Centre for Behavioural and Experimental Social Science

University of East Anglia

2020-06

Depth of Reasoning Models with Sophisticated Agents.

Peter G. Moffatt
Ganna Pogrebna
Graciela Zevallos-Porles

www.uea.ac.uk/cbess

Depth of Reasoning Models with Sophisticated Agents

Peter G. Moffatt*
School of Economics
University of East Anglia

Ganna Pogrebna†
The Alan Turing Institute
The University of Sydney

Graciela Zevallos-Porles‡
School of Economics
University of East Anglia

July 23, 2020

Abstract

In the context of guessing games, we propose the Sophisticated Reasoning Model (SRM) which includes a "sophisticated" type. A parameter p_s represents the proportion of sophisticated players in the population. A sophisticated player is one who forms a belief ($\widetilde{p}_s$) of the proportion of the population who are sophisticated (following the same cognitive process as themselves) and best responds to this belief. The model nests the standard Level-k and cognitive hierarchy models (when $\widetilde{p}_s = 0$) and also Nash behaviour (when $\widetilde{p}_s = 1$). Moreover, a sophisticated player with correct beliefs ($\widetilde{p}_s = p_s$) has best response equal to the winning guess. The model is extended to allow heterogeneity in beliefs. When applied to field data from a guessing game, only 9% of players are estimated to be sophisticated, but these players greatly over-estimate the proportion who are of the same type. This is interpreted as a manifestation of the Dunning-Kruger effect.

Keywords: Beauty contest game; Sophisticated reasoning model; Level k-model; Cognitive hierarchy model; Dunning-Kruger effect.

*Corresponding author. E-mail addresses: p.moffatt@uea.ac.uk
†E-mail addresses: gpogrebna@turing.ac.uk
‡E-mail addresses: graciela.zevallos@uea.ac.uk

1 Introduction

Stahl and Wilson (1994) proposed a class of models which used depth of reasoning to explain strategic thinking and applied it to symmetric 3×3 games. They proposed that in strategic games a population of players may have a hierarchical structure, where *boundedly rational* players with more advanced levels of reasoning best-respond to their beliefs about behavior of players with less advanced levels of reasoning. Level-0 type (henceforth, L0) picks her action in a strategic game at random from the uniform distribution (also known as *uniform random* play); level-1 type (henceforth, L1) assumes that all players in the population are of L0 type and best-responds to L0's actions; and level-2 type (henceforth, L2) believes that all players in the population are of L1 type and best-responds to L1's actions. They also include a type whose behaviour corresponds to the Nash equilibrium prediction, sometimes referred to as the *(naïve) Nash* type.[1]

Stahl and Wilson (1995) extended their original (1994) model to include unboundedly rational types in addition to the *Nash* type. One is the *worldly* type, who best responds to the belief that the population consists of lower-level types and Nash types. Another is the *rational expectations (RE)* type, who best responds to the belief that the population consists of lower-level types, Nash types, *and other RE types* (i.e. they take account of the possibility that other players follow the same cognitive process as themselves). Although Stahl and Wilson (1995) find evidence of the existence of worldly types, they find no evidence of RE types in the context of symmetric 3×3 games.

In this paper, we reconsider the possibility of the existence of Stahl and Wilson's (1995) RE types, but in a different context. We will refer to these types as *sophisticated*.[2] Similarly to Stahl and Wilson (1995), our model includes both boundedly rational and unboundedly rational types. Boundedly rational types include L0 (uniform random), L1 (optimizing on L0), ..., and LK (optimizing on L(K-1)) types, where K represents the highest level among boundedly rational types. There will be only one unboundedly rational type: the sophisticated type; for reasons to be explained, any other unboundedly rational types will automatically be subsumed into the classification of "sophisticated".

The context in which we will conduct our analysis is the class of games known as *guessing* games. Nagel (1995) proposed to test behavior in strategic games against predictions of the game theory (Nash, 1951) using a simple guessing game which is also known as *ω-beauty contest* game. In this game, each player i of N players

[1] We will avoid the adjective *naïve* in describing the Nash type - we prefer to interpret their behaviour as sophisticated with incorrect beliefs.

[2] We prefer the label "sophisticated" to "rational", because the latter is often interpreted as implying Nash behaviour.

picks a number y_i between 0 and 100 and is told that the winner is the player whose guess is closest to $\omega \cdot \frac{1}{N}\sum_{i=1}^{N} y_i$ where $\omega \in [0,1]$.[3] The prize received by the winner is typically a money amount, M. If more than one player makes the winning guess, the prize is divided equally amongst these players.

The subgame perfect Nash equilibrium predicts that an unboundedly rational utility-maximizing agent should pick 0, and hence that all N players should do this, and share the prize M, each receiving a payoff of $\frac{M}{N}$. Yet, much empirical evidence (e.g., Nagel (1995); Stahl and Wilson (1995); Bosch-Domènech et al. (2010)) suggests that players in the guessing game fail to play in accordance with the Nash equilibrium; most pick numbers greater than zero.

Several models have been proposed to explain this experimental regularity. In this paper, we focus on the two most widely used models: the level-k model (e.g., Nagel (1995); Crawford and Iriberri (2007)) and the cognitive hierarchy model (e.g., Camerer et al. (2004)). Much in the spirit of Stahl and Wilson (1994, 1995), both of these models are built on the assumption that each player forms a belief about the levels of reasoning (and, therefore the strategies) of other players, and adopts a strategy which represents the best response to this belief.

When the level-k model is applied to the guessing game, a key assumption relates to the behaviour of L0. The current convention is to follow Stahl and Wilson (1995) and assume that L0 randomly chooses her action from the uniform $(0, 100)$ distribution, resulting in an average of $y_{L0} = 50$. From this assumption, it follows that L1 chooses $\omega \cdot y_{L0}$ or 50ω; L2 chooses $\omega^2 \cdot y_{L0}$ or $50\omega^2$, and so on.

According to the cognitive hierarchy model (CH, Camerer et al. (2004)), players are assumed to be distributed between levels according to a discrete probability distribution (usually the Poisson), and each player is assumed to compute their best response on the assumption that other players are distributed between levels lower than their own level according to a truncated Poisson distribution.

Econometrically, both the level-k and CH models turn out to be natural applications of the finite mixture approach. The "subject types" are simply different levels of reasoning. We also require a stochastic element that allows the observed decision to differ from the "best response" for a player of a given type. In developing mixture models of this type, we are following the approaches of Bosch-Domènech et al. (2010) and Runco (2013).

Despite their popularity, both level-k and CH models suffer from a serious drawback: both models are built on the assumption that each player (of higher level than L0) assumes all other players to have a lower level of reasoning than their own. Apart from being intuitively questionable, this assumption presents particular problems in the case of CH, because it results in a lower bound to the guess: guesses can never fall below this lower bound however high the level of reasoning. This

[3] In most experiments ω is set to $2/3$.

leads to a logical problem if the winning guess happens to fall below this lower bound. Our principal goal in this paper is to rectify this drawback by proposing a hybrid model which we will call the sophisticated reasoning model (SRM).

The sophisticated reasoning model (SRM), is also a finite mixture model, with one of the components of the mixture consisting of "sophisticated" players. Sophisticated players are defined as those who compute a best response taking into account the presence of low-level types, and also of other members of their own sophisticated type; that is, they allow for the possibility that a proportion of other players follow the same cognitive process as they do themselves. The SRM has a free parameter $\widetilde{p}_s$, representing sophisticated agents' beliefs concerning the proportion of other sophisticated agents in the population. The SRM nests the standard level-k model and the cognitive hierarchy model (when $\widetilde{p}_s = 0$) as well as the Nash equilibrium prediction (when $\widetilde{p}_s = 1$) as special cases. Furthermore, if a sophisticated agent's belief ($\widetilde{p}_s$) happens to coincide with the actual proportion of sophisticated agents in the population p_s, we classify this agent as *clairvoyant sophisticated*, since in this situation their best response is the winning response. The SRM therefore represents a more general approach than existing models. Moreover, the model includes various stochastic elements which allow observed behaviour to deviate from the best-responses predicted by the theoretical version of the model.

The purpose of this paper is twofold: (i) to understand the implications of introducing sophisticated players for modelling strategic behavior in games; and (ii) to test the sophisticated reasoning model against other depth of reasoning models. We will re-evaluate existing experimental evidence in the form of field data previously analysed by Bosch-Domènech et al. (2010). We will estimate the SRM on this data set and provide evidence, in the form of in-sample and out-of-sample predictive performance, that it can capture behavior better than level-k or cognitive hierarchy on their own.

The remainder of this paper is organized as follows. Section 2 introduces the sophisticated reasoning model and formulates predictions for a typical guessing game. Section 3 considers the problem of econometric estimation of the level-k, CH, and SRM models, deriving the likelihood functions for each. Section 4 reports on estimation of all three models using experimental data from Bosch-Domènech et al. (2010) and compares the models using a number of criteria, including in-sample and out-of-sample predictive performance. Finally Section 5 concludes.

2 The Sophisticated Reasoning Model

As explained in Section 1, our analysis is entirely in the context of the ω-beauty contest game. We will assume that L0 types choose randomly from a uniform (0,100) distribution, and that higher level types anchor on the belief that L0 types behave in this way.

The way in which we modify the level-k model is based on the following intuition. Consider the common case in which $\omega = 2/3$. If a player has succeeded in reaching the realisation that the best response for (say) a level-2 player is 22.2, and is considering the next step of the process, it is reasonable to suppose that such a player will detect the downward trajectory of best responses, and will be able to take the cognitive jump to the end of the process: the best response of zero. However, such a player may also realise that only a proportion of other players will reach this stage of the reasoning process, and hence that the winning guess will be somewhat higher than zero. Such a player will be labelled "sophisticated". Note that, essentially, we are removing types $3 \ldots \infty$ from the level-k model, and replacing them with this single type labelled "sophisticated". We could of course consider a different number of boundedly rational types, and the preferred number becomes an empirical question. For the moment, let us assume that the population of players is made up of L0, L1, L2,..., LK (where K is the maximum level of reasoning for boundedly rational types) and sophisticated types, in proportions $p_0, p_1, p_2, ..., p_K$, and p_s, respectively.

The central problem for a sophisticated player is to form a belief about the proportion of players in the population who are also sophisticated (i.e., of the same type as themselves). For the moment, let us assume that all sophisticated players have the same belief of $\widetilde{p}_s$. Although we are labeling these players as *sophisticated*, this belief may or may not be correct. That is, it may be that $\widetilde{p}_s \neq p_s$.

Although sophisticated players do not know the proportion p_s, we assume that they do know how the population divides among boundedly rational types (L0, L1, L2,..., LK). This means that while sophisticated players do not know the exact proportion of, for example, L0 players in the population (p_0), they do know the proportion of boundedly rational players who are L0, this proportion being $\frac{p_0}{1-p_s}$.

Based on these assumptions, the best guess of a sophisticated player (denoted y_s) is given implicitly by:

$$y_s = \omega \left[\widetilde{p}_s y_s + (1 - \widetilde{p}_s) \cdot \frac{p_0 y_{L0} + \omega p_1 y_{L0} + \omega^2 p_2 y_{L0} + ... + \omega^K p_K y_{L0}}{1 - p_s} \right]. \tag{1}$$

where y_{L0} is the expected guess for L0 (50 under the assumption of a uniform distribution for L0). Note that this best guess, y_s, appears on both sides of (1). Hence we need to find the fixed-point solution for y_s that produces equality in (1). This is achieved by rearranging (1) so as to obtain an explicit expression for y_s as a function of the belief $\widetilde{p}_s$:

$$y_s(\widetilde{p}_s) = \omega \left[\frac{1-\widetilde{p}_s}{1-\omega\widetilde{p}_s} \cdot \frac{p_0 y_{L0} + \omega p_1 y_{L0} + \omega^2 p_2 y_{L0} + ... + \omega^K p_K y_{L0}}{1-p_s} \right]. \tag{2}$$

It is useful to consider special cases of (2). First, consider the situation in which a sophisticated player assumes that *no* other sophisticated players exist, so that $\widetilde{p}_s = 0$. In this situation, the sophisticated player becomes essentially an L(K+1) player in a cognitive-hierarchy model; that is, a player who computes a best response on the assumption that all other players are distributed (with known probabilities) between L0, L1, L2,..., and LK. Next, consider the situation in which a sophisticated player assumes that *all* other players are sophisticated, so that $\widetilde{p}_s = 1$. In this situation, (2) yields a best response of zero - the Nash prediction. Hence our model is seen to include the special case of Nash equilibrium behaviour. Finally, consider the situation in which sophisticated players form *correct* beliefs about the proportion of sophisticated players in the population; that is, $\widetilde{p}_s = p_s$. In this situation, (2) can be rearranged to give:

$$y_s = \omega \left(p_s y_s + p_0 x_{L0} + \omega p_1 y_{L0} + \omega^2 p_2 y_{L0} + ... + \omega^K p_K y_{L0} \right). \tag{3}$$

Equation (3) is, of course, an expression for the winning guess, since it is ω multiplied by the *true* expected guess, with expectation taken over all $K+2$ types in the population. Thus we see that if sophisticated players correctly anticipate the proportion of sophisticated players, their best response will win the game, and they will share the prize. In this situation, it is reasonable to label sophisticated players as *clairvoyant sophisticated*.

A numerical example is useful at this point. Assume that $\omega = 2/3$ and $K = 2$, and the true type proportions in the population are as in Table 1 below.

Player Type	Description	Mixing Proportion Symbol	Mixing Proportion Value
L0	Level 0	p_0	0.10
L1	Level 1	p_1	0.30
L2	Level 2	p_2	0.20
S	Sophisticated	p_s	0.40
Total	All players		1.00

Table 1: Numerical example: The distribution of types in the population.

If the sophisticated players form beliefs correctly, so that $\widetilde{p}_s = 0.40$, their best response will be (from Equation (2)):

$$y_s = \frac{2}{3}\left[\frac{1-0.40}{1-\frac{2}{3}\times 0.40} \times \frac{50 \times 0.10 + 33.3 \times 0.30 + 22.2 \times 0.20}{1-0.40}\right] = \underline{17.66} \tag{4}$$

and they will win the game by implementing this best response.[4] However, if they seriously overestimate the proportion of sophisticated players in the population, so that, e.g., $\widetilde{p}_s = 0.80$, their best response will be (again from (2)):

$$y_s = \frac{2}{3}\left[\frac{1-0.80}{1-\frac{2}{3}\times 0.80} \times \frac{50 \times 0.10 + 33.3 \times 0.30 + 22.2 \times 0.20}{1-0.40}\right] = \underline{9.25} \tag{5}$$

Hence the over-estimation of p_s results in a guess that is too low to win the prize. Figure 1 shows the sophisticated best response plotted against their belief of the proportion of sophisticated players, for the example just used.

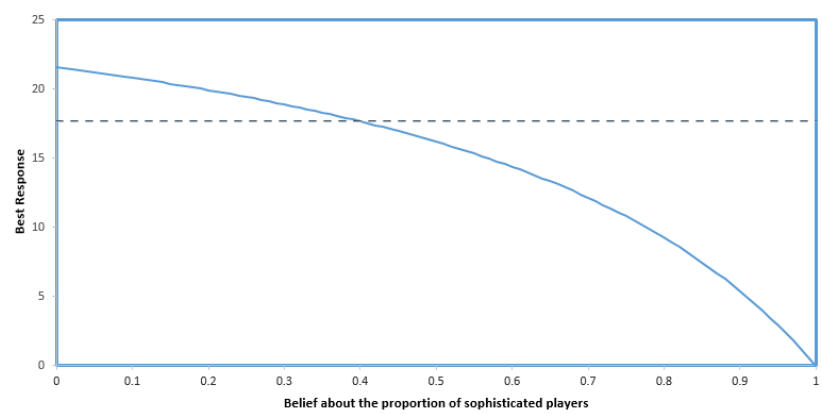

Figure 1: Solid curve: sophisticated players' best response $y_s(\widetilde{p}_s)$ against belief about the proportion of sophisticated players $\widetilde{p}_s$, assuming true type proportions shown in Table 1. Dashed line: winning guess (17.66).

[4]To verify that 17.66 is the winning guess, we can compute the average guess, based on this winning guess together with the proportions in Table 1, as $0.10 \times 50 + 0.30 \times 33.3 + 0.20 \times 22.2 + 0.40 \times 17.66 = 26.49$. Taking two-thirds of this average brings us back to the winning guess of 17.66.

3 Econometric Modelling

The first model we estimate is the standard level-k model. We follow a mixture modelling approach similar to that used by Bosch-Domènech et al. (2010) and Runco (2013). We assume that there are a finite number $(J+1)$ of types, and the "Nash" players being type J. Apart from Level-0 reasoners, who are assumed to choose from a uniform distribution, we assume that an individual's choice is the best guess for an individual of their type, plus a random normally distributed error with mean zero. That is, we assume that if y_j^* is the best guess for type j, then the actual guess (y) will be determined by:

$$(y|\text{type j}) = y_j^* + \epsilon_j, \qquad \epsilon \sim N\left(0, \sigma_j^2\right) \quad j = 1, ..., J. \tag{6}$$

Note that it is assumed that the variance of guesses around the best guess is assumed to differ between types. These assumptions give us the conditional density functions for each type:

$$\begin{aligned} f\left(y|L_0\right) &= 1/100, & 0 \leqslant y \leqslant 100 \\ f\left(y|L_j\right) &= \frac{1}{\sigma_j}\phi\left(\frac{y-y_j^*}{\sigma_j}\right) I_{y>0} + \Phi\left(-\frac{y_j^*}{\sigma_j}\right) I_{y=0} & 0 \leqslant y \leqslant 100 \quad j = 1, ..., J. \end{aligned} \tag{7}$$

where $I_{(.)}$ is the indicator function. Note that we are assuming censoring at zero, since zero is the lowest permissible guess.[5]

We also assume that the population is made up of the $J+1$ types with mixing proportions $p_0, p_1, \cdots, p_J$. Combining the mixing proportions with the conditional densities (7) gives us the sample log-likelihood (for a sample of guesses y_i, $i = 1, ..., n$):

$$\begin{aligned} &\text{LogL}\left(p_0, p_1, \ldots, p_J, \sigma_1, \ldots, \sigma_J\right) = \\ &\qquad \sum_{i=1}^{n} \ln\left[\frac{p_0}{100} + \sum_{j=1}^{J} p_j \left[\frac{1}{\sigma_j}\phi\left(\frac{y_i - y_j^*}{\sigma_j}\right) I_{y_i>0} + \Phi\left(-\frac{y_j^*}{\sigma_j}\right) I_{y_i=0}\right]\right]. \end{aligned} \tag{8}$$

We set $J = 5$, and the "best guesses" are $y_1^* = 33.3$, $y_2^* = 22.2$, $y_3^* = 14.8$, $y_4^* = 9.9$, $y_5^* = 0$. Note that the best guess for level 5 is zero, because, as noted above, we assign level J to "Nash" players.

Having estimated the standard level-k model by maximising (8), we will

[5] Strictly speaking, we should also censor at the upper limit of 100. However, the best guesses are all so far from 100 that this makes no difference to the results.

generate the predicted density function of guesses. This is given by:

$$\hat{f}\left(y;\hat{p}_0,\hat{p}_1,\ldots,\hat{p}_J,\hat{\sigma}_1,\ldots,\hat{\sigma}_J\right)= \frac{\hat{p}_0}{100}+\sum_{j=1}^{J}\hat{p}_j\left[\frac{1}{\hat{\sigma}_j}\phi\left(\frac{y-y_j^*}{\hat{\sigma}_j}\right)I_{y>0}+\Phi\left(-\frac{y_j^*}{\hat{\sigma}_j}\right)I_{y=0}\right],\qquad y\in[0,100]\,. \tag{9}$$

The second post-estimation task we conduct is to compute the posterior type probabilities using Bayes' Rule. The posterior probability of a subject being of each type $k=0,\ldots,J$ is given by:

$$\hat{P}\left(\text{Type }0|y_i\right)=\frac{\frac{\hat{p}_0}{100}}{\frac{\hat{p}_0}{100}+\sum_{j=1}^{J}\hat{p}_j\left[\frac{1}{\hat{\sigma}_j}\phi\left(\frac{y_i-y_j^*}{\hat{\sigma}_j}\right)I_{y_i>0}+\Phi\left(-\frac{y_j^*}{\hat{\sigma}_j}\right)I_{y_i=0}\right]}$$

$$\hat{P}\left(\text{Type k}|y_i\right)=\frac{\hat{p}_k\left[\frac{1}{\hat{\sigma}_k}\phi\left(\frac{y_i-y_k^*}{\hat{\sigma}_k}\right)I_{y_i>0}+\Phi\left(-\frac{y_k^*}{\hat{\sigma}_k}\right)I_{y_i=0}\right]}{\frac{\hat{p}_0}{100}+\sum_{j=1}^{J}\hat{p}_j\left[\frac{1}{\hat{\sigma}_j}\phi\left(\frac{y_i-y_j^*}{\hat{\sigma}_j}\right)I_{y_i>0}+\Phi\left(-\frac{y_j^*}{\hat{\sigma}_j}\right)I_{y_i=0}\right]} \tag{10}$$

$$k=1,\ldots,J.$$

The second model we estimate is the cognitive hierarchy (CH) model, developed by Camerer et al. (2004). This model assumes that the distribution of the population over reasoning levels is Poisson (τ), and that a subject of type k believes other members of the population to be distributed between types $0,\ldots,(k-1)$ according to an upper truncated Poisson (τ) distribution. That is, the (true) distribution of types over the population is given by:

$$p\left(j\right)=\left(\frac{e^{-\tau}\tau^j}{j!}\right),\qquad j=0,1,2,\ldots \tag{11}$$

and a subject of type k believes that the probability distribution of other subjects between types is:

$$p_k\left(j\right)=\left(\frac{e^{-\tau}\tau^j}{j!}\right)\Bigg/\left(\sum_{m=0}^{k-1}\frac{e^{-\tau}\tau^m}{m!}\right),\qquad j=0,\ldots,k-1. \tag{12}$$

On the basis of the probabilities (12), the "best guesses" for each type may be computed recursively as:

$$
\begin{array}{rcl}
y_1^* & = & 0.67 \times 0.5 = 0.33 \\
y_2^* & = & 0.67[p_2(1) \times y_1^* + p_2(0) \times 0.5] \\
y_3^* & = & 0.67[p_3(2) \times y_2^* + p_3(1) \times y_1^* + p_3(0) \times 0.5] \\
y_4^* & = & 0.67[p_4(3) \times y_3^* + p_4(2) \times y_2^* + p_4(1) \times y_1^* + p_4(0) \times 0.5] \\
y_5^* & = & 0.
\end{array} \tag{13}
$$

Note that, for practical purposes, we assume reasoning levels up to 4. We also assume that type 5 is the "Nash" type whose "best guess" is zero.

The log-likelihood function is then constructed by combining the observed guesses (y) with the type probabilities $(p(j))$ and the best guesses, y_j^* given in (13):

$$
\begin{aligned}
&\text{LogL}\,(\lambda, \sigma_1, \ldots, \sigma_J) = \\
&\qquad \sum_{i=1}^{n} \ln\left[\frac{p(0)}{100} + \sum_{j=1}^{J} p(j)\left[\frac{1}{\sigma_j}\phi\left(\frac{y_i - y_j^*}{\sigma_j}\right) I_{y_i>0} + \Phi\left(-\frac{y_j^*}{\sigma_j}\right) I_{y_i=0}\right]\right].
\end{aligned} \tag{14}
$$

As with the standard level-k model, we assume that zero guesses are censored zeros. Having estimated the CH model, the predicted density and posterior probabilities may be obtained using formulae similar to (9) and (10) respectively.

The third model that we estimate, and the focus of the paper, is the sophisticated reasoning model (SRM) described in Section 2. Recall that this model is an extension of the standard level-k model that includes a "sophisticated" type with mixing proportion p_s, and assumes that sophisticated players have belief $\widetilde{p}_s$ of the proportion of sophisticated players in the population. The first version of this model that we estimate is the homogeneous version, in which it is assumed that all sophisticated players have the same belief $\widetilde{p}_s$, so that $\widetilde{p}_s$ is a single parameter to be estimated. The best guess for this representative sophisticated player is given as $y_s(\widetilde{p}_s)$ in (2) above. As with other types, we assume that sophisticated players' guesses are normally distributed around this best guess with standard deviation σ_s and we allow for censoring at zero. The log-likelihood for this model is:

$$
\begin{aligned}
&\text{LogL}\,(p_0, p_1, \ldots, p_J, p_s, \sigma_1, \ldots, \sigma_J, \sigma_s, \widetilde{p}_s) = \\
&\qquad \sum_{i=1}^{n} \ln\Bigg(p_s\left[\frac{1}{\sigma_s}\phi\left(\frac{y_i - y_s(\widetilde{p}_s)}{\sigma_s}\right) I_{y_i>0} + \Phi\left(-\frac{y_s(\widetilde{p}_s)}{\sigma_s}\right) I_{y_i=0}\right] \\
&\quad + (1-p_s)\left[\frac{p_0}{100} + \sum_{j=1}^{J} p_j\left[\frac{1}{\sigma_j}\phi\left(\frac{y_i - y_j^*}{\sigma_j}\right) I_{y_i>0} + \Phi\left(-\frac{y_j^*}{\sigma_j}\right) I_{y_i=0}\right]\right]\Bigg).
\end{aligned} \tag{15}
$$

A superior model would be one that assumes that sophisticated agents each have their own belief of the proportion of sophisticated agents. Hence $\widetilde{p}_s$ becomes a random variable. Given the requirement that $\widetilde{p}_s \in [0,1]$, a natural distribution to assume is the beta distribution: $\widetilde{p}_s \sim B(\alpha, \beta)$. With this assumption, the log-likelihood becomes:

$$
\begin{aligned}
&\mathrm{LogL}\left(p_0, p_1, \ldots, p_J, p_s, \sigma_1, \ldots, \sigma_J, \sigma_s, \alpha, \beta\right) = \\
&\sum_{i=1}^{n} \ln\left(p_s \int_0^1 \left[\frac{1}{\sigma_s}\phi\left(\frac{y_i - y_s\left(\widetilde{p}_s\right)}{\sigma_s}\right) I_{y_i>0} + \Phi\left(-\frac{y_s\left(\widetilde{p}_s\right)}{\sigma_s}\right) I_{y_i=0}\right] f\left(\widetilde{p}_s; \alpha, \beta\right) d\widetilde{p}_s \right. \\
&\quad \left. + \left(1 - p_s\right)\left[\frac{p_0}{100} + \sum_{j=1}^{J} p_j \left[\frac{1}{\sigma_j}\phi\left(\frac{y_i - y_j^*}{\sigma_j}\right) I_{y_i>0} + \Phi\left(-\frac{y_j^*}{\sigma_j}\right) I_{y_i=0}\right]\right]\right). \qquad (16)
\end{aligned}
$$

where $f(\widetilde{p}_s; \alpha, \beta)$ is the $B(\alpha, \beta)$ density function evaluated at $\widetilde{p}_s$. Note that the scalar parameter $\widetilde{p}_s$ estimated using (15) is being replaced by the two parameters α and β in (16). We will refer to (16) as the heterogeneous sophisticated reasoning model. Estimation is performed using the method of maximum simulated likelihood (MSL), with the integral in (16) being evaluated using Halton draws (see Train (2009)).

4 Data and Results

We will use the same data set as used by Bosch-Domènech et al. (2010). The data consists of 7,892 guesses from three different 2/3 guessing games played by newspaper readers.[6] The distribution of guesses is shown in Figure 2. We see that the entire range of possible guesses is populated, although there is a clear skew towards the lower end of the distribution. We also see modes consistent with the predictions of the level-k model, at 33, 22, 15, and zero. None of our models attempt to explain the (less prominent) modes at 50, 66 and 100.

Four different models have been estimated: (i) level-k; (ii) cognitive hierarchy; (iii) sophisticated reasoning model with homogeneity; and (iv) sophisticated reasoning model with heterogeneity. As mentioned earlier, the appropriate number of types to assume in these models is an empirical question. We address this by estimating each model with different numbers of types, and then comparing model performance in various ways.

[6] For further details concerning the data, see Bosch-Domènech et al. (2010).

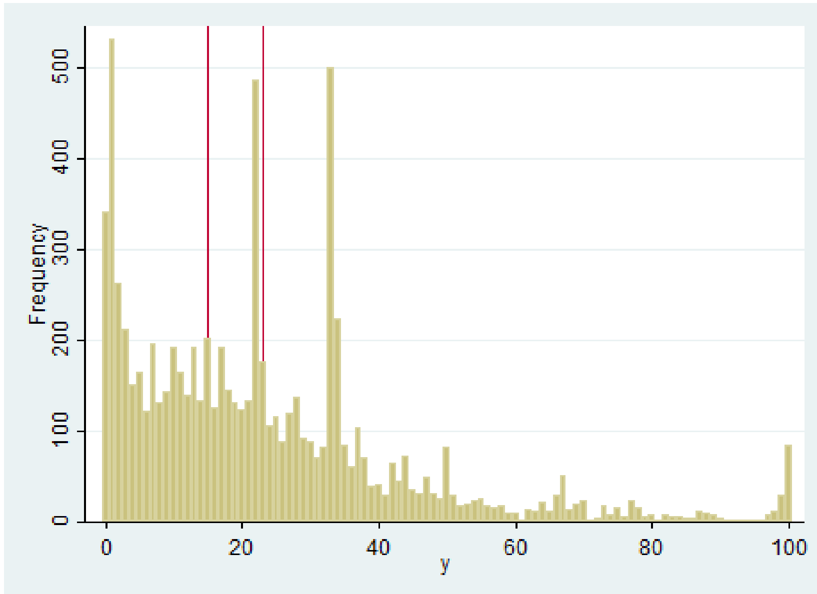

Figure 2: The distribution of guesses in the 2/3 guessing game from three different newspaper experiments, pooled. Sample size: 7,892. Vertical lines shown at mean (23.08) and winning guess (15.39). Source: Bosch-Domènech et al. (2010)

The results from estimation of the level-k model are shown in Table 2. Estimation has been performed by maximising the log-likelihood function specified in (8). The model has been estimated for J=3,4,5. Recall that type J is assumed to be the Nash type (with best guess equal to zero) and J-1 is the maximum reasoning level assumed apart from the Nash type. We see that according to the AIC, the model with J=5 is the best-fitting model. Estimates from this model tell us that level-2 is a dominant type, with mixing proportion 0.44, but this type also displays a large spread around the best guess, with standard deviation estimated to be 12.536. We also see that 10.3% of the population are estimated to be the "Nash" type.

A very useful way of presenting the results of mixture models is the posterior probability plot. The posterior probabilities for the level-k model are obtained using (10) above. For the model with J=5, the posterior probabilities are plotted against the guess in Figure 3. The dashed curve is the posterior probability of being L0, we see that, if a player's guess is above 60, the probability of being L0 is, as expected, very close to 1. The solid curves are the posterior probabilities for the other levels of reasoning. We see in particular that players with guesses close to zero have a reasonably high probability of being Nash-types.

Parameter	Level-K (J=3)	Level-K (J=4)	Level-K (J=5)
p_0	0.168(0.007)	0.172(0.008)	0.173(0.007)
p_1	0.062(0.003)	0.073(0.003)	0.070(0.003)
p_2	0.634(0.010)	0.356(0.018)	0.441(0.021)
p_3		0.291(0.018)	0.061(0.039)
p_4			0.151(0.026)
p_{Nash}	0.136(0.006)	0.109(0.006)	0.103(0.006)
σ_1	0.459(0.018)	0.482(0.020)	0.475(0.020)
σ_2	12.457(0.199)	13.978(0.352)	12.536(0.367)
σ_3		8.379(0.267)	7.994(0.610)
σ_4			6.299(0.322)
σ_{Nash}	2.452(0.132)	1.986(0.102)	1.861(0.096)
N	7,892	7,892	7,892
k	6	8	10
LogL	-32,424.59	-32,242.84	-32,217.31
AIC	64,861.18	64,501.68	64,454.62
In-sample KS (critical 0.015)	0.097	0.047	0.040
Out-of-sample KS (critical 0.023)	0.093	0.041	0.034

Table 2: Maximum likelihood estimates of level-k model with J=3, 4, 5. Asymptotic standard errors in parentheses. Type J is the Nash type. Final two rows contain Kolmogorov-Smirnoff (KS) test statistics. These should be compared to the critical value shown in the leftmost column.

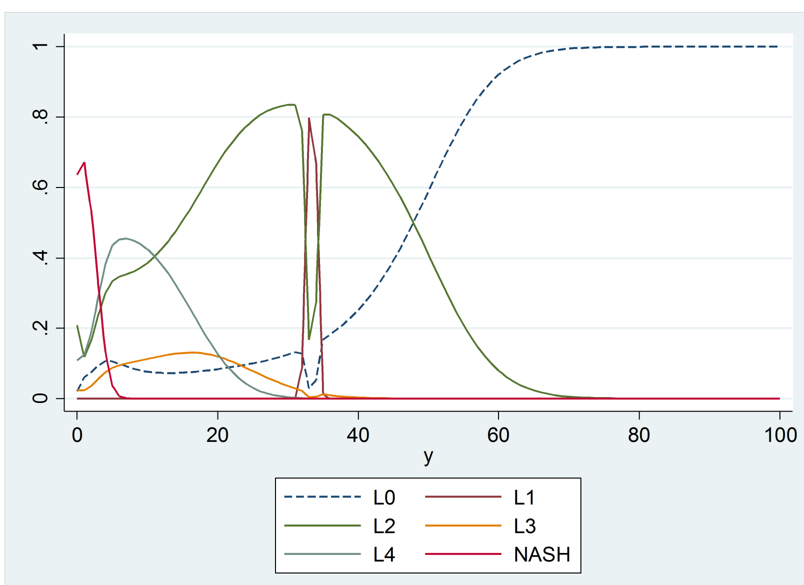

Figure 3: Posterior type probabilities from the level-k model with J=5.

As noted, the AIC indicates that the model with J=5 provides the best fit of the data. We also provide measures of predictive performance, both in-sample and out-of-sample. The measure we use for this purpose is the Kolmogorov-Smirnov (KS) test statistic. To understand how this test statistic is computed, it is useful to plot the fitted distribution functions from each model (obtained via (9)), and the (in-sample) empirical cumulative distribution, over the range of possible guesses. This plot is shown in Figure 4. We see that the fitted distribution from the model with J=5 is closest to the actual distribution (F_emp). The KS test statistic is computed as the maximum distance between the fitted and empirical distributions. These are the (in-sample) KS test statistics reported in Table 2. We see that the model with J=5 has the smallest (in-sample) KS statistic, confirming superior (in-sample) fit.

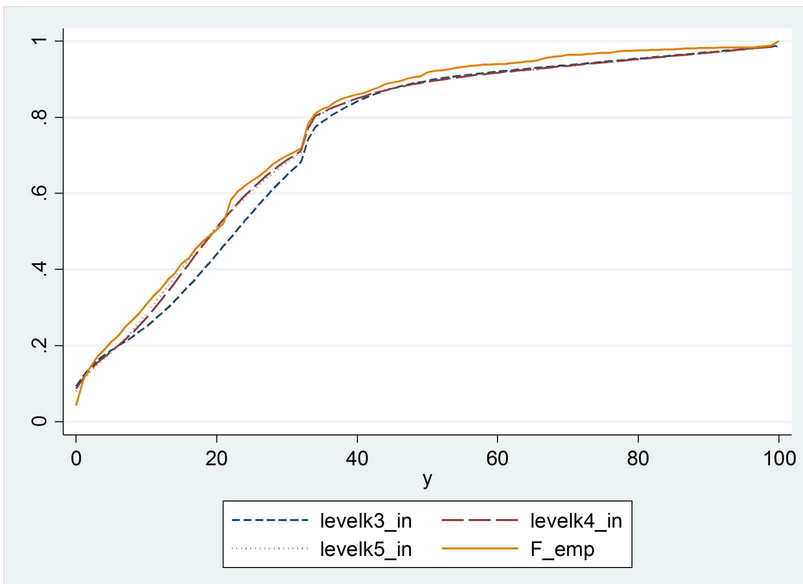

Figure 4: Fitted probability distributions from the level-k models compared to (in-sample) empirical cumulative distribution.

We also perform a test of out-of sample predictive performance. To do this, we randomly select exactly 50% of the original sample (resulting in a sample of 3,946 observations) and perform estimation on this reduced sample. We then compare the fitted distributions resulting from this estimation to the empirical distribution of the 50% of the sample not included in estimation, again using the KS-test. These are the out-of-sample KS tests shown in Table 2. Based on the out-of-sample test, the level-k model with J=5 again appears to have the best predictive performance of the three estimated level-k models.

The second model is the Cognitive Hierarchy (CH) model. Estimation has been performed by maximising the log-likelihood function specified in (14). The model is estimated for J=3,4,5, and the results are presented in Table 3. Here J-1 is the maximum reasoning level assumed, and type J is the "Nash" type. We see that the estimates of the Poisson mean (τ) are between 1.6 and 2.7, and this is in broad agreement with Camerer et al. (2004). On all three measures of model fit (AIC, in-sample KS, out-of-sample KS), the CH model with J=5 is seen to perform best. However, it must also be said that on all three measures the best of the CH models falls short of the performance of the best of the level-k models whose results appear in Table 2.

Parameter	CH (J=3)	CH (J=4)	CH (J=5)
τ	1.603(0.026)	2.225(0.024)	2.724(0.026)
σ_1	16.372(0.333)	18.103(0.451)	19.865(0.605)
σ_2	10.999(0.260)	12.057(0.287)	12.913(0.353)
σ_3	2.625(0.182)	10.232(0.266)	10.996(0.290)
σ_4		1.937(0.086)	9.903(0.317)
σ_5			1.490(0.061)
N	7,892	7,892	7,892
k	4	5	6
LogL	-33,960.94	-33,429.437	-33,196.01
AIC (=2k-2LogL)	67,929.88	66,868.88	66,404.02
In-sample KS (critical: 0.015)	0.257	0.187	0.131
Out-of-sample KS (critical: 0.023)	0.254	0.185	0.129

Table 3: Maximum likelihood estimates of cognitive hierarchy model with J = 3, 4, 5. Asymptotic standard errors in parentheses. Type J is the Nash type. Final two rows contain Kolmogorov-Smirnoff (KS) test statistics. These should be compared to the critical value shown in the left-most column.

The third model is the sophisticated reasoning model (SRM) with homogeneity. Estimation has been performed by maximising the log-likelihood function specified in (15). The model is estimated for J=2,3,4, and the results are presented in Table 4. Here J is the maximum reasoning level assumed for the boundedly rational types, but recall that this model includes a "sophisticated" type. On all three measures of model fit (AIC, in-sample KS, out-of-sample KS), the model with J=4 is seen to perform best. On the basis of the AIC, this model also performs better than the best of the level-k models (Table 2). However, in both in-sample and out-of-sample prediction, the performance of the SRM with homogeneity is seen to be inferior to that of the level-k model with J=5.

Considering the results from the SRM with homogeneity with J=4 (final column of Table 4) we see that the estimate of p_s is 0.110, indicating that 11% of the population are sophisticated. Furthermore, the estimate of $\widetilde{p_s}$ of 0.969 indicates that these sophisticated players seriously over-estimate the proportion of players who are sophisticated – they apparently assume that nearly all other players are sophisticated.

The quantity $\widetilde{p}_s - p_s$ reported in Table 4 (with standard error obtained using the delta method) amounts to a formal test of the correctness of sophisticated agents' beliefs, and its strong significance amounts to a strong rejection. This is clearly the model's attempt to account for the mode at zero; sophisticated players who believe that nearly all other players are sophisticated behave very similarly to Nash types. It is therefore not surprising that the estimates of p_s appearing in Table 4 are very close to the estimates of p_{Nash} appearing in Table 2.

Parameter	SRM-HOM (J=2)	SRM-HOM (J=3)	SRM-HOM (J=4)
p_0	0.144(0.007)	0.144(0.007)	0.145(0.007)
p_1	0.061(0.003)	0.062(0.003)	0.063(0.003)
p_2	0.689(0.009)	0.677(0.009)	0.654(0.011)
p_3		0.011(0.002)	0.011(0.002)
p_4			0.017(0.002)
p_s	0.105(0.006)	0.105(0.006)	0.110(0.007)
σ_1	0.457(0.018)	0.458(0.018)	0.460(0.018)
σ_2	13.911(0.200)	14.087(0.205)	14.034(0.242)
σ_3		0.190(0.015)	0.190(0.014)
σ_4			0.120(0.007)
σ_s	0.863(0.077)	0.874(0.080)	0.942(0.115)
$\widetilde{p}_s$	0.970(0.001)	0.970(0.001)	0.969(0.001)
$\widetilde{p}_s - p_s$	0.865**(0.006)	0.864**(0.006)	0.859**(0.008)
N	7,892	7,892	7,892
k	7	9	11
LogL	-32,150.13	-32,123.40	-32,024.84
AIC (=2k-2LogL)	64,314.26	64,264.80	64,071.68
In-sample KS (critical:0.015)	0.106	0.098	0.071
Out-of-sample KS (critical:0.023)	0.092	0.083	0.051

Table 4: Maximum likelihood estimates of sophisticated reasoning model (SRM) with homogeneity, and with J=2,3,4. Asymptotic standard errors in parentheses. Final two rows contain Kolmogorov-Smirnoff (KS) test statistics. These should be compared to the critical value shown in the leftmost column.

The final model we estimate is the SRM with heterogeneity. Estimation has been performed by maximising the log-likelihood function specified in (16). The model is estimated for J=2,3,4, and the results are presented in Table 5. Note that, instead of estimating a single value of $\widetilde{p}_s$ as in Table 4, we are now estimating the parameters of a beta distribution (α and β) representing the heterogeneity in sophisticated players' beliefs. We will interpret the estimates of these parameters shortly. We first note that on the basis of the three measures of model fit (AIC, in-sample KS, out-of-sample KS), the model with J=2 appears inferior to the other two, but there is very little to choose between the models with J=3 and J=4. This suggests that allowing for heterogeneity in the beliefs of sophisticated agents reduces the need to add boundedly rational types to the model. Secondly, note that on all three of the performance measures, the models with J=3 and J=4 outperform all of the other models previously estimated.

Parameter	SRM-HET (J=2)	SRM-HET (J=3)	SRM-HET (J=4)
p_0	0.167(0.007)	0.148(0.007)	0.147(0.007)
p_1	0.067(0.003)	0.073(0.003)	0.072(0.003)
p_2	0.529(0.013)	0.381(0.020)	0.396(0.021)
p_3		0.301(0.018)	0.265(0.026)
p_4			0.021(0.012)
p_s	0.237(0.009)	0.098(0.004)	0.098(0.004)
σ_1	0.467(0.019)	0.482(0.020)	0.481(0.020)
σ_2	12.061(0.272)	16.178(0.387)	15.927(0.398)
σ_3		8.722(0.321)	8.784(0.363)
σ_4			4.017(1.161)
σ_s	0.970(0.068)	0.089(0.022)	0.089(0.022)
$\widetilde{p}_s$			
α	3.754(0.425)	92.160(3.252)	92.100(3.284)
β	0.487(0.029)	3.485(0.136)	3.484(0.136)
$E(\widetilde{p}_s)$	0.885(0.008)	0.964(0.001)	0.964(0.001)
$E(\widetilde{p}_s) - p_s$	0.648**(0.015)	0.866**(0.004)	0.866**(0.004)
N	7,892	7,892	7,892
k	8	10	12
LogL	-32,063.20	-31,884.56	-31,882.266
AIC (=2k-2LogL)	64,142.4	63,789.12	63,788.532
In-sample KS (critical: 0.015)	0.059	0.028	0.028
Out-of-sample KS (critical: 0.023)	0.054	0.025	0.025

Table 5: Maximum likelihood estimates of sophisticated reasoning model (SRM) with heterogeneity, and with J=2,3,4. Asymptotic standard errors in parentheses.

Turning to the interpretation of parameters, we will focus on the SRM-HET model with J=4 (final column of Table 5). We see that the proportion of sophisticated players is here estimated to be only 9%. The distribution of sophisticated players' beliefs is the beta distribution defined by the estimates of α and β appearing in the final column of Table 5. This beta distribution is shown in Figure 5. Again we see that (estimated) beliefs are a very long way above the (estimated) proportion of sophisticated players in the population. Again we perform a test of correct beliefs, this time based on the quantity $E(\widetilde{p}_s) - p_s$, and again we reject strongly.

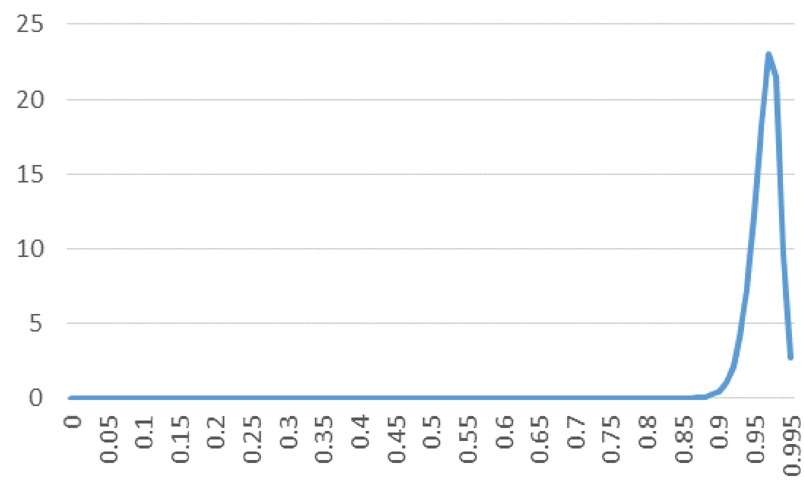

Figure 5: Beta distribution (for sophisticated players' beliefs) with α = 92.1 and β = 3.48, estimated from SRM-HET with J=4.

For the best performing model, SRM-HET with J=4, we plot the posterior type probabilities in Figure 6. We see that a player's guess needs to be very low (<10) for the model to classify them as sophisticated with reasonable probability.

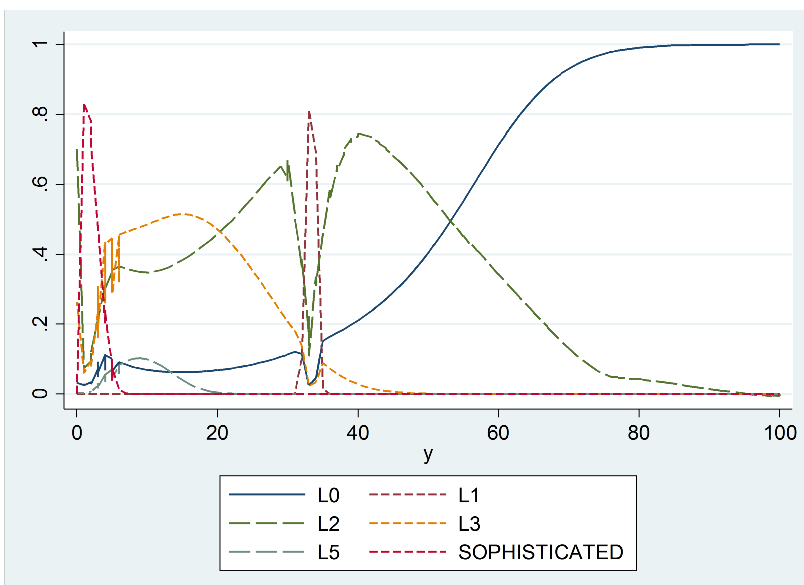

Figure 6: Posterior type probabilities from the SRM-HET with J=4.

It may seem incongruous that we have defined a class of "sophisticated" agents on the basis that they fully understand the behaviour of other players, but then find that they are so inaccurate in their beliefs about the proportion of players who are like themselves. However there is a behavioural anomaly known as the Dunning-Kruger effect which may partly explain this. This is the tendency for: unskilled individuals to *overestimate* their own ability; and experts to *underestimate* their own ability. Kruger and Dunning (1999, p.1127) have postulated that the effect is the result of internal illusion in those of low ability, and external misperception in those of high ability: "The miscalibration of the incompetent stems from an error about the self, whereas the miscalibration of the highly competent stems from an error about others."

It appears that the SRM incorporates both predictions of the Dunning-Kruger effect. Firstly, the prediction that unskilled individuals overestimate their own ability is consistent with the assumption, implicit in the level-k component of the model, that players of low reasoning levels assume that all other players are of a lower level than themselves.[7] Secondly, the prediction that experts underestimate their own ability (relative to others) is consistent with our finding that sophisticated players over-estimate the proportion of others who, like them, have a complete understanding of the situation.

5 Conclusion

Models involving depth of reasoning are widely used in the economics literature to explain human behavior in strategic games, such as "beauty contest" guessing games. We have applied a number of depth-of-reasoning models to an existing data set on guesses in such a game. The focus has been on testing for the presence of "sophisticated" players, and we have found that allowing for the existence of sophisticated players greatly improves explanatory power. Explanatory power again improves when sophisticated players are assumed to vary in their belief about the proportion of the population who are sophisticated. We have assumed a beta distribution for this distribution of beliefs. The estimation of this model reveals that sophisticated players tend to over-estimate the proportion of the population who are sophisticated, with average belief around 96%. This is significantly larger than the estimated proportion of sophisticated players which is only around 9%. This discrepancy has been explained as a manifestation of the Dunning-Kruger effect: highly able players underestimate their abilities relative to others (and therefore overestimate the abilities of others).

The sophisticated player whose belief is correct ($\widetilde{p}_s = 0.09$) guesses $y_s^*(0.09) \approx 17$

[7] Another phenomenon that would be consistent with the Level-k component is the self-selection of entrants to newspaper contests on the basis of over-confidence. See Camerer and Lovallo's (1999) discussion of "reference group neglect".

(in accordance with (2)). This guess is very close to the winning guess of 15 (see Figure 2). The reason why this player's guess is not exactly the winning guess is because she is assuming that all other sophisticated players have the same belief as her own. If, instead, she correctly recognised that the beliefs of sophisticated players tend to be biased to the extent seen in Figure 5, and built these biased beliefs into her calculation, then her best guess would come down to 15 and she would win the prize. Such a player would certainly deserve the accolade of "clairvoyant sophisticated". The capacity to predict the winning guess is, in our view, another attractive feature of the sophisticated reasoning model.

Acknowledgements

The authors are grateful to Nadine Chlaß for helpful comments.

References

Bosch-Domènech, A., J. G. Montalvo, R. Nagel, and A. Satorra, 2010, A finite mixture analysis of beauty-contest data using generalized beta distributions, *Experimental economics* 13, 461–475.

Camerer, C. and D. Lovallo, 1999, Overconfidence and excess entry: An experimental approach, *American Economic Review* 89, 306–318.

Camerer, C. F., T.-H. Ho, and J.-K. Chong, 2004, A cognitive hierarchy model of games, *The Quarterly Journal of Economics* , 861–898.

Crawford, V. P. and N. Iriberri, 2007, Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions?, *Econometrica* , 1721–1770.

Kruger, J. and D. Dunning, 1999, Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-assessments., *Journal of personality and social psychology* 77, 1121.

Nagel, R., 1995, Unraveling in guessing games: An experimental study, *American Economic Review* , 1313–1326.

Nash, J., 1951, Non-cooperative games, *Annals of mathematics* , 286–295.

Runco, M., 2013, Estimating depth of reasoning in a repeated guessing game with no feedback, *Experimental Economics* 16, 402–413.

Stahl, D. O. and P. W. Wilson, 1994, Experimental evidence on players' models of other players, *Journal of Economic Behavior & Organization* 25, 309–327.

Stahl, D. O. and P. W. Wilson, 1995, On players models of other players: Theory and experimental evidence, *Games and Economic Behavior* 10, 218–254.

Train, K. E., 2009, *Discrete choice methods with simulation* (Cambridge University Press).