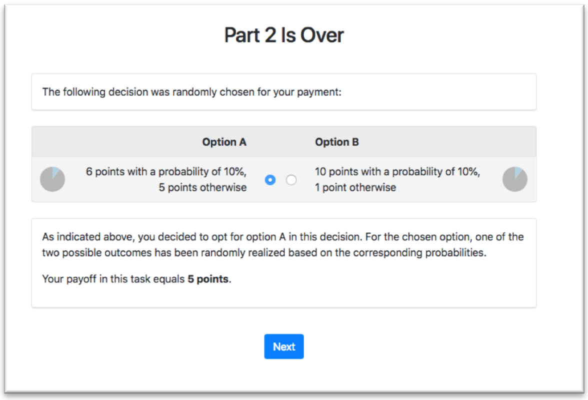

At the end of the experiment, the points you have earned in the selected decision will be converted to Euros at an exchange rate of 1/6 Euros per point.

PART 3

In this part, you are asked to answer three questions. For each correct answer, you will receive two points. After you have submitted your answers, you will see the correct answers and the amount of points you have earned. At the end of the experiment, the points you have earned in this part will be converted to Euros at an exchange rate of 1/6 Euros per point.

PART 4

In this part, you are asked to provide some information about yourself (your sex and your undergraduate major). As stated before, your responses are completely confidential and anonymous.

F.2 *BEL*

This appendix reproduces the instructions for the *BEL* treatment of the experiment.

INSTRUCTIONS

You are about to participate in an experiment in the economics of decision-making. If you follow these instructions carefully, you can earn an amount of money which will be paid to you in cash at the end of the experiment.

Your computer screen will display useful information. Remember that the information on your computer screen is private. To insure best results for yourself and accurate data for the experimenters, please do not communicate with the other participants at any point during the experiment. If you have any questions, or need assistance of any kind, raise your hand and one of the experimenters will come.

PARTS and PAYMENTS

This experiment will consist of four parts. At the end of the experiment, you will be paid 5€, plus earnings based on the points you have earned during the experiment. Your points will be converted to Euros at an exchange rate of 1/6 Euro per point. To sum up, your final payment in Euros is given by the following formula:

$$Your\ points/6 + 5 \quad €$$

PART 1

THE BASIC IDEA

In this part, you have to decide whether to form a partnership with other players. The benefit of forming a partnership with another player will depend on the other player's type, which is randomly assigned by the computer. During the experiment, you will not know the type of your potential partners. However, you will receive clues about their types before you decide to form a partnership. If you form a partnership, you receive an amount of points that depends on your partner's type. If you do not form a partnership, you receive an amount of points that depends only on your own type.

ROUNDS

The experiment will be divided into **60 rounds**. In each round, you are randomly paired with another player. Decisions and points you make in one round do not affect other rounds.

Each round you will be **randomly assigned either an active role or a passive role**. When you are assigned a passive role, you cannot make any decision. An active player is always paired with a passive player.

PLAYERS AND TYPES

In each round, every participant is assigned a type. A participant's type can be X, Y or Z. The type of active players is always Y while the type of passive players will be randomly determined. This means that, in each of your active rounds, you are **equally likely** to be paired with an X, Y or Z type player.

CLUES

When you are active, you will not know the type of the participant you are paired with. However, you will receive a clue about your potential partner's type. Clues are determined in the following way. The **computer will digitally draw a random ball from a box containing 24 balls of different colors**. Each ball can be either blue, yellow or red. The number of blue, yellow and red balls in **the box depends on your partner's type**. The boxes used in the experiment are illustrated in the figure below.

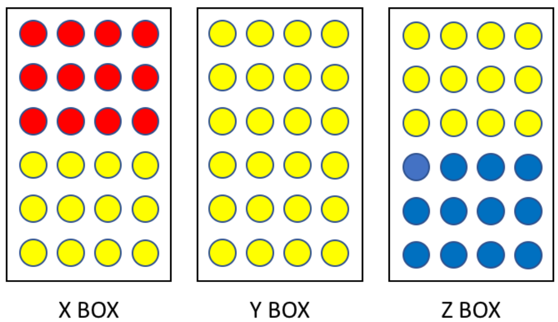

If your partner's type is X, the box contains no blue balls, 12 yellow balls and 12 red balls. If your partner's type is Y, the box contains 24 yellow balls but no blue or red balls. If your partner's type is Z, the box contains 12 blue balls, 12 yellow balls and no red balls. To give you a clue about the type of your potential partner, the computer will first determine which box to use given your partner's type. Then it will digitally draw a random ball: each single ball in the box has the same probability of being selected, equal to 1/24. The clue you receive is the color of this randomly drawn ball.

FORMING A PARTNERSHIP

When you are an active player, after you have received a clue about your potential partner's type, you can decide whether you want to form a partnership or not.

The computer will decide on behalf of the passive player whether he agrees to form a partnership or not. Specifically, **if the passive player's type is Z or Y, the computer will always agree to form a partnership. If the type of the passive player is X, the computer will agree to form a partnership with some probability**, for example with 50% probability. There are five versions of this game, called **A, B, C, D** and **E**. There will be 12 rounds for each game but the exact sequence of games will be random. The following table reports the actual probability that the computer will propose to form a partnership on behalf of an X type player in different games.

	Game				
	A	B	C	D	E
Probability computer agrees to form a partnership when passive type is X.	100%	75%	50%	25%	0%
Probability computer agrees to form a partnership when passive type is Y.	100%	100%	100%	100%	100%
Probability computer agrees to form a partnership when type is Z.	100%	100%	100%	100%	100%

To determine the choice for the passive player, the computer will digitally draw a random number between 0 and 100. If the random number is below the probability given in the table above, then the computer will agree to form a partnership. If the random number is above the probability, then the computer will not agree to form a partnership. For example, assume you are playing game D, the passive player's type is X and the computer randomly draws a 40. Then, the computer will not agree to form a partnership (because 40>25).

Only if the active player proposes to form a partnership and the computer acting on the passive player's behalf agrees, then a partnership is formed. For example, if you are an active player and propose to form a partnership but the computer of the passive player does not agree, the partnership is not formed.

POINTS

Every time you form a partnership, you earn an amount of points that depends only on the type of the other player. Whenever you do not form a partnership, you earn an amount of point that depends only on your own type. The actual points you can earn are reported in the following table. If you form a partnership with an X type you earn 160 points, if you form a partnership with a Y type you earn 80 points and if you form a partnership with a Z type you earn 40 points. If you do not form a partnership and your type is X, you earn 100 points. If you do not form a partnership and your type is Y, you earn 75 points. If you do not form a partnership and your type is Z, you earn 25 points.

Points if you **form** a partnership:			
Partner's type	X	Y	Z
Your points	160	80	40
Points if you do **not form** a partnership:			
Your type	X	Y	Z
Your points	100	75	25

PAYMENT

At the end of the experiment, your payment will depend on the points you have earned in this part. In particular, the computer will randomly select one round out of the 60 rounds in part 1. At the end of the experiment, the points you have earned in the selected round will be converted to Euros at an exchange rate of 1/6 Euros per point.

PART 1 APP PAGES

The experiment app will show you several pages, described below. Between a page and the next you may have to wait for other participants to make their choices. At the beginning of Part 1 you will see a page like this.

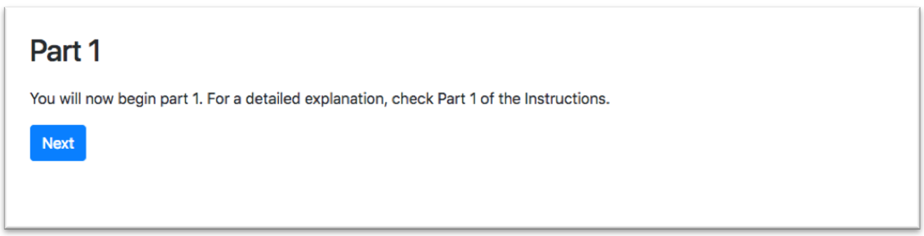

Part 1

You will now begin part 1. For a detailed explanation, check Part 1 of the Instructions.

Next

At the beginning of a new round, you will see a page informing you of whether you are playing game A, B, C, D or E.

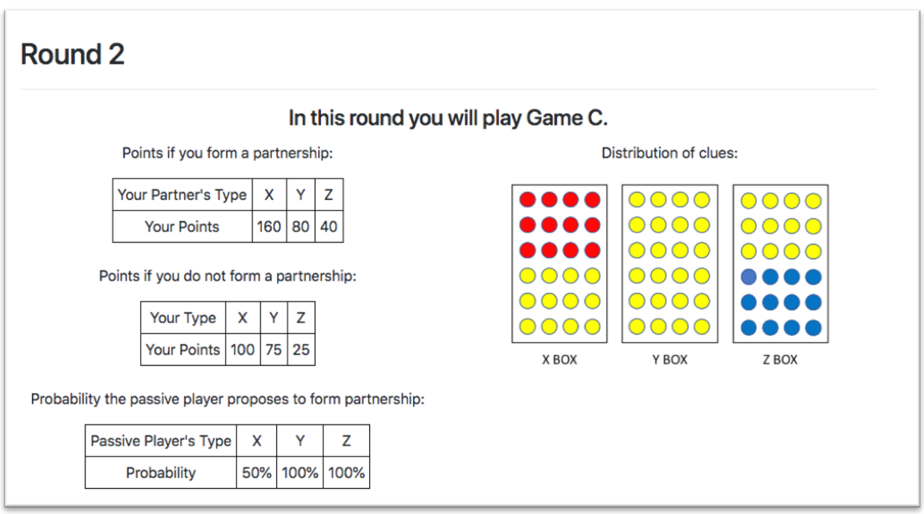

Round 2

In this round you will play Game C.

Points if you form a partnership:

Your Partner's Type	X	Y	Z
Your Points	160	80	40

Points if you do not form a partnership:

Your Type	X	Y	Z
Your Points	100	75	25

Probability the passive player proposes to form partnership:

Passive Player's Type	X	Y	Z
Probability	50%	100%	100%

Distribution of clues:

In this page, you are told your type and whether you are an active player or passive player.

In this round you are passive.

You are an X type (if you do not form a partnership, you earn 100 points).

Next

If you are an active player in this round, you are also given a clue about your potential partner's type and you decide whether you want to form a partnership or not.

In this round you are active: make a choice.

You are a Y type (if you do not form a partnership, you earn 75 points).

You have received a red clue.

Do you want to form a partnership?

○ Yes ○ No

Next

At the bottom of this page you can see a table with a history of all the previous rounds.

History Table

Round	Game	Your Role	Your Type	Your Clue	Your Choice	Partner's Choice	Partner's Type	Your points
1	A	active	Y	Blue	No	Yes	Z	75
2	C	active	Y	Red	-	-	-	-

If you do not click the "next" button on this page, after one minute the app will move you to the next page. In the next page you are told your points in this round and the actual type of your partner.

Round 2

You did not form a partnership. Your partner's type was: X Your points are: 75.

Next

When you have played the last round of Part 1, you will see a page informing you that this part is over. In this page you will find out the Part 1 paying round. At the bottom of this page you can see the history table summarizing all the rounds.

Part 1 Is Over

Part 1 of the experiment is over. The paying round was round 2. Your payoff in that round is: 75 points.

Next

History Table

Round	Game	Your Role	Your Type	Your Clue	Your Choice	Partner's Choice	Partner's Type	Your points
1	A	active	Y	Blue	No	Yes	Z	75
2	C	active	Y	Red	Yes	No	X	75
3	E	passive	Z	-	Yes	No	Y	25

PART 2

In this part, you will face 10 decisions listed on your screen. In each decision you have to choose between "Option A" and "Option B". If you choose Option A, you will earn either 6 or 5 points. If you choose Option B, you will earn either 10 points or 1 point. After you choose one option, whether you earn the higher payoff or the lower payoff is randomly determined by the computer. Before making a choice, you will know the exact probability of earning the higher payoff rather than the lower payoff in each option. For example, in one decision Option A will give you 6 points with a probability of 30% and 5 points otherwise, while Option B will give you 10 points with a probability of 30% and 1 point otherwise.

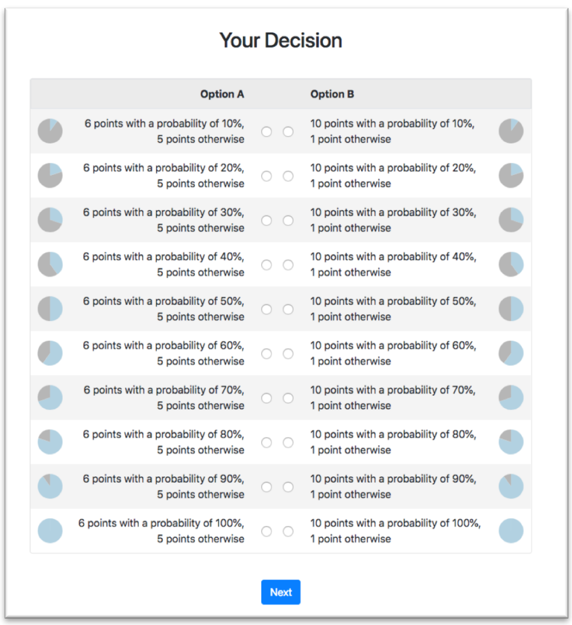

Your Decision

Option A	Option B
6 points with a probability of 10%, 5 points otherwise	10 points with a probability of 10%, 1 point otherwise
6 points with a probability of 20%, 5 points otherwise	10 points with a probability of 20%, 1 point otherwise
6 points with a probability of 30%, 5 points otherwise	10 points with a probability of 30%, 1 point otherwise
6 points with a probability of 40%, 5 points otherwise	10 points with a probability of 40%, 1 point otherwise
6 points with a probability of 50%, 5 points otherwise	10 points with a probability of 50%, 1 point otherwise
6 points with a probability of 60%, 5 points otherwise	10 points with a probability of 60%, 1 point otherwise
6 points with a probability of 70%, 5 points otherwise	10 points with a probability of 70%, 1 point otherwise
6 points with a probability of 80%, 5 points otherwise	10 points with a probability of 80%, 1 point otherwise
6 points with a probability of 90%, 5 points otherwise	10 points with a probability of 90%, 1 point otherwise
6 points with a probability of 100%, 5 points otherwise	10 points with a probability of 100%, 1 point otherwise

Next

As in this example, in any one of the 10 decisions, the probability you will earn the higher payoff (6 if Option A is chosen or 10 if Option B is chosen) is the same between option A and option B. In the first decision, at the top of the list, the probability you will earn the higher payoff is 10%. As you move down the table, the chances of the higher payoff for each option increase. In fact, for decision 10 in the bottom row, each option pays the highest payoff for sure. So, your choice in decision 10 is simply between 6 points (Option A) or 10 points (Option B).

For each of the ten decisions, you will be asked to choose Option A or Option B by clicking on the appropriate button. The computer will ensure that you switch at most once from Option A to Option B. If you choose Option A in one decision, the computer will automatically select Option A for all the previous decisions. If you choose Option B in one decision, the computer will automatically select Option B for all the following decisions. Once you have made a choice in all decisions, you can click on the Next button to submit your choices.

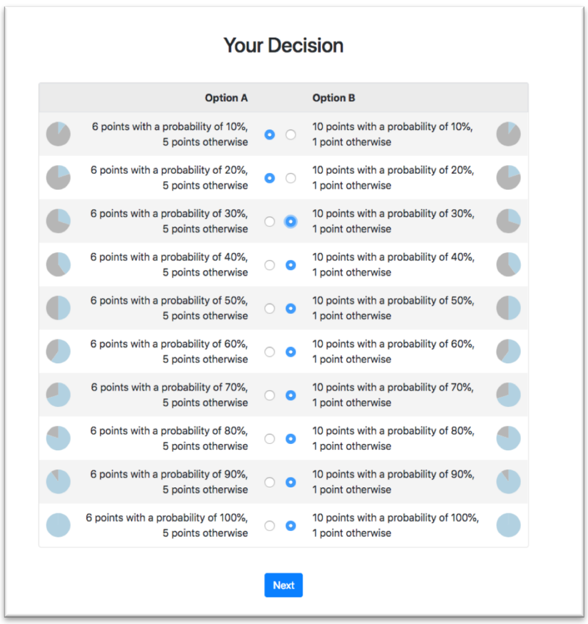

After you have submitted your choices, one of the 10 decisions will be randomly chosen for your payment. For the option you chose, A or B, in this decision, the computer will randomly determine whether you earn the higher or lower payoff. To determine the outcome of your choice, the computer will digitally draw a random number between 0 and 100. If the random number is below the probability of earning the higher payoff, then you receive the higher payoff. If the random number is above the probability, then you receive the lower payoff. For example, assume you chose Option A in the first decision and this decision is selected for payment. If the computer randomly draws a 60, you will earn 5 points.

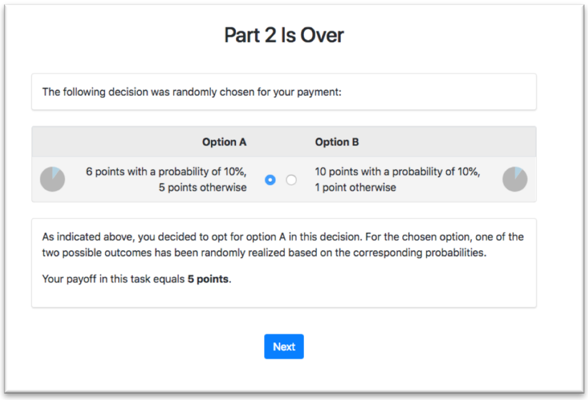

At the end of the experiment, the points you have earned in the selected decision will be converted to Euros at an exchange rate of 1/6 Euros per point.

PART 3

In this part, you are asked to answer three questions. For each correct answer, you will receive two points. After you have submitted your answers, you will see the correct answers and the amount of points you have earned. At the end of the experiment, the points you have earned in this part will be converted to Euros at an exchange rate of 1/6 Euros per point.

PART 4

In this part, you are asked to provide some information about yourself (your sex and your undergraduate major). As stated before, your responses are completely confidential and anonymous.

F.3 *COND*

This appendix reproduces the instructions for the *COND* treatment of the experiment.

INSTRUCTIONS

You are about to participate in an experiment in the economics of decision-making. If you follow these instructions carefully, you can earn an amount of money which will be paid to you in cash at the end of the experiment.

Your computer screen will display useful information. Remember that the information on your computer screen is private. To insure best results for yourself and accurate data for the experimenters, please do not communicate with the other participants at any point during the experiment. If you have any questions, or need assistance of any kind, raise your hand and one of the experimenters will come.

PARTS and PAYMENTS

This experiment will consist of four parts. At the end of the experiment, you will be paid 5€, plus earnings based on the points you have earned during the experiment. Your points will be converted to Euros at an exchange rate of 1/6 Euro per point. To sum up, your final payment in Euros is given by the following formula:

$$Your\ points/6 + 5 \quad €$$

PART 1

THE BASIC IDEA

In this part, you have to decide whether to form a partnership with other players. The benefit of forming a partnership with another player will depend on the other player's type, which is randomly assigned by the computer. During the experiment, you will not know the type of your potential partners. However, you will receive clues about their types before you decide to form a partnership. If you form a partnership, you receive an amount of points that depends on your partner's type. If you do not form a partnership, you receive an amount of points that depends only on your own type.

ROUNDS

The experiment will be divided into **60 rounds**. In each round, you are randomly paired with another player. Decisions and points you make in one round do not affect other rounds.

Each round you will be assigned **either an active role or a passive role**. When you are assigned a passive role, you cannot make any decision. An active player is always paired with a passive player.

PLAYERS AND TYPES

In each round, every participant is assigned a type. A participant's type can be X, Y or Z. The type of active players is always Y while the type of passive players will be randomly determined. This means that, in each of your active rounds, you are **equally likely** to be paired with an X, Y or Z type player.

CLUES

When you are active, you will not know the type of the participant you are paired with. However, you will receive a clue about your potential partner's type. Clues are determined in the following way. The **computer will digitally draw a random ball from a box containing 24 balls of different colors**. Each ball can be either blue, yellow or red. The number of blue, yellow and red balls in **the box depends on your partner's type**. For example, consider the boxes illustrated in the figure below.

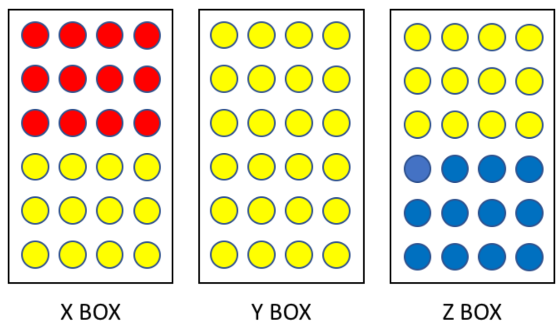

In this example, if your partner's type is X, the box contains no blue balls, 12 yellow balls and 12 red balls. If your partner's type is Y, the box contains 24 yellow balls but no blue or red balls. If your partner's type is Z, the box contains 12 blue balls, 12 yellow balls and no red balls. To give you a clue about the type of your potential partner, the computer will first determine which box to use given your partner's type. Then it will digitally draw a random ball: each single ball in the box has the same probability of being selected, equal to 1/24. The clue you receive is the color of this randomly drawn ball.

The exact composition of the X box will change in each round. There are five versions of this game, called **A, B, C, D** and **E**. There will be 12 rounds for each game but the exact sequence of games will be random. Each game has a different number of yellow and red balls in the X box, as summarized in the following table and illustrated in the figure below.

	Game				
	A	B	C	D	E
Red balls in X box	12	15	18	21	24
Yellow balls in X box	12	9	6	3	0
Blue balls in X box	0	0	0	0	0

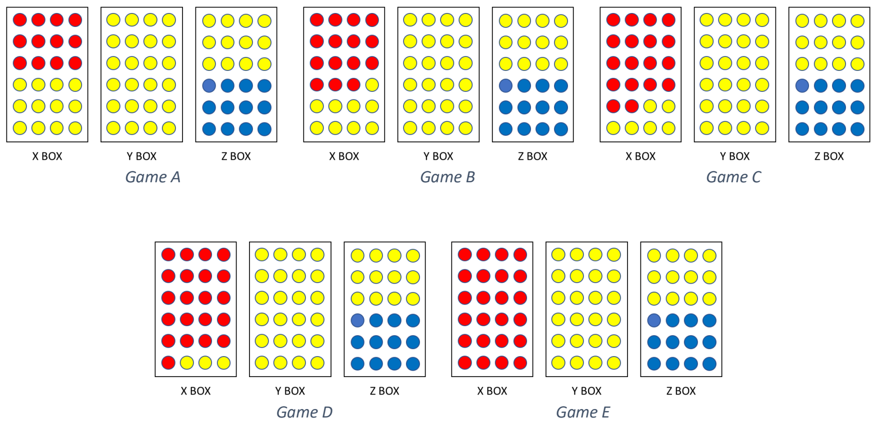
X BOX
Y BOX
Z BOX
Game A
X BOX
Y BOX
Z BOX
Game B
X BOX
Y BOX
Z BOX
Game C
X BOX
Y BOX
Z BOX
Game D
X BOX
Y BOX
Z BOX
Game E

4

FORMING A PARTNERSHIP

When you are an active player, after you have received a clue about your potential partner's type, you can decide whether you want to form a partnership or not. Note that passive players cannot make any choice. Thus, **a partnership is formed whenever the active player in the pair decides so**.

POINTS

Every time you form a partnership, you earn an amount of points that depends only on the type of the other player. Whenever you do not form a partnership, you earn an amount of point that depends only on your own type. The actual points you can earn are reported in the following table. If you form a partnership with an X type you earn 160 points, if you form a partnership with a Y type you earn 80 points and if you form a partnership with a Z type you earn 40 points. If you do not form a partnership and your type is X, you earn 100 points. If you do not form a partnership and your type is Y, you earn 75 points. If you do not form a partnership and your type is Z, you earn 25 points.

Points if you **form** a partnership:			
Partner's type	X	Y	Z
Your points	160	80	40
Points if you do **not form** a partnership:			
Your type	X	Y	Z
Your points	100	75	25

PAYMENT

At the end of the experiment, your payment will depend on the points you have earned in this part. In particular, the computer will randomly select one round out of the 60 rounds in part 1. At the end of the experiment, the points you have earned in the selected round will be converted to Euros at an exchange rate of 1/6 Euros per point.

PART 1 APP PAGES

The experiment app will show you several pages, described below. Between a page and the next you may have to wait for other participants to make their choices. At the beginning of Part 1 you will see a page like this.

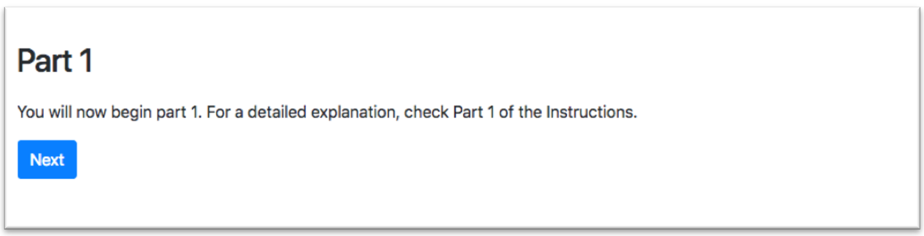

At the beginning of a new round, you will see a page informing you of whether you are playing game A, B, C, D or E.

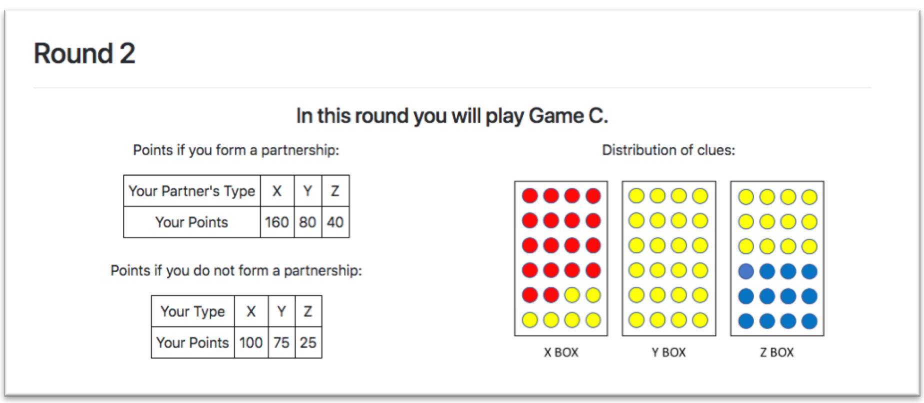

In this page, you are told your type and whether you are an active player or passive player.

7

In this round you are passive.

You are an X type (if you do not form a partnership, you earn 100 points).

Next

If you are an active player in this round, you are also given a clue about your potential partner's type and you decide whether you want to form a partnership or not.

In this round you are active: make a choice.

You are a Y type (if you do not form a partnership, you earn 75 points).

You have received a yellow clue.

Do you want to form a partnership?

○ Yes ○ No

Next

At the bottom of this page you can see a table with a history of all the previous rounds.

History Table

Round	Game	Your Role	Your Type	Your Clue	Your Choice	Partner's Choice	Partner's Type	Your points
1	A	active	Y	Red	Yes	-	X	160
2	C	active	Y	Yellow	-	-	-	-

If you do not click the "next" button on this page, after one minute the app will move you to the next page. In the next page you are told your points in this round and the actual type of your partner.

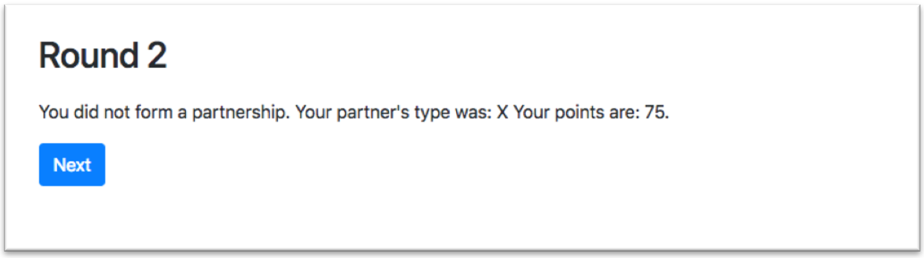

When you have played the last round of Part 1, you will see a page informing you that this part is over. In this page you will find out the Part 1 paying round. At the bottom of this page you can see the history table summarizing all the rounds.

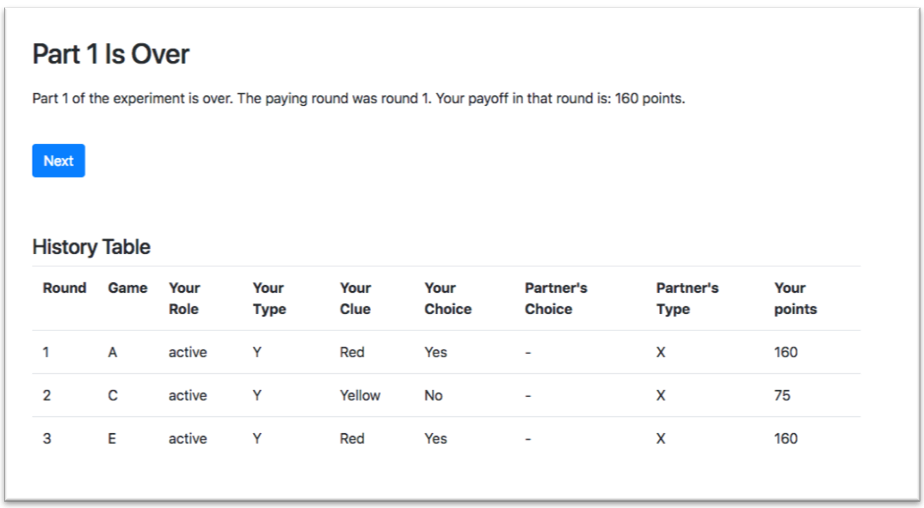

Round	Game	Your Role	Your Type	Your Clue	Your Choice	Partner's Choice	Partner's Type	Your points
1	A	active	Y	Red	Yes	-	X	160
2	C	active	Y	Yellow	No	-	X	75
3	E	active	Y	Red	Yes	-	X	160

PART 2

In this part, you will face 10 decisions listed on your screen. In each decision you have to choose between "Option A" and "Option B". If you choose Option A, you will earn either 6 or 5 points. If you choose Option B, you will earn either 10 points or 1 point. After you choose one option, whether you earn the higher payoff or the lower payoff is randomly determined by the computer. Before making a choice, you will know the exact probability of earning the higher payoff rather than the lower payoff in each option. For example, in one decision Option A will give you 6 points with a probability of 30% and 5 points otherwise, while Option B will give you 10 points with a probability of 30% and 1 point otherwise.

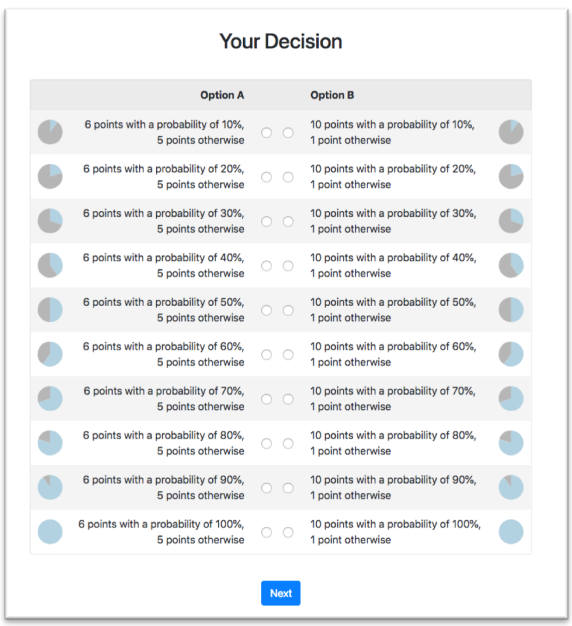
Your Decision

Option A	Option B
6 points with a probability of 10%, 5 points otherwise	10 points with a probability of 10%, 1 point otherwise
6 points with a probability of 20%, 5 points otherwise	10 points with a probability of 20%, 1 point otherwise
6 points with a probability of 30%, 5 points otherwise	10 points with a probability of 30%, 1 point otherwise
6 points with a probability of 40%, 5 points otherwise	10 points with a probability of 40%, 1 point otherwise
6 points with a probability of 50%, 5 points otherwise	10 points with a probability of 50%, 1 point otherwise
6 points with a probability of 60%, 5 points otherwise	10 points with a probability of 60%, 1 point otherwise
6 points with a probability of 70%, 5 points otherwise	10 points with a probability of 70%, 1 point otherwise
6 points with a probability of 80%, 5 points otherwise	10 points with a probability of 80%, 1 point otherwise
6 points with a probability of 90%, 5 points otherwise	10 points with a probability of 90%, 1 point otherwise
6 points with a probability of 100%, 5 points otherwise	10 points with a probability of 100%, 1 point otherwise

Next

As in this example, in any one of the 10 decisions, the probability you will earn the higher payoff (6 if Option A is chosen or 10 if Option B is chosen) is the same between option A and option B. In the first decision, at the top of the list, the probability you will earn the higher payoff is 10%. As you move down the table, the chances of the higher payoff for each option increase. In fact, for decision 10 in the bottom row, each option pays the highest payoff for sure. So, your choice in decision 10 is simply between 6 points (Option A) or 10 points (Option B).

For each of the ten decisions, you will be asked to choose Option A or Option B by clicking on the appropriate button. The computer will ensure that you switch at most once from Option A to Option B. If you choose Option A in one decision, the computer will automatically select Option A for all the previous decisions. If you choose Option B in one decision, the computer will automatically select Option B for all the following decisions. Once you have made a choice in all decisions, you can click on the Next button to submit your choices.

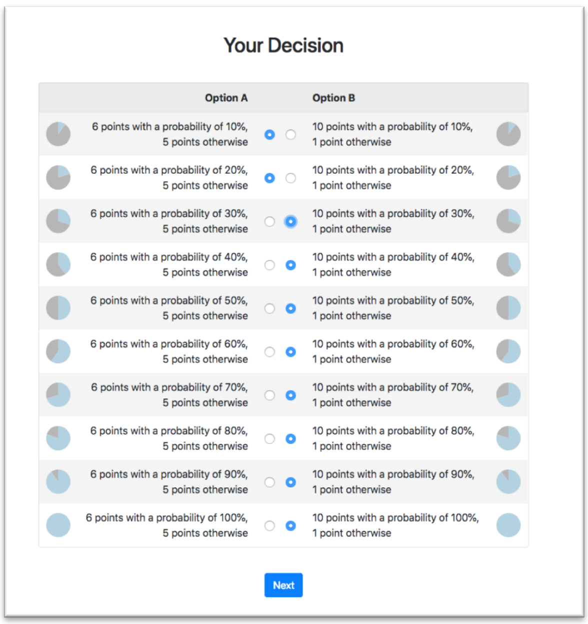

After you have submitted your choices, one of the 10 decisions will be randomly chosen for your payment. For the option you chose, A or B, in this decision, the computer will randomly determine whether you earn the higher or lower payoff. To determine the outcome of your choice, the computer will digitally draw a random number between 0 and 100. If the random number is below the probability of earning the higher payoff, then you receive the higher payoff. If the random number is above the probability, then you receive the lower payoff. For example, assume you chose Option A in the first decision and this decision is selected for payment. If the computer randomly draws a 60, you will earn 5 points.

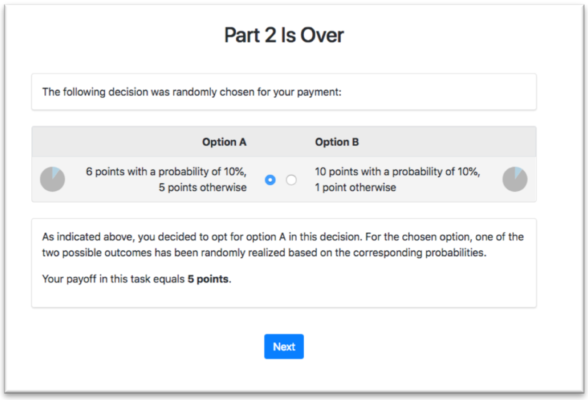

At the end of the experiment, the points you have earned in the selected decision will be converted to Euros at an exchange rate of 1/6 Euros per point.

PART 3

In this part, you are asked to answer three questions. For each correct answer, you will receive two points. After you have submitted your answers, you will see the correct answers and the amount of points you have earned. At the end of the experiment, the points you have earned in this part will be converted to Euros at an exchange rate of 1/6 Euros per point.

PART 4

In this part, you are asked to provide some information about yourself (your sex and your undergraduate major). As stated before, your responses are completely confidential and anonymous.