A Window of Cognition: Eyetracking the Reasoning Process in Spatial Beauty Contest Games

Chun-Ting Chen, Chen-Ying Huang and Joseph Tao-yi Wang*

June 14, 2018

Abstract

We study the reasoning process in an environment where final choices are well understood and the associated theory is procedural by introducing two-person beauty contest games played spatially on two-dimensional grid maps. Players choose locations and are rewarded by hitting targets dependent on opponents' choice locations. By tracking subjects' eye movements (lookups), we infer their reasoning process and classify subjects into various levels. More than a half of the subjects' classifications coincides with their classifications using final choices, supporting a literal interpretation of the level-k model for subject's reasoning process. Lookup analyses reveal that the center area is where most subjects initially look at. This sheds light on the level-0 belief. Moreover, learning lookups of a trial on average could increase payoffs of that trial by roughly 60%, indicating how valuable lookups can help predict choices.

Keywords beauty contest game, level-k model, best response hierarchy, guessing game, cognitive hierarchy

JEL C91, C72, D87

*Department of Economics, National Taiwan University, 1 Roosevelt Road, Section 4, Taipei 106, Taiwan. Chen: chuntingchen@ntu.edu.tw; Huang: chenying@ntu.edu.tw; Wang: josephw@ntu.edu.tw (corresponding author). Acknowledgement: Research support was provided by the Ministry of Science and Technology of Taiwan (grant 96-2415-H-002-006). Joseph thanks the advice, guidance and support of Colin F. Camerer. We thank Ming Hsu for valuable suggestions that direct us to eyetracking. We thank comments from Vincent Crawford, Rosemarie Nagel, Matthew Shum, Yi-Ting Chen, Shih-Hsun Hsu, Ching-Kang Ing, Chung-Ming Kuan, and the audience of the ESA 2008 International Meeting and North American Region Meeting, TEA 2008 Annual Meeting, 2009 Stony Brook Workshop on Behavioral Game Theory, AEA 2010 Annual Meeting, 2010 KEEL conference, the 12th BDRM conference and the 2010 World Congress of the Econometric Society.

1 Introduction

Since Nash (1950) defined equilibrium as mutual best responses, game theory has been focusing on interpreting observed choices as the outcome of strategic best responses to consistent beliefs. This strategic approach has achieved tremendous success by simply assuming utility optimization and belief consistency. The focus on explaining final choices is widespread. Nonetheless, this emphasis on final choices should not exclude the possibility of analyzing the decision-making process prior to final decision if the process is observable and contains information about cognition, potentially hard to extract from observing choices alone.

In this paper, we study the reasoning process as well as final choices in a game-theoretic environment. We extend Costa-Gomes and Crawford (2006) and design the two-person spatial beauty contest game, which is a graphical simplification of the two-person guessing game played on two dimensional grid maps.[1] It is known that initial responses in the p-beauty contest games can be well explained by theories of heterogeneous levels of rationality such as the level-k model.[2] Since the level-k model predicts choices well, it is plausible to come up with satisfactory hypotheses on the reasoning processes.

Taking the level-k model procedurally, we propose a natural hypothesis regarding the reasoning process of the spatial beauty contest game. A key in the level-k model is that players of higher levels of rationality best respond to players of lower levels, who in turn best respond to players of even lower levels and so on. This best response hierarchy is the perfect candidate for procedurally modeling the reasoning process of a subject prior to making the final choice.[3] As an example, a level-2 subject would first focus on what a level-0 subject would choose since her opponent thinks of her as a level-0. She would next consider what a level-1 opponent would choose since her opponent would best respond to a level-0. Finally she would think about her level-2 choice since she would best respond to her level-1 opponent. Since the graphical representation may induce the subject to go through this hierarchical procedure of best responses by counting on the computer screen, we trace subject's eye-movements with video-based eyetracking. While the subject could go through the level-k reasoning process entirely in her mind, we hypothesize that she

[1]Nagel (1995) and Ho et al. (1998) studied the p-beauty contest game. Variants of two-person guessing games are studied by Costa-Gomes and Crawford (2006) and Grosskopf and Nagel (2008). In designing our games, we attempt to mimic the Bertrand competition in price and quality. Imagine firms compete by either cutting the price or increasing the quality of service. Suppose there is a common marginal cost and a best obtainable quality of service. If one firm typically competes by cutting the price while the other by increasing the quality of service, then the equilibrium is where price equals the marginal cost and the best service is obtained.

[2]Level-k models are developed by Stahl and Wilson (1995), Nagel (1995), Costa-Gomes and Crawford (2006), and Ho and Su (2013). Camerer et al. (2004) proposes the related cognitive hierarchy model.

[3]To avoid confusion, the subject is denoted by her while her opponent is denoted by him.

would count on the screen since this reduces memory load. Then, the locations being looked up reflect how a subject reasons before making her final choice. We eyetrack the entire sequence of every location a subject has ever fixated at in the experiment real-time. Following the convention, we call this real-time fixation data "lookups" even though there is no hidden payoff-relevant information to be looked up.

We estimate a level-k lookup model as follows. When a subject reasons through a particular best response hierarchy designated by her level-k type, each step of thinking is characterized as a "state." To describe changes between the thinking states, we construct a constrained Markov-switching model. We classify subjects into various level-k types based on maximum likelihood estimation using individual lookup data; Vuong's test is employed to ensure separation. Among the seventeen subjects we tracked, one follows the level-0 ($L0$) best response hierarchy the closest with her lookups, six follow the level-1 ($L1$) hierarchy, four follow the level-2 ($L2$) hierarchy, another four follow the level-3 ($L3$) hierarchy, and the remaining two follow the equilibrium (EQ) hierarchy, which coincides with level-4 ($L4$) hierarchy in most games. The average thinking step is 2.00 (if EQ is viewed as having 4 thinking steps), in line with results of other p-beauty contest games.

To see whether the lookup data indeed aligns well with choice data, we classify subjects by using their final choice data only by following Costa-Gomes and Crawford (2006). After all, if a subject reasons in her mind, one should not expect lookups being informative about final choices. On the other hand, if she counts on the screen to go through the best response hierarchy as we hypothesize, the estimated level based on lookups may coincide with that based on choices since the level reflects her strategic sophistication. We find that lookup-based and choice-based classifications are pretty consistent, classifying ten of the seventeen subjects as the same type. This suggests that if a subject's lookups are classified as a particular level-k type, her final choices follow the prediction of that level-k type as well, supporting a literal interpretation of the level-k model. This lends support to the level-k model as a procedural theory in addition to a theory of final choices.

Since the level-k model explains both final choices and lookups well for more than a half of our subjects, one might wonder what could lookup data tells us beyond which can be learned from final choices. To infer empirical level-0 beliefs, we use initial lookups to identify the starting point of reasoning as mostly the center. The top-left and the top-center, though less likely, are also possible. Since the center and the top-left (due to the reading habit in English) are salient, it hints on how salience may be important in determining level-0 belief. We also find that relying on lookups of a trial to predict the choice of that same trial is roughly as good as relying on choices of all other trials. Both can increase payoffs of about 60%, indicating how informative lookups are.

In the related literature, procedural data are used to infer the reasoning process and identify the level of subjects. Closest to our work are Costa-Gomes and Crawford (2006) and Costa-Gomes et al. (2001). Costa-Gomes and Crawford (2006) employs the mouse-tracking technology "mouselab " to study two-person guessing games and explicitly derive the procedural implication of level-k model by tracking how subjects click on payoff-relevant information. Burchardi and Penczynski (2014) and Penczynski (2016) use within-team text messages in the first trial of the p-beauty contest game to identify empirical level-0 beliefs and the number of steps a subject reasons. These approaches are complementary and confirm different aspects of the procedural process to be consistent with the level-k model.[4]

The remaining of the paper is structured as follows: Section 2.1 describes the spatial beauty contest game and its theoretical predictions; Section 2.2 describes details of the experiment; Section 3 reports aggregate statistics on lookups; Section 4 reports classification results based on lookups; Section 5 compares classification results based on lookups with those based on final choices. Section 6 investigates additional insights of lookups, and Section 7 concludes.

2 The Experiment

2.1 The Spatial Beauty Contest Game

We now introduce our design, the equilibrium prediction, the prediction by the level-k model and formulate the hypotheses which will be tested. To create a spatial version of the p-beauty contest game, we reduce the number of players to two, so that we can display the choices of all players on the computer screen visually. Players choose locations (instead of numbers) simultaneously on a 2-dimensional grid map attempting to hit one's target location determined by the opponent's choice. The target location is defined as a relative location to the other player's choice of location by a pair of coordinates (x, y). We use the standard Euclidean coordinate system. For instance, $(0, -2)$, means the target location of a player is "two steps below the opponent," and $(-4, 0)$ means the target location of a player is "four steps to the left of the opponent." These targets are common knowledge to the players. Payoffs are determined by how "far" (the sum of horizontal distance and vertical distance) a player is away from the target. The larger this distance is, the lower her payoff is. Players can only choose locations on a given grid map, though one's target may fall outside if the opponent is close to or on the boundary.[5] For example, consider

[4]Brocas, Carrillo, Wang and Camerer (2014) use mouselab to observe failure of acquiring payoff-relevant information necessary to find an equilibrium. Agranov, Caplin and Tergiman (2015) create an incentivized choice process protocol with random termination to identify level-0 behaviors.

[5]Similar designs could also be found in Kuo et al. (2009). They addressed different issues.

the 7×7 grid map in Figure I. For the purpose of illustration, suppose a player's opponent has chosen the center location labeled O ($(0,0)$) and the player's target is $(-4,0)$. Then to hit her target, she has to choose location $(-4,0)$. But location $(-4,0)$ is not on the grid map, while choosing location $(-3,0)$ is optimal among all 49 feasible choices because location $(-3,0)$ is the only feasible location that is one step from location $(-4,0)$.[6]

The spatial beauty contest game is essentially a spatial version of Costa-Gomes and Crawford (2006)'s asymmetric two-person guessing games, in which one subject would like to choose α of her opponent's choice and her opponent would like to choose β of her choice. Hence, similar to Costa-Gomes and Crawford (2006), the equilibrium prediction of this spatial beauty contest game is determined by the targets of both players. For example, if the targets of the two players are $(0,2)$ and $(4,0)$ respectively, the equilibrium consists of both players choosing the top-right corner of the grid map. This conceptually coincides with a player hitting the lower bound in the two-person guessing game of Costa-Gomes and Crawford (2006) if $\alpha\beta$ is less than 1, or all choosing zero in the p-beauty contest game where p is less than 1.[7] Note that in general the equilibrium needs not be at the corner since targets can have opposite signs. For example, when the targets are $(4,-2)$ and $(-2,4)$ played on a 7×7 grid map, the equilibrium locations for the two players are both two steps away from the corner (labeled as $\mathbf{E}_1$ and $\mathbf{E}_2$ for the two players respectively in Figure I).

Supplementary Appendix A1 derives the equilibrium predictions for general spatial beauty contest games, while Supplementary Appendix A2 develops the predictions of the level-k model and shows that all level-k types with k above a threshold level $\overline{k}$ coincide with equilibrium. In Table I we list all the 24 spatial beauty contest games used in the experiment, their various level-k predictions, equilibrium predictions and the thresholds $\overline{k}$. Notice that the first 12 games are "easy games" where targets of each player are 1-dimensional, while the last 12 games are "hard games " where targets are 2-dimensional. Also, Games $(2m-1)$ and $(2m)$ (where $m = 1, 2, \ldots, 12$) are the same but with reversed roles of the two players, so for instance, Games 1 and 2 are the same, Games 3 and 4 are the same, etc.

The spatial beauty contest game extends the one dimensional two-person guessing games in Costa-Gomes and Crawford (2006) to two dimensions. Extending to two dimensions allows us to separate choices and the reasoning of two players better. In easy games where targets are 1-dimensional, we let two players' target be on different dimensions so that the dimension corresponding to one's target is presumably more salient for that

[6]For instance, to go from location $(-3,1)$ to $(-4,0)$, one has to travel one step left and one step down and hence the distance is 2.

[7]However, choosing the top-right corner is *not* a dominant strategy in our design, unlike in the symmetric two-person guessing game analyzed by Grosskopf and Nagel (2008).

player (and differs from that for the other player). This may lead to more distinct reasoning processes and final choices, providing a better chance to identify when a player is reasoning for herself instead of reasoning for her opponent. In hard games where targets of both players are 2-dimensional, equilibrium choices do not coincide with the corners. This separates equilibrium reasoning from corner-choosing heuristics.

The thresholds $\bar{k}$ for our 24 games are almost always 4, but some are 3 (Games 1, 10, 17), 5 (Games 5, 11, 12) or 6 (Game 6). This indicates that as long as we include level-k types with k up to 3 and the equilibrium type, we will not miss the higher level-k types much since higher types coincide with the equilibrium most of the time. Moreover, as evident in Table I, different levels make different predictions. In other words, various levels are strongly separated on the grid map.[8] The level-k model predicts what final choices are made for each level k. We assume that a subject is of a particular level-k type in all games. This is formulated in Hypothesis 1.

Hypothesis 1 (Final Choice Data) *Consider a series of one-shot spatial beauty contest games without feedback, a subject's final choices for games $n = 1, 2, \ldots, N$ follows the prediction of a particular level-k type where k is constant across games.*

Since our games are spatial, players can literally count using their eyes how many steps on the grid map they have to move to hit their targets. Thus, a natural way to use lookups is to *take the level-k reasoning processes literally* assuming subjects look through the following best response hierarchy: An $L1$ best responds to an $L0$, an $L2$ best responds to an $L1$, …, and an Lk best responds to an $L(k-1)$. Though this process could be carried out solely in one's mind, counting on the map reduces memory load and is more straightforward. Hence, we formulate Hypothesis 2 and base our econometric analysis of lookups on this.

Hypothesis 2 (Lookup Data) *Consider a series of one-shot spatial beauty contest games without feedback where subjects are assumed to carry out the reasoning process on the grid map. A subject's lookup sequence for games $n = 1, 2, \ldots, N$ should follow a particular level-k type where k is constant across games, and:*

(a) **(Duration of Lookups):** *Focus on the associated level-k best response hierarchy and fixate longer than random at locations of $L0$ player's choices, …, own $L(k-2)$ player's choices, opponent $L(k-1)$ player's choices, and own Lk player's choices.*

(b) **(Sequence of Lookups):** *Have adjacent fixations (from level i to $i+1$) that correspond to steps of the associated level-k best response hierarchy.*

[8]The only exceptions are $L3$ and EQ in Games 1, 10, 17, $L2$ and $L3$ in Games 2, 6, 9, and $L2$ and EQ in Game 18. See the underlined predictions in Table I.

2.2 Experimental Procedure

We conduct 24 spatial beauty contest games (with various targets and map sizes) randomly ordered without feedback at the Social Science Experimental Laboratory (SSEL), California Institute of Technology.[9] In addition to recording subjects' final choices (Figure II), we also employ Eyelink II eyetrackers (SR-research Inc.) to track the entire decision process before the final choice is made. The experiment is programmed using the Psychophysics Toolbox of Matlab (Brainard, 1997), which includes the Video Toolbox (Pelli, 1997) and the Eyelink Toolbox (Cornelissen et al., 2002). For every 4 milliseconds, the eyetracker records the location one's eyes are looking at on the screen and one's pupil sizes. Location accuracy is guaranteed by first calibrating subjects' eyetracking patterns (video images and cornea reflections of the eyes) when they fixate at certain locations on the screen (typically 9 points), interpolating this calibration to all possible locations, and validating it with another set of similar locations. Since there is no hidden information in this game, the main goal of eyetracking is not to record information search. Instead, the goal is to capture how subjects carry out reasoning before making their decisions and to test whether they think through the best response hierarchy implied by a literal interpretation of the level-k model.

Before each game, a drift correction is performed in which subjects fixate at the center of the screen and hit a button (or space bar). This realigns the calibration at the center of the screen. During each game, when subjects use their eyes to fixate at a location, the eyetracker sends the current location back to the display computer, and the display computer lights up the location (real time) in red (as Figure II shows). Seeing this red location, if subjects decide to choose that location, they could hit the space bar. Subjects are then asked to confirm their choices ("Are you sure?"). then have a chance to confirm their choice ("YES") by looking at the bottom left corner of the screen, or restart the process ("NO") by looking at the bottom right corner of the screen. In each session, two students at Caltech were recruited through the SSEL website to be eyetracked. Since there was no feedback, each subject was eyetracked in a separate room individually and their results were matched with the other subject's at end of the experiment. Three trials were randomly drawn from the 48 trials played to be paid. Average payment is US$15.24 plus a show-up fee of US$20. A sample of the instructions can be found in Supplementary Appendix A8. A quiz was administered after the instructions were read out to make sure subjects understood the instructions (which all of them passed). Due to insufficient show-up of eligible subjects, three sessions were conducted with only one subject eyetracked,

[9] Each game is played twice with two different presentations that are mathematically identical. Since the results are similar, we focus on the presentation shown in Figure II that allows us to trace the decision-making process through observing the lookups.

and their results were matched with a subject's from a different session. Hence, we have eyetracking data for 17 subjects.

3 Lookup Summary Statistics

We first summarize subjects' lookups to test Hypothesis 2a, namely, subjects do look at and count on the grid map during their reasoning process. Here are two examples of the raw data. Figure III shows the lookups of subject 2 in trial 14, acting as a Member B. The diameter of each fixation circle is proportional to the length of each lookup. Note that these circles fall almost exclusively on the best response hierarchy of an $L2$, which is exactly her level-k type. Figure IV shows noisier lookups of an $L3$ type (subject 1) in trial 8 acting as member A.

We present the aggregate data regarding empirical lookups for all 24 Spatial Beauty Contest games in Supplementary Figures 1 through 24. For each game, we calculate the percentage of time a subject spent on each location. The radius of the circle is proportional to the average percentage of time spent on each location, so bigger circles indicate longer time spent. The level-k choice predictions are labeled as L0, L1, L2, L3, E. Consistent with Hypothesis 2a, subjects indeed spend more time at locations corresponding to the thinking steps of a particular best response hierarchy, so the empirical lookups concentrate on locations predicted by the level-k best response hierarchy. However, many *other* locations are also looked up.

We attempt to quantify this concentration of attention game-by-game. First, we define Hit area for every level-k type as the minimal convex set enveloping the locations predicted by this level-k type's best response hierarchy in game n. Figure V shows an example of Hit areas for various level-k types in a 7×7 spatial beauty contest game with target $(4, -2)$ and the opponent's target $(-2, 4)$ (Game 16). An $L2$ player 1 has a best response hierarchy consisting of locations O, $\mathbf{L1}_2$, $\mathbf{L2}_1$. Thus we can construct a minimal convex set enveloping these three locations (see the dashed line area in Figure V). For each game, we then take the union of Hit areas of all level-k types and see if the aggregate lookups of all subjects are indeed within the union.

We define the empirical percentage of time spent on an area as hit time, denoted as h_t, and define the size percentage of an area as hit area size, denoted as h_{as}. We calculate the difference between hit time and hit area size, $h_t - h_{as}$, to correct for the contribution of hit area size, and report it for the union of Hit areas for various level-k types as the $LookupScore$ in Figure VI. Note that this is essentially Selten (1991)'s linear "difference measure of predicted success." In fact, if subjects scan randomly over the grid map, her $LookupScore$ for any area is expected to be zero, since the percentage of time she spends on an area will roughly equal the hit area size of that area. By subtracting the hit area size, we can evaluate how high hit time is compared with what random scanning over the grid map would imply. These measures are all positive (except for Game 22), strongly rejecting the null hypothesis of random lookups. The p-value of one sample t-test is 0.0001, suggesting that subjects indeed spend a disproportionately long time on the union of Hit areas for various level-k types.

To sum up, the aggregate result is largely consistent with Hypothesis 2a that subjects look at locations of the level-k best response hierarchy longer than random scanning would imply, although the data is noisy. We next turn to test Hypothesis 2b and consider whether individual lookup data can be used to classify subjects into various level-k types.

4 Markov-Switching Level-k Reasoning

We now analyze subjects' lookups with a constrained Markov-switching model to classify them into various level-k types to test Hypothesis 2b. As a part of the estimation, we employ Vuong's test for non-nested but overlapping models to ensure separation between competing types. When performing this estimation, we use the entire sequence of lookups on the grid map for each trial. We summarize the estimation method here and provide the details in Supplementary Appendix A3.

We define *each stage of the reasoning process* as a *state*. The states are in the mind of a subject. If she is a level-2, there are three states according to the best response hierarchy of reasoning. For example, in Game 16 shown in Figure I, the three states are $s = 0$ (she thinks her opponent thinks she is a level-0), $s = -1$ (she thinks her opponent is a level-1), and $s = 2$ (she is a level-2). Note that to distinguish a state regarding beliefs about self from beliefs about the opponent, if a state is about the opponent, we indicate it by a minus sign.[10]

To account for the transitions of states within a subject's mind, we employ a Markov-switching model by Hamilton (1989) and characterize the transition using a Markov transition matrix. Instead of requiring a level-k subject to "strictly" obey a monotonic order of level-k thinking going from lower states to higher states, we allow subjects to move back from higher states to lower states. This is to account for the possibilities that subjects may go back to double check as may be typical in experiments. However, since a level-k

[10]We hasten to point out that these states are in the mind of a subject. It is not the level of a player. Take a level-2 subject as an example. Her level, according to the level-k model, is 2. But there are three states, $s = 0$, $s = -1$, and $s = 2$, in her mind. Which state she is in depends on what she is currently reasoning about. A level-2 subject could be at state $s = -1$ because at that point of time, she is thinking about what her opponent would choose, who is a level-1 according to the best response hierarchy. However, this state $s = -1$ is not to be confused with a level-1 subject (whose $k = 1$ and states of thinking consist of $s = -0$ and $s = 1$).

player best responds to a level-$(k-1)$ opponent, it is difficult to imagine a subject directly jumping from the reasoning state of say $s = (k-2)$ to that of $s = k$ without first going through the reasoning state of $s = -(k-1)$. Thus, we restrict the probabilities for all transitions that involve a jump to higher states to be zero.

When a subject is in a particular state, her reasoning will be reflected in the lookups which we can track. We now describe this mapping between the particular state a subject is in and her lookup. For instance, if a level-2 player with target $(4, -2)$ in Game 16 (player 1 as shown in Figure I) is at state $s = 0$ at a point of time, her state-to-lookup mapping would give us the location $(0, 0)$ which a level-0 player would choose (O in Figure I) since at this particular point of time, she is thinking about what her opponent thinks she would choose as a level-0. Similarly, if a level-2 player is in state -1, the mapping would give us the location $(-2, 3)$ which a level-1 opponent would choose ($\mathbf{L1}_2$ in Figure I) since at this particular point of time, she is thinking about what her opponent would choose as a level-1. Finally, if a level-2 player 1 is in state 2, then the mapping would give us the location $(2, 1)$ which a level-2 subject would choose ($\mathbf{L2}_1$ in Figure I) since at this particular point of time, she is thinking about her choice as a level-2.

For a level-0 player in state $s = 0$ thinking about choosing the center, the state-to-lookup mapping predicts her lookup should fall exactly on the location $(0, 0)$. If her lookup is not on that location, we interpret this as an error. We assume a logistic distribution of error so that looking at locations farther away from $(0, 0)$ is less likely, and estimate the precision parameter λ of the logistic distribution of error. When $\lambda \to +\infty$, subjects look at exactly $(0, 0)$. When $\lambda \to 0$, subjects look at all locations randomly with equal probability.

To summarize, for each level k, we estimate a state transition matrix and a precision parameter for the logistic distribution of error. Thus, for a given initial distribution of the states, we can infer the probability distribution of states at any point of time using the state transition matrix. Moreover, at any point of time, the mapping from the state to the lookup gives us the lookup location corresponding to any state when there is no error. Coupled with the error structure, we can calculate the probability distribution of various errors and therefore the distribution of predicted lookup locations. We then maximize the likelihood to explain the entire observed sequence of lookups. We do this for various level-k types. The final step is to select the k in various level-k types to best explain the observed sequence of lookups for each subject.

We caution that the above econometric model may be plagued by an overfitting problem because higher level-k types have more states and hence more parameters. It is not surprising if one discovers that models with more parameters fit better.[11] Hence, we need

[11] In particular, the Markov-switching model for level-k has $(k+1)$ states with a $(k+1) \times (k+1)$

to make sure our estimation does not select higher levels merely because it contains more states and more parameters. However, usual tests for model restrictions may not apply, since the parameters involved in different level-k types could be non-nested. In particular, the state space of a level-2 subject $\{0, -1, 2\}$ and the states of a level-1 subject $\{0, 1\}$ are not nested. Yet, the state space of a level-1 type, $\{0, 1\}$, is nested in the state space of a level-3 type, $\{0, 1, -2, 3\}$. In order to evaluate the classification, we use Vuong's test for non-nested but overlapping models (Vuong, 1989).[12]

Let Lk^* be the type which has the largest likelihood based on lookups. Let Lk^a be an alternative type having the next largest likelihood among all lower level types based on lookups.[13] If according to Vuong's test, Lk^* is a better model than Lk^a, we can be assured that the maximum likelihood criterion does not pick up the reported type by mere chance. Thus, we conclude that the lookup-based type is Lk^*. If instead we find that according to Vuong's test, Lk^* and Lk^a are equally good, then we conservatively classify the subject as the second largest lower type Lk^a.

Table II shows the results of maximum likelihood estimation and Vuong's test for each subject. For each subject, we list her Lk^* type, her Lk^a type, her Vuong's test statistic, and her lookup-based type according to Vuong's test in order. Six of the seventeen subjects (subjects 1, 5, 6, 8, 11, 13) pass Vuong's test and have their lookup-based type as Lk^*. The remaining eleven subjects are conservatively classified as Lk^a. The overall results are summarized in column (A) of Table III. After employing Vuong's test, the type distribution for $(L0, L1, L2, L3, EQ)$ is $(1, 6, 4, 4, 2)$. The distribution is in line with typical type distributions reported in previous studies. Treating the EQ type as having a thinking step of 4, we find that the average number of thinking steps is 2.00.

Up to now, we have shown that lookups do fall on the hotspots of the best response hierarchy (Hypothesis 2a). Classifying subjects based on lookups (Hypothesis 2b) gives us a reasonable level of sophistication as argued above. However, one might still wonder whether the results reported in Table II are merely a misspecification, as many assumptions are required for Hypothesis 2b to hold. In the next section, we take up this issue by matching subject lookup results with their final choices. Our argument is that if we take the level-k theory literally to interpret the underlying reasoning process, the classification

transition matrix. This gives the model $\frac{k(k+3)}{2}$ parameters in the transition matrix alone: Since each row sums up to one and elements with the column index greater than the row index plus one are zero, we have in total $(k+1)(k+1) - (k+1) - k(k-1)/2 = k(k+3)/2$ parameters. For example, a level-2 subject has 3 states and 5 parameters, but a level-1 subject has only 2 states and 2 parameters.

[12]See Supplementary Appendix A4 for the details of Vuong's test for non-nested but overlapping models. Note that this is the generalized version of the well-known "nested" Vuong's test.

[13]Recall that the reason why we use Vuong's test is to avoid overfitting. Hence, if the alternative type has a larger transition matrix (more parameters) but a lower likelihood, there is no point to perform a test, since Lk^* has fewer parameters but a higher likelihood. This leads us to consider only lower level types as the alternative type.

based on lookups should match well with the classification using final choices since the level k reflects a player's sophistication.

5 Matching Up with Final Choices

Following the literature, we classify individual subjects into various level-k types based on final choices. In particular, similar to Costa-Gomes and Crawford (2006), we perform a maximum likelihood estimation to classify each individual subject into a specific level-k type with beliefs that a level-0 chooses (on average) the center. Subjects are modeled as following a constant level-k and playing quantal response using a logistic error structure. Supplementary Appendix A5 provides the details of maximum likelihood estimation. The aggregate distribution of types is reported in row (B) of Table III. The type distribution for $(L0, L1, L2, L3, EQ)$ is $(2, 4, 4, 4, 3)$. The average number of thinking steps is 2.12, close to that based on lookup classifications.

If we consider the classification results on a subject-by-subject basis, the similarity between the two classifications are even more evident. Table III compares the lookup-based and choice-based classification results. For ten out of the seventeen subjects, their lookup-based types and their choice-based types are the same. In other words, for those subjects, when their choices reflect a particular level of sophistication, their lookup data suggests the same level of sophistication. This supports a literal interpretation of the level-k model—When a subject's choice data indicates a particular level, her lookups suggest that the best response hierarchy of that level is carried out when she reasons. We also report the average response time for each level-k in Table III. It is mostly increasing in k for both choice-based and lookup-based classifications. In other words, if anything, we find subjects of higher levels of sophistication take longer to make choices.

As a robustness check, to examine whether there are clusters of subjects whose choices resemble each other's and thus predict other's choices in the cluster better than the pre-specified level-k types, we conduct the pseudotype test of Costa-Gomes and Crawford (2006) and report the results in Supplementary Table 1(a).[14] We find only one cluster of pseudo-17 types, consisting of subject 3 and subject 17, indicating at most a small cluster of subjects that are not explained well by the predefined level-k model.[15]

Since the classification based on lookups and that based on choices align for more than a half of the subjects, we next turn to discuss the subtle differences between them. We evaluate the robustness of individual choice-based classification by performing bootstrap (Efron, 1979; Efron and Tibshirani, 1994; Salmon, 2001), as maximum likelihood esti-

[14]The idea of pseudotypes is to treat each subject's choices as a possible type. Since we have 17 subjects, we include 17 pseudotypes, each constructed from one of our subject's choices in 24 trials.

[15]The type distribution with pseudotypes is very similar and reported in row (C) of Table III.

mation may not have enough power to distinguish between various types. For example, reading from Supplementary Table 1(a), for subject 14, the log likelihood is −98.89 for $L0$, −84.17 for $L1$, −96.99 for $L2$, −76.67 for $L3$, and −74.45 for EQ. Maximum likelihood estimation classifies her as EQ, although the likelihood of $L3$ is also close.

To bootstrap, suppose a subject is classified as a particular level-k type with the logistic precision parameter λ_k from the maximum likelihood estimation. We then draw (with replacement) 24 new trials out of the original dataset and re-estimate her k and λ_k. We do this 1000 times to generate the discrete distribution of k (and corresponding λ_k), and evaluate the robustness of k by looking at the distribution of k. Each level-k type estimated from a re-sampled dataset that is not the same as her original level-k type is viewed as a "misclassification," and counted against the original classification k. By calculating the total misclassification rate (out of 1000 re-samples), we can measure the robustness of the original classification.

The results of this bootstrap procedure are listed in Table IV. For each subject, we report the bootstrap distribution of k (the number of times a subject is classified into $L0$, $L1$, $L2$, $L3$ or EQ in the 1000 re-sampled datasets). The bootstrap misclassification rate (percentage of times classifying the subject as a type different from her original type) is listed in the last column. For example, subject 14 is originally classified as EQ, but is only re-classified as EQ 587 times during the bootstrap procedure. She is instead classified as $L3$ 228 times and as $L1$ 185 times. Hence, the distribution on the number of times that subject 14 is classified into $L0$, $L1$, $L2$, $L3$ or EQ in the 1000 re-sampled datasets is $(0, 185, 0, 228, 587)$ and the corresponding misclassification rate is 0.413.

The bootstrap results align surprisingly well with whether the lookup-based classifications match their choice-based types. In particular, for the ten subjects whose two classifications match, all but three of them have (choice-based) bootstrap misclassification rates lower than 0.05, suggesting that their classifications are truly sharp.[16] In contrast, for six of the remaining seven subjects whose two classifications do not match, their choice-based type have bootstrap misclassification rates higher than 18.4%. The difference is significant, having a p-value of 0.0123 according to Mann-Whitney-Wilcoxon rank sum test. There may be some reasons why the two classifications sometimes disagree and why their choices seem noisy as we discover in the bootstrap procedure. We attempt a more systematic analysis using lookup and choice data next.

[16]One of these three subjects (subject 17) fails the pseudotype test and is unlikely to resemble any of the level-k types. The remaining two subjects (subjects 2 and 4) have misclassification rates of 0.076 and 0.110, respectively. These are marginally higher than 0.05.

6 Extracting Information From Lookups and Choices

To prepare us for a more systematic analysis, we first watch raw videos animating the entire lookup sequences trial-by-trial to gain some insights. We briefly summarize our observations though we caution that they are highly conjectural. More details are provided in Supplementary Appendix A6. Two out of the seventeen subjects can be deemed as textbook literal level-k types as their lookups follow the best response hierarchy very precisely. Most level-1 subjects do not look at the opponent's goal even once in many trials, suggesting that whether the minimal knowledge of the opponent is looked up may be the first criterion for judging a subject's level of strategic sophistication. We also discover some alternative ways which may have been used to simplify the reasoning process. These include breaking the two-dimensional games into two one-dimensional games to reason in order, adopting choosing-the-corner heuristic, and utilizing a short-cut by summing up the targets of the subject and the opponent. Though these alternative ways are harder to reconcile with the best response hierarchy we use, it remains to be seen how prevalent they are before alternative procedural assumptions can be made. We leave this for future research. Finally, one subject seems to be jumping between level-2 and level-4, while another skips reasoning in some trials. Recall there is no feedback in the experiment. Hence, alternating between different levels or skipping reasoning sometimes poses interesting challenges to the usual assumption of treating each trial as a one-shot game. This is beyond the scope of this paper.

We next turn to address what lookups can do whereas choices cannot. We explore two possibilities. First, we narrow down a subject's level-0 belief by analyzing where she initially looks at in every game. Second, we attempt to predict the final choice of any trial using only lookups of that particular trial. Choice data of several trials can be used to predict the choice of some other trial. Yet because there is only one choice in each trial, attempting to predict the choice using only the choice within a trial is impossible. We will see how informative the within-trial lookups prior to choice are. Finally, on examining whether subjects are literal constant level-k players, note that our lookups and choice classifications so far are based on three implicit assumptions. First, subjects are characterized by the level-k theory. Second, they have a constant k. Third, they are literal. Hence, when their lookups follow a particular level, their choices will follow that level as well. We will also examine these assumptions in a more systematic way.

6.1 Starting Point for Level-k Reasoning

One possibility lookups can help where final choices cannot is to narrow down level-0 belief. If a subject carries out the best response hierarchy, her initial lookups may reflect the level-0 play in her mind. To look into that, a natural way is to look at where initial lookups distribute. Since the grid map of each game is different, we need a way to summarize how initial lookups distribute over maps of different sizes.

Hence, for every game, we partition both dimensions of the map into three equal-sized bins. This way we divide each map into $3 \times 3 = 9$ equal-sized areas. Most salient areas for level-0 belief arguably are the center (O) and the top-left (TL), with the latter being focal because of the reading habit in English. However, for completeness, we also include the top-center (TC), the top-right (TR), the middle-left (ML), the middle-right (MR), the bottom-left (BL), the bottom-center (BC) and the bottom-right (BR) of the map. When subjects are free to look at any place in the map, we record which area their initial lookups lie. In particular, we consider the first 1% of the time spent on the grid map and count the percentage time spent in each area. The subject-by-subject percentage distribution of the initial lookups is reported in Table V. We further illustrate the aggregate percentage distribution over all subjects in Figure VII.

We find that O, TL, TC together account for 71% of the initial lookups. If subjects scan on the map uniformly, these three areas should account for only 1/3 of the time. Reading from the last column of Table V, eight out of the seventeen subjects have their modal initial lookups at O whereas four at TL and the other four at TC. This is broadly in support of using the center as the level-0 belief since indeed it is looked up most often initially. Moreover, subject 3 has her modal initial lookups at TL, in agreement with the observation that she starts her reasoning from TL reported in Supplementary Appendix A6. The three subjects (8, 11, 15) suspected to first perform reasoning regarding the horizontal dimension and then the vertical dimension in Supplementary Appendix A6, as well as pseudotype subject 17, have modal initial lookups at TC. This is possible if they start reasoning near the top row of the grid map. Finally, subject 14 is the only one whose modal initial lookups are not at O, TL or TC but at BL. However, as we point out in Supplementary Appendix A6, she has very few lookups and quickly chooses a corner. Her initial lookups fall on the four corner areas (TL, TR, BL, BR) quite evenly, totaling 82.6% of the time. This may reflect her quick final choices closely since choosing the corner directly does not rely on starting the reasoning process from any fixed location.

To summarize, we find evidence from initial lookups that the center area may most often be where subjects start reasoning. The top-left and the top-center may also be important but not as much as the center. Since the center and the top-left are salient, it is consistent with the importance of salience in determining level-0 belief (Burchardi and Penczynski, 2014).

Based on these results, we consider level-k types starting from the top-left. The last column of Table IV report alternative choice-based level-k types if they yield maximum likelihood (across all level-k types starting from C and TL). Only four subjects have maximum likelihood with level-0 belief of TL (subjects 3, 4, 7 and 17 as $L1$, $L3$, $L2$, and $L1$ via TL, respectively), and they indeed have modal or large fractions of initial lookups at TL.[17] Including level-k types starting from TL in the lookup-based classification results in subjects 3, 4, 7, 15, 17 being re-classified as $L3$ via TL, and all have modal or large fractions of initial lookups at TL.

6.2 Trial-by-Trial Lookup Estimation

Another possibility that lookups can do whereas final choices cannot is a trial-by-trial out-of-sample prediction. To this end, for each trial, we use lookups of only that trial to predict the final choice of the same trial. Relying on choice data alone cannot make such predictions because by definition each trial has only one choice data point, i.e. the choice itself, giving no further choice data to base predictions on. Supplementary Appendix A7 describes in detail how we use lookups of a trial to predict the final choice of that trial (shorthanded as the "1-trial lookup model"). In essence, taking the reasoning process of a level-k type starting from the center literally would imply specific locations to be looked at most often. We assume a subject looks uniformly over these locations, but conditional on each location her lookups follow the same logistic distribution over the grid map. Since there are $k+1$ locations, her lookup distribution will be a mixture of $k+1$ logistic distributions. We drop the last lookup since it is highly correlated with the final choice we want to predict and take the subject's lookup duration heatmap of the entire trial as our empirical distribution, and classify her into the type which minimizes the mean absolute difference between the mixture of logistic distributions of each level-k type and the empirical duration heatmap. Her choice of that trial is then predicted to be the choice of that classified level. In short, we classify every trial into a level based on only lookups of that trial before choice. The subject is then predicted to make a final choice of that level. We caution that such within-trial prediction may be noisy because we constrain ourselves to use only lookups of a single trial. However, this arguably tests whether lookups contain valuable information despite of the noisiness.

We use the economic value (EV) as our judgement criterion. EV is a widely-used measure to indicate how well a model performs. It is normalized so that EV = 0% means the model prediction leads to the same expected payoff of the actual subjects. On the other hand, EV = 100% implies the model prediction leads to the highest possible payoff as if playing the best response. Hence EV is interpreted as percentage gain from the prediction of a model, treating actual payoffs as the baseline. Rightly because so, if the

[17] Further including level-k types starting from the top-center only reclassifies subject 9 as $L2$ via TC, but she is reported to often skip reasoning in Supplementary Appendix A6.

model performs worse than the actual subjects, EV could be negative.[18] Compared with hit or miss of a model prediction, EV has the advantage of distinguishing near-misses (with EV close to 100%) from predictions that are much worse (with lower or even negative EV).

The level-k theory imposes the maximal prediction power of the 1-trial lookup model. Take Figure I as an example. The level-k theory predicts that player 1 chooses O, $\mathbf{L1}_1$, $\mathbf{L2}_1$, $\mathbf{L3}_1$ or $\mathbf{E}_1$ if her level is 0, 1, 2, 3 or 4 correspondingly. Accordingly, the 1-trial lookup model must eventually predict her choice to be from O, $\mathbf{L1}_1$, $\mathbf{L2}_1$, $\mathbf{L3}_1$ or $\mathbf{E}_1$. Therefore the prediction power of the 1-trial lookup model cannot exceed that of the best-predicting level among all possible k's. Hence, we first calculate the EV for each level k and find the maximum over all possible k's for each trial. We do this trial-by-trial, allowing the best level to differ from trial to trial. We then average this maximum EV over all trials. This is the maximal prediction power of the level-k theory with a trial-by-trial dependent level k. It is reported in column 2 of Table VI as "Maximum Level-k EV." Averaging over all subjects, the maximum EV imposed by the level-k theory is 83.7%. This is quite a significant amount, indicating that if the various best predicting level of every trial is known, payoffs can be increased substantially. The question is whether the lookups before the choice can help obtain this valuable information trial-by-trial.

To evaluate performance of the 1-trial lookup model, for each trial, we divide the EV based on the prediction of the 1-trial lookup model by the maximum EV to obtain the EV ratio of that trial. This ratio reflects the fraction of possible EV realized by the 1-trial lookup model in that trial. Averaging over all trials we construct a measure of how well the 1-trial lookup model performs, compared with the upper bound imposed by the level-k theory. This is reported in column 3 as "1-Trial Lookup." Averaging over all subjects, the EV ratio of the 1-trial lookup model is 0.71, roughly indicating that 0.71 of the 83.7% gain can indeed be realized by relying on the lookups of a trial to make a prediction. Hence, if an opponent knows a subject's lookups of a trial before making a choice, his payoffs can be increased roughly by 60%. This supports that lookups contain valuable information for making choices optimally.

For comparison, we also look into how valuable choice information is. Since there are 24 trials in the experiment, we conduct the leave-1-choice-out model. This assumes a stable level for every 23 trials, and relies on choices of these 23 trials to classify a subject into a level. She is then predicted to make a final choice of that level in the remaining left-out trial. This is similar to what we did in Section 5 except we do it for every 23

[18] Precisely, EV= $\dfrac{\pi^{\text{Follow}} - \pi^{\text{Actual}}}{\pi^{\text{BR}} - \pi^{\text{Actual}}}$. To illustrate, suppose a model predicts subject 1 to be level-3. Then π^{Follow} is her opponent's payoff should he follow the model prediction and best respond to level-3. π^{Actual} is his actual payoff. π^{BR} is her opponent's payoff should he best respond to subject 1's choice. Hence EV is the percentage of the gain should a model be followed, compared with the maximum possible gain implied by playing best response to subject 1's choice.

trials and hence 24 times. For each trial, we again divide the EV based on the prediction of the leave-1-choice-out model by the maximum EV to obtain the EV ratio of that trial. Averaging over all trials we construct a measure of how well the leave-1-choice-out model performs, compared with the upper bound imposed by the level-k theory. This is reported in column 5 of Table VI as "Leave-1-Choice-Out." Averaging over all subjects, the EV ratio of the leave-1-choice-out model is 0.72, similar to the EV ratio of the 1-trial lookup model. In words, this means, in terms of playing optimally, learning the lookups of a trial is as useful as learning choices of 23 trials. This is important in situations where the entire history of choices is unavailable, in which case the 1-trial lookup model would be a good substitute to make a prediction.

6.3 Literal Constant Level-k

We next turn to address whether the constant level-k reasoning process is carried out literally. If a player is a literal level-k whose k is constant throughout, we argue three criteria have to be met. First, the maximum EV the level-k theory imposes on has to be high enough. After all, if the maximum EV is low, the level-k theory does not help reaching optimal plays and some alternative theory may perform better. Second, the EV ratio of the leave-1-choice-out model has to be high as well. Since the leave-1-choice-out model assumes a stable level for every 23 trials, a low EV ratio could suggest a violation of the stability of levels, questioning the constant level assumption. Third, the EV ratio of the 1-trial lookup model has to be high too. Otherwise, one may doubt whether the literal level-k reasoning process is carried out before the choice is made, suggesting alternative reasoning processes that might be considered to reach the choice.

We use the criteria that the maximum EV has to be at least 70% and the two EV ratios at least 0.7. Among all subjects, eleven pass the criteria. This suggests that for most subjects, we cannot reject that the literal, constant level-k theory predicts optimal plays quite well, even though we cannot directly prove that subjects are indeed carrying out the constant level-k reasoning literally. This lends support to the level-k theory as coming up with a good prediction as far as EV is concerned. This is in contrast to Georganas, Healy and Weber (2015) which assigns the final choice to the closest level-k type and finds individual type not persistent within the family of guessing games.[19] Moreover, this demonstrates that the lookup and choice analyses in Sections 4 and 5 are well-founded because had the opponent assumed the literal constant level-k reasoning process of the subject, this prediction serves him well by making his EV quite close to the best response

[19] Georganas, Healy and Weber (2015) cannot estimate both k and the logistic precision parameter λ with one data point. Nonetheless, they also find persistence individual types within the family of undercutting games. Arad and Rubinstein (2012) correlate level-k reasoning in 91–100 games and iterative reasoning in a Blotto game.

to the actual play of the subject.

The remaining six subjects, 14, 9, 4, 3, 7, 17 fail the criterion. We now look into why their EV or EV ratio is not as high and how lookup information may help.

Subjects 3, 7 and 17 have the lowest maximum EV. While other subjects have a maximum EV around 80% or higher, they all have maximum EV below 70%. This indicates they likely do not follow the level-k model some way or the other. In fact, the pseudotype test of Section 5 reveals that subject 3 and 17 are pseudotypes to each other and subject 7 is classified as level-0, or close to random, both suggesting that the level-k theory may not work well for them. Not surprisingly, the alternative estimation in Section 6.1 allowing for level-k types starting from the top-left corner shows that they indeed follow level-k reasoning with the alternative starting point of the top-left corner instead of the center. We thus expand the 1-trial lookup model by adding alternative types which start the reasoning process from the top-left corner. If we classify every trial of a subject into a level of this expanded set based on only lookups of that trial, the EV ratios of the 1-trial lookup model in column 3 of Table VI are substantially increased from -0.08 to 2.36 for subject 3, from 0.74 to 0.98 for subject 7 and from 0.94 to 1.33 for subject 17. The fact that the EV ratio is higher than 1 for subjects 3 and 17 implies that allowing the reasoning process from the top-left corner makes the 1-trial lookup model even more informative than the maximum level-k model. We thus are even more confident that subjects 3 and 17 might have started their reasoning process from the top-left corner as this makes their lookups become so informative presumably because starting from the top-left corner fits their lookups well.

Subjects 9 and 4 have maximum EV around 80%, but especially low EV ratios (0.47 and 0.36 respectively) for the leave-1-choice-out model. This suggests that we may question the stability of a fixed level for them. As we find in Supplementary Appendix A6, subject 9 initially behaves like level-2, but eventually has very few lookups (indicating not reasoning) and chooses the center as level-0 in the later trials. Subject 4 seems to jump around initially, but eventually settles down as level-1 and hits the level-1 choice perfectly in the last 8 trials. Hence the stability of levels may indeed be questionable for them. Relying on their lookups instead, their 1-trial lookup model has somewhat higher EV ratios (0.71 for subject 9 and 0.57 for subject 4) presumably because the 1-trial lookup model does not assume a fixed level. Hence the trade-offs of relying on the 1-trial lookup model are that within-trial predictions by lookups potentially could be noisy. Yet the noisy prediction has the flexibility to predict better when levels are not constant.

Subject 14 has a maximum EV of 87.1%, but a low EV ratio of 0.30 for the 1-trial lookup model. This suggests that she likely does not follow level-k reasoning literally. As we indicate in Section 6.1, subject 14 has very few lookups and seems to jump straight to a corner. We conjecture she follows the corner heuristics by choosing the equilibrium corner in easy games, but the corner in the direction of her goals in hard games. In fact, if we include this particular corner heuristic, her EV ratio is significantly improved from 0.30 to 1.29. Hence, the low EV ratio of the 1-trial lookup model helps indicate that this subject might perform an alternative reasoning process to arrive at the same choice.

Overall, in terms of EV, for eleven subjects, we find support that the literal, constant level-k theory predicts optimal plays quite well. This echoes our finding in Section 5 that subjects' lookup-based types and their choice-based types are quite consistent. For the remaining subjects, lookup information can help us confirm whether they start the reasoning process instead from the top-left corner (for subjects 3 and 17), they might not have a stable level (subjects 9 and 4) or they may not go through the best response hierarchy literally (subject 14).

7 Conclusion

We introduce the spatial beauty contest game in which the process of reasoning can be tracked, and provide theoretical predictions together with a procedural interpretation of the level-k theory. This procedural interpretation yields a plausible hypothesis on the decision-making process. We then conduct a laboratory experiment using video-based eyetracking technology to test this hypothesis, and fit the eyetracking data on lookups using a constrained Markov-switching model of level-k reasoning. Results show that based on lookups, subjects' lookup sequences could be classified into following various level-k best response hierarchies, which for more than a half of them coincide with levels that they are classified into using final choices. Finally, initial lookup data and trial-by-trial lookup estimation indicate that most subjects indeed follow the constant level-k reasoning literally on the grid map starting from the center. In fact, lookups of a trial contain valuable information to predict the choice of that same trial well.

Analyzing reasoning processes is a hard task. The spatial beauty contest game is designed to fully exploit the structure of the p-beauty contest so that subjects are induced to literally count on the grid map to carry out their reasoning as implied by the best response hierarchy of a level-k theory. The high percentage of subjects whose classifications based on lookups and those based on choices align could be read as a support to the level-k model as a complete theory of reasoning and choice altogether in the spatial beauty contest game. Whether this holds true for more general games remains to be seen. Nevertheless, the paper adds on the literature and points out a possibility of analyzing reasoning before arriving at choices. To best utilize the procedural data, a design which suits the tracking technology used is indispensable.

References

Agranov, M., Caplin, A. and Tergiman, C. (2015), 'Naive play and the process of choice in guessing games', *Journal of the Economic Science Association* **1**(2), 146–157.

Arad, A. and Rubinstein, A. (2012), 'Multi-dimensional iterative reasoning in action: The case of the Colonel Blotto game', *Journal of Economic Behavior & Organization* **84**(2), 571 – 585.

Brainard, D. H. (1997), 'The Psychophysics Toolbox', *Spatial Vision* **10**, 433–436.

Brocas, I., Carrillo, J. D., Wang, S. W. and Camerer, C. F. (2014), 'Imperfect Choice or Imperfect Attention? Understanding Strategic Thinking in Private Information Games', *Review of Economic Studies* **81**(3), 944–970.

Burchardi, K. B. and Penczynski, S. P. (2014), 'Out Of Your Mind: Eliciting Individual Reasoning in One Shot Games', *Games and Economic Behavior* **84**, 39–57.

Camerer, C. F., Ho, T.-H. and Chong, J.-K. (2004), 'A cognitive hierarchy model of games', *Quarterly Journal of Economics* **119**(3), 861–898.

Cornelissen, F. W., Peters, E. M. and Palmer, J. (2002), 'The Eyelink Toolbox: Eye tracking with MATLAB and the Psychophysics Toolbox', *Behavior Research Methods, Instruments and Computers* **34**, 613–617.

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

Costa-Gomes, M., Crawford, V. P. and Broseta, B. (2001), 'Cognition and behavior in normal-form games: An experimental study', *Econometrica* **69**(5), 1193–1235.

Efron, B. (1979), 'Bootstrap Methods: Another Look at the Jackknife', *The Annals of Statistics* **7**(1), 1–26.

Efron, B. and Tibshirani, R. J. (1994), *An Introduction to the Bootstrap*, Chapman and Hall/CRC Monographs on Statistics and Applied Probability, Chapman and Hall/CRC.

Georganas, S., Healy, P. J. and Weber, R. A. (2015), 'On the persistence of strategic sophistication', *Journal of Economic Theory* **159**, 369 – 400.

Grosskopf, B. and Nagel, R. (2008), 'The two-person beauty contest', *Games and Economic Behavior* **62**(1), 93–99.

Hamilton, J. D. (1989), 'A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle', *Econometrica* **57**(2), 357–384.

Ho, T. H., Camerer, C. F. and Weigelt, K. (1998), 'Iterated dominance and iterated best response in experimental "p-beauty contests"', *American Economic Review* **88**(4), 947–969.

Ho, T.-H. and Su, X. (2013), 'A Dynamic Level-k Model in Sequential Games', *Management Science* **59**(2), 452–469.

Kuo, W.-J., Sjostrom, T., Chen, Y.-P., Wang, Y.-H. and Huang, C.-Y. (2009), 'Intuition and Deliberation: Two Systems for Strategizing in the Brain', *Science* **324**(5926), 519–522.

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

Nash, J. F. (1950), 'Equilibrium points in n-person games', *Proceedings of the National Academy of Sciences* **36**(1), 48–49.

Pelli, D. G. (1997), 'The VideoToolbox software for visual psychophysics: Transforming numbers into movies', *Spatial Vision* **10**, 437–442.

Penczynski, S. P. (2016), 'Strategic thinking: The influence of the game', *Journal of Economic Behavior & Organization* **128**(Supplement C), 72 – 84.

Salmon, T. C. (2001), 'An Evaluation of Econometric Models of Adaptive Learning', *Econometrica* **69**(6), 1597–1628.

Selten, R. (1991), 'Properties of a measure of predictive success', *Mathematical Social Sciences* **21**(2), 153–167.

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

Vuong, Q. (1989), 'Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses', *Econometrica* **57**(2), 307–333.