Level-k Reasoning in a Generalized Beauty Contest*

Dmitry Shapiro Xianwen Shi Artie Zillante

First Version: October 2009
This Version: March 2014

Abstract

We study how the predictive power of level-k models changes as we perturb the classical beauty contest setting along two dimensions: the strength of the coordination motive and the information symmetry. We use a variation of the Morris and Shin (2002) model as the unified framework for our study, and find that the predictive power of level-k models varies considerably along these two dimensions. Level-k models are successful in predicting subject behavior in settings with symmetric information and a strong coordination motive. Their predictive power weakens significantly when either private information is introduced or the importance of the coordination motive is decreased.

1 Introduction

The experimental literature on beauty contests and related guessing games has documented substantial evidence that individuals tend to have a limited degree of strategic sophistication, especially in settings where the strategic reasoning is not straightforward. The "p-beauty contest," in which participants choose a number between 0 and 100 and whoever picks the number closest to a multiple p of the group average wins a prize, best illustrates this limited degree of strategic sophistication. The p-beauty contest can be solved by iterated elimination of weakly dominated strategies, and the unique equilibrium occurs when every player chooses 0. In order to reach this equilibrium subjects need to go through a large number of rounds of elimination of dominated strategies. The experimental literature on beauty contests, however, shows that subjects usually perform one to three rounds of elimination and that their behavior is consistently different from the equilibrium prediction.

*Shapiro and Zillante: Belk College of Business, University of North Carolina Charlotte, 9201 University City Boulevard, Charlotte NC 28223-0001, USA; email address: dashapir@uncc.edu, azillant@uncc.edu. Shi: Department of Economics, University of Toronto, 150 St. George Street, Toronto, Ontario M5S 3G7, Canada; email address: xianwen.shi@utoronto.ca. We would like to thank Vincent Crawford, an advisory editor, three anonymous referees, Stephen Morris, Kirill Evdokimov and Pei-Yu Lo for helpful comments and suggestions. We are also grateful to Rosemarie Nagel for providing us data from her papers and to Yuanyuan Wan for sharing his expertise on maximum likelihood estimation. This research was supported by funds provided by a Childress Klein Research Grant, the Belk College of Business, the Connaught start-up fund at University of Toronto and a SIG award from SSHRC of Canada.

Stahl and Wilson (1995) and Nagel (1995) are seminal proponents of the theory of level-k reasoning. Ho et al. (1998), Costa-Gomes et al. (2001), Bosch-Domenech et al. (2002), Costa-Gomes and Crawford (2006), and Crawford and Iriberri (2007a,b), among others, have further developed and applied level-k models to beauty contests and related settings. The level-k model is based on the presumption that subjects' behavior can be classified into different levels of reasoning. The zero level of reasoning, $L0$, corresponds to non-strategic behavior when strategies are selected at random without forming any beliefs about opponents' behavior. In the literature $L0$ is typically considered to be a person's model of others rather than an actual person. Level-1 players, $L1$, believe that all other players are $L0$ and play a best response to this belief. Level-2 players, $L2$, play the best response to the belief that all other players are $L1$ and so on. For example, when p is equal to 2/3 in the beauty contest, level-1 players choose 33 and level-2 players choose 22. As is shown in Nagel (1995), Duffy and Nagel (1997), and many other papers, there is indeed a salient pattern of levels of reasoning in the beauty contest setting.

While level-k thinking is not particularly unique to the beauty contest (see e.g. Costa-Gomes and Crawford, 2006), the structure of the game and its simplicity are very conducive to this type of behavior. Success in the beauty contest largely depends on a person's ability to correctly predict the average choice made by others which explicitly forces individuals to think about the decisions of other players. Moreover, the symmetry of information makes this task relatively simple, which can further encourage participants to focus on the behavior of others.

Most of the existing literature focuses on games with complete information. An early exception is Crawford and Iriberri (2007a), who applied level-k reasoning to first- and second-price auctions. In many real applications, however, market participants often have access to both public and private information on the underlying fundamentals, and choose actions that are not only responsive to peer action choices but also appropriate to the fundamentals. A natural question then arises: how will level-k models perform beyond the classical beauty contest setting?

To answer this question, we introduce a framework that generalizes the classical beauty contest setting along two dimensions. First, it allows players to have private information that is relevant for their action choice. Second, it allows the importance of coordination to change so that the ability of correctly guessing other players' actions can have a different impact on players' payoffs. We then analyze how the predictive power of the level-k models varies along these two dimensions.

The generalized framework for our study is a modification of the Morris and Shin (2002, hereafter MS) model on the social value of public information.[1] In our setting, just as in MS, the agent's payoff is determined by two criteria: how well an agent's action matches an unknown state of the world and how well it matches the average actions of other agents. The relative importance of both factors can be varied within the model. In particular, as the latter becomes more important it makes the coordination motive of the game stronger. Agents in our model receive two signals

[1] Subsequently, Angeletos and Pavan (2007) generalize the MS analysis of the social value of information by allowing both strategic complementarity and strategic substitutability among agents' actions. The MS framework has been applied to many different settings including asset pricing (Allen et al., 2006, Bacchetta and Wincoop, 2005), venture capital (Angeletos et al., 2007) and political science (Dewan and Myatt, 2007, 2008).

about the (unknown) underlying state. If both signals are public the information is symmetric. If one signal is public and the other is private (as in the original MS setting) then the information is asymmetric and, in particular, different participants have different information.

Based on this framework we design several experimental treatments that differ from each other in the symmetry of information and the importance of predicting the average action of other players. Our main findings are as follows. First, in aggregate subjects place less weight on the public signal than the MS model predicts, a result that is consistent with the theoretical prediction of level-k models. An important implication is that, if agents have limited cognitive ability, the detrimental effect of increased public disclosure on social welfare may not be as strong as the MS model predicts. The MS setting is also used by Cornand and Heinemann (2014), who conduct two-player experiments within the MS framework and also find that subjects put less weight on the public signal than the theory predicts. While similarities exist between our paper and theirs, the two papers were developed contemporaneously and independently and differ substantially in both experimental design and research focus. In particular, they exclusively focus on the welfare implications of public disclosure, whereas our main focus is to test the performance of level-k models across settings with different information and payoff structures.

Second, we use the data to determine which (expected) value of level-0 maximizes the predictive power of the level-k model. For most of the treatments the value is very close to the average of the two signals, which would correspond exactly to 50 in the standard beauty contest. There is, however, an important difference between the public information treatment with the strongest coordination motive (the one closest to the beauty contest) and the others. In the former the share of level-0 players is negligible as is usually assumed. In other treatments many subjects routinely choose an action equal to the average of their two signals, as if they were ignoring the strategic aspect of the game entirely and simply focusing on matching the underlying state of the world. Typically these players would be classified as level-0 players, and there are a non-negligible number of them. Ignoring this behavior that looks like a level-0 action in the standard setting substantially decreases the performance of the level-k model in those treatments. Two recent experimental studies on level-k models, Agranov, Caplin and Tergiman (2013) and Burchardi and Penczynski (2014), find that there are a substantial fraction of subjects who are $L0$ players in their experiments, and the average action of $L0$ subjects are around 50.

Third, after determining the appropriate $L0$ definition for the experimental treatments, we compare the performance of the level-k model in explaining subjects' behavior. Treatments with public information in which available actions are restricted to be between the two signals are the most consistent with level-k predictions. For the restricted public treatments with the strongest coordination motive, the level-k model can explain more than 60% of the data. For the restricted public treatments with lower coordination motives the level-k model explains more than 40% of the data in all treatments other than the one with the lowest coordination motive. The performance of the level-k model is slightly worse in treatments with public information when the restriction on the action set is removed, save for the treatment with strong coordination motives. When private information is introduced, the performance of level-k model is much worse than its performance in corresponding public information treatments, with the exception of treatments with the lowest coordination motive.

Finally, because level-k behavior is viewed as an appropriate framework for subjects' initial behavior before any learning takes place, we redo our analysis using only first round data. In treatments with public information and a restricted action set, the level-k model performed better for the first round only data than for the full data set. In particular, the values of $L0$ that can best explain the data are very close to the average between two signals (equivalent to 50 in the beauty contest) and, with the exception of the treatment with the lowest coordination motive, the success rates varied between 56% and 76%. In all other treatments, however, the results were considerably weaker. The level-k model often explained the entire data better than the first round data.

Overall, the data suggest that subject behavior in public information treatments with a restricted action set (i.e. the setting closest to the classical beauty contest) is the most consistent with level-k predictions. The average of the two signals serves as a focal point for subjects' initial beliefs, a large share of chosen actions can be explained by the level-k models, and in the first round the explanatory power of the level-k model is higher. Furthermore, by any criterion that we use in the paper, the level-k models perform the best in the treatment with the strongest coordination motive and public information.

Our analysis highlights the strengths and limitations of level-k models. The modified MS framework used in our study is considerably more complicated than those typically used in the level-k literature. Despite this complexity level-k models are very successful in predicting subjects' behavior in settings that are close to the classical beauty contest, such as when the coordination motive is strong and information is symmetric. At the same time the predictive power of level-k models diminishes as we move away from the classical setting by either weakening the coordination motive or by introducing private information.

Our experimental findings also have important policy implications. The key insight in the analysis of Morris and Shin (2002) is that in equilibrium players often place too much weight on the public signal relative to the weight that would be used by the social planner. Therefore, individual information aggregation is not socially efficient and enhanced public disclosure could hurt social welfare. However, our theoretical analysis of level-k reasoning shows that limited cognitive ability, either due to limited level of reasoning or incapability of Bayesian updating, necessarily leads to subjects underweighting the public signal compared to the equilibrium prediction. In our experiment subjects indeed put less weight on the public signal than the theory predicts, implying that limited cognitive ability can limit the detrimental effect of increased public disclosure.

The rest of the paper is organized as follows. In Section 2 we provide a theoretical background for our study which is largely based on the MS model. We derive the prediction of level-k models in this setting and show that subjects with limited cognitive ability will put less weight on the public signal than the equilibrium predicts. Section 3 provides details of our experimental design and various treatments. Our experimental results are reported in Section 4 and alternative models are considered in Section 5. Section 6 concludes. The experimental instructions are given in the Appendix.

2 Theoretical Background

This section provides a theoretical background for our study. The primary goal of our paper is to analyze performance of level-k reasoning in a setting that is similar to the classical beauty contest yet allows us to vary the importance of the coordination motive and information structure. For this purpose we use the MS framework as a basis for our experimental analysis. We cannot implement the original MS model directly in the lab, however, because it uses assumptions such as a continuum of agents and an improper uniform distribution. In section 2.1 we modify the original MS model to adapt it to an experimental environment. In section 2.2 we use the modified MS framework to derive predictions of the level-k model, and in section 2.3 we discuss different options for specifying $L0$.

2.1 Modified Morris-Shin Model

There are I ex-ante identical agents, $i = 1, ..., I$. Agent i chooses an action $a_i \in \mathbb{R}$. The payoff function for agent i is given by

$$u_i\left(a_i, \overline{a}_{-i}, \theta\right) = C - (1-r)\left(a_i - \theta\right)^2 - r\left(a_i - \lambda \overline{a}_{-i}\right)^2, \tag{1}$$

where C is a constant, θ represents the underlying state, r and λ are constants between 0 and 1, and $\overline{a}_{-i}$ is the average action of i's opponents: $\overline{a}_{-i} = \frac{1}{I-1}\sum_{j \neq i} a_j$.

The payoff function has three terms. The first term is a constant C and is the highest payoff the individual can possibly receive. The second term reflects the loss from mismatching the underlying state θ and is simply the square of the distance between θ and a_i. The third term is the "beauty contest" term. It measures the loss from mismatching the average action of opponents, $\overline{a}_{-i}$, which is scaled by λ. The parameter r measures the relative importance of coordinating with opponents' actions versus matching the underlying state. When $\lambda = 1$ and $C = 0$ the game becomes the coordination game specified in MS. When $r = 1$ and $\lambda < 1$ the game becomes similar to the beauty contest in the sense that subjects only need to match λ times the average of other players' actions. Unlike the beauty contest, however, everyone, not just the player whose guess is the closest to the target, receives a non-negative payoff.

Our payoff function differs from the MS one in three ways. First, we consider a setting with a finite number of players while in MS there is a continuum of players. Second, we introduce the term λ inside the payoff function to match the classical p-beauty contest. Third, the payoff function in MS is always negative, which is difficult to implement in the laboratory. By adding a positive constant C to the original payoff function we allow participants' payoffs to be positive without altering equilibrium predictions.

As in MS, before taking actions, agent i will receive two signals about θ and we assume that both signals have the same precision α. The first signal y is always public and is given by

$$y = \theta + \eta, \;\; \eta \sim N\left(0, 1/\alpha\right). \tag{2}$$

As for the second signal, x_i, it can be either public or private. If it is private, then

$$x_i = \theta + \varepsilon_i,\ \varepsilon_i \sim N\left(0, 1/\alpha\right), \tag{3}$$

where η and ε_i are independent. If it is public, then it is the same across agents and is given by

$$x_i = \theta + \varepsilon,\ \varepsilon \sim N\left(0, 1/\alpha\right).$$

Again η and ε are independent. After receiving x_i and y, agent i chooses action a_i. MS assume that θ is distributed with the improper uniform distribution over the real line in which case the expected value of θ given x_i and y is

$$\mathbb{E}_i(\theta|x_i, y) = \frac{y + x_i}{2}. \tag{4}$$

Following the same procedure as in MS we can show that when x_i is private the unique equilibrium is linear and is given by

$$a_i\left(y, x_i\right) = \frac{1-r}{2-\lambda r}x_i + \frac{1-r}{(2-\lambda r)(1-\lambda r)}y. \tag{5}$$

When signal x_i is public, the unique Nash equilibrium (NE) is

$$a_i\left(y, x_i\right) = \frac{1-r}{2-2\lambda r}x_i + \frac{1-r}{2-2\lambda r}y. \tag{6}$$

Notice, in particular, that when $\lambda < 1$ and $r = 1$ the NE is 0, as in the beauty contest.

A major difficulty of implementing the MS setup in the lab is to generate θ according to the improper uniform distribution. To deal with this problem we adopted the following strategy. We generate θ using the uniform distribution on interval $[a, b]$ and then given θ we generate the signals y and x_i according to (2) and (3). We then normalize state θ and signals (x_i, y) by subtracting y from each of them, so that $\theta^* = \theta - y$, $x_i^* = x_i - y$ and $y^* = 0$.

Because the prior of θ has a bounded support, the formula (4) to obtain $E(\theta|x_i, y)$ may not be valid, and thus the NE would no longer be given by (5) and (6). However, normalized signals are immune to this problem. By the definition of y, we have $\theta^* = -\eta$ and $x_i^* = \varepsilon_i - \eta$. As both $-\eta$ and $\varepsilon_i - \eta$ are normally distributed, by the standard formula for the conditional distribution of normally distributed random variables we have

$$\mathbb{E}(\theta^*|x_i^*) = \mathbb{E}(-\eta|\varepsilon_i - \eta) = \frac{1}{2}(\varepsilon_i - \eta) = \frac{x_i^*}{2}.$$

Given $y^* = 0$, this is the same as (4). Therefore, when agents observe normalized signals the MS logic and the equilibrium derivations remain valid. In the experimental design section we provide more details on how the normalization is implemented.

2.2 Calculating Levels of Reasoning

Within the setting introduced in the previous section we derive actions that correspond to different levels of reasoning. From now on we assume that signals and the state are normalized, and with slight abuse of notation, will use θ, x_i, and $y(=0)$ to denote the normalized state and signals. Note that the posterior estimate of the state θ for player i with private signal x_i is $\mathbb{E}_i[\theta] = \frac{1}{2}x_i$.

Player i chooses a_i to maximize (1). It follows from the first-order condition that the best response is

$$a_i^* = (1-r)\,\mathbb{E}_i[\theta] + r\lambda \mathbb{E}_i[\bar{a}_{-i}].$$

Except for the non-strategic $L0$ type, agents with different levels of reasoning will form different beliefs about $\mathbb{E}_i[\bar{a}_{-i}]$ and will choose an action accordingly.

The first step in calculating Lk actions is to define the behavior of $L0$. In the literature type $L0$ is usually viewed as the starting point of a player's analysis of others' actions, so it should be unsophisticated and non-strategic (see e.g. Crawford and Iriberri, 2007). According to the standard level-k model, an $L1$ agent expects that other players are $L0$ players. Different $L0$ specifications will affect $L1$ players' beliefs about the average action of other players.

Fix an $L0$ specification. Suppose, under this $L0$ specification, an $L1$ player (say player i, who by definition thinks all other players are $L0$ players) believes that the average action of player i's opponents is a fraction π of the average of the signals received by i's opponents:

$$\bar{a}_{-i} = \pi \overline{x}_{-i}, \tag{7}$$

where $\overline{x}_{-i}$ denotes the average of signals received by player i's opponents. Player i's estimate about the average action of other players is then $\mathbb{E}_i[\bar{a}_{-i}] = \pi \mathbb{E}_i[\overline{x}_{-i}]$. For example, if player i believes that her opponents' actions are uniformly distributed between 0 and their private signals, then $\pi = \frac{1}{2}$ and $\mathbb{E}_i[\bar{a}_{-i}] = \frac{1}{2}\mathbb{E}_i[\overline{x}_{-i}]$.

In the settings where x_i is private, $\mathbb{E}_i[\overline{x}_{-i}] = \mathbb{E}_i[\theta]$. It follows that

$$\mathbb{E}_i[\bar{a}_{-i}] = \pi \mathbb{E}_i[\overline{x}_{-i}] = \pi\left(\frac{1}{2}x_i\right) = \frac{1}{2}\pi x_i.$$

Therefore, an $L1$ player in the settings with private signals will play

$$a_{L1} = (1-r)\frac{1}{2}x_i + r\lambda\frac{1}{2}\pi x_i = \left[(1-r)\frac{1}{2} + \frac{1}{2}r\lambda\pi\right]x_i.$$

We use induction to derive the action choice of a level-n agent. Let a_{Ln} denote the action taken by an Ln player with private signal x_i. Then it takes the following linear form: $a_{Ln} = \beta_n x_i$, where β_n is a coefficient depending on r, λ, and π. In particular

$$\beta_0 = \pi \text{ and } \beta_1 = (1-r)\frac{1}{2} + \frac{1}{2}r\lambda\pi. \tag{8}$$

Now consider an $L(n+1)$ player with private signal x_i. Then she expects that other players are Ln players and

$$\mathbb{E}_i[\bar{a}_{-i}] = \mathbb{E}_i[\beta_n \overline{x}_{-i}] = \beta_n \frac{1}{2}x_i.$$

Therefore,

$$a_{Ln+1} = (1-r)\,\mathbb{E}_i[\theta] + r\lambda\mathbb{E}_i[\bar{a}_{-i}] = \frac{1}{2}(1-r)\,x_i + r\lambda\frac{1}{2}\beta_n x_i.$$

It follows that

$$\beta_{n+1} = \frac{1}{2}(1-r) + \frac{1}{2} r\lambda\beta_n,$$

which implies the following difference equation:

$$(\beta_{n+1} - \beta_n) = \frac{1}{2} r\lambda (\beta_n - \beta_{n-1}).$$

Using the initial condition (8), we can solve[2]

$$\beta_n = \frac{1-r}{2-r\lambda} + \frac{(1-\lambda\pi)r + 2\pi - 1}{2-r\lambda}\left(\frac{1}{2}r\lambda\right)^n. \tag{9}$$

When signal x_i is public, by following a similar procedure we can show that an Ln agent with signal $(x_i, 0)$ will choose action $\widetilde{\beta}_n x_i$, where $\widetilde{\beta}_n$ is given by

$$\widetilde{\beta}_n = \frac{1-r}{2-2r\lambda} + \frac{(1-2\lambda\pi)r + 2\pi - 1}{2-2r\lambda}(r\lambda)^n. \tag{10}$$

It follows from (5) and (9) that, if x_i is private and $\pi \geq \frac{1-r}{2-r\lambda}$, β_n is decreasing in n and converges to the NE prediction in MS as $n \to \infty$. Similarly, for all $\pi \geq \frac{1-r}{2-2r\lambda}$, with both signals public, $\widetilde{\beta}_n$ is decreasing in n and converges to the NE prediction given by (6) as $n \to \infty$. Therefore, we have proven the following result:

Proposition 1 *In our modified MS model with private signals, all level-k players choose higher actions than the NE prediction for any L0 specification with* $\pi \geq \frac{1-r}{2-r\lambda}$. *In the model with both signals public, all level-k players choose higher actions than the NE prediction for any L0 specification with* $\pi \geq \frac{1-r}{2-2r\lambda}$.

When $\lambda = 1$ the weights on the public and private signals sum to 1. Therefore, it follows from Proposition 1 that when $\lambda = 1$, level-k agents put less weight on the public signal when compared to the theoretical prediction in the MS model. This result has an important policy implication. MS argue that the coordination motive forces players to place too much weight on the public signal relative to the weight that would be chosen by the social planner. Consequently, information is not aggregated efficiently and public disclosure of more information could be detrimental to the social welfare. Our Proposition 1, however, shows that the detrimental effect of public disclosure may be smaller than predicted by MS if agents have limited levels of reasoning. In particular, all level-k players with $\pi \geq 1/2$ put a higher weight on the private signal – and consequently a lower weight on the public signal – than MS predicts.

[2] From the difference equation, we obtain

$$\begin{aligned}\beta_n - \beta_0 &= \frac{1}{2}r\lambda(\beta_{n-1} - \beta_{n-2}) + ... + \frac{1}{2}r\lambda(\beta_1 - \beta_0) + (\beta_1 - \beta_0)\\ &= \left(\left(\frac{1}{2}r\lambda\right)^{n-1} + ... + \frac{1}{2}r\lambda + 1\right)(\beta_1 - \beta_0)\\ &= \frac{1 - \left(\frac{1}{2}r\lambda\right)^n}{1 - \frac{1}{2}r\lambda}(\beta_1 - \beta_0)\end{aligned}$$

Equation (9) then follows by substituting β_0 and β_1 from (8).

2.3 $L0$ Specifications

As Crawford and Iriberri (2007) mention, the specification of $L0$ is the key to the explanatory power of level-k models. The natural candidate, which we call "random $L0$", assumes that $L0$'s actions are uniformly distributed between the two signals, which implies $\pi = 1/2$. Under this assumption, $L0$'s behavior is unsophisticated and serves as a natural focal point for higher level players to start their reasoning (see discussion in Crawford, Costa-Gomes, and Iriberri, 2013). Furthermore, this $L0$ assumption is also directly related to the $L0$-specification in the standard beauty contest. In particular, when $r = 1$ and signals are public, our game is reduced to a beauty contest game and the two $L0$ definitions coincide.

Another $L0$-specification, which we call "non-strategic $L0$" and is related to "truthful $L0$" in Crawford and Iriberri (2007), assumes that the $L0$ type ignores all strategic aspects of the game (guessing other players' actions) and focuses solely on the nonstrategic aspect of the game (guessing the state). Again, under this specification, $\pi = 1/2$. The behavior of the random-$L0$ type and the non-strategic $L0$ type is observationally different: observed choices for random $L0$ types would be uniformly distributed between their two signals while observed choices for the non-strategic $L0$ types are always halfway between their two signals. Nonetheless, because both $L0$-specifications imply $\pi = 1/2$, these two specifications yield the same prediction for the behavior of higher types. Therefore, we will label both $L0$ specifications as "$L050$". Other authors, notably Cornand and Heinemann (2014) and Baeriswyl and Cornand (2014), argue that ignoring the strategic aspect of the game and focusing on guessing the state can be considered $L1$ behavior. This notational difference is primarily a manner of semantics when considering our "non-strategic $L0$" model as it will only change the level of classification of the players.

These two $L0$-specifications are standard in the literature on beauty contests. Our generalized beauty contest game, however, differs from the standard beauty contest game in several aspects. First, subjects have private information in several of our treatments. Second, subjects have to choose actions to match not only the average actions or a fraction of the average actions, but also the fundamentals. Third, but not least, unlike the standard beauty contest in which subjects are restricted to choose a number in an interval, in the MS framework agents can choose any real number as their action. Therefore, it is not clear that our subjects will anchor their beliefs in the same way as in the much simpler beauty contest game.

Later we introduce several different alternatives to the $L050$ specification, including one estimated from the data. However, all $L0$-specifications considered in this paper have the property of $\pi \geq 1/2$ and, therefore, Proposition 1 holds.

3 Experimental Design

The design of all treatments in our study is based on the modified MS framework as described above. This section contains our experimental implementation of the MS framework as well as similarities and differences across treatments.

3.1 Payoff Function and Signals

In all treatments the payoff function of subject i is given by

$$u_i(a_i, a_{-i}) = 2000 - (1-r)(a_i - \theta)^2 - r(a_i - \lambda \overline{a}_{-i})^2, \tag{11}$$

where a_i is the action of subject i, θ is the true state of the world, $\overline{a}_{-i}$ is the average of all other subjects' actions, $\lambda \in [0,1]$ is the weight on $\overline{a}_{-i}$, and $r \in [0,1]$ is the relative importance of matching the weighted average of other players' actions. Note that negative values of $u_i(a_i, a_{-i})$ are possible and so we publicly announce to participants that negative payoffs count as 0. Otherwise, subjects may incur a large loss in a single period of the experiment that would be impossible to recover even if they receive the maximum of 2000 each period afterwards.[3]

To ensure participants' understanding of the payoff structure we took advantage of the fact that each term had a very simple and intuitive interpretation. We began by verbally explaining that three factors determine the payoff: mismatching the underlying state, mismatching $\lambda \overline{a}_{-i}$, and their relative importance r. After these factors are explained, we present the actual mathematical form, explain the meaning of each term, and go through several numerical examples. Finally, during the actual experiment at the end of each period the second and third terms in (11) are calculated and displayed together with a_i, θ, and $\lambda \overline{a}_{-i}$.

The information available to subject i is given by two signals: y and x_i. The signals and state θ are generated prior to the experiment according to the following procedure. For each round t, state θ is generated randomly according to $U[400, 700]$. Given θ, the signals are independently drawn from $N(\theta, 3600)$. Signal y is public and the same for all subjects. Signal x_i can be public or private. In treatments with private x_i different subjects in a group observe different signals. In treatments with public x_i all subjects observe the same signal. Signals and the state are generated in such a way so that each period *all* groups of subjects receive the same signals and the underlying state is the same. If, say, members of group 1 received private signals 105, 72, 41, and 36 then in all other groups there would be a member who receives signal 105, a (different) member with signal 72 and so on.

After the state and signals are generated, we normalize them by subtracting y from each of them so the triple (θ, x_i, y) becomes $(\theta - y, x_i - y, 0)$ and the normalized signal y, therefore, is always 0. Both normalized signals are then displayed on the computer screens and the payoffs are calculated using the normalized state value, $\theta - y$. Note that normalized x-signals and the normalized state could be negative. While the main reason for using the normalization is theoretical and is explained in Section 2 there are additional benefits. First, it simplifies the environment as it is easier to make a decision with signals 0 and 43 than with signals 529 and 572. Second, this guarantees that subjects know that y is indeed a public signal. Third, it makes our setting similar to the standard beauty contest setting.

[3]This can potentially affect the equilibrium prediction because when the maximum of (11) is negative the agent would be indifferent between all actions. However, this happens only when the two signals are very far apart. In our experiment this happened in approximately 0.1% of all observations.

To keep matters simple subjects are not informed about the distributions used for state and signal generation. Subjects are told, however, that the best guess for the state is the average of the two signals (see instructions in the Appendix for the exact wording). Section 2.2 shows that derivations of levels of reasoning and the equilibrium action do not require knowledge of the distribution as long as one knows how to estimate the state given the two signals.

3.2 Treatment and Session Description

There are four aspects in which the MS model differs from the classical beauty contest. First, in the MS model there is private information because agents receive private signals. Second, the goal is divided between guessing λ-average and guessing fundamentals. Third, the action domain is unrestricted. Finally, in the standard MS model $\lambda = 1$ and in the classical beauty contest $\lambda < 1$. The treatments designed for this paper will reflect these differences.

It is convenient to classify each treatment based on the information structure and value of λ. In total, there are three groups of treatments. In the first group signal x_i is private and $\lambda = 1$. We label this group **Pr-A** as the non-zero signal was *pr*ivate and the participants must match the *a*verage action of other subjects. This environment is directly related to the MS model, especially when the domain is unrestricted. In the second group we set $\lambda = 1/2$ so that subjects need to match θ and *one-half* of the average action of the other subjects in their groups. The latter consideration makes the game related to the p-beauty contest with $p = 1/2$, but some information is private. We label the group **Pr-H** where the H represents that individuals must now match one-*h*alf of the average action. In our third group $\lambda = 1/2$ as in Pr-H but both signals are public. As such only two signals are drawn every period, and it is common knowledge that both signals are public. We label this treatment **Pu-H** as the non-zero signal is now a *pu*blic signal and subjects need to match one-*h*alf of the average action. Pu-H is directly related to the beauty contest especially when the domain is restricted. For a fixed information structure and value of λ we will vary values of r from 0.15 to 0.95 with higher r corresponding to a higher coordination motive. Finally, for every information structure and (r, λ) pair there is a treatment in which the strategy choice is bounded by the two signals, like in the beauty contest, and one where it is unrestricted, like in MS.

Given the goal of this paper it is instructive to be more precise regarding the relationship between Pu-H and the beauty contest. First, similar to the beauty contest, Pu-H is a game with perfect information. Second, when the choices are restricted to $[0, x_t]$, it makes Pu-H dominance solvable.[4] Finally, as r approaches 1, the state, θ, becomes irrelevant because the only remaining goal is to match $\lambda\overline{a}_{-i}$. One notable difference from the beauty contest is that here *all* subjects, not just the player who is closest to one-half of the average, are paid. However, the tournament aspect is still retained in that subjects with actions closer to $\lambda\overline{a}_{-i}$ in Pu-H receive higher payoffs

[4] To see this, recall that the best response is given by $a_i = (1-r)\, x_i + r\lambda\overline{a}_{-i}$. Without loss of generality we can assume $x_i > 0$. Because subjects are restricted to choose actions between $[0, x_i]$, we can first eliminate actions outside of the interval $[(1-r)\, x_i/2, (1-r)\, x_i/2 + r\lambda x_i]$. Once we do that, we can further eliminate actions outside of the interval $\left[(1+r\lambda)\,(1-r)\, x_i/2, (1+r\lambda)\,(1-r)\, x_i/2 + r^2\lambda^2 x_i\right]$ and so on. By repeating this procedure we will get a sequence of intervals with length $r^k\lambda^k x_i$, and this sequence will shrink to a point, which is NE.

than those farther away.

To sum up there is a group of treatments that is close to the MS model, Pr-A; a group of treatments that is close to the beauty contest, Pu-H; and a group of treatments that is between the two, Pr-H. Table 1 summarizes the information about the treatments, their mnemonic names, and the number of subjects in each treatment. The total number of subjects is 90.

	y	x_i	λ	Unrestricted Domain	Restricted Domain
Pr-A	0	private	1	19	8
Pr-H	0	private	$\frac{1}{2}$	13	8
Pu-H	0	public	$\frac{1}{2}$	17	25

Table 1: Description of experimental sessions and the number of subjects.

Sessions are based on one of the three treatments described. Within each session the information structure, value of λ, and restrictions on the domain remain the same. The only variation is due to changes in r. Each session consists of 6 phases with 10 rounds in each phase, for a total of 60 rounds.[5] Within each phase the value of r is fixed but r differs across phases. We use six values of r: 0.15, 0.3, 0.5, 0.65, 0.8 and 0.95. For each session we use the following order of r across phases: 0.15, 0.5, 0.8, 0.95, 0.3, and 0.65. Thus, in the first phase (first 10 rounds) subjects make decisions with $r = 0.15$, while in the second phase (rounds 11-20) subjects make decisions with $r = 0.5$ and so on. We start with a low value of r, gradually increase r until phase four, decrease r between the fourth and fifth phases, and then increase it again. The choice of a non-monotone sequence of r's can help us separate the effect of r from the effect of learning. For example, if subjects' behavior is similar in phases with $r = 0.15$ (the first 10 rounds) and $r = 0.3$ (the fifth ten rounds) then it suggests that this behavior is caused by low r and not by lack of subject's experience with the environment.

Overall, our design enables us to vary the standard beauty contest setting in two directions. First, by changing r we vary the strength of the coordination motive. This change is interesting because games in which the importance of coordination varies can capture a wide range of economic applications such as monetary policy (Morris and Shin, 2002), asset pricing (Allen, Morris and Shin, 2003, Bacchetta and Wincoop, 2005), venture capital (Angeletos, Lorenzoni and Pavan, 2007) and political campaigns (Dewan and Myatt, 2007, 2008). While levels of reasoning are well-defined for any value of r, one would expect that subjects will focus less on the actions of others as the coordination motive weakens. If this conjecture is correct it suggests that in games where coordination is less important or its effect is less obvious subjects will be less likely to follow level-k reasoning.

Second, we introduce private information into the game by making the second signal x_i private. Private information is prevalent in many economic applications and therefore it is important to understand how well level-k models can explain the data in settings with private information.

[5] Due to time constraints the Pr-H treatment with restricted domain had only five phases. The last phase (the one with $r = 0.65$) was not conducted.

Indeed, level-k reasonings have been applied to classical settings with private information, such as the winner's curse in common value auctions and overbidding in private value auctions (see Crawford and Iriberri, 2007a). However, the comparison of level-k model performance between the complete and private information settings, both in absolute and relative terms, has not been studied yet.

3.3 Procedures

Sessions were conducted at UNC Charlotte between 2008 and 2010. Subjects were typically undergraduate students, primarily recruited from the business school but not exclusively. Subjects were seated at visually isolated carrels and were forbidden to communicate with other subjects throughout the duration of the experiment. Instructions were read aloud to subjects, and a few minutes were spent discussing how different values of r could impact the subjects' loss from mismatching the state θ (i.e. the term $-(1-r)(a_i-\theta)^2$) and the loss from mismatching the decisions of other investors (i.e. the term $-r(a_i-\lambda\overline{a}_{-i})^2$). To reinforce this distinction in the actual experiment after each round a payoff screen displayed the loss from mismatching each of these two terms as well as the total payoff.

All subjects were divided into four-person groups which were re-assigned in the beginning of each period. In some sessions we had a number of subjects that was not divisible by 4. In those instances we used the following procedure. First, the computer would form as many groups as possible. The remaining subjects would form an incomplete group that was completed by the decisions of a subject(s) from fully completed groups. When relevant the subject(s) chosen from the fully formed group was the one who observed the private signal different from those observed by members of incomplete group. For instance, if the private signals in a fully completed Pr-H were 105, 72, 41, and 36, and the private signals of an incomplete group were 105, 72, and 41, then a decision from a subject who saw a private signal of 36 would be used to complete the incomplete group. Even though the decision of this randomly chosen subject is used for two groups, that subject will only receive the payoff based on the outcome within her fully formed group.

At the beginning of each round, subjects were shown signals and were asked to submit a decision for a_i. Depending on the treatment, subjects were informed that either both signals were public or one was public and the other was private and only observable to that specific subject. When all decisions were submitted, $\overline{a}_{-i}$ and profit were calculated for each agent. At the end of each round subjects were shown a screen containing their own action choice, a_i, the true state, θ, the average opponent action, $\overline{a}_{-i}$ and their payoff for the current round.

A subject's cash payment is determined as follows. At the end of the experiment one of the six phases is randomly chosen. A subject's total payoff during the chosen phase is calculated and converted it into USD by multiplying it by .001. Thus, if a subject earned 10500 during the chosen phase it will become \$10.50. This is in addition to the \$5 show-up fee that all subjects received. The average payment to subjects, including the show-up fee, was \$15 for a 75-90 minute session.

4 Results

In this section we analyze subjects' behavior and study how well it matches NE and level-k (Lk) predictions. Given that NE and Lk actions are linear combinations of a random non-zero signal x and zero signal y, they will vary each period even when the treatment and the value of r are fixed. To make results comparable across periods and treatments we normalize the non-zero signal to 100 and adjust subjects' actions and NE and Lk predictions accordingly. For example, given action a and non-zero signal x, the normalized action is $a_n = 100 \cdot a/x$ so that action $a = x/2$ is normalized to 50 and action $a = x$ is normalized to 100. The interpretation of normalized values is that they represent the percentage weight a particular action or a prediction puts on a non-zero signal.

We begin by comparing subject behavior in the aggregate to NE predictions. Next we consider behavior at the individual level under various definitions of the $L0$ type. We consider three different specifications, including one estimated from the data, to compare level-k performance across different treatments. The comparison is done based on the full data as well as only first round data. Finally, we estimate the frequency of different levels of reasoning in our data, which helps us determine which levels of reasoning are most commonly followed and also serves as a robustness check of our earlier findings.

Throughout this section all the calculations are done using normalized values of actions and levels of reasoning.

4.1 Preliminary Tests

We begin by comparing subjects' behavior with NE predictions. For each treatment and value of r we calculate the average normalized action a_n. Figure 1 shows these a_n and the normalized NE prediction. In all three treatments subjects' actions are higher than NE predicts, meaning that compared to the NE prediction, subjects tend to overweight the non-zero signal (which can be private or public). Earlier we established that overweighting the non-zero signal is consistent with level-k reasoning. Using a non-parametric signed rank test we find that the difference between observed behavior and NE is significant in two-thirds of the treatments. In Pr-A the difference is significant for every $r \neq 0.15$, and in Pu-H for every $r \neq 0.65$. As for Pr-H the difference is significant for $r = 0.5$, $r = 0.8$ and $r = 0.95$.[6] Importantly, the difference is significant in all three treatments for $r = 0.8$ and $r = 0.95$.

Table 2 presents the average absolute deviation of actions from NE (in normalized units). Deviations are quite substantial in all treatments and especially in treatments with private signals. Notably, the observed behavior is closest to NE when $r = 0.65$ regardless of the value of λ and the information structure. The most likely reason is that the $r = 0.65$ phase is last in each session and subjects' learning could bring them closer to NE. As level-k reasoning is usually thought of as the

[6] The test uses individual level data. Because the observations within one experimental session are not independent, as a robustness check we also apply the signed rank test to the session level data. The session-level data is independent and we can apply the signed rank-test to compare the session level average action with NE. Based on the session-level data the signed rank test rejects the null hypothesis that average action is equal to NE in every treatment.

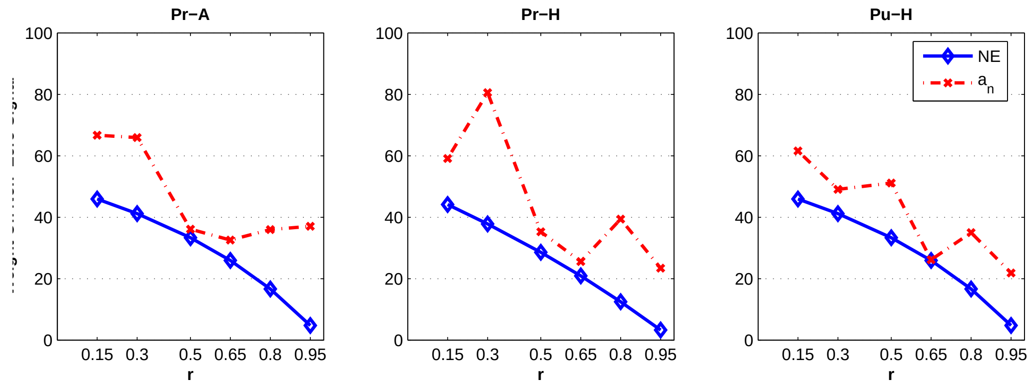

Figure 1: Subjects' behavior and NE in all treatments. On the y-axis is the average weight that subjects put on the non-zero signal which is public in Pu-H and private otherwise. Solid line is NE; and dash-dotted line, a_n, is the average over normalized actions.

framework that describes people's behavior in the beginning of experiments, a better performance of NE in the final stage of experimental sessions is not surprising.

r	0.15	0.30	0.50	0.65	0.80	0.95
Pr-A	67.20	53.86	36.88	30.08	46.23	44.76
Pr-H	58.57	86.97	37.14	23.28	44.22	36.93
Pu-H	30.63	16.46	22.43	10.71	18.91	17.59

Table 2: Average absolute deviation of observed behavior from NE across different treatments and phases. The deviation is calculated based on normalized data with non-zero signal normalized to 100. Higher r means stronger coordination motive.

Result 1: *Consistent with level-k behavior, subjects tend to put a higher weight on the non-zero signal than NE predicts. Overall, NE performs best in the last phase of the study with* $r = 0.65$.

Figure 2 plots a histogram of subjects' choices. We use normalized actions and exclude observations when the chosen actions are outside of the signal range. We also exclude observations when the distance between signals is less than 15 to reduce the sensitivity of normalized actions to small changes in subjects' behavior when the distance between signals is small. In total, 19.72% of all observations are excluded. Figure 2 shows a histogram for all three treatments: Pr-A, Pr-H and Pu-H; and the two most extreme values of r, 0.15 and 0.95. The distance between signals varies from 15 to 226. When $r = 0.95$ we mark the normalized levels of reasoning based on the assumption $L0 = 50$. When $r = 0.15$ all normalized levels of reasoning are close to each other and all belong to the interval $[44, 47]$.

When $r = 0.15$ the normalized action of 50 is the most popular choice in treatments with private signals. In Pu-H the normalized action of 50 is also the most popular choice but the distribution of actions is more dispersed. When $r = 0.95$ the two most prominent spikes occur in the public treatment. The spikes correspond to the normalized $L1$ and $L2$ actions. The modal choice in Pr-A

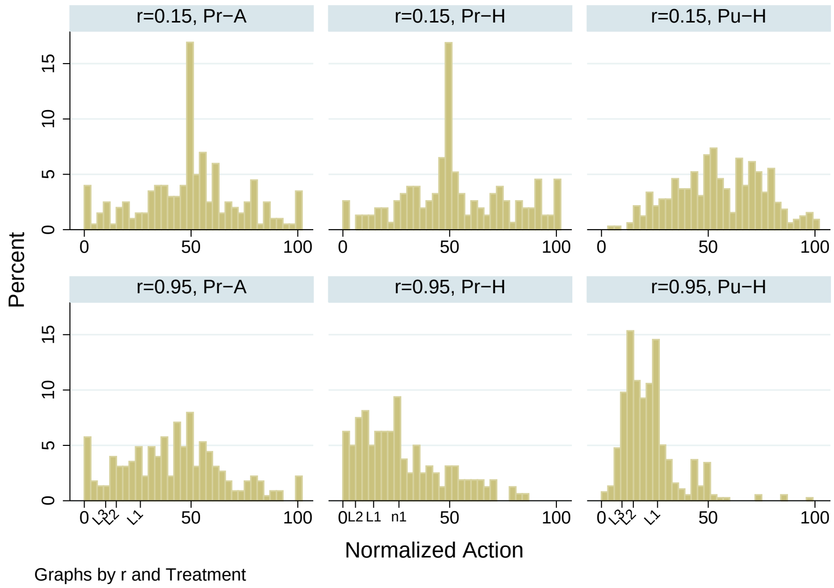

Figure 2: Histogram of normalized actions in the cases of $r = 0.15$ and $r = 0.95$. Actions outside of the signal range are excluded. Actions when the distance between the two signals is less than 15 are excluded. Levels of reasoning are calculated based on $L0 = 50$. Label $n1$ corresponds to the first naive levels of reasoning as defined in Section 5.2. For $r = 0.15$ all levels of reasoning belong to the interval $[44, 47]$ (not marked on the plot).

remains 50, though its frequency is lower than in the $r = 0.15$ case. In Pr-H the modal choice does not correspond to any level of reasoning. In Section 5.2 we introduce "naive" levels of reasoning. The most frequent choice in Pr-H corresponds to the first naive level of reasoning, labeled $n1$.

4.2 Explanatory Power of the Level-k Model

In this section we evaluate and compare performance of the level-k model under three different $L0$ specifications. The comparison is done for the full data as well as for first round data only.

4.2.1 Determining $L0$. Performance of the Level-k Model with Full Data

A key element of the level-k model is the initial level, $L0$, that people use to start their reasoning process. Behavior of the $L0$ type is usually assumed to be as simple and non-strategic as possible. Normally, it is not meant to describe actual choices but rather serves as a starting point in subjects' logic. In section 2.3 we mention that a natural candidate for $L0$ is either a "random $L0$" or a "truthful $L0$," both of which have an expected normalized action choice of 50. Under "random $L0$," level-0 players pick a random action between their signals, while under "truthful $L0$" level-0 players non-strategically attempt to match fundamentals. Both specifications are often used to explain p-beauty contest behavior. Hereafter, we use $L050$ to denote $L0$ specifications with an expected normalized action of 50.

Our setup includes treatments with private signals and/or unrestricted domains, making it richer than the standard beauty contest, and plausible that subjects use some value other than 50 as the starting point in their reasoning.[7] In this section we use the data to estimate the value of $L0$ that maximizes the explanatory power of the level-k model. In a way, by estimating an optimal $L0$, we are determining the upper bound of level-k explanatory power. We then compare the performance of the $L050$ specification with the optimal $L0$ model.

We use the following estimation procedure. For a given value of $L0$ and parameters of the treatment, we calculate $L1$, $L2$, $L3$, and NE.[8] For each treatment we calculate the (expected) value of $L0$ that maximizes what we call the *success rate* of the level-k model. The *success rate* is defined as the number of actions that are within 3 normalized units of at least one level of reasoning divided by the number of all actions in that treatment. To see how the error of 3 normalized units is related to the non-normalized values, we note that the average distance between two signals in our experiments is 72. The 3 normalized units in this case mean that the non-normalized action has to be at most 2.16 away from a given level. For instance, in treatment Pu-H with $r = 0.95$ and $L050$ the normalized value of $L1$ is 26.25 which corresponds to a non-normalized value of 18.9. The action would be classified as consistent with $L1$ if it is between 16.8 and 21.[9]

It is possible, especially for low values of r, that ± 3 normalized unit intervals around Lk overlap. For instance, in treatment Pu-H with $r = 0.15$, when $L0 = 50$ the normalized values of the first three levels of reasoning and NE are approximately 46.25, 45.97, 45.95, and 45.95. Our approach in this case means that any normalized action between 42.95 and 49.25 is considered to be consistent with level-k. In comparison, in the Pu-H treatment with $r = 0.95$ the normalized values are $26.25, 14.97, 9.61$ and 4.76. By our approach any action from the intervals $[1.76, 17.97]$ and $[23.25, 29.25]$ will be consistent with the level-k model.

In general, levels of reasoning are more dispersed and the set of actions consistent with level-k is larger in treatments with higher r. It means the level-k model generates a sharper prediction in treatments with low r's. It also means that using the success rate to compare performance of level-k models across treatments with different r will favor high r phases, where more actions are consistent with level-k. To address the latter, as a benchmark we calculate the success rate of the level-k model in explaining randomly generated data.

The values of optimal $L0$ as well as the success rate generated by optimal $L0$ are presented in Table 3. For a fixed measure of goodness-of-fit – in our case the success rate ± 3 normalized units – this value provides the upper bound for level-k performance.[10]

[7] Burchardi and Penczynski (2014) have detailed data on subjects' reasoning in a timed p-beauty contest game and show that subjects differ in where they start their reasoning.

[8] Levels of reasoning higher than three are very close to the NE, so we simply keep track of the first three levels and the NE itself.

[9] We allow actions to be different from the precise values for level-k action because of subjects' natural propensity to round or choose actions that correspond to "nice" numbers. In the beauty contest with $p = 2/3$ the action of 33 is more likely to be chosen than the correct $33\frac{1}{3}$ value for $L1$.

[10] The optimal $L0$ is not reported in Table 3 for the restricted Pr-H treatment with $r = 0.65$ because time constraints kept us from conducting that phase.

r	Restricted			Unrestricted		
	L0	**rate**	**Implied L1**	**L0**	**rate**	**Implied L1**
			Pr-A			
0.15	132	*0.29*	52.40	135	*0.24*	52.63
0.3	92	*0.25*	48.80	99	*0.24*	49.85
0.5	99	*0.38*	49.75	91	*0.26*	47.75
0.65	105	*0.41*	51.63	98	*0.43*	49.35
0.8	94	*0.41*	47.60	94	*0.32*	47.60
0.95	117	*0.30*	58.08	94	*0.30*	47.15
			Pr-H			
0.15	154	*0.31*	48.28	*174*	*0.21*	49.03
0.3	192	*0.26*	49.40	*116*	*0.15*	43.70
0.5	142	*0.24*	42.75	*176*	*0.25*	47.00
0.65	*	*	*	*182*	*0.31*	47.08
0.8	198	*0.53*	49.60	*187*	*0.23*	47.40
0.95	169	*0.36*	42.64	*83*	*0.31*	22.21
			Pu-H			
0.15	131	*0.22*	52.33	139	*0.17*	52.93
0.3	98	*0.42*	49.70	197	*0.30*	64.55
0.5	102	*0.44*	50.50	188	*0.36*	72.00
0.65	117	*0.52*	55.53	142	*0.47*	63.65
0.8	107	*0.61*	52.80	82	*0.64*	42.80
0.95	50	*0.62*	26.25	53	*0.72*	27.68

Table 3: Columns labeled $L0$ report the value of $L0$ that maximizes the success rate of level-k models in treatments with restricted and unrestricted domains respectively. Columns labeled *rate* show share of observations explained by level-k; columns labeled $L1$ report the first of level reasoning that corresponds to the optimal $L0$.

At first look, the optimal (expected) values of $L0$ seem to be extremely high. In treatments Pr-A and Pu-H they are near 100, meaning that level-1 subjects believe the other players' average actions will be equal to their non-zero signal. This result holds for treatments with unrestricted *and* restricted domains. For treatments with restricted domain such an initial belief seems to be particularly counterintuitive. In Pr-H optimal $L0$ is even higher and is often near 200. That implies that a level-1 subject believes that everyone else picks an action equal to twice their non-zero signal, which is impossible in treatments with restricted domain.

However, these unrealistically high values of optimal $L0$ can be explained if the value of $L1$ implied by optimal $L0$ is calculated (reported in columns 4 and 7). Generally the implied value of $L1$ is very close to 50 and it stays within 3 normalized units of 50 for 14 of the 17 treatments with restricted domain and for 12 of the 18 treatments with unrestricted domain. We interpret this result as follows. The literature typically assumes that the fraction of $L0$ players is negligible and $L0$ is simply the initial step for subjects to structure their reasoning. Consistent with this interpretation our calculations of success rate only include $L1$, $L2$, $L3$, and NE and do not include actions that are close to $L0$. However, given the large number of treatments in which $L1$ is close to 50, our findings suggest that in our experiment there is a non-negligible number of non-strategic subjects playing 50, which would normally be classified as level-0 actions. To put it differently, it is not really the case that in Pr-H level-1 subjects choose 50 because they believe that level-0 subjects will on average choose 200, but rather there is a non-negligible number of non-strategic subjects who consistently choose 50 and then there are higher level subjects who best respond to 50 and so on. There is one notable exception to the result that the implied $L1$ is close to 50. The exception is the Pu-H treatment with $r = 0.95$ – the closest treatment to the beauty contest setting – where the optimal $L0$ is exactly 50 in the case of restricted domain and 53 in the case of unrestricted domain.

Based on this result, we evaluate and compare three "$L0$ specifications". The first specification, $L050$, assumes that the average action of $L0$ is equal to 50 and the share of non-strategic subjects is negligible. In other words the success rate is calculated based on values of $L1$, $L2$, $L3$, and NE. The second specification assumes that there is a non-negligible share of non-strategic subjects who play 50 and higher levels are calculated as usual. We label this specification as $L050_{NS}$, where NS indicates that non-strategic players are included. The success rate for $L050_{NS}$ is calculated based on values of 50, $L1$, $L2$, and NE.[11] The third specification is based on optimal $L0$, that is the $L0$ specification that generates the highest success rate, no matter how unrealistically high it is. Similar to the $L050$ specification the success rate is based on $L1$, $L2$, $L3$, and NE.

The comparison of $L050$ with $L050_{NS}$ shows the share of non-strategic players. The comparison of $L050_{NS}$ and $L050$ with optimal $L0$ indicates how much explanatory power is lost if we use one parsimonious and intuitive value of $L0$ instead of separately estimating it for each treatment. The comparison of $L050_{NS}$ and optimal $L0$ is particularly interesting because Table 3 shows the $L1$ under the optimal $L0$ approach is often close to 50 and, therefore, $L2$ under optimal $L0$ is close to $L1$ under $L050_{NS}$ and so on. If the success rates of the two specifications are close it would be strong evidence in favor of $L050_{NS}$, which has the advantage of being more intuitive and consistent with the approach used in the literature.

The success rates for each of the three specifications are presented in Table 4.[12] Table 4 also presents the success rate of the level-k model in explaining random data generated with $U[0, 100]$. The evidence provides stronger support to level-k behavior if the success rate with experimental data is higher than the success rate with random data.

First, optimal $L0$ performs better than $L050_{NS}$ and $L050$, as it is supposed to by construction. However, the decline in success rate between optimal $L0$ and $L050_{NS}$ is rather modest. The major exceptions occur when $r = 0.95$ (in both Pu-H and Pr-H with unrestricted domain, and in Pr-H with restricted domain) and when r takes relatively low values in Pu-H with unrestricted domain.

[11] We exclude $L3$ from $L050_{NS}$ so that $L050$ and $L050_{NS}$ use the same number of levels and results are comparable.

[12] As a comparison benchmark we use data from Nagel (1995) and Bosch-Domenech et al. (2002). The levels of reasoning are calculated based on $L050$ and the success rate, using our approach, is 56%. The success rate in explaining the $U[0, 100]$ randomly generated behavior is 21%. While this is a useful benchmark, especially for our public treatments with $r = 0.95$, it is worth noting several differences between their settings and ours. First, they use $\lambda = 2/3$ while we use $\lambda = 1/2$. Second, the support of allowed actions is fixed to $[0, 100]$ in their experiments and is varied in our setup. Third, they have $r = 1$ while our maximum r is 0.95. Finally, our subjects play treatments with $r = 0.15$, 0.5, and 0.8 before they play the treatment with $r = 0.95$.

	Optimal L0		$L050_{NS}$			L050		
r	R-ed	Unr-ed	Random	R-ed	Unr-ed	Random	R-ed	Unr-ed
				Pr-A				
0.15	*0.29*	*0.24*	0.10	*0.24*	*0.21*	0.06	*0.09*	*0.08*
0.3	*0.25*	*0.24*	0.13	*0.23*	*0.24*	0.08	*0.09*	*0.08*
0.5	*0.38*	*0.26*	0.16	*0.36*	*0.24*	0.10	*0.13*	*0.08*
0.65	*0.41*	*0.43*	0.20	*0.38*	*0.42*	0.14	*0.23*	*0.23*
0.8	*0.41*	*0.32*	0.24	*0.38*	*0.30*	0.17	*0.23*	*0.19*
0.95	*0.30*	*0.30*	0.24	*0.15*	*0.28*	0.22	*0.14*	*0.17*
				Pr-H				
0.15	*0.31*	*0.21*	0.12	*0.31*	*0.21*	0.06	*0.09*	*0.07*
0.3	*0.26*	*0.15*	0.13	*0.24*	*0.13*	0.07	*0.08*	*0.10*
0.5	*0.24*	*0.25*	0.15	*0.14*	*0.22*	0.09	*0.06*	*0.09*
0.65	*	*0.31*	0.17	*	*0.29*	0.11	*	*0.27*
0.8	*0.53*	*0.23*	0.20	*0.53*	*0.22*	0.13	*0.44*	*0.14*
0.95	*0.36*	*0.31*	0.20	*0.36*	*0.20*	0.15	*0.30*	*0.15*
				Pu-H				
0.15	*0.22*	*0.17*	0.10	*0.18*	*0.13*	0.06	*0.05*	*0.08*
0.3	*0.42*	*0.30*	0.13	*0.41*	*0.20*	0.08	*0.21*	*0.13*
0.5	*0.44*	*0.36*	0.16	*0.44*	*0.27*	0.10	*0.25*	*0.15*
0.65	*0.52*	*0.47*	0.20	*0.48*	*0.38*	0.14	*0.44*	*0.36*
0.8	*0.61*	*0.64*	0.24	*0.60*	*0.58*	0.17	*0.45*	*0.56*
0.95	*0.62*	*0.72*	0.24	*0.60*	*0.52*	0.22	*0.62*	*0.69*

Table 4: Columns 2 and 3 show the success rate under optimal $L0$ for restricted and unrestricted treatments. Columns 4 to 6 (7 to 9) show the success rate of level-k models when $L050_{NS}$ ($L050$) in explaining random data and the data in restricted and unrestricted treatments.

Second, the success rate of the $L050$ model is lower than that of $L050_{NS}$, except for Pu-H with $r = 0.95$. Overall it suggests that the normalized action of 50 serves as an initial point of subjects' reasoning and that in all treatments but Pu-H with $r = 0.95$ there is a non-negligible share of subjects who play a non-strategic action of 50. This result is consistent with recent experimental findings (Agranov, Caplin, and Tergiman, 2013, and Burchardi and Penczynski, 2014) that in level-k experiments non-strategic subjects exist.

Next we compare the success rate with the random benchmark and across treatments. In treatments with public information, level-k based on $L050_{NS}$ is considerably better at explaining subjects' behavior than random data. This result holds for treatments with restricted and unrestricted domain, although the success rates of the former are higher.[13] With private information, the result differs for Pr-A and Pr-H treatments. In Pr-A the success rate of the $L050_{NS}$ model explains experimental data better than random data with the the exception of $r = 0.95$ in restricted treatments. In Pr-H, the $L050_{NS}$ performance is the weakest as there are three incidents with success rates less than or equal to those of the random benchmark.

[13]The treatments with restricted domain have higher success rate partly because with restricted domain subjects cannot choose actions outside of the bounds which are not consistent with any level-k predictions.

The highest success rates, whether in absolute terms or as compared to the random benchmark, are achieved in treatments with public information and with $r = 0.8$ or 0.95. Under these circumstances the coordination motive is the strongest and, in addition, the public information makes it easier for subjects to predict the actions of other players. Overall, 60% or more of the data is consistent with level-k predictions and the level-k models considerably outperform the random benchmark. Also, with the exception of $r = 0.15$, treatments with public information tend to have higher success rates than the corresponding treatments with private information. The data suggest that level-k behavior is more prevalent in treatments with public information than in treatments with private information.

Result 2: *For most values of r, and regardless of the information type and the restrictions (or lack thereof) on domain, we find evidence that the level-k model based on $L050_{NS}$ explains the data better than the level-k model based on* $L050$, *and almost as well as the level-k model based on optimal* $L0$. *Treatment Pu-H with $r = 0.95$ is a notable exception where it is* $L050$ *that performs better than* $L050_{NS}$.

Result 3: *With few exceptions, most notable being the case of $r = 0.95$ in treatments with private information, level-k models are better in explaining observed actions than in explaining randomly generated actions.*

Result 4: *The highest success rates of the level-k model are in treatments with public information and high values of r. Furthermore, in treatments with public information the success rate of the level-k model is higher than in corresponding treatments with private information.*

Result 5: *The success rate of treatments with the restricted domain tends to be higher than in corresponding treatments with unrestricted domains.*

4.2.2 Performance of the level-k Model with the First Round Data

A usual motivation for the level-k approach is that it characterizes how subjects think at the beginning of the game before any learning takes place. In this section we calculate the success rate of the level-k model using first round data and compare it with the success rate based on the full data. As with earlier success rate calculations an action needs to be within 3 normalized units of a particular level of reasoning. Results are presented in Table 5.

First, comparing the success rates based on optimal $L0$, with few exceptions level-k models explain first round data much better than the entire data. This result is true for both restricted and unrestricted domains, though the improvement in restricted domain treatments is particularly notable, with the success rate increasing by a factor of two or more. For instance, in the Pr-H treatment with restricted domain, the success rate increases from 0.36 to 0.75 when $r = 0.95$ and from 0.24 to 0.63 when $r = 0.5$. Overall, with restricted domain, the average success rate is 0.48 in treatment Pr-A, 0.5 in Pr-H, and 0.59 in Pu-H. In treatments with unrestricted domain the increase in success rate is also present, though slightly less pronounced. Furthermore, in the unrestricted Pu-H treatment with $r = 0.95$ the success rate decreases from 0.72 to 0.65.

r	Restricted L1	Restricted rate	Unrestricted L1	Unrestricted rate	Optimal L0 (Full) R-ed	Optimal L0 (Full) Unr-ed	$\mathbf{L050_{NS}}$ (1st Round) R-ed	$\mathbf{L050_{NS}}$ (1st Round) Unr-ed	$\mathbf{L050_{NS}}$ (Full) R-ed	$\mathbf{L050_{NS}}$ (Full) Unr-ed
					Pr-A					
0.15	*51.50*	*0.50*	47.00	0.11	*0.29*	0.24	*0.38*	0.11	*0.24*	0.21
0.3	*56.75*	*0.63*	49.10	0.42	*0.25*	0.24	*0.25*	0.37	*0.23*	0.24
0.5	*46.50*	*0.50*	47.75	0.37	*0.38*	0.26	*0.25*	0.32	*0.36*	0.24
0.65	*62.35*	*0.38*	47.73	0.47	*0.41*	0.43	*0.13*	0.32	*0.38*	0.42
0.8	*47.20*	*0.50*	71.20	0.37	*0.41*	0.32	*0.50*	0.26	*0.38*	0.30
0.95	*41.93*	*0.38*	51.43	0.42	*0.30*	0.30	*0.25*	0.37	*0.15*	0.28
					Pr-H					
0.15	*47.00*	*0.38*	42.54	0.08	*0.31*	0.21	*0.38*	0.00	*0.31*	0.21
0.3	*49.40*	*0.25*	49.40	0.23	*0.26*	0.15	*0.25*	0.23	*0.24*	0.13
0.5	*42.50*	*0.63*	47.75	0.31	*0.24*	0.25	*0.13*	0.31	*0.14*	0.22
0.65	*	*	25.95	0.31	*	0.31	*	0.23	*	0.29
0.8	*48.40*	*0.50*	44.80	0.31	*0.53*	0.23	*0.50*	0.31	*0.53*	0.22
0.95	*45.49*	*0.75*	2.74	0.38	*0.36*	0.31	*0.38*	0.31	*0.36*	0.20
					Pu-H					
0.15	*47.00*	*0.28*	47.00	0.12	*0.22*	0.17	*0.28*	0.12	*0.18*	0.13
0.3	*48.35*	*0.56*	61.25	0.47	*0.42*	0.30	*0.52*	0.12	*0.41*	0.20
0.5	*48.25*	*0.60*	48.25	0.47	*0.44*	0.36	*0.48*	0.35	*0.44*	0.27
0.65	*51.95*	*0.64*	38.30	0.76	*0.52*	0.47	*0.56*	0.47	*0.48*	0.38
0.8	*52.00*	*0.76*	42.80	0.47	*0.61*	0.64	*0.72*	0.29	*0.60*	0.58
0.95	*51.43*	*0.68*	33.38	0.65	*0.62*	0.72	*0.60*	0.47	*0.60*	0.52

Table 5: Columns 2 through 5 show the implied $L1$ and success rate based on the first round data. For brevity we omit the actual value of $L0$. Columns 6 and 7 show the earlier reported success rate based on the full data and optimal $L0$. The last four columns compare success rate under the $L050_{NS}$ assumption between the first round data and full data. Numbers in italics are for treatments with restricted domain.

Second, except for restricted Pu-H, there is considerable volatility in the $L1$ implied by the optimal $L0$. In restricted Pu-H the implied $L1$ is within 3 normalized units of 50 for every value of r. In other treatments, there does not appear to be a parsimonious value of optimal $L0$ that explains the data in every treatment, which is why we calculate the first round success rate under $L050_{NS}$ and compare it with the full data success rate under $L050_{NS}$. Results are presented in the last four columns of Table 5. Using the $L050_{NS}$ specification, the first round success rates are lower, sometimes considerably, than "optimal $L0$" success rates. Furthermore, the comparison of $L050_{NS}$ first round success rates with the full data success rates is less clean. With the exception of restricted Pu-H, all other treatments have an average success rate below 30%. The average success rate for restricted Pu-H is the highest at 53%.

With the exception of restricted Pu-H the evidence is inconclusive. On one hand, using the optimal $L0$ approach the first round data is explained quite well by level-k models. However, it is difficult to explain why subjects would coordinate on 71.20 in unrestricted Pr-A with $r = 0.8$ and then suddenly on 51.43 when $r = 0.95$. The increased success rate under optimal $L0$ could just have a mechanical nature, meaning that first round data have less data points and allowing variation in $L0$ provides an extra degree of freedom to explain a smaller data set.

Result 6: *In the Pu-H treatment with restricted domain the level-k model has a considerably higher success rate in explaining the first round data than in explaining full data. With the exception of* $r = 0.15$ *the success rate is 56% or higher. Furthermore, the success rate of level-k model only slightly decreases when we switch from optimal* $L0$ *to* $L050_{NS}$.

Result 7: *In other treatments the first round success rate also tends to be higher than the full data success rate. However, there is no parsimonious value of implied* $L1$ *that can consistently explain the first round data well. In this sense, for treatments other than restricted Pu-H, we do not have conclusive support to the hypothesis that level-k reasoning is more prevalent in the first round.*

Given our findings thus far and for the sake of uniformity we conduct most of the remaining tests and comparisons using the level-k framework that is based on $L050_{NS}$, which is equivalent to the framework with $L050$ with a non-negligible number of level-0 subjects. Additionally, given Result 2 we report results based on the $L050$ assumption for treatment Pu-H with $r = 0.95$.

4.3 Shares of Level-k Types

In this subsection we estimate the distribution of subjects across different level-k types. This estimation can illustrate which levels of reasoning are more or less commonly used and can also serve as a robustness check to our earlier findings. We first identify shares of level-k subjects who consistently followed a particular level of reasoning. We then use maximum likelihood to estimate the distribution of types based on the success rate measure as defined in the previous section.

4.3.1 Consistency of Subjects' Behavior

The success rate approach may capture a significant number of false positives, because, as reported in Table 4, the level-k model may explain 10% to 24% randomly generated data. As a robustness check, we apply a more stringent criterion where a subject is labeled as following a particular level of reasoning only when the subject's actions are close to a given level for several consecutive rounds. Because each phase consists of ten rounds, we split each phase into two five-round halves: first five and second five. We then calculate the average absolute deviation of subjects' normalized actions from normalized levels of reasoning (NS, $L1$, $L2$ and NE for $L050_{NS}$; and $L1$, $L2$, $L3$ and NE for $L050$) during the first five and the second five rounds of each phase.

We apply a gradual classification of subjects' behavior into different level-k categories. We count a subject as a certain level-k type with higher probability if his normalized action choices are closer to the corresponding normalized value of the k^{th}-level of reasoning, Lk. Formally, let a subject's normalized actions in five consecutive rounds be $\{a_t\}_{t=1}^{5}$, or for brevity $\{a_t\}$. Then we classify a subject as a level-k subject with probability $\Pr(Lk|\{a_t\})$ defined as

$$\Pr(Lk|\{a_t\}) = \begin{cases} 100\% & \text{if} \quad \frac{1}{5}\sum_{t=1}^{5}|a_t - Lk| \leq 3 \\ 50\% & \text{if} \quad 3 < \frac{1}{5}\sum_{t=1}^{5}|a_t - Lk| \leq 5 \\ 20\% & \text{if} \quad 5 < \frac{1}{5}\sum_{t=1}^{5}|a_t - Lk| \leq 7 \\ 0 & \text{if} \quad \frac{1}{5}\sum_{t=1}^{5}|a_t - Lk| > 7 \end{cases}. \tag{12}$$

That is, if the five-round average absolute deviation $\frac{1}{5}\sum_{t=1}^{5}|a_t - Lk|$ is within 3 units, we assign probability 100% to a subject being level-k; if it is between 3 and 5 units, we assign probability 50%; if it is between 5 and 7 units, we assign probability 20%; and if it exceeds 7 units, we conclude the subject does not play level-k.

Applying classification criterion (12) can result in a subject being classified into several levels of reasoning with positive probabilities and the sum of the probabilities may exceed 1. This multiple classification is especially possible for treatments with low values of r and could overestimate the predictive power of the level-k model. To avoid double counting, we scale down the "crude" probabilities specified in (12) proportionally by the sum of the "crude" probabilities. For instance, for the $L050$ specification the final probability that we classify a subject with action choices $\{a_t\}$ as a level-k subject is defined as

$$\frac{\Pr(Lk|\{a_t\})}{\Pr(L1|\{a_t\}) + \Pr(L2|\{a_t\}) + \Pr(L3|\{a_t\}) + \Pr(NE|\{a_t\})}\Pr(Lk|\{a_t\}). \tag{13}$$

The final probability for $L050_{NS}$ is defined similarly except that the term $\Pr(NS|\{a_t\})$ would replace $\Pr(L3|\{a_t\})$ in the denominator of (13).

To illustrate, suppose $\frac{1}{5}\sum_{t=1}^{5}|a_t - Lk| \leq 3$ and $5 < \frac{1}{5}\sum_{t=1}^{5}|a_t - Lk'| \leq 7$, so by (12) we have $\Pr(Lk|\{a_t\}) = 100\%$ and $\Pr(Lk'|\{a_t\}) = 20\%$. In addition, suppose the average absolute deviations from other levels all exceed 7. Then the final probability assigned to Lk is

$$\frac{100\%}{100\% + 20\%} \times 100\% = \frac{5}{6},$$

and the final probability assigned to Lk' is

$$\frac{20\%}{100\% + 20\%} \times 20\% = \frac{1}{30}.$$

The advantage of this approach is two-fold.[14] First, it greatly reduces the chance of false positives as compared to the success rate approach. Simulations show that when actions are generated using $U[0,100]$ then the average weight assigned to a particular level, say, $L1$ is as low as 0.00037. Second, it allows us to identify which levels of reasoning subjects follow.

Table 6 shows the frequencies (sum of probabilities) assigned to each particular level. The summation is across all subjects and for each subject we sum the probabilities in the first five rounds and the second five rounds. For the sake of brevity we pool restricted and unrestricted treatments. Levels of reasoning are calculated based on $L050_{NS}$ and, in addition, for $r = 0.95$ we reports results based on $L050$. Overall level-k behavior can occur for almost any combination of model parameters and regardless of whether the information is public or private. However, subjects' propensity to use level-k varies considerably. It is most commonly used in treatments with public information and high values of r (0.8 and 0.95). In Pr-A the most common level of reasoning is NS and in Pr-H it is $L1$. In Pu-H, especially for high values of r, higher levels of reasoning are common.

4.3.2 Maximum likelihood estimation

Lastly we use the maximum likelihood (ML) approach to estimate the shares of different level-k types in our data. The likelihood function is defined as follows. The estimation parameters are μ_{NS}, μ_{L1}, μ_{L2}, and μ_{NE}, where μ_{Lk} is the share of subjects playing according to Lk, μ_{NE} is the share of subjects playing according to NE and μ_{NS} is the share of non-strategic subjects. Let Lk denote the normalized value of the k^{th} level of reasoning. We assume that a level-k subject plays a normalized action according to the uniform distribution in the interval $[Lk - 3, Lk + 3]$.[15] We also introduce the type *other* with share $\mu_{other} = 1 - \mu_{NS} - \mu_{L1} - \mu_{L2} - \mu_{NE}$. We assume that the type *other* picks a normalized action according to $U\,[0,100]$.

Given these assumptions, conditional on a vector of estimated parameters μ the likelihood of a particular normalized action, a_i, is calculated as follows. The likelihood that action a_i comes from

[14] In addition to the criterion reported in the paper we also examine several alternatives. We perform the calculations using only the first five periods of each phase (as in Crawford and Iriberri, 2007a), pooling the data from all ten periods of each phase, and consider different thresholds. We also use actual values instead of normalized ones. Qualitatively, the results appear to be quite robust. Level-k models perform better in treatments with public signals and in phases with high r. Quantitatively, numbers change depending on whether the criterion is more or less stringent. When it is less stringent, say because of a higher threshold, then all numbers are higher. If it is more stringent, say, because of a lower threshold or because we consider all 10 periods of the phase instead of two five-period intervals, then numbers are lower.

[15] In addition to the uniform distribution we also considered an alternative where level-k subjects choose normalized actions according to a truncated normal distribution on the interval $[Lk-3, Lk+3]$ with standard deviation 2.5. The truncated normal distribution assigns higher likelihood to actions closer to Lk. For instance, the probability that a level-k subject chooses an action between $[Lk-1, Lk+1]$ is 0.4 and not 1/3 as in the case of the uniform distribution. The results of ML estimation based on the truncated normal distribution are qualitatively similar and not reported.

	Pr-A					**Pr-H**					**Pu-H**				
	NS	**L1**	**L2**	**L3**	**NE**	**NS**	**L1**	**L2**	**L3**	**NE**	**NS**	**L1**	**L2**	**L3**	**NE**
0.15	0.7	0.2	0.2	*n/a*	0.2	0.9	0.1	0.1	*n/a*	0.1	2.0	0.4	0.4	*n/a*	0.4
0.3	2.2	-	-	*n/a*	-	0.2	-	-	*n/a*	-	6.9	2.1	1.9	*n/a*	1.6
0.5	3.2	0.2	-	*n/a*	-	1.0	0.2	0.2	*n/a*	0.2	3.6	2.8	2.3	*n/a*	1.7
0.65	4.5	0.7	0.3	*n/a*	0.3	-	2.1	0.8	*n/a*	0.7	0.7	2.2	4.2	*n/a*	6.9
0.8	0.7	0.2	-	*n/a*	-	0.2	2.8	0.8	*n/a*	0.4	6.1	8.0	9.1	*n/a*	3.6
0.95 ($\mathbf{L050_{NS}}$)	2.4	0.2	0.2	*n/a*	-	0.2	1.7	0.9	*n/a*	0.5	2.2	14.1	17.6	*n/a*	2.9
0.95 (L050)	*n/a*	0.2	0.2	-	-	*n/a*	1.7	0.7	0.3	0.3	*n/a*	14.1	12.7	7.9	1.4

Table 6: The frequencies assigned to a particular levels of reasoning. For a given five-round window (either from periods 1 to 5 or 6 to 10) and a subject the probabilities are calculated according to (13). Then probabilities are summed across two five-round windows and subjects. For all r except for 0.95 levels are calculated based on $L050_{NS}$ specification. For $r = 0.95$ calculations are made based on $L050_{NS}$ and $L050$ specifications. Because the $L050_{NS}$ specification is estimated using types $NS, L1, L2$ and NE and $L050$ is estimated using types $L1, L2, L3$ and NE we use label n/a for $L3$ in the $L050_{NS}$ specification and for NS in the $L050$ specification. Label '-' means no subject follows a given level of reasoning. Treatments with restricted and unrestricted domains are pooled together. The number of subjects in Pr-A is 27, the number of subjects in Pr-H is 21, the number of subjects in Pu-H is 42.

type *other* is equal to $\mu_{other} \cdot \frac{1}{100}$, where $1/100$ is the density of $U[0, 100]$. The likelihood that this same action a_i comes from type k is equal to zero if $|a_i - Lk| > 3$ and is equal to $\mu_k \cdot \frac{1}{6}$, otherwise. Here $1/6$ is the density of the uniform distribution over $[Lk - 3, Lk + 3]$. Thus, the likelihood of an action a_i is

$$\mathcal{L}(\mu|a_i) = \mu_{other} \cdot \frac{1}{100} + \sum_{k \in \{NS,1,2,NE\}} \frac{\mu_k}{6} \cdot \mathbf{1}_{|a_i - Lk| \geq 3}.$$

Given all the actions observed in a particular treatment the log-likelihood function is

$$\log \mathcal{L}(\mu|a_1, \ldots, a_n) = \sum_i \log \mathcal{L}(\mu|a_i),$$

which we maximize numerically.[16] In addition, in the case of $r = 0.95$ we estimate the ML under the $L050$ assumption. The difference in the likelihood function is that instead of estimating μ_{NS} we estimate μ_{L3} that corresponds to the share of $L3$ players.

The results of the ML estimation are reported in Table 7. We report point estimates of the share of different levels of reasoning as well as the values of the log-likelihood function and p-values of testing the joint null hypothesis $H0 : \mu_{NS} = \mu_{L1} = \mu_{L2} = \mu_{NE} = 0$ ($H0 : \mu_{L1} = \mu_{L2} = \mu_{L3} = \mu_{NE} = 0$ in the case of $L050$).[17] Under $H0$ the subjects' behavior is random. If $H0$ is rejected it means that the effect of level-k types is statistically significant.

[16]As it is evident from the definition of the likelihood functions we treat all observations as independent. This is a simplifying assumption, because otherwise we would have to make some ad hoc assumptions regarding the dependence of observations within a session. An expected effect of this assumption is when each observation is treated as independent its informativeness is overstated and the estimated standard errors are smaller than the correct ones.

[17]Given the non-negativity restriction on shares of level-k reasoning one cannot use the normal distribution for the

First, Table 7 demonstrates that $H0$ is rejected in every public treatment except the one with unrestricted domain and $r = 0.15$.[18] Second, the point estimates become larger as r increases. Third, the $L050$ specification is better at explaining subjects' behavior in the public treatments with $r = 0.95$, where the share of non-strategic types, as estimated under the $L050_{NS}$ specification, is low and the log-likelihood is higher under $L050$. Finally, in treatments with private information, there is no particularly discernible pattern. In Pr-A the $H0$ is typically rejected. Based on point estimates the rejection comes from a relatively large share of non-strategic types, NS. The unrestricted Pr-H treatments have the lowest number of $H0$ rejections.

Result 8: *Maximum likelihood estimation confirms that level-k reasoning is statistically significant in public treatments for all but one treatment. In Pr-A, the null is rejected in all but one case (three if* $L050$ *is included). The largest point estimates in Pr-A correspond to the non-strategic level,* NS.

hypothesis testing. Instead, one can apply Lemma 21.1 in Gourieroux and Monfort (1995) to show that, under the null hypothesis, the limiting distribution of the likelihood ratio $T = 2(L_n(\hat{\mu}_{NS}, \hat{\mu}_{L1}, \hat{\mu}_{L2}, \hat{\mu}_{NE}) - L_n(0,0,0,0))$ is the same as that of the random variable

$$\zeta = \min_{\{\gamma : \mathcal{I}_0^{1/2}\gamma \geq \vec{0}\}} ||\gamma - \hat{\gamma}||^2,$$

where $\hat{\gamma}$ is the four-dimensional standard normal random variable and $\mathcal{I}_0$ is the information matrix. We simulate the distribution of ζ using 1000 draws from the multivariate standard normal distribution and then for each $\hat{\gamma}$ solve the minimization problem above. As an estimate for the information matrix we use the Hessian calculated at the point estimates. We are grateful to Yuanyuan Wan for his help on this issue.

[18] If instead of estimating $(\mu_{NS}, \mu_{L1}, \mu_{L2}, \mu_{NE})$ we exclude the non-strategic type and estimate $(\mu_{L1}, \mu_{L2}, \mu_{L3}, \mu_{NE})$ then the null is rejected whenever $r > 0.15$ in the public treatments.

	Restricted ($\mathbf{L050_{NS}}$)						L050		Unrestricted ($\mathbf{L050_{NS}}$)						L050	
Pr-A	**0.15**	**0.3**	**0.5**	**0.65**	**0.8**	**0.95**		**0.95**	**0.15**	**0.3**	**0.5**	**0.65**	**0.8**	**0.95**		**0.95**
NS	0.119	0.084	0.191	0.102	0.106	0.000	**L1**	0.000	0.114	0.103	0.104	0.139	0.057	0.076	**L1**	0.002
L1	0.000	0.000	0.053	0.058	0.093	0.001	**L2**	0.000	0.000	0.024	0.013	0.102	0.046	0.007	**L2**	0.013
L2	0.013	0.000	0.000	0.035	0.029	0.000	**L3**	0.000	0.000	0.000	0.000	0.029	0.000	0.018	**L3**	0.000
NE	0.006	0.045	0.000	0.028	0.000	0.000	**NE**	0.000	0.000	0.000	0.000	0.000	0.000	0.000	**NE**	0.000
LL	*-362*	*-365*	*-354*	*-361*	*-351*	*-369*	**LL**	*-368*	*-862*	*-863*	*-864*	*-845*	*-851*	*-868*	**LL**	*-878*
p-value	*0.001*	*0.004*	*0.000*	*0.000*	*0.000*	*0.947*	**p-value**	0.934	*0.000*	*0.000*	*0.000*	*0.000*	*0.005*	*0.002*	**p-value**	*0.693*
Pr-H																
NS	0.189	0.110	0.016	*	0.055	0.013	**L1**	0.110	0.084	0.000	0.076	0.000	0.029	0.000	**L1**	0.037
L1	0.032	0.027	0.000	*	0.235	0.101	**L2**	0.073	0.008	0.011	0.004	0.132	0.000	0.003	**L2**	0.037
L2	0.000	0.000	0.000	*	0.120	0.063	**L3**	0.000	0.005	0.000	0.003	0.000	0.008	0.026	**L3**	0.005
NE	0.000	0.000	0.000	*	0.017	0.000	**NE**	0.019	0.002	0.011	0.002	0.039	0.008	0.000	**NE**	0.000
LL	*-354*	*-363*	*-368*	*	*-332*	*-362*	**LL**	*-363*	*-593*	*-598*	*-594*	*-585*	*-584*	*-598*	**LL**	*-598*
p-value	*0.000*	*0.001*	*0.703*	*	*0.000*	*0.002*	**p-value**	*0.001*	*0.003*	*0.447*	*0.005*	*0.000*	*0.334*	*0.454*	**p-value**	*0.405*
Pu-H																
NS	0.094	0.159	0.143	0.005	0.117	0.029	**L1**	0.256	0.024	0.015	0.072	0.000	0.000	0.000	**L1**	0.123
L1	0.000	0.088	0.076	0.054	0.125	0.253	**L2**	0.189	0.014	0.031	0.000	0.054	0.197	0.111	**L2**	0.227
L2	0.000	0.000	0.105	0.084	0.185	0.201	**L3**	0.093	0.000	0.000	0.065	0.040	0.180	0.228	**L3**	0.208
NE	0.000	0.070	0.000	0.206	0.051	0.000	**NE**	0.000	0.000	0.031	0.000	0.153	0.087	0.057	**NE**	0.038
LL	*-1139*	*-1092*	*-1095*	*-1077*	*-1068*	*-1032*	**LL**	*-1019*	*-782*	*-779*	*-773*	*-755*	*-717*	*-729*	**LL**	*-688*
p-value	*0.000*	*0.000*	*0.000*	*0.000*	*0.000*	*0.000*	**p-value**	*0.000*	*0.196*	*0.012*	*0.000*	*0.000*	*0.000*	*0.000*	**p-value**	*0.000*

Table 7: The ML estimate of shares of level-k types. NS stands for non-strategic type. The reported p-value is for the test $H0: \mu_{NS} = \mu_{L1} = \mu_{L2} = \mu_{NE} = 0$ ($H0: \mu_{L1} = \mu_{L2} = \mu_{L3} = \mu_{NE} = 0$ in the case of $L050$). The number of observations for a given value of r is: 80 in Restricted Pr-A and 190 in unrestricted Pr-A; 80 in Restricted Pr-H and 130 in unrestricted Pr-H; 250 in restricted Pu-H and 170 in unrestricted Pu-H.

5 Alternative Models

In this section we examine the performance of two alternative models to the standard level-k model. We first study the performance of the cognitive hierarchy model (CH) introduced by Camerer et al. (2004). We then study an alternative level-k model where subjects are assumed to update their beliefs in a naïve manner.

5.1 Cognitive Hierarchy

	Restricted				Unrestricted			
r	**L0**	τ	**rate**	**Implied L1**	**L0**	τ	**rate**	**Implied L1**
				Pr-A				
0.15	197	1.0	0.36	57.28	194.00	1.0	0.28	57.05
0.3	188	1.0	0.39	63.20	135.00	1.0	0.31	55.25
0.5	157	2.5	0.49	64.25	138.00	1.5	0.34	59.50
0.65	102	2.5	0.48	50.65	98.00	3.5	0.48	49.35
0.8	82	7.0	0.43	42.80	94.00	2.0	0.32	47.60
0.95	115	1.5	0.36	57.13	94.00	1.5	0.32	47.15
				Pr-H				
0.15	154	0.5	*0.31*	48.28	*174.00*	0.5	*0.21*	49.03
0.3	192	1.0	*0.31*	49.40	*192.00*	1.0	*0.18*	49.40
0.5	142	1.0	*0.34*	42.75	*182.00*	1.5	*0.31*	47.75
0.65	*	*	*	*	*182.00*	2.5	*0.35*	47.08
0.8	127	2.0	*0.60*	35.40	*187.00*	2.5	*0.30*	47.40
0.95	164	4.5	*0.50*	41.45	*83.00*	2.0	*0.35*	22.21
				Pu-H				
0.15	194	1.0	*0.29*	57.05	194.00	1.0	0.23	57.05
0.3	140	1.0	*0.55*	56.00	199.00	2.0	0.37	64.85
0.5	103	2.0	*0.53*	50.75	188.00	2.0	0.46	72.00
0.65	95	3.5	*0.57*	48.38	106.00	3.5	0.51	51.95
0.8	103	10.0	*0.66*	51.20	82.00	4.5	0.70	42.80
0.95	49	4.0	*0.68*	25.78	46.00	5.0	0.79	24.35

Table 8: The table reports the success rate of the CH model for treatments with restricted and unrestricted domains. Value of $L0$ and τ were determined to maximize the success rate within a given treatment. Columns "Implied $L1$" report the level of $L1$ that corresponds to the optimal level of $L0$.

The underlying idea of the CH model is that higher types believe that the population of other players is a mixture of lower types. For example, type $L2$ believes that some players are $L1$ and others are $L0$ and best responds accordingly. Camerer et al. (2004) assume that types are distributed according to the Poisson distribution with parameter τ so that $\Pr(Lk) = f(k) = exp(-\tau)\tau^k/k!$. Each type does not realize that there are players of the same or higher types but it correctly estimates relative proportions of lower types. For example, type $L2$ will believe that the share of $L0$ is $f(0)/(f(0) + f(1))$ and the share of $L1$ is $f(1)/(f(0) + f(1))$.

We calculate the success rate of the CH model in a similar fashion as we did for the level-k

model. We vary $L0$ and τ to find which combination generates the highest success rate. The precision used for calculating the success rate of the CH-model is 3 normalized units as before. Results are reported in Table 8.

Qualitatively speaking, the results of the CH success rate calculations are somewhat similar to level-k results. First, the success rate is the highest in treatments with public information and treatments with restricted domain have higher success rate than treatments with unrestricted domain. Second, the Pu-H treatment with $r = 0.95$ is the only treatment where optimal $L0$ is close to 50.

In general, the success rate of the CH model is higher than that of level-k model, but the value of implied $L1$ is more volatile and is not as uniformly close to 50 as in the previous section.[19]

5.2 Naïve update

	Restricted			Unrestricted		
r	**L0n50**	**L0n50$_{\mathbf{NS}}$**	**L050$_{\mathbf{NS}}$**	**L0n50**	**L0n50$_{\mathbf{NS}}$**	**L050$_{\mathbf{NS}}$**
			Pr-A			
0.15	0.18	0.18	0.24	0.17	0.17	0.21
0.30	0.14	0.14	0.23	0.16	0.16	0.24
0.50	0.24	0.24	0.36	0.16	0.16	0.24
0.65	0.15	0.15	0.38	0.18	0.18	0.42
0.80	0.15	0.15	0.38	0.11	0.11	0.30
0.95	0.04	0.04	0.15	0.13	0.13	0.28
			Pr-H			
0.15	0.11	0.30	0.31	0.07	0.19	0.21
0.30	0.08	0.24	0.24	0.08	0.12	0.13
0.50	0.18	0.25	0.14	0.07	0.20	0.22
0.65	*	*	*	0.24	0.26	0.29
0.80	0.50	0.59	0.53	0.21	0.29	0.22
0.95	0.45	0.38	0.36	0.34	0.34	0.20

Table 9: The success rate level-k models based on naïve update $L0n50$ and $L0n50_{NS}$ as well as the success rate of level-k model based on the correct update, $L050_{NS}$.

The second alternative model considered is the level-k model with naïve update. This model modifies the standard level-k model that is applicable to treatments with private information. Under naïve update, we assume that subjects take their own (private) signal as a proxy for private signals that other subjects observe. The derivation of level-k reasoning with naïve update, therefore, coincides with the derivation of level-k reasoning in treatments with the public information as was done in Section 2.2.

Table 9 shows the success rate of naïve level-k models based on $L0n50$ and $L0n50_{NS}$, where the letter n signifies the naïve update. For comparison, it also shows the success rate under $L050_{NS}$

[19]For the CH model, we also compute model performance using maximum likelihood, and model performance using first round data only. The results are qualitatively similar to the standard level-k models we analyze earlier. For the sake of brevity, we do not report them here. They are available upon request.

with the Bayesian update. Notice that for Pr-A the success rates for $L0n50$ and $L0n50_{NS}$ are the same. This is because in both cases under naïve update all levels of reasonings are simply equal to 50.

Comparing naïve update and Bayesian update, in Pr-A Bayesian update tends to perform better than naïve update, and that holds for every r and regardless of whether the domain is restricted or unrestricted. As for Pr-H, naïve update outperforms Bayesian update in restricted domain treatments and in unrestricted domain treatments when values of r are high. Overall, using naïve update improves the performance of level-k models in explaining Pr-H data but not in explaining Pr-A data.

	Pr-A		**Pr-H**				
	Ln	**NE**	**NS**	**L1n**	**L2n**	**L3n**	**NE**
0.15	1.4	0.1	0.6	0.3	0.3	n/a	0.1
0.3	2.2	-	0.2	-	-	n/a	-
0.5	3.2	-	1.0	0.1	0.1	n/a	0.6
0.65	4.5	0.5	-	-	1.0	n/a	1.3
0.8	0.7	-	0.2	-	2.6	n/a	1.0
0.95 ($\mathbf{L0n50_{NS}}$)	2.4	-	0.2	2.0	2.4	n/a	1.0
0.95 (L0n50)	2.4	-	n/a	2.0	1.9	0.9	0.6

Table 10: The sum of frequencies assigned to particular levels of reasoning. For a given five-round window (either from periods 1 to 5 or 6 to 10) and a subject, the frequencies are calculated according to (13). Then probabilities are summed across two five-round windows and subjects. For all r except for 0.95 levels are calculated based on $L0n50_{NS}$ specification. For $r = 0.95$ calculations are made based on $L0n50_{NS}$ and $L0n50$ specifications. Because $L0n50_{NS}$ specification is estimated using types $NS, L1n, L2n$ and NE and $L0n50$ is estimated using types $L1n, L2n, L3n$ and NE we use label n/a for $L3n$ in the $L0n50_{NS}$ specification and for NS in the $L0n50$ specification. Label '-' means no subject follows a given level of reasoning. Treatments with restricted and unrestricted domains are pooled together. The number of subjects in Pr-A is 27.

Table 10 shows the frequency with which a particular level of reasoning is followed. Frequencies are calculated based on (13) as described in Section 4.3.1. With the $L0n50_{NS}$ (or $L0n50$) specification all levels of reasoning in Pr-A are simply 50 and we label them as Ln. Comparing the results with Table 6, there is virtually no difference between Pr-A and Pr-A with naïve update. When $r = 0.95$ in Pr-H, the frequency of naive levels of reasoning is greater than in Table 6. Nonetheless, the numbers are still considerably smaller than those in Pu-H with high values of r.

6 Concluding Remarks

The goal of this paper is to investigate the strengths and limits of level-k models. To do that, we use a modified MS framework to generalize the classical beauty contest setting. The MS framework allows us to introduce private information and vary the strength of the coordination motive. We find that level-k models, though successful in organizing data in public information treatments with high coordination motive, are much less successful in explaining subject behavior in other treatments. We conjecture that the reason for these results is that when the coordination motive weakens, the behavior of other players becomes less important, so the incentives for subjects to predict it become lower.

The introduction of private information into the model weakens level-k behavior even further because the task of predicting the beliefs and actions of opponents becomes considerably more complex. For example, in the p-beauty contest with $p = 1/2$, $L1$ logic can be summarized in the following simple phrase: people will just pick actions randomly between 0 and 100 so the average action will be 50 and so I should play 25. In contrast, in the setting with private information the same $L1$ logic becomes more complicated because subjects do not know the range from which others are choosing and have to estimate it. Given the increased complexity of level-k reasoning, participants may rely on a different rule of thumb in settings with private information. For example, we find that some subjects may become completely non-strategic and pick their actions simply to match the underlying state, while other subjects may take a naïve approach regarding their belief about how other agents use their information. The identification of the exact rule of thumb subjects used in the experiment is an important research question, and we leave it for future research.

7 Appendix. Instructions for Treatment Pr-H

Welcome to a decision-making study!

Introduction

Thank you for participating in today's study in economic decision-making. These instructions describe the procedures of the study, so please read them carefully. If you have any questions while reading these instructions or at any time during the study, please raise your hand. *At this time I ask that you refrain from talking to any of the other participants.*

General Description

This study consists of 60 rounds, time permitting. In each round all participants (including you) have the role of investors. All participants are divided into groups with 4 investors in each group. The division is random and will be re-done in the beginning of each round. You and the 3 other investors in your group can invest some amount of experimental currency in a particular project. Your task is to decide how much you would like to invest into this project. Returns on your investment will be determined by the amount that you invest (a_{you}) and by the following two factors:

- the project's quality q;
- one-half of the average investments made by others: $\frac{1}{2} \cdot a_{average} = \frac{1}{2} \cdot \frac{a_1 + a_2 + a_3}{3}$;

  Example: Assume that the other three investors in your group invested 150, 200 and 250. The average amount invested by the others is $a_{average} = 200$. One-half of the average then is $\frac{1}{2} \cdot 200 = 100$.

At the time when you make decisions you will **NOT** know either of these two factors. You will not know one half of the average amount invested by others, $\frac{1}{2} \cdot a_{average}$, because other participants are making their decisions at the same time as you. You will not know q because you must make your investment decision before q is revealed. Therefore, you will need to decide how much to invest based on the information that will be made available to you.

Information. Signals.

In the beginning of each round you and all other investors in your group will receive two signals that will provide you with information about the project's quality. Both signals are randomly drawn given the project's quality q. Because signals are randomly drawn it is impossible to precisely predict q given the signals. However, they will give you an idea of a range where q might be. The Table below shows to you how signals should be interpreted.

First, to make calculations easier for you one signal is always set equal to 0. Second, given the two signals that you will see the best guess of q will be simply the average of the two signals. Because of the randomness it is unlikely that q will ever be precisely equal to the average of the two signals. The last two columns in the table give you an idea of how precise your guess is. You see that in two cases out of three, i.e. with probability 2/3, the quality, q, will be at most 40 away from the average and with probability 95% the quality will be at most 80 away from the average.

Signal 1	**Signal 2**	**The best guess of q**	**With prob. 2/3 q will be in**	**With prob. 95% q will be in**
0	s	(0+s)/2	(0+s)/2±40	$(0+s)/2 \pm 80$

Example 1: Assume that you received two signals 0 and 100. Then the best guess of the project quality would be $(0 + 100)/2 = 50$. With probability of 2/3 you can conclude that the project quality will be between $10(= 50 - 40)$ and $90(= 50 + 40)$ and with probability 95% the project quality will be between -30 and 130. In the remaining 5% of the cases the quality will be outside of the $[-30, 130]$ interval.

Guessing one-half of the average

In the previous section we explained how to guess q given the information that you will receive (the two signals). However, your profit will also depend on how well you can guess one-half of the average amount invested by other investors in your group. The decisions of other investors are decisions made by humans and therefore there is no precise theory that will tell you where one-half of the average will be.

Therefore, your best option would be to try to predict how much the other investors are going to invest given their information. Here is what you know and what you don't know about the information available to other investors in your group:

- They receive two signals, just like you do;
- You know the first signal that everyone receives. It is 0. All investors in your group will have 0 as the first signal.
- You do NOT the second signal that they receive. The second signal is a private signal. It means that you cannot see private signals received by other investors. It also means that they cannot see the private signal that you receive.
- You DO know that private signals of other investors are generated in the same way as your private signal. Most importantly that they are also centered around the project's quality q.

Use your knowledge about the information that other investors have to predict how much they will invest. Based on that you can form your guess of one-half of the average investment.

Your Profit and Cash Payments

Your profit will be calculated as follows. In the beginning of each round you will be given 2000 experimental points. From this amount we will deduct points when your action does not match the project's quality. We will also deduct points when your action does not match ***one-half*** of the average investments made by others. Your final profit will be calculated by the following formula:

$$Payoff = 2000 - (1-r)(a_{you} - q)^2 - r\left(a_{you} - \frac{1}{2}a_{average}\right)^2.$$

The first term says that your investment will bring you at most 2000. The second term determines your loss from mismatching the project's quality q. The third term determines your loss from mismatching ***one-half*** of the average investments made by others.

It is possible that the project quality and one-half of the average investment will be two different numbers. In this case parameter r measures the **relative importance** of matching the investments of others versus matching the quality. A lower r means matching the quality is more important. Relative importance will be changed every 10 rounds.

The following two examples are used to illustrate how r impacts your payoff. While you will submit decisions for these two examples they are for illustrative purposes and will not impact your payment.

Example: Let $r = 0.15$ so that is it is *more important to match the quality*. Let quality, q, be 10, and $a_{average}$ be 120. At your computer terminal, please submit an action of 30 now. If your action, a_{you}, is 30 then your loss from mismatching the quality is $(1-0.15)\cdot(30-10)^2 = 340$. Your loss from mismatching one-half of the average investments is $0.15\cdot(30-60)^2 = 135$. You see that your mismatch of the average investment is larger than the mismatch of quality, but your losses from mismatching the quality are higher. Your total profit is $2000 - 340 - 135 = 1525$.

Example: Now assume that $r = 0.8$ so that is it is *more important to match the investments of others*. As before assume that $q = 10$ and $a_{average} = 120$. Thus everything is the same as in the example above except for r. Again, please submit an action of 30 now. Your loss from mismatching the quality is $(1-0.8)\cdot(30-10)^2 = 80$ and your loss from mismatching the average investment is much higher and is equal to $0.8\cdot(30-60)^2 = 720$. Your total profit is $2000 - 80 - 720 = 1200$.

The profit that you made in each round will be converted into cash by the following procedure. The study lasts for 60 rounds. In the end of the study we will openly and randomly choose a sequence of 10 rounds: either from round 1 to round 10, or from round 11 to round 20 and so on. Your cash earnings will be equal to the total profit that you earned during these 10 rounds times 0.001. This is in addition to the $5 that you receive as a show-up fee. For example, if round 21 to 30 is chosen and you earned 10000 during these rounds your cash payoff will be: $10000\cdot0.001+5 = \$15$. If in a particular round you make a negative profit it will count as 0.

Summary

The study consists of 60 rounds, time permitting. In the beginning of each round, the computer will generate the project quality q and randomly determine 3 other investors who will be in your group. Computer will also generate two signals for each participant. The first signal — zero — will be the same among all participants. The second signal will be private. It means that you cannot see the signals received by other investors, and they cannot see the second signal received by you.

Your task is to submit an amount that you would like to invest. After you and all other members of your group enter their decisions, the computer will calculate and display your profit in that particular round. Your profit will be determined based on how well you guessed the project's quality and how well you guessed one-half of the average investment made by others. In the end of the study we will take the profit you made in a randomly chosen sequence of 10 rounds and will convert it into cash payment.

References

[1] Agranov, M., Caplin, A., Tergiman, C. 2013. Naive Play and the Process of Choice in Guessing Games. Working paper, California Institute of Technology.

[2] Allen, F., Morris, S., Shin, H.S. 2006. Beauty Contests and Iterated Expectations. Rev. Financ. Stud. 19 (3), 161-177.

[3] Angeletos, G.-M., Lorenzoni, G., Pavan, A. 2007. Wall Street and Silicon Valley: A Delicate Interaction. NBER Working Paper No. 13475.

[4] Angeletos, G.-M., Pavan, A. 2007. Efficient Use of Information and Social Value of Information. Econometrica 75 (4), 1103–1142.

[5] Bacchetta, P., van Wincoop, E. 2005. A Theory of the Currency Denomination of International Trade. J. Int. Econ. 67 (2), 295-319.

[6] Baeriswyl, R., Cornand, C. 2014. Reducing Overreaction to Central Banks' Disclosures: Theory and Experiment. J. Eur. Econ. Assoc. forthcoming.

[7] Bosch-Domenech, A., Montalvo, J.G., Nagel, R., Satorra, A. 2002. One, Two, (Three), Infinity, . . . : Newspaper and Lab Beauty-Contest Experiments. Amer. Econ. Rev. 92 (5), 1687–1701.

[8] Burchardi, K.B., Penczynski, S.P. 2014. Out of Your Mind: Eliciting Individual Reasoning in One Shot Games. Games Econ. Behav. 84 (1), 39-57.

[9] Camerer, C.F., Ho, T.-H., Chong, J.-K. 2004. A Cognitive Hierarchy Model of Games. Quart. J. Econ. 119 (3), 861–898.

[10] Cornand, C., Heinemann, F. 2014. Measuring agents' Overreaction to Public Information in Games with Strategic Complementarities. Exp. Econ. 17 (1), 61-77.

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

[12] Costa-Gomes, M.A., Craword, V.P. 2006. Cognition and Behavior in Two-Person Guessing Games: An Experimental Study. Amer. Econ. Rev. 96 (5), 1737-1768.

[13] Costa-Gomes, M.A., Crawford, V.P., Iriberri, N. 2009. Comparing Models of Strategic Thinking in Van Huyck, Battalio, and Beil's Coordination Games. J. Eur. Econ. Assoc. 7 (2-3), 365-376.

[14] Crawford, V.P., Iriberri, N. 2007a. Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions? Econometrica 75 (6), 1721–1770.

[15] Crawford, V.P., Iriberri, N. 2007b. Fatal Attraction: Salience, naïvete, and Sophistication in Experimental Hide-and-Seek Games. Amer. Econ. Rev. 97 (5), 1731-1750.

[16] Crawford, V.P., Costa-Gomes, M.A., Iriberri, N. 2013. Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications. J. Econ. Lit. 51 (1), 5-62.

[17] Dewan, T., Myatt, D.P. 2007. Leading the Party: Coordination, Direction and Communication. Am. Polit. Sci. Rev. 101 (4), 825–843.

[18] Dewan, T., Myatt, D.P. 2008. The Qualities of Leadership: Communication, Direction and Obfuscation. Am. Polit. Sci. Rev. 102 (3), 351–368.

[19] Duffy, J., Nagel, R. 1997. On the Robustness of Behaviour in Experimental 'Beauty Contest' Games. Econ. J. 107 (445), 1684-1700.

[20] Gneezy, U. 2005. Step-Level Reasoning and Bidding in Auctions. Manage. Sci. 51 (11), 1633–1642.

[21] Gourieroux, C., Monfort, A. 1995. Statistics and Econometric Models, volume 2. Cambridge University Press, Cambridge.

[22] Ho, T.-H., Camerer, C., Weigelt, K. 1998. Iterated Dominance and Iterated Best Response in Experimental 'p-Beauty Contests'. Amer. Econ. Rev. 88 (4), 947-969.

[23] Morris, S., Shin, H.S. 2002. The Social Value of Public Information. Amer. Econ. Rev. 92 (5), 1521-1534.

[24] Nagel, R. 1995. Unraveling in Guessing Games: An Experimental Study. Amer. Econ. Rev. 85 (5), 1313-1326.

[25] Stahl, D.O., Wilson, P.W. 1995. On Players' Models of Other Players: Theory and Experimental Evidence. Games Econ. Behav. 10 (1), 218-254.

[26] Stahl, D.O., Wilson, P.W. 1994. Experimental Evidence on Players' Models of Other Players. J. Econ. Behav. Org. 25 (3), 309–327.