Out Of Your Mind: Eliciting Individual Reasoning in One Shot Games *

Konrad B. Burchardi†
London School of Economics

Stefan P. Penczynski‡
Cornell University

June 2011

Abstract

We experimentally investigate the fundamental element of the level-k model of reasoning, the non-strategic level-0. We use data from an experimental design that allows to obtain incentivised written accounts of individuals' reasoning. We show that around one third of the participants play non-strategically. The non-strategic level-0 actions are not uniformly distributed. They might be shaped by salience considerations. Higher level players correctly anticipate the non-uniform distribution of non-strategic actions.

JEL Classification: C91, D83

Keywords: individual reasoning, one shot games, level-k model

*We are indebted to Madhav Aney, Ayala Arad, Dan Benjamin, Colin Camerer, Mariano Cena, Syngjoo Choi, Vincent Crawford, Ziad Daoud, Erik Eyster, Nagore Iriberri, Gilat Levy, Rosemarie Nagel, Michele Piccione, Mireia Raluy i Callado, Ariel Rubinstein, Reinhard Selten, and in particular Georg Weizsäcker for helpful discussions and comments on earlier drafts. We thank seminar audiences at Cornell, ESA International Meeting 2009, ESA North America Meeting 2009, GfeW Meeting 2008, LSE, Nordic Behavioural and Experimental Conference 2008, Universitat Pompeu Fabra, and University of East Anglia for comments. Kara Contreary, Jan Zapal, and Piotr Zurawski provided excellent research assistance. We are grateful for the financial support of the Michio Morishima fund (STICERD), the London School of Economics and the German Association of Experimental Economists (GfeW, Heinz-Sauermann Förderpreis 2007). This work was partly accomplished at the Universitat Pompeu Fabra in Barcelona, Spain, and benefitted from its hospitality. Stefan Penczynski gratefully acknowledges the financial support of the German Academic Exchange Service (DAAD Doktorandenstipendium Nr. D/05/44596) during his stay in Barcelona.

†STICERD, LSE, Houghton Street, London WC2A 2AE, United Kingdom (K.B.Burchardi@lse.ac.uk)

‡446 Uris Hall, Cornell University, Ithaca, NY 14850, USA (spp45@cornell.edu)

1 Introduction

Equilibrium concepts have proven to be a powerful tool in economics, political science, international relations and other fields when trying to understand and predict strategic behaviour of individuals, firms, countries, and other entities. However, manifold experimental studies of human behaviour have shown that the equilibrium concept does poorly in predicting outcomes of one shot games, even when a unique equilibrium exists (see for example Nagel, 1995, Rubinstein, Tversky, and Heller, 1996 or Camerer, 2003 for an overview). This class of games reflects *strategic situations without precedent* which are faced, for example, by consumers who – otherwise price-takers – buy a house and bargain with the seller about its price.

A leading explanation for the failure of equilibrium concepts to explain and predict behaviour in one shot games is that players may not believe that other players choose an equilibrium strategy. This explanation appears particularly pertinent in unprecedented strategic situations, where learning cannot cause a convergence of beliefs and strategies to equilibrium.[1] The level-k model of reasoning, as first proposed by Nagel (1995) and Stahl and Wilson (1995), postulates the following alternative belief structure: There exist so-called level-0 players, who do not play strategically. The model defines level-1 players to best respond to what they believe level-0 players do. Level-2 players form a belief about the fractions and strategies of lower level players and best respond to this. This process continues for higher level players. Hence the model assumes a hierarchy of types who best respond to non-equilibrium beliefs, referred to as *levels*. They distinguish themselves by the number of iterated best responses to the distribution of level-0 actions.[2]

Although there is extensive empirical work in support of the level-k model, little is known about its anchoring elements: Do non-strategic *level-0* players exist? How do they choose their actions? And do strategic players correctly anticipate the actions of non-strategic players? Answers to these questions are required for the model to have predictive power.

The lack of empirical evidence is no coincidence. Firstly, the prominent theoretical papers typically assume level-0 actions to be uniformly distributed over the action space and the beliefs of higher level players to be non-heterogeneous and consistent with the actions of level-0 players. These assumptions naturally exclude any investigation of the level-0 actions

[1] A distinct explanation could be that players fail to choose the best response to their belief about the other players' strategies. This is the idea behind the Quantal Response Equilibrium (QRE) proposed by McKelvey and Palfrey (1995) in which actions of higher expected payoff are more likely.

[2] Applications of the level-k model include dominance solvable games (Nagel, 1995; Costa-Gomes and Crawford, 2006), normal-form games (Stahl and Wilson, 1995; Costa-Gomes, Crawford, and Broseta, 2001), two-person zero-sum games with non-neutral framing (Crawford and Iriberri, 2007a), common-value auctions (Crawford and Iriberri, 2007b), and coordination games (Crawford, Gneezy, and Rottenstreich, 2008). For a more complete overview see Camerer, Ho, and Chong (2004) and Crawford, Costa-Gomes, and Iriberri (2010).

and beliefs. We generalise standard versions of the level-k model by allowing for (i) a non-uniform level-0 action distribution, (ii) heterogeneous level-0 beliefs and (iii) the distribution of level-0 beliefs to be independent of the distribution of level-0 actions. These are necessary generalisations to answer the questions of interest. To fix ideas, we present a formal version of this generalised model in Section 2.

Secondly, there are important empirical obstacles to investigating these questions: choice data alone can typically be explained by various reasoning patterns – different levels of reasoning and associated level-0 beliefs or actions – and hence are not informative. In particular, non-strategic behaviour cannot be robustly identified and the level-0 beliefs are difficult to uncover. Moreover, in repeated games individuals are undergoing a learning process, which renders the use of a sequence of observations difficult.

We present two methods which allow to investigate the above questions. Both rely on an experimental design that give access to *incentivised* written accounts of individual reasoning which are stated at the time of the decision-making. These accounts are obtained through a particular team communication protocol which can be applied to any one shot game. We played a standard 'beauty contest' game: each participant is asked to state a number between 0 and 100 and the participant whose number is closest to 2/3 of the average of all numbers wins a prize. The unique Nash equilibrium of this game is to play 0 for all participants.

The first empirical strategy is to obtain a non-parametric estimate of the distribution of level-0 beliefs and level-0 actions from the written accounts of reasoning. We detect non-strategic reasoners and obtain an estimate of the distribution of level-0 actions. We find that at least 20% of the participants play non-strategically, in the sense of not attempting to best respond to any belief. The mode and median of the actions of these players are close to 60, and the mean is significantly higher than 50. This is consistent with the interpretation that salience due to the multiplier of 2/3 shapes the actions of non-strategic players. It suggests that level-0 players can be understood as having no stronger reasons for choice other than salience. Further, we analyse stated level-0 beliefs of higher level players. While the majority of them starts their reasoning at exactly 50, there is substantial heterogeneity. We cannot reject the hypothesis that the beliefs are, on average, correctly anticipating the mean of the level-0 actions.

The second empirical strategy is to estimate the structural parameters of our generalised level-k model with a maximum likelihood estimation. In the estimation we make use of information about the players' sophistication from the written accounts. We estimate about one third of participants to be playing non-strategically. The distributions of level-0 actions and beliefs have an estimated mean of 58 and 54, respectively. These are close to the non-parametric estimates obtained with the first empirical strategy. The belief distribution is more concentrated than the action distribution. However, it still has substantial variance.

The estimation of these structural parameters is only possible under the generalised version of the model. Therefore, acknowledging heterogeneity in the beliefs does not only seem correct, but it as well enables novel estimations of the level-k model's structural parameters.

The two empirical strategies provide consistent evidence on the questions this paper seeks to answer. We show that around one third of the participants play non-strategically. The non-strategic level-0 actions are not uniformly distributed. They might be shaped by salience considerations. Higher level players correctly anticipate the non-uniform distribution of non-strategic actions.

After presenting the generalised level-k model in the next section, this paper proceeds as follows: In section 3 we discuss the existing empirical literature on the level-k model. We then present the experimental design and procedures in sections 4. Results from the two empirical strategies are presented in sections 5 and 6. The final section offers concluding comments.

2 Generalised Level-k Model

Consider a game in which N players simultaneously choose an action $x \in X$. Level-k type models deviate from equilibrium models in that they allow for heterogeneity in the belief about other players' actions. A player is characterised by a level of reasoning k, $k = 0, 1, 2, 3, \ldots$. Denote the fraction of level-k types in the population with l_k and define $\boldsymbol{l} := \{l_0, l_1, l_2, ...\}'$. Level-0 players are defined as playing non-strategically in the sense that their action does not depend on a belief over other players' actions. We denote the distribution of their actions by $g^0(x \,|\, \boldsymbol{\theta}^0)$, where $\boldsymbol{\theta}^0$ is a parameter vector characterising the distribution, and refer to it as 'level-0 distribution'. Players of a higher level of reasoning ($k \geq 1$) are thought to best respond to the actions of lower level players. The belief about the distribution of those actions is derived from (i) a belief about the level-0 distribution and (ii) a belief about the relative proportions of lower-level players. We refer to the belief about the level-0 distribution as 'level-0 belief'. Note that in most games, the player only needs to form a belief about one or several statistics of the level-0 distribution, in the sense that conditional on knowledge of these statistics, information about other aspects of the level-0 distribution leaves the best response unchanged. For example in the 'beauty contest' game a level-k player, $k \geq 1$, only needs to form a belief about the mean of the level-0 distribution. We allow the level-0 belief about the vector of statistics $\boldsymbol{d}$ to be heterogeneous across individuals. Let its distribution be characterised by the probability density function $g^b(\boldsymbol{d} \,|\, \boldsymbol{\theta}^b)$, where $\boldsymbol{\theta}^b$ is again a vector of parameters. Denote the belief of a level-k player with $k \geq 1$ about the proportion of level-i players in the population as $b_k(i)$, $i \leq k-1$. Given this belief a level-2 player can calculate the distribution of actions of lower level players and best respond to it. Higher level players find their best response analogously.[3] In conclusion, the strategy of a level-k player, $k \geq 1$, is found as probability distribution over

$$s_k(b_k(\cdot), g^b(\cdot)) = \underset{x \in X}{\operatorname{argmax}}\, u(x; b_k(\cdot), g^b(\cdot)). \tag{1}$$

Given an assumption on the anchoring level-0 distribution, the level-0 belief, the true level-k distribution, and a specification for $b_k(\cdot)$, the level-k model makes a probabilistic prediction about the frequencies of actions.

This general level-k model nests the models by Nagel (1995), Stahl and Wilson (1995), Costa-Gomes, Crawford, and Broseta (2001), Camerer, Ho, and Chong (2004) and Costa-Gomes and Crawford (2006) as special cases.[4] All of these models assume that the level-0 actions are uniformly distributed over the action space, that higher level players' level-0 beliefs are non-heterogeneous and that they correctly anticipate the actions of level-0 players.

3 Literature

The paper relates to an extensive empirical literature on the level-k model of reasoning. Most of this literature is concerned with estimating the distribution of level-k types in a given population. For example, Camerer, Ho, and Chong (2004) use a GMM estimator, essentially choosing the parameter of a one-parameter version of the level-k model (as well referred to as 'cognitive hierarchy model') to match the mean of the action data. This is an elegant methodology to estimate the level-k distribution that assumes certain level-0 action and belief distributions. An approach which does not rely on a structural model is proposed in Bosch-Domènech, Montalvo, Nagel, and Satorra (2004), who use a large dataset to fit a

[3]Note that for players with $k \geq 2$ higher-order beliefs need to be specified. In particular, a player needs to form a belief over the lower-level players' beliefs of the level-0 distribution. Similarly, a $k \geq 3$ player needs to form a belief about the population beliefs of player i, $2 \leq i \leq k-1$, i.e. a level-3 player needs to know what a level-2 player believes the relative proportions of level-0 and level-1 players are. Here and in the literature, these higher-order beliefs are assumed to be consistent with the beliefs of the lower-level player.

[4]Nagel (1995) assumes the players' population beliefs for $k > 0$ to be degenerate on $k-1$: $b_k(i) = 1$ if $i = k-1$ and $b_k(i) = 0$ otherwise. She assumes – as all later versions of the level-k model – the level-0 distribution to be uniform and every player's level-0 belief to be consistent, which makes $g^b(\cdot)$ degenerate. Stahl and Wilson (1995) study normal form games and present a version of the level-k model where best responses are calculated with error. The uniform level-0 distribution, $g^0(\cdot)$, reflects a fully imprecise best-response. Levels 1 and 2 best respond with error, playing actions with a higher expected payoff with a higher probability. Higher levels than $k = 2$ are not considered and the population distribution belief is not restricted to be degenerate. Costa-Gomes, Crawford, and Broseta (2001) and and Costa-Gomes and Crawford (2006) model the level-k types in a similar fashion as Nagel (1995), but assume level-0 players not to exist. Further they consider 'equilibrium' types that play the Nash equilibrium strategy and call 'sophisticated' types those players that best respond to the actual distribution of others' responses. Camerer, Ho, and Chong (2004) introduced the 'cognitive hierarchy' model where the players' beliefs reflect the true relative frequencies of lower level types and the true distribution of types $l(k)$ follows a Poisson distribution with parameter τ. Formally, for all $k > 0$, $b_k(i) = l(i;\tau)/\sum_{m=0}^{k-1} l(m;\tau)$, where $l(k;\tau) = \tau^k e^{-\tau}/(k!)$.

set of normal distributions to action data from various 'beauty contest' games. Some of the super-imposed normal distributions are then associated to underlying levels of reasoning, which provides both estimates of the type distribution and the choices for individual types.

In contrast, we estimate the various parameters of a structural model using maximum likelihood estimation in our second empirical strategy. This is typically not possible, since the level-k model makes stark choice predictions when the distribution of beliefs is assumed to be non-heterogeneous. For example, a uniform level-0 distribution in the 'beauty contest' game, together with non-heterogeneous, consistent beliefs about the mean of level-0 actions, imply a degenerate belief at 50. Any action slightly off 33, 22, etc. will be attributed to level-0 play. A maximum likelihood estimation then typically yields almost all players to be level-0 players and fits the level-0 distribution to the full-sample action distribution. We instead explicitly model heterogeneity in the level-0 belief, which causes heterogeneity in actions of players of the same level-k, $k > 0$, and allows us to separate this from level-0 play. In this sense the generalisation of the model is key to our estimation strategy.

A central challenge faced by the empirical literature is that the level-k model can typically explain a given *action* in a one-shot game with several reasoning patterns.[5] Like we do, other studies try to circumvent this problem by obtaining additional information which allows to conclude about the underlying reasoning. Costa-Gomes, Crawford, and Broseta (2001) and Costa-Gomes and Crawford (2006) obtain for each player multiple choices made in subsequently played variants of a game without feedback. Under the assumption of a constant reasoning level, they then match a player's sequence of choices to a typical 'fingerprint' of, say, a level-1 player. This provides estimates of individual levels of reasoning. In addition to the 'fingerprint', they used information search data to identify types of players.[6] Tracking both choices and response times, Rubinstein (2007) associates longer response times with more cognitive effort, differentiating between cognitive, instinctive and reasonless choices. In a similar vein, Agranov, Caplin, and Tergiman (2010) incentivise and observe provisional choices over time in order to get insights in the choice process, including initial naive considerations. In her original study, Nagel (1994) asked the participants after the experiment to verbally state the reason for their chosen action and then classified and analysed comments of participants. A descriptive analysis of optionally given verbal comments received in the context of a newspaper experiment is presented by Bosch-Domènech, Montalvo, Nagel, and Satorra (2002). These studies do not analyse the non-strategic actions and beliefs thereof.

Our experimental design analyses written accounts of reasoning. We carefully ensure

[5] For example, a player who chooses 33 in the standard 'beauty-contest' game might do so because he is a level-1 reasoner who believes that level-0 reasoners play on average 50. But he might just was well be a level-0 reasoner who has chosen the number at random or a level-2 reasoner who beliefs that the population is composed of a combination of level-1 and level-0 reasoners who on average choose 50.

[6] This method was introduced by Camerer, Johnson, Sen, and Rymon (1993).

that these are written during the actual thinking period and provide an incentive to the participants to state their reasoning fully and clearly.[7] This is central to our empirical strategy, since only with such incentives we can be sure that the reasoning process is fully represented in the written accounts. In the absence of such incentives one is likely to underestimate the level of reasoning.

At the heart of our experimental design lies the use of team communication as a means of observing individual reasoning. The experimental literature has used team setups on various occasions to obtain insights into the reasoning process of participants and to investigate the performance of teams as opposed to individuals. In this respect, our experimental design is related to the innovative study by Cooper and Kagel (2005) who were the first to let team players communicate via an instant messenger, allowing the experimenters to observe the speed of learning in strategic play. We use a communication protocol in which the message is written prior to any team interaction, so that it purely reflects individual reasoning in a one-shot game.

4 Observing Individual Reasoning

We present two strategies to estimate the distribution of level-0 actions and beliefs thereof. Both empirical strategies rely on an experimental design which allows us to obtain incentivised accounts of individual reasoning. This section explains our design. It emphasises how our design ensures that individual participants have an incentive to state their reasoning fully and clearly. It as well describes how we extract information from these written accounts by classifying the reasoning pattern along the lines of a general level-k model. We corroborate the robustness of our findings and the replicability of the classification procedure.

4.1 Experimental Design

In order to elicit the reasoning underlying a player's action, we designed the following game structure: Individuals are randomly assigned in teams of two players. Their payoff in the game depends on a joint 'team action'. To determine this, both players are given the chance to choose an action – which we call the 'final decision'. Then one player's decision is chosen randomly, with probability one half as the team action. Consequently, each player faces a 50% chance of having her partner's final decision determining the team action. The players hence have an incentive to ensure that their team partner's final decision is as sound as

[7] Indeed, the standard theme in protocol analysis, the research field in psychology concerned with methods of eliciting verbal accounts from participants, is that "[t]he closest connection between thinking and verbal reports should be found when participants were instructed to focus on the task while verbalising their ongoing thoughts" (Ericsson, 2002, p. 983).

possible. Importantly, the players are given the possibility to convince their partner of the optimal team action. In particular, players are allowed to write one message to their team partner, which consists of a 'suggested decision' and a justifying text. This text is unlimited in size and its writing is not limited in time. The messages are exchanged *simultaneously* once both players have entered their message, and thereafter the players take their final decision individually.

The simultaneous exchange of a single message ensures that every explanatory statement and suggested decision in the first round is written without any previous communication with the team partner, hence reflecting an individual's reasoning. It is therefore this first message and suggested decision which we will analyse in this paper in order to understand individual decision making in situations of strategic interaction.[8] Importantly, under the reasonable assumption that the best way to convince one's team partner is by explaining him the own reasoning, this design gives an incentive to write down the reasoning process as fully and clearly as possible. It is generally applicable in one shot games and we believe it constitutes an important methodological contribution of our paper.

The design has two potential caveats. First, the suggested decision is taken while knowing that the opponents are teams of 2 players. This leads, if anything, to a population belief with more weight on higher levels due to the team reasoning. Second, the decision is taken while justifying it in a message. This might lead to a more thoughtful decision. However, we cannot exclude the possibility that the requirement to explain one's decision can cause stress, leading to a lower level of reasoning. In any case, if these distortions were present, this should show up in the action data. It is comforting that the distribution of suggested decisions is similar to other non-communication treatments in the literature, e. g. Nagel (1995).

4.2 Experimental Procedures

We conducted the experiment in the Experimental Economics Laboratory of the Department of Economics in Royal Holloway (University of London). In 6 sessions we played three rounds of the 'beauty-contest' game. Since in this paper we are interested in individual reasoning, we only analyse the first round suggested decision and the accompanying first message, i. e. with the activities that took place before any interaction of the team players.

At the start of each session the participants were made familiar with the structure of the experiment and the messaging system in two practice rounds. We used the same software as above for the practice rounds, but asked the teams to find the answer to two unrelated questions. Since we wanted to avoid any pre-treatment sensitisation to strategic considerations, we asked them to provide the year of two historic events. The questions in the test

[8] The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007).

round were chosen to be relatively difficult to stimulate the use of the messaging system. The participants of the experiment were paid a show-up fee of £5 and the winning team won a prize of £20 (£10 per team player).

A total of 84 individuals participated in our experiment. Sessions had 12, 14 or 16 participants. The participants were mainly undergraduate students in Royal Holloway and all of them were recruited by the host institution. Out of the 84 students 15 were studying Economics, 13 of them being in their first year of studies, one in the second year, one being in the third year. 16 of the 84 students had received some form of training in game theory, but only 5 had been confronted with the 'beauty contest' game. The majority of students had participated in an economic experiment before.

4.3 Classification of Communication Transcripts

We use the information elicited on individual reasoning in the following way: First, we identify the cases in which no level-k type of reasoning is undertaken. Then, conditional on a level-k reasoning being applied by the player, which includes non-strategic play, we uncover two sets of data: (a) the maximum and minimum number of steps of reasoning which can be interpreted into the message, including possibly 0, and (b) the level-0 belief which a player states, if any.

In particular, two research assistants read the messages and classified the type of reasoning with the following procedure:

1. To investigate the prevalence of reasoning patterns different from level-k reasoning we asked the RAs to indicate whether the player puts forward equilibrium reasoning. For this it was not necessary that the player actually played the unique equilibrium strategy.[9] The RAs were further instructed to denote whether the player applied an iterated elimination of dominated strategies. For this it is necessary that *first* some actions are excluded and then a strategy is formed for the remaining action space.[10]

2a. If any level-k reasoning was explained in the message, we asked the RAs to indicate if the argument contained a belief about others' play that served as a starting point for best responses, but was in itself not derived by choosing a best response. If so, we asked them to denote this level-0 belief.[11]

[9] We call a player a 'sophisticated' type, if she recognised the equilibrium, but played an action different from the unique equilibrium action.

[10] For example the statement "Everybody plays on average 50 so I should not play higher than 34" is not an iterated elimination of dominated strategies.

[11] For completeness, we also asked whether an argument revealed a population belief distribution. If so, the classifiers were asked to indicate whether it was degenerate or non-degenerate.

2b. Lastly, we were interested in how many steps of reasoning the player applies. When designing the classification procedure we were worried that in some cases it might not be possible to identify from the communication *exactly* how many steps of reasoning were applied.[12] We therefore asked the classifiers to only indicate the lowest level of reasoning which is clearly stated and the highest level of reasoning which could possibly be interpreted into the messages.[13] We refer to these as 'lower bound' and 'upper bound', respectively. We instructed the classifiers to consider as level-0 a player whose message does not exhibit "any strategic reasoning whatsoever". This might arise as a result of choosing a number randomly or based on non-strategic considerations such as taste. We emphasised that for this classification to be chosen, it was important that the player was not in any way best responding to what he thought others would play.

When designing the classification procedure we intended to avoid two potential concerns: First, the classifiers might try to extract more information than the messages actually contain. We therefore instructed the classifiers to only enter information when it was clearly contained in the message.[14] Second, we were concerned that in the 'beauty-contest' game in the case of an ambiguous statement relatively low suggested decisions might lead the classifiers to indicate a higher lower-bound on the level of reasoning than was clearly exhibited. In contrast when indicating the upper-bound, knowing about low choices should, if anything, lead the classifiers to indicate a higher upper-bound. We therefore split the classification of the messages into two parts. We did not reveal the choice data to the classifiers when asking for the lower-bound but revealed it subsequently when asking for the upper bound.[15]

The classification was undertaken by two Ph.D. students in the Department of Economics at LSE. First they classified the transcripts individually. After this phase their classification of the lower bound coincided for 77% of all participants and the classification of the upper bound coincided in 76% of all cases. Then the two RAs met to reconcile their judgements

[12] Think for example of the imaginary statement: "I presume everybody else will play 33, so let us play 22." This clearly exhibits one step of reasoning. But it seems possible, too, that the player skipped the first step of his reasoning when writing down his argument.

[13] The instructions specified that the classifiers, after writing down the lower bounds should be able to say to themselves: "It seems impossible that the players' level of reasoning is below this number!", and after writing down the upper bounds: "Although maybe not clearly communicated, this statement could be an expression of this level. If the player reasoned higher than this number, this was not expressed in the statement!"

[14] The instructions were self-contained and were not complemented by verbal comments. The instructions were written by the two authors, of whom one had taken a look at the communication transcripts beforehand. The instructions can be obtained from the authors upon request. Remaining questions were answered via an e-mail list that included all four persons involved and which can be obtained from the authors.

[15] Other studies applying a classification use a similar procedure in order to avoid any unconscious alignment of the classification with the choice data that might result from implicit assumptions (for example Rydval, Ortmann, and Ostatnicky, 2009).

and provide a joint classification, if possible. We only use data on which they could agree in the reconciliation.

Later we asked further 6 RAs to again classify the lower bounds. Table 6 in appendix A.7 shows that for 70 out of 78 messages (∼90%), 6 or more of the 8 classifiers agreed on exactly one level.[16] We take this as comforting evidence that our method of classification is robust, provides informative insights about individual reasoning and can easily be replicated.

5 Experimental Results

5.1 Action Data

Table 1 presents aggregate summary statistics for all 6 sessions. Figure 1 shows histograms of the suggested decision and the final decisions aggregated over all sessions. The suggested decision is comparable to the first period's decision of other experiments with individual participants. However, the final decision is not, since the participants have, at the time of taking this decision, already received a message from their team partner. For the subsequent analysis and classification of the individual reasoning we will exclusively use the suggested decision.

Table 1: Summary statistics

	Mean	Std. Dev.	Median	Min.	Max.	N
PANEL A: Full Sample						
Suggested Decision	43.93	21.14	40	0	100	84
Final Decision	39.73	18.75	35	0	100	84
Team Action[a]	40.02	18.98	35.5	16	100	84
PANEL B: Level-0 players						
Level-0 Action	62.35	22.39	60	16	100	17
PANEL C: Non-level-0 players with stated belief						
Level-0 Belief	55.26	12.33	50	40	100	36

Notes: [a] The team action is a random draw of the two final decisions.

The data on the suggested decision is similar to data generated in other comparable experiments in having similar means and a high fraction of choices between 20 and 50. The original study by Nagel (1995) had an average of 36.6, which is slightly lower than ours.[17] A concern with our design is that having to communicate to the team partner might increase

[16] The 6 subjects that did not write a message are dropped from this exposition.

[17] The spike at 40 might be unusual. The communication data reveals that this arises mainly as a result

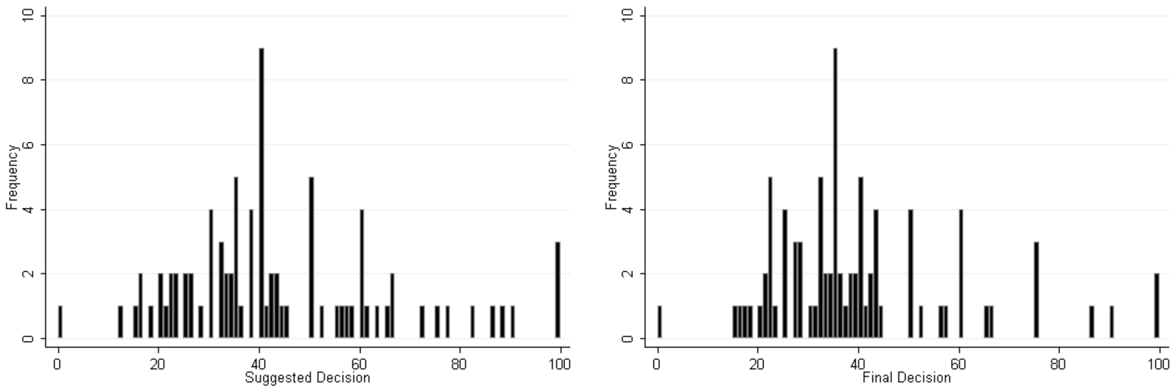

(a) Suggested decisions.

(b) Final decisions.

Figure 1: Individual decisions.

the participants' level of reasoning, e. g. because the participants would examine the task at hand more thoroughly in order to state sensible arguments in the communication. The fact that our data exhibits – if anything – slightly higher action choices is reassuring in this respect.

The final decision of the participants has a lower mean and median than the suggested decision. This supports the intuition that the group-decision making process increases the level of reasoning and hence, in the 'beauty contest' game, leads to lower chosen numbers on average. Moreover the standard deviation of decisions across participants decreases after the exchange of the message and outliers become fewer, consistent with the idea that level-0 players who potentially chose the outliers, become fewer.[18]

5.2 Level-k Bounds

Table 2 presents the lower and upper bounds on the level of reasoning of individuals. For 70 participants both a lower and an upper bound was indicated. Eight participants have a non-classified upper bound. Another 6 participants did not make any statement and could therefore not be classified.

For 50 of the 84 participants the lower and the upper bounds coincide ($\approx 60\%$) and hence the classification fully determines their level of reasoning. These are 17 level-0, 26 level-1, 6 level-2 and 1 level-3 players, corresponding to the diagonal of Table 2. For further 20 players the classification restricts the level of reasoning to be one of two possibilities. Only for one

of two factors: some level-0 players chose 40 and some level-1 players chose 40 as they held the belief that the level-0 mean would equal 60. This insight shows how relaxing the dependence on action data allows for an analysis of the structural characteristics of reasoning.

[18]An analysis of this process of persuasion is presented in Penczynski (2010b).

Table 2: Level classification results

		Level upper bounds					
		0	1	2	3	NA	Total
	0	17	11	1	0	6	35
Level	1		26	3	0	2	31
lower	2			6	5	0	11
bounds	3				1	0	1
	NA					6	6
	Total	17	37	10	6	14	84

participant we have an interval between 0 and 2. None of those participants for whom both a lower and upper bound is indicated was classified as potentially reasoning higher than level 3.[19]

Two players identified the Nash equilibrium, one of them being an 'equilibrium'-type that suggested playing 0, another one being 'sophisticated' in the sense that she imitated a level-2 player and suggested 20. No upper bound was assigned to those players who identified the equilibrium. Further two participants were found to apply elimination of dominated strategies.

For our purposes, the main take-away from the data is that at least 20% of the subjects are non-strategic reasoners, since 17 out of 84 participants are identified as level-0 reasoners. This is, however, only a lower bound on the fraction of level-0 reasoners, since for an additional 24 participants we cannot exclude the possibility that they are level-0 reasoners. When estimating the level-k model, we estimate over one third of the population to be level-0 reasoners (37%, see section 6.3).

5.3 Level-0 Action

We analyse the actions chosen by those with a lower and upper bound of 0 in order to provide an estimate of the distribution of level-0 play. Panel B of table 1 shows summary statistics of the suggested decisions of level-0 players and figure 2 presents the corresponding histogram. A one sample Kolmogorov-Smirnov test suggests that the distribution is significantly different from uniform (p-value$= 0.038$). Level-0 players choose on average 62 and their median choice is 60. The hypothesis that the mean of their choices is equal to 50 is rejected when tested against the alternative hypothesis of a higher mean (t-test, p-value$= 0.019$).

We find a higher frequency of level-0 actions around 50, 66 and 100. These can be seen as focal points in the spirit of Schelling (1960): 50 and 100 are focal due to the action space being integers between 0 and 100 and 66 may be focal due to the multiplier in the

[19] The data by subject can be obtained from the authors upon request.

game being 2/3. This adds to the evidence provided by Bacharach and Stahl (2000) and Crawford and Iriberri (2007a), who argue that salience considerations importantly shape the action distribution of non-strategic players.[20] However, the games they analyse display purposely framed actions. This is not true for the 'beauty contest' game. The fact that we still find an action distribution markedly different from uniform and centered around certain numbers, is consistent with an idea presented by Lewis (1969, p. 35) on the role of salience: he hypothesises that individuals "tend to pick the salient as a last resort, when they have no stronger ground for choice." Level-0 players can be understood as having no stronger reasons for choice than salience.

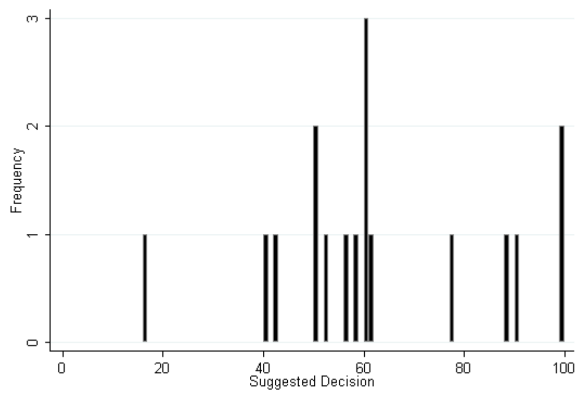

Figure 2: Suggested decisions of level-0 players

5.4 Level-0 Beliefs

The analysis of the written accounts of individual reasoning further allows us to analyse the players' beliefs about the average action of level-0 players. In 36 of the messages the players stated a non-derived belief about the average action of other players. Table 1 shows summary statistics of the stated level-0 beliefs. The mean belief is significantly higher than 50 (one-sided t-test, p-value$= 0.007$). Figure 3 presents the distribution of the level-0 beliefs.

More than 20 participants started reasoning with a level-0 belief of exactly 50. However, 11 players have a level-0 belief between 55 and 66. We interpret this as evidence that level-0 beliefs are indeed heterogeneous. This provides an explanation for why the actions of higher level, same-level players are heterogeneous.

Secondly, and to us surprisingly, the data suggests that level-0 beliefs are closely linked to the distribution of level-0 actions. In particular, they might pick up the same salience considerations which are found in the level-0 actions. Further, we fail to reject the hypothesis

[20]Penczynski (2010a) uses the present paper's method to illuminate the level-0 distribution and belief in the context of the 'hide and seek' game.

that the mean of the level-0-action and -belief distributions as shown in figures 2 and 3 are the same (two sample t-test with unequal variances, p-value: 0.235).

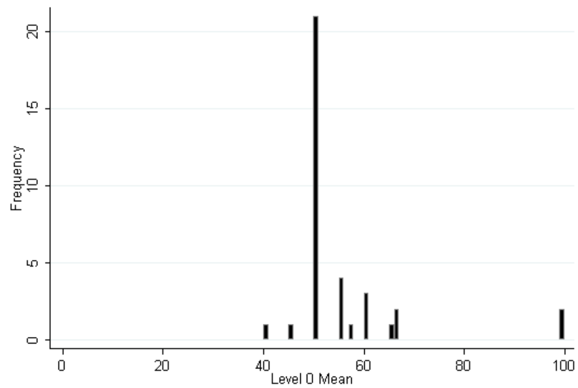

Figure 3: LEVEL-0 BELIEFS FROM COMMUNICATION TRANSCRIPTS

6 Estimation Of Structural Level-k Model

This section discusses how a generalised level-k model, which allows for heterogeneous level-0 beliefs, can be estimated. We will use the suggested decision together with the upper and lower bounds obtained from the classification to estimate the model's structural parameters, in particular the distribution of level-k types and the distributions of level-0 actions and beliefs. Among others this allows us to estimate the fraction of level-0 players.

6.1 An Estimable Model

The level-k model outlined in section 2 makes a probabilistic prediction about the observed actions. Let $f_j(x \,|\, \boldsymbol{\theta}_j)$ denote the probability mass function over the actions of a level j player. Then the unconditional probability mass function of the action of some player i can be written as

$$p(x_i\,;\,\boldsymbol{\psi}) = \sum_{j=1}^{k} l_j f_j(x_i \,|\, \boldsymbol{\theta}_j) \tag{2}$$

where $\boldsymbol{\psi} = (\boldsymbol{l}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2, ..., , \boldsymbol{\theta}_n)$, $l_j \geq 0$ for all j and $\sum_{j=1}^{k} l_j = 1$. This is a convex combination of component densities denoted in the statistics literature as 'finite mixture distribution'. The l_j's give the weight of each component distribution, and the $\boldsymbol{\theta}_j$'s are the parameters which characterize each component distribution. The distribution of all actions is hence a finite mixture of the action distributions of level-0, level-1, level-2 players and so on. We will outline in section 6.2 how to consistently estimate the parameter vector $\boldsymbol{\psi}$. In the following we will first describe which form $f_j(x \mid \boldsymbol{\theta}_j)$ takes, how it depends on the level-0 action and belief distributions and how we parameterise these.

Action distributions in the level-k model The action distribution of a level-0 player, $f_0(x \mid \boldsymbol{\theta}_0)$, is simply the level-0 distribution $g^0(x \mid \boldsymbol{\theta}^0)$. The actions of higher level players, $f_j(x \mid \boldsymbol{\theta}_j)$ with $j \geq 1$, are derived from their level-0 belief, $g^b(x \mid \boldsymbol{\theta}^b)$, taking into account the players' population belief. As the level-0 belief is a random variable, the action distribution of a level-k player can be understood as the distribution of a transformed random variable.[21]

For a general form of population beliefs, this might be a complicated transformation. This is simplified by the fact that we predominantly found degenerate population beliefs in the messages. We will therefore assume degenerate population beliefs throughout.[22] The action of a level-k player then follows from his level-0 belief, say b^*, as $p^k b^*$ or in our case $(2/3)^k b^*$.

Parametric assumptions We parameterise $g^0(x \mid \boldsymbol{\theta}^0)$ as bounded normal distribution defined on the interval $[0, 100]$ and characterised by the mean and standard deviation $(\mu^0, \sigma^0) = \boldsymbol{\theta}^0$. This allows for a concentration of the level-0 actions as well as an approximate uniform distribution when σ^0 is large. We parameterise the distribution of the level-0 belief as well as bounded normal defined on the interval $[0, 100]$ and characterised by the mean and standard deviation $(\mu^b, \sigma^b) = \boldsymbol{\theta}^b$. As a special case this allows for the level-0 belief to be concentrated on the mean of the level-0 action distribution. To find the action distribution of level-k players that follows from this level-0 belief distribution, Lemma 1 in the appendix is useful. It states that a random variable which is distributed as bounded normal on $[0, 100]$ with parameters (μ, σ) will – when applying a multiplicative transformation using a factor a – be distributed as bounded normal with parameters $(a\mu, a\sigma)$ and support $[0, a100]$. Therefore we find the action distribution $f_k(x \mid \boldsymbol{\theta}_k)$, $k \geq 1$, as a bounded normal distribution with $\boldsymbol{\theta}_k = (2/3)^k \boldsymbol{\theta}^b$ and support $[0, (2/3)^k 100]$.[23]

We estimate the fraction of level reasoners for levels 0 to 3. No parametric structure is imposed on the distribution of level-reasoners, but we assume the highest level of reasoning observed to be 3. We do not find any communication that hints to a level of reasoning higher than that.

[21] Note that a non-degenerate level-0 belief distribution will imply a non-degenerate action distribution for higher levels. This is in contrast to other models, where only if a player exactly matches the point prediction of the level-k model he would be classified as a higher level player.

[22] 9 out of 12 players exhibited a degenerate population belief.

[23] Note that in equation 2 the component density f_j is indexed, allowing for the possibility that they are of different parametric families. In the case of the 'beauty contest' we can omit the subscript.

6.2 Estimator and Identification

Likelihood Function Given the probability mass function in equation 2 for the action of a player of unknown level, we can write the log-likelihood of the data as

$$L(\boldsymbol{x};\boldsymbol{\psi}) = \sum_{i=1}^{n} \log p(x_i;\, \boldsymbol{\psi}) \tag{3}$$

where $\boldsymbol{x}$ is the vector of actions observed. We use the information on the bounds by imposing $l_j = 0$ in $p(x_i;\boldsymbol{\psi})$ when the classification information is such that individual i is certainly not of level j.

Identification For the model to be identified, the mixture densities need to be linearly independent for all mixture probabilities $l_j \neq 0$. In the 'beauty contest' game with a parameterisation of the level-0 action and belief distribution as bounded normal this will necessarily be satisfied, irrespective of how many levels are estimated.[24]

Maximum Likelihood Estimator We estimate the parameter vector $\boldsymbol{\psi}$ with the maximum-likelihood estimator, denoted $\hat{\psi}_{MLE}$. The log-likelihood function is thrice differentiable and the expectation of the third partial derivative is finite. We are unable to calculate the true information matrix, but we calculate an estimate of it and verify that it is positive definite. The MLE for this estimation problem is hence consistent, asymptotically normal and with asymptotic variance given by the inverse of the Fisher information matrix. We find the global maximiser of the log-likelihood function numerically, since the likelihood equations cannot be solved analytically. Details of this procedure are given in appendix A.2.

In order to ensure that the estimator is unbiased in a small sample, we ran Monte Carlo studies of sample size $N = 84$. For these we have generated data with the model given in equation 2, using $\hat{\boldsymbol{\psi}}_{MLE}$ for $\boldsymbol{\psi}$. The results give no reason to believe that our estimator is biased. Details are given in appendix A.4.

6.3 Estimation results

Level-k distribution Table 3 shows the estimation results for the level-k distribution. We estimate 47% of the participants to be level-1 reasoners, and 15% to be level-2 reasoners. We estimate only 1% of the participants to be level-3 reasoners. This is similar to the classification results in section 5.2.[25] Crucially, we estimate that more than one third of the players (37%) are level-0 players.

[24] Unless in the special case where $\boldsymbol{\theta}^0 = (2/3)^k\boldsymbol{\theta}^b$ for some $k \geq 1$. For for any non-zero level-0 action mean, this will not be true if the mean of the level-0 belief is at the mean of the level-0 action.

[25] For computational reasons, the maximum level in the 'beauty contest' estimation is level-3.

The estimates in Table 3 on the relative frequencies of level-k reasoners, conditional on $k \geq 1$, resemble the relative fractions of level-1, level-2 and level-3 reasoners found in the literature for various games (see Camerer, Ho, and Chong (2004), Costa-Gomes and Crawford (2006), Crawford and Iriberri (2007a) etc.). This is reassuring and adds to the earlier evidence that the level-k model can have predictive power for the distribution of actions as a function of population and game characteristics.

Table 3: Estimated level-k distribution

Parameter	l_0	l_1	l_2	l_3
Estimate	0.37	0.47	0.15	0.01
	(0.057)	(0.058)	(0.042)	(0.016)

Notes: The table presents the results from a maximum likelihood estimation of the structural model as outlined in section 6.1. This table only presents the results for the level-k distribution, but the level-0 action and belief distribution were estimated simultaneously. Those results are reported in table 4. Bootstrapped standard errors are given in brackets. These are obtained from 200 iterations of our estimation when sampling 84 observations from our data.

However, our estimate of the fraction of non-strategic reasoners is substantially higher than previously estimated. Nagel (1995) associates certain actions with level-0 play and estimates between 2% and 17% of the population to be level-0 reasoners. Camerer, Ho, and Chong (2004) parameterise the type distribution as Poisson and estimate the mean to be roughly 1.5, corresponding to roughly 22% non-strategic play. Note that our estimation procedure does not conflate level-0 play and errors of higher-level players, since we de-couple the level-0 action and belief distribution. Another way of interpreting the heterogeneity in the level-0 beliefs in the estimation is to think of it as errors of higher level players.[26]

Some of the level-k literature suggests that level-0 players only exist in the heads of other players. The evidence presented here sheds substantial doubts on this assumption. Of course, studies differ in the amount of testing that is done before the experiment, which might influence the capability of players to play strategically. In our study, we took care not to train or hint towards any strategic consideration. In our view, this should be the approach for studying one shot games.

Level-0 actions and beliefs Table 4 presents our estimates of the parameters characterising the distribution of actions of level-0 players and the beliefs of higher level players regarding their play.

We estimate the mean of the level-0 action distribution to be 58.38 and the mean of the level-0 belief distribution to be 54.01. These are close to the estimates we obtained non-

[26]However, we want to emphasise, that we indeed find substantial heterogeneity in the stated level-0 beliefs in the classification procedure.

Table 4: Estimated level-0 actions and beliefs

Parameter	μ^0	σ^0	μ^b	σ^b
Estimate	58.38	19.73	54.01	16.28
	(7.09)	(3.45)	(2.49)	(2.41)

Notes: The table presents the results from a maximum likelihood estimation of the structural model as outlined in section 6.1. This table only presents the results for the level-0 action and belief distribution, but the level-k distribution was estimated simultaneously. Those results are reported in table 3. Bootstrapped standard errors are given in brackets. These are obtained from 200 iterations of our estimation when sampling 84 observations from our data.

parametrically in Section 5. We estimate the level-0 action distribution to have a variance of 19.73 and the level-0 belief distribution to have a variance of 16.28. These are as well similar to the non-parametric estimates obtained before. Hence, we again find the belief distribution to be more concentrated than the action distribution, but – in contrast to the assumption made in the literature[27] – we estimate the belief distribution to have substantial variance.[28]

6.4 Quantifying the classification information

We can quantify the benefit of the additional information from our design by comparing the variance of our estimator with the variance of the equivalent estimator which does not use classification information.[29] Under suitable regularity conditions the maximum-likelihood estimator $\hat{\boldsymbol{\psi}}_{MLE}$ will have an asymptotic distribution with variance matrix $\boldsymbol{I}^{-1}$, where $\boldsymbol{I}$ is the Fisher information matrix. The finite sample distribution then has an approximate variance matrix $\boldsymbol{I}^{-1}/n$. Let $\boldsymbol{I_0}^{-1}$ denote the information matrix of the estimator using unclassified data and $\boldsymbol{I_c}^{-1}$ the information matrix of our estimator which uses the partly classified data. Then $\boldsymbol{I_0}^{-1}/n$ and $\boldsymbol{I_c}^{-1}/n$ are the approximate finite variance matrices of these estimators, respectively. We can calculate how many more unclassified data points one would need to achieve the same efficiency as our design. As the variances are matrices, one needs to choose some real-valued summary statistic of the matrix to calculate an exact ratio. It is a common procedure in the literature on optimal experimental design to compare the traces of the variance matrices, also known as the 'total sum of variances'. As a bottom line, to estimate μ^0, σ^0, μ^b and σ^b with the same sum of variances, one would need 13 times as many

[27]Recall that the literature assumes the level-0 action to be uniformly distributed and the level-0 belief to be degenerate at 50.

[28]When estimating the level-0 action distribution with a beta distribution that nests the uniform distribution, the estimated distribution is very similar in shape to the one found here.

[29]For references on the literature of optimal experimental design with mixture densities, check for example Hosmer and Dick (1977).

choice observations as we have classified observations. Furthermore to estimate the l_j's with the same precision one would need 17 times as many observations. In our case, this would correspond to more than 1000 and 1400 observations respectively.[30]

7 Concluding Comments

In this study we present a novel experimental design which allows to investigate the anchoring element of the level-k model, the level-0 actions and beliefs. In a standard 'beauty contest' game, we find that one third of the participants are playing non-strategically. The level-0 beliefs are mostly 50, but significant upward deviations occur. On average, the beliefs are consistent with the non-uniform level-0 actions, which are concentrated around points 50, 66 and 100. These numbers suggest that one plausible determinant for level-0 actions is salience.

The novel features of our design are that the participants state their individual reasoning, (i) in close temporal proximity to the reasoning process itself, and (ii) are supplied an incentive to state their reasoning as fully and clearly as possible.

The written accounts provide an immediate insight into the individual's reasoning in addition to the participant's action in the game. This relaxes the need to make assumptions on aspects of the reasoning in order to interpret the player's action. The level of reasoning, for example, can be determined without particular assumptions on the level-0 belief or population belief distributions. In addition, it is possible to directly observe level-0 play or level-0 beliefs without any preconception of their distributions.

Obtaining an incentivised written account of individual reasoning also relaxes the need to design complex games with a view to drawing inference from actions alone. It allows us to learn about individual reasoning in games which are economically interesting even when actions alone are not informative. The design is generally applicable to one shot games and we believe can prove useful for other purposes in experimental economics.

[30] Appendix A.5 gives details on the estimation and estimates of $\boldsymbol{I_0}^{-1}$ and $\boldsymbol{I_c}^{-1}$.

References

Agranov, M., A. Caplin, and C. Tergiman (2010): "The Process of Choice in Guessing Games," Working paper, California Institute of Technology.

Bacharach, M., and D. O. Stahl (2000): "Variable-Frame Level-n Theory," *Games and Economic Behavior*, 32(2), 220–246.

Berndt, E. K., B. H. Hall, R. E. Hall, and J. A. Hausman (1974): "Estimation and Inference in Nonlinear Structural Models," *Annals of Economic and Social Measurement*, 3/4, 653–665.

Bosch-Domènech, A., J. G. Montalvo, R. Nagel, and A. Satorra (2002): "One, Two, (Three), Infinity, . . . : Newspaper and Lab Beauty-Contest Experiments," *American Economic Review*, 92(5), 1687–1701.

——— (2004): "Finite Mixture Analysis of Beauty-Contest Data from Multiple Samples," Economics Working Papers 737, Department of Economics and Business, Universitat Pompeu Fabra.

Camerer, C. (2003): *Behavioral Game Theory*. Princeton University Press.

Camerer, C., E. J. Johnson, S. Sen, and T. Rymon (1993): "Cognition and framing in sequential bargaining for gains and losses," in *Frontiers of Game Theory*, ed. by A. K. Ken Binmore, and P. Tani, pp. 27–48. MIT Press.

Camerer, C. F., T.-H. Ho, and J.-K. Chong (2004): "A Cognitive Hierarchy Model of Games," *The Quarterly Journal of Economics*, 119(3), 861–898.

Cooper, D. J., and J. H. Kagel (2005): "Are Two Heads Better than One? Team versus Individual Play in Signaling Games," *American Economic Review*, 95(3), 477–509.

Costa-Gomes, M. A., and V. P. Crawford (2006): "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," *American Economic Review*, 96(5), 1737–1768.

Costa-Gomes, M. A., V. P. Crawford, and B. Broseta (2001): "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica*, 69(5), 1193–1235.

Crawford, V. P., M. Costa-Gomes, and N. Iriberri (2010): "Strategic Thinking," Discussion paper, University of Oxford.

Crawford, V. P., U. Gneezy, and Y. Rottenstreich (2008): "The Power of Focal Points Is Limited: Even Minute Payoff Asymmetry May Yields Large Coordination Failures," *American Economic Review*, 98(4), 1443–1458.

Crawford, V. P., and N. Iriberri (2007a): "Fatal Attraction: Focality, Naïveté and Sophistication in Experimental 'Hide and Seek' Games," *American Economic Review*, 97(5), 1731–1750.

——— (2007b): "Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?," *Econometrica*, 75(6), 1721–1770.

Ericsson, K. (2002): "Towards a procedure for eliciting verbal expression of non-verbal experience without reactivity: interpreting the verbal overshadowing effect within the theoretical framework for protocol analysis," *Applied Cognitive Psychology*, 16(8), 981–987.

Fischbacher, U. (2007): "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 10(2), 171–178.

Hosmer, D. W., and N. P. Dick (1977): "Information and mixtures of two normal distributions," *Journal of Statistical Computation and Simulation*, 6(2), 137–148.

Lewis, D. K. (1969): *Convention: A Philosophical Study*. Harvard University Press.

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

Nagel, R. (1994): "Reasoning and Learning in Guessing Games and Ultimatum Games With Incomplete Information: An Experimental Investigation," Ph.D. thesis, Universität Bonn.

——— (1995): "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85(5), 1313–1326.

Penczynski, S. P. (2010a): "Focality, Tasks and the Beginning of Level-k Reasoning," Working paper, Cornell University,.

——— (2010b): "Persuasion and Strategic Thinking," Working paper, Cornell University.

Rubinstein, A. (2007): "Instinctive and Cognitive Reasoning: A Study of Response Times," *Economic Journal*, 117, 1243–1259.

Rubinstein, A., A. Tversky, and D. Heller (1996): “Naive Strategies in Zero-Sum Games,” in *Understanding Strategic Interaction–Essays in Honor of Reinhard Selten*, ed. by W. Albers, W. Güth, P. Hammerstein, B. Molduvano, and E. van Damme, pp. 394–402. Springer Verlag.

Rydval, O., A. Ortmann, and M. Ostatnicky (2009): “Three very simple games and what it takes to solve them,” *Journal of Economic Behavior & Organization*, 72(1), 589–601.

Schelling, T. C. (1960): *The Strategy of Conflict*. Harvard University Press.

Stahl, D. O., and P. W. Wilson (1995): “On Players’ Models of Other Players: Theory and Experimental Evidence,” *Games and Economic Behavior*, 10(1), 218–254.

A Appendix

A.1 Transformation of Bounded Normal

Lemma 1 *A random variable which is distributed as bounded normal on* $[0,100]$ *with parameters* (μ,σ) *will still be distributed as bounded normal when applying a multiplicative transformation with a factor* a*, with parameters* $(a\mu, a\sigma)$ *and support* $[0, a100]$.

Proof 1 *For a monotonic, differentiable transformation of a random variable* B*, say* $X = g(B)$*, the distribution of* X *is found as* $f_X(x) = f_B(g^{-1}(x)) \cdot |g^{-1\prime}(x)|$ *where the support needs to be suitably changed. In our case* $g(b) = p^k b$, $g^{-1}(x) = x/p^k$ *and* $g^{-1\prime}(x) = 1/p^k$

$$f(x;\bullet) = \frac{\frac{1}{p^k\sigma^b}\phi(x/p^k;\boldsymbol{\theta}^b)}{\Phi(100;\boldsymbol{\theta}^b) - \Phi(0;\boldsymbol{\theta}^b)} = \frac{\frac{1}{p^k\sigma^b}\phi(y;p^k\boldsymbol{\theta}^b)}{\Phi(p^k100;p^k\boldsymbol{\theta}^b) - \Phi(0;p^k\boldsymbol{\theta}^b)}$$

where the equality of the nominators follows by straight-forward manipulation of $\phi()$ *and the denominator follows from*

$$\Phi(T;\boldsymbol{\theta}^b) = \int_{-\infty}^{T} \frac{1}{\sigma^b} \cdot \phi\left(\frac{x-\mu^b}{\sigma^b}\right) dz = \int_{-\infty}^{p^kT} \frac{1}{p^k\sigma^b} \cdot \phi\left(\frac{w - p^k\mu^b}{p^k\sigma^b}\right) dw = \Phi\left(p^kT; p^k\boldsymbol{\theta}^b\right)$$

after a change of variable $w = p^k z$.

A.2 Details of the numerical estimation

In order to find the maximum likelihood estimates, we numerically maximise the log-likelihood. We start out searching over a grid covering the the full range of parameter values.[31] Once the vector of parameter values is found which gives the highest likelihood, a new, finer grid is defined around this set. The new grid includes the parameter values that were neighbors of the maximising set in the previous, coarser grid. The calculation is iterated until the mesh size is sufficiently small.

A.3 Bootstrapping

The standard errors for the maximum likelihood estimates have been bootstrapped. Estimating 200 samples that consist of 84 draws with replacement from the original dataset, we obtain a measure of the estimates's variance. Figure 4 shows the histograms of the 200 estimates for the 8 estimated parameters.

[31] All sets of possible parameter values are bounded except for the standard deviation of the bounded normal distribution. Using a beta distribution allows us to cover distribution shapes that are equivalent to standard deviations up to infinity in the bounded normal case.

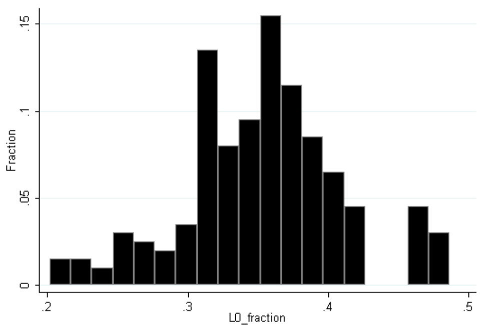

(a) l_0

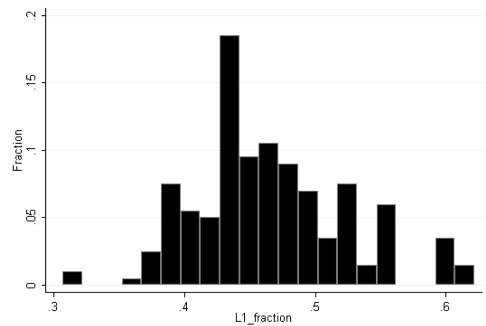

(b) l_1

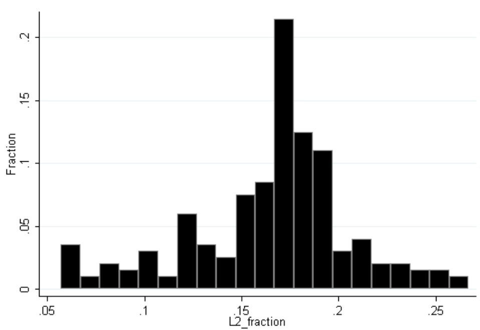

(c) l_2

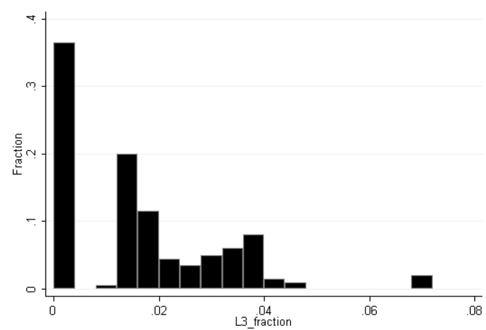

(d) l_3

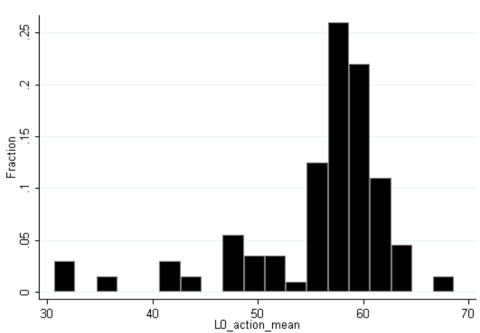

(e) μ^0

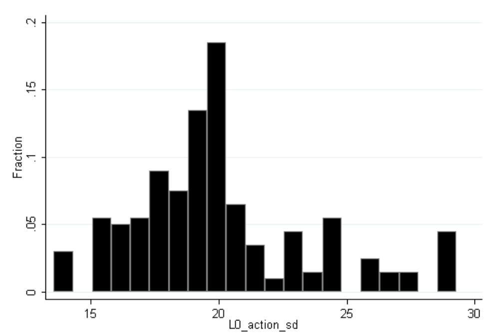

(f) σ^0

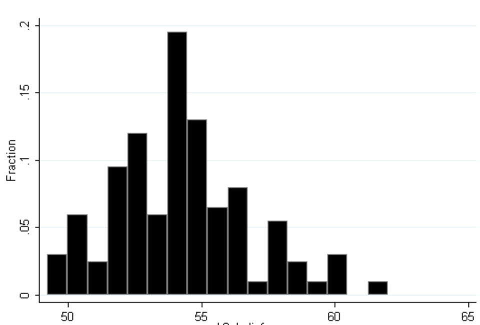

(g) μ^b

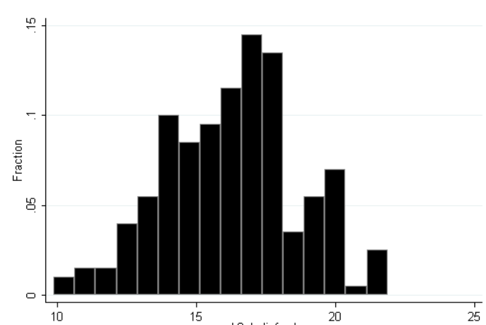

(h) σ^b

Figure 4: Bootstrapped distributions of MLE estimators

A.4 Monte Carlo studies

We use Monte Carlo studies in order to analyse the properties of our maximum likelihood estimators. We use the estimable model outlined in section 6.1 to generate data for 84 participants, reflecting the number of data points in the 'beauty contest' game. By design, the Monte Carlo results are specific to the parameters for which data is generated. We use the parameter estimates presented in section 6 as parameters in the data generating process.

From the generated levels of reasoning we generate level bounds. Firstly, there are participants whose communication does not lend itself to any classification. We reflect this by not giving bounds for 17% of the generated observations, making this value comparable to the 16.7% in the experiment. For the remaining data points we generate an upper bound which is one above the true level with probability 0.2 and similarly (and independently) for the lower bound. In our sample, 27.1% of the level classifications have an interval of two possible level values and 1.4% of three possible levels. The above data generating process gives comparable, but slightly less precise values of 32% and 4%, respectively.

Table 5 presents the results based on 50 runs. The average estimate from the MC studies is very close to the data generating parameter and the data generating parameter is always well within one standard deviation of estimates from the MC studies.

Table 5: MONTE CARLO RESULTS 'BEAUTY CONTEST'.

	Level				*Level-0 belief*		*Level-0 action*	
	0	*1*	*2*	*3*	*mean*	*st.dev.*	*mean*	*st.dev.*
true DGP	0.36	0.47	0.15	0.02	50.00	13.16	60.00	21.00
mean of MC estimates	0.39	0.46	0.14	0.01	50.13	13.07	56.45	21.17
st. dev. of MC estimates	0.05	0.06	0.04	0.01	1.95	1.68	5.45	4.80

A.5 Estimators of the information matrices

We estimate the Fisher information matrices with the outer-product of the score vector estimator introduced by Berndt, Hall, Hall, and Hausman (1974). Recall that $\boldsymbol{\psi} = (l_0, l_1, l_2, \mu^0, \sigma^0, \mu^b, \sigma^b)$ and let $\boldsymbol{z}_i$ be a vector of dimension 1×4 with element z_{ij} generally equal to 1 and set to 0 if the classification information is such that individual i is certainly not of level j. We outline below only the estimation of $\boldsymbol{I}_c$. The estimator of $\boldsymbol{I}_0$ is obtained similarly by setting all elements of $\boldsymbol{z}_i$ equal to 1. Define $\boldsymbol{v}_i(\boldsymbol{\psi}) \equiv \nabla_\psi p(x_i; \boldsymbol{\psi}, \boldsymbol{z}_i)$. Then the estimator of the information matrix $\boldsymbol{I}_c$ is $\sum_i \boldsymbol{v}_i(\boldsymbol{\psi})\boldsymbol{v}_i(\boldsymbol{\psi})'$. The score vector $\boldsymbol{v}_i(\boldsymbol{\psi})$ has elements

$$
\begin{aligned}
\frac{\partial p_i}{\partial l_0} &= \frac{1}{p(x_i;\boldsymbol{\psi},\boldsymbol{z}_i)}\left[z_{i0}f(x_i;\boldsymbol{\theta}^0) - z_{i3}f(x_i;p^3\boldsymbol{\theta}^b)\right] \\
\frac{\partial p_i}{\partial l_1} &= \frac{1}{p(x_i;\boldsymbol{\psi},\boldsymbol{z}_i)}\left[z_{i1}f(x_i;p^1\boldsymbol{\theta}^b) - z_{i3}f(x_i;p^3\boldsymbol{\theta}^b)\right] \\
\frac{\partial p_i}{\partial l_2} &= \frac{1}{p(x_i;\boldsymbol{\psi},\boldsymbol{z}_i)}\left[z_{i2}f(x_i;p^2\boldsymbol{\theta}^b) - z_{i3}f(x_i;p^3\boldsymbol{\theta}^b)\right] \\
\frac{\partial p_i}{\partial \mu^0} &= \frac{1}{p(x_i;\boldsymbol{\psi},\boldsymbol{z}_i)}\left[z_{i0}l_0\frac{\partial f(x_i;\boldsymbol{\theta}^0)}{\partial \mu^0}\right] \\
\frac{\partial p_i}{\partial \sigma^0} &= \frac{1}{p(x_i;\boldsymbol{\psi},\boldsymbol{z}_i)}\left[z_{i0}l_0\frac{\partial f(x_i;\boldsymbol{\theta}^0)}{\partial \sigma^0}\right] \\
\frac{\partial p_i}{\partial \mu^b} &= \frac{1}{p(x_i;\boldsymbol{\psi},\boldsymbol{z}_i)}\left[z_{i1}l_1\frac{\partial f(x_i;p^1\boldsymbol{\theta}^b)}{\partial \mu^b} + z_{i2}l_2\frac{\partial f(x_i;p^2\boldsymbol{\theta}^b)}{\partial \mu^b} + z_{i3}(1-l_0-l_1-l_2)\frac{\partial f(x_i;p^3\boldsymbol{\theta}^b)}{\partial \mu^b}\right] \\
\frac{\partial p_i}{\partial \sigma^b} &= \frac{1}{p(x_i;\boldsymbol{\psi},\boldsymbol{z}_i)}\left[z_{i1}l_1\frac{\partial f(x_i;p^1\boldsymbol{\theta}^b)}{\partial \sigma^b} + z_{i2}l_2\frac{\partial f(x_i;p^2\boldsymbol{\theta}^b)}{\partial \sigma^b} + z_{i3}(1-l_0-l_1-l_2)\frac{\partial f(x_i;p^3\boldsymbol{\theta}^b)}{\partial \sigma^b}\right].
\end{aligned}
$$

We find

$$
\begin{aligned}
\frac{\partial f(x;a\boldsymbol{\theta})}{\partial \mu} &= \frac{\frac{1}{(a\sigma)}\phi(x,a\boldsymbol{\theta})\left[\frac{x-a\mu}{a\sigma^2}[\Phi(a100;a\boldsymbol{\theta})-\Phi(0;a\boldsymbol{\theta})] - \frac{1}{\sigma}[\phi(0;a\boldsymbol{\theta})-\phi(a100,a\boldsymbol{\theta})]\right]}{[\Phi(a100;a\boldsymbol{\theta})-\Phi(0;a\boldsymbol{\theta})]^2} \\
&= f(x;a\theta)\frac{1}{\sigma}\left[\frac{x-a\mu}{a\sigma} + a\sigma(f(a100;a\theta)-f(0;a\theta))\right]
\end{aligned}
$$

where the second step follows by the definition of $f(x;a\theta)$. Secondly we find

$$
\frac{\partial f(x;a\boldsymbol{\theta})}{\partial \sigma} = \frac{\frac{\partial \frac{1}{a\sigma}\phi(x,a\theta)}{\partial \sigma}}{[\Phi(a100;a\theta)-\Phi(0;a\theta)]} - \frac{\frac{1}{a\sigma}\phi(x,a\theta)\frac{\partial[\Xi_0]}{\partial \sigma}}{[\Phi(a100;a\theta)-\Phi(0;a\theta)]^2}
$$

where $\Xi_0 = \Phi(a100;a\theta) - \Phi(0;a\theta)$. We can calculate

$$
\frac{\partial \frac{1}{a\sigma}\phi(x,a\theta)}{\partial \sigma} = \frac{1}{\sigma}\frac{1}{a\sigma}\phi(x;a\theta)\left[\frac{(x-a\mu)^2}{a^2\sigma^2} - 1\right]
$$

and

$$
\begin{aligned}
\frac{\partial[\Phi(a100;a\theta)-\Phi(0;a\theta)]}{\partial\sigma} &= \frac{\partial\int_0^{a100}\frac{1}{a\sigma}\phi(x;a\theta)\mathrm{dx}}{\partial\sigma} \\
&= \int_0^{a100}\frac{\partial\frac{1}{a\sigma}\phi(x;a\theta)}{\partial\sigma}\mathrm{dx} \\
&= \int_0^{a100}\frac{1}{\sigma}\frac{1}{a\sigma}\phi(x;a\theta)\left[\frac{(x-a\mu)^2}{a^2\sigma^2}-1\right]\mathrm{dx} \\
&= \frac{1}{\sigma}\int_0^{a100}\frac{(x-a\mu)^2}{(a\sigma)^2}\frac{1}{a\sigma}\phi(x;a\theta)\mathrm{dx}-\frac{1}{\sigma}[\Phi(a100;a\theta)-\Phi(0;a\theta)]
\end{aligned}
$$

Divide by $\Phi(a100;a\theta)-\Phi(0;a\theta)$ to find

$$
\begin{aligned}
&\frac{1}{\sigma^3a^2}\int_0^{a100}(x^2-2a\mu x+a^2\mu^2)f(x;a\theta)\mathrm{dx}-\frac{1}{\sigma} \\
=\ &\frac{1}{\sigma^3a^2}\underbrace{\int_0^{a100}x^2f(x;a\theta)\mathrm{dx}}_{\Xi_2}-\frac{2a\mu}{\sigma^3a^2}\underbrace{\int_0^{a100}xf(x;a\theta)\mathrm{dx}}_{\Xi_1}+\frac{a^2\mu^2}{\sigma^3a^2}\underbrace{\int_0^{a100}f(x;a\theta)\mathrm{dx}}_{=1}-\frac{1}{\sigma}
\end{aligned}
$$

where Ξ_1 and Ξ_2 are the first and second central moments of the bounded normal. From the m.g.f. of the bounded normal they are found as

$$
\begin{aligned}
\Xi_1 &= a\mu-a\sigma\frac{\phi(a100;a\theta)-\phi(0;a\theta)}{\Phi(a100;a\theta)-\Phi(0;a\theta)} \\
\Xi_2 &= a^2\mu^2+a^2\sigma^2+a^2\sigma^2\frac{\phi'(a100;a\theta)-\phi'(0;a\theta)}{\Phi(a100;a\theta)-\Phi(0;a\theta)}-2a^2\mu\sigma\frac{\phi(a100;a\theta)-\phi(0;a\theta)}{\Phi(a100;a\theta)-\Phi(0,a\theta)}
\end{aligned}
$$

Note that $\phi'(a100;a\theta)=-\frac{a100-a\mu}{a\sigma}\phi(a100;a\theta)$ and $\phi'(0;a\theta)=-\frac{0-a\mu}{a\sigma}\phi(0;a\theta)$. Collecting terms we find

$$
\frac{\partial\left[\Phi(a100;a\theta)-\Phi(0;a\theta)\right]/\partial\sigma}{\Phi(a100;a\theta)-\Phi(0;a\theta)}=-\frac{1}{\sigma}\frac{\frac{a100-a\mu}{a\sigma}\phi(a100;a\theta)-\frac{0-a\mu}{a\sigma}\phi(0;a\theta)}{\Phi(a100;a\theta)-\Phi(0;a\theta)}
$$

Collecting terms and substituting in $f(x;a\theta)$ we can write

$$
\frac{\partial f(x;a\theta)}{\partial\sigma} = f(x;a\theta)\frac{1}{\sigma}\left[\left(\frac{x-a\mu}{a\sigma}\right)^2-1+(a100-a\mu)\cdot f(a100,a\theta)-(0-a\mu)\cdot f(0,a\theta)\right]
$$

A.6 Estimates of information matrices

Our estimate of the information matrix when not using information on the upper and lower bounds of reasoning is

$$\hat{\boldsymbol{I}}_0 = \begin{bmatrix} 0.428 & -0.595 & 0.049 & -29.234 & 17.102 & 6.466 & -12.630 \\ -0.595 & 0.896 & -0.127 & 41.296 & -25.289 & -11.914 & 17.787 \\ 0.049 & -0.127 & 0.064 & -3.800 & 3.095 & 3.129 & -1.504 \\ -29.234 & 41.296 & -3.800 & 2041.117 & -1184.958 & -469.305 & 869.482 \\ 17.102 & -25.289 & 3.095 & -1184.958 & 751.216 & 334.845 & -515.823 \\ 6.466 & -11.914 & 3.129 & -469.305 & 334.845 & 241.982 & -200.424 \\ -12.630 & 17.787 & -1.504 & 869.482 & -515.823 & -200.424 & 382.519 \end{bmatrix}$$

and our estimate of the information matrix when using this information is

$$\hat{\boldsymbol{I}}_c = \begin{bmatrix} 0.018 & -0.007 & -0.004 & -0.684 & 0.341 & -0.309 & -0.064 \\ -0.007 & 0.021 & -0.010 & 0.417 & -0.148 & -0.021 & 0.096 \\ -0.004 & -0.010 & 0.042 & 0.162 & -0.136 & 0.191 & 0.030 \\ -0.684 & 0.417 & 0.162 & 116.136 & -17.443 & 15.078 & 3.438 \\ 0.341 & -0.148 & -0.136 & -17.443 & 83.307 & 2.554 & 3.215 \\ -0.309 & -0.021 & 0.191 & 15.078 & 2.554 & 40.129 & 9.249 \\ -0.064 & 0.096 & 0.030 & 3.438 & 3.215 & 9.249 & 8.512 \end{bmatrix}.$$

A.7 Classification agreement

Table 6: Classification Agreement

Subject	*Level-0*	*Level-1*	*Level-2*	*Level-3*	*Coinciding*
1	8				8
2			8		8
3	8				8
4	8				8
5	8				8
6		8			8
7	8				8
8	8				8
9	8				8
10		8			8
11	8				8
12		8			8
13		8			8
14	8				8
15		8			8
16	8				8
17	8				8
18	8				8
19			8		8
20	8				8
21	8				8
22	8				8
23	8				8
24		8			8
25		8			8
26	8				8
27	8				8
28	8				8
29	8				8
30		8			8
31	8				8
32	8				8

Continued on next page

Table 6 – continued from previous page

Subject	*Level-0*	*Level-1*	*Level-2*	*Level-3*	*Coinciding*
33		8			8
34				8	8
35	8				8
36		8			8
37		8			8
38		8			8
39		8			8
40	8				8
41	1	7			7
42			7	1	7
43	1	7			7
44	7	1			7
45			7	1	7
46		7	1		7
47	7	1			7
48	7	1			7
49	7	1			7
50	7	1			7
51		1	7		7
52	1	7			7
53	1	7			7
54	7	1			7
55	1	7			7
56	6	2			6
57	2	6			6
58	6	2			6
59	6	2			6
60	6	2			6
61		6	2		6
62	6	2			6
63	6	2			6
64	6	2			6
65		2	6		6
66		6	2		6

Continued on next page

Table 6 – CONTINUED FROM PREVIOUS PAGE

Subject	*Level-0*	*Level-1*	*Level-2*	*Level-3*	*Coinciding*
67	2	6			6
68	6	1	1		6
69		1	6	1	6
70	1	6	1		6
71	3	5			5
72	1	5	2		5
73		1	5	2	5
74	4	4			4
75		4	4		4
76	4	4			4
77	4	2	2		4
78	1	2	4	1	4

Notes: The table presents the number of classifiers (out of 8 classifiers) who picked the indicated lower bound, by subject. The subjects are ordered by the maximum number of coinciding choices.