

Robustness of Level- k Reasoning in Generalized Beauty Contest Games*

Dmitry Shapiro Xianwen Shi Artie Zillante

October 28, 2009

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 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. When we introduce private information or weaken the strength of the coordination motive, the predictive power of level- k models decreases significantly.

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. This is best illustrated by 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. When p is less than one the game can be solved by iterative elimination of strictly dominated strategies, and the unique equilibrium is where 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, 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 Stephen Morris for many helpful comments and constant feedback. We are also grateful to Danielle Catambay, Vincent Crawford and Pei-Yu Lo for helpful comments and suggestions. This research was supported by funds provided by a Childress Klein Research Grant, the Belk College of Business, UNC Charlotte, and the Connaught start-up fund at University of Toronto.

The theory of level- k reasoning, first proposed by Stahl and Wilson (1995) and Nagel (1995) with further extensions by Ho, Camerer and Weigelt (1998), Costa-Gomes, Crawford and Broseta (2001) and Costa-Gomes and Crawford (2006), can be used to rationalize subject behavior in the p -beauty contest. 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, L_0 , corresponds to non-strategic behavior when strategies are selected at random without forming any beliefs about opponents' behavior. In the literature L_0 is typically considered to be a person's model of others rather than an actual person. Level-1 players, L_1 , believe that all their opponents are L_0 and play a best response to this belief. Level-2 players, L_2 , play the best response to the belief that all their opponents are L_1 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 was shown in Nagel (1995) 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 entirely depends on a person's ability to correctly predict the average choice made by others. This explicitly forces individuals to think about 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.

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 which 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 may 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 that we use for our study is a modification of the Morris and Shin (2002) model (hereafter MS) on the social value of public information. In our setting, just as in Morris and Shin, the agents' payoff is determined by two criteria: how well an agent's action matches an unknown state of the world and how well his action 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 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 Morris and Shin 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 in how important it is to predict the average action of other players. Our main findings are as follows. First, in aggregate we find that subjects place less

weight on the public signal than the MS model predicts. We show that this 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.

Next, we compare individual subjects' behavior with level- k predictions and show that treatments with public information *and* a strong coordination motive are the only treatments where the percentage of individuals playing according to level- k models is high. Furthermore, in treatments with public signals and a strong coordination motive, subjects who followed a particular level of reasoning in the beginning would either continue to follow the same level or move to a higher one. We interpret such a consistent pattern as evidence that subjects actually *did* play according to the level- k logic rather than choosing level- k actions by coincidence. In treatments with private information very few subjects would play according to level- k predictions, even when the coordination motive was strong. Moreover, those who followed level- k reasoning in early periods often stopped following any level of reasoning or regressed in their level- k thinking in later periods.

Finally, we perform maximum likelihood (ML) estimation of two level- k models: the standard level- k model and the cognitive hierarchy (CH) model introduced in Camerer, Ho, and Chong (2004). We find that both level- k models perform better as the coordination motive gets stronger. In particular, in treatments with public information and a stronger coordination motive, the log-likelihood is higher and the estimated shares of strategic types (i.e. the types that are not level-0) are higher and more significant. With few exceptions both level- k models predict subject behavior better than Nash equilibrium (NE), especially in treatments with public information. The performance of the standard level- k model and the CH model is comparable, although the CH model has fewer parameters (e.g. in treatments with public information it has only one parameter) and fits the data better than the standard level- k model in terms of log-likelihood.¹

Overall, our analysis highlights the strengths and limitations of level- k models. The Morris and Shin framework used in our study is considerably more complicated than those typically used in 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. That is, when the coordination motive is strong and information is symmetric. At the same time we find that the predictive power of level- k models diminishes quickly as we move away from the classical setting and either weaken the coordination motive or introduce 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. We

¹Gneezy (2005) applies the framework of cognitive hierarchy to analyze first-price and second-price common value auctions with complete information, and finds evidence supporting the CH theory.

find in our experiment that subjects indeed put less weight on the public signal than the theory predicts. This finding is also documented independently in a recent experimental study by Cornand and Heinemann (2009). This implies 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 discuss relevant literature on level- k thinking and the MS model. In Section 3 we provide a theoretical background for our study. In Section 3.1 we develop the theoretical framework for our experiments which is largely based on the MS model. In Section 3.2 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 4 provides details of our experimental design and various treatments. Our experimental results are reported in Section 5. Section 6 then concludes and the experimental instructions are given in the Appendix.

2 Literature Review

Our experimental study contributes to the existing literature on the classical beauty contest beginning with Nagel (1995), who first documents the clear pattern of level- k thinking in subject behavior. Ho, Camerer, and Weigelt (1998); Bosch-Domenech, Montalvo, Nagel, and Satorra (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. However, most of the existing literature focuses on games with complete information. Notable exceptions are Crawford and Iriberri (2007a), who applied level- k reasoning to first- and second-price auctions, and a recent independent work by Cornand and Heinemann (2009), which is closely related to our paper.

Cornand and Heinemann (2009) conduct experiments within the framework of the MS model and find that subjects put less weight on the public signal than the theory predicts. By assuming that all subjects use a common level of reasoning, they find that subject behavior is consistent with the second level of reasoning ($L2$). They further argue that, if all subjects behave according to $L2$, the welfare result in Morris and Shin does not hold: increasing the precision of public information is always beneficial. Their paper and ours share the same theoretical framework and both find that the increased disclosure of public information is less detrimental than the theory predicts. But there are important differences. 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. Moreover, we assume that the population consists of a mixture of different levels, whereas for the purposes of Cornand and Heinemann it was sufficient to assume a common level.

The framework underlying our experimental study is first developed by Morris and Shin (2002) to evaluate the value of public information on social welfare in a coordination environment. Subsequently, Angeletos and Pavan (2007) generalize their analysis of the social value of information by allowing both strategic complementarity and strategic substitutability among agents' actions. The Morris and Shin framework has been applied to many different settings including asset pric-

ing (Allen, Morris and Shin, 2006, Bacchetta and Wincoop, 2005), venture capital (Angeletos, Lorenzoni and Pavan, 2007) and political science (Dewan and Myatt, 2007, 2008).

3 Theoretical Background

This section provides a theoretical background for our study. First, based on the MS model we develop the framework that will be used in our experimental design. Then within this framework we derive level- k predictions.

3.1 Modified Morris-Shin Model

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

$$u_i(a_i, \bar{a}_{-i}, \theta) = C - (1 - r)(a_i - \theta)^2 - r(a_i - \lambda \bar{a}_{-i})^2, \quad (1)$$

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

The payoff function has three terms. The first one 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 \bar{a}_{-i} multiplied 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.

As in MS, before taking actions, agent i will receive two signals about θ . The first signal y with precision α is always public and given by

$$y = \theta + \eta, \quad \eta \sim N(0, 1/\alpha).$$

We allow the second signal x_i with precision β to be either public or private. If it is private, then

$$x_i = \theta + \varepsilon_i, \quad \varepsilon_i \sim N(0, 1/\beta),$$

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

$$x_i = \theta + \varepsilon, \quad \varepsilon \sim N(0, 1/\beta).$$

Again η and ε are independent. After receiving x_i and y , agent i chooses action a_i .

In what follows we will be specifically interested in a special case of this framework when $\alpha = \beta$. When signal x_i is private, we can follow the procedure in MS to show that the unique equilibrium is linear and is given by

$$a_i(y, x_i) = \frac{1-r}{2-\lambda r}x_i + \frac{1+r-\lambda r}{2-\lambda r}y \quad (2)$$

When signal x_i is public, the unique Nash equilibrium is

$$a_i(y, x_i) = \frac{1-r}{2-2\lambda r}x_i + \frac{1+r-2\lambda r}{2-\lambda r}y \quad (3)$$

Our model differs from MS 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.

3.2 Calculating Levels of Reasoning

Within the setting introduced in the previous section we derive actions that correspond to different levels of reasoning. Given that the formulas for level- k actions quickly become large and cumbersome all the calculations in the text are presented for the case when $y = 0$.

It is convenient to introduce a variable μ which measures the relative precision of the two signals. Recall that α denote the precision of the public signal and β denote the precision of the private signal. We define μ as $\beta/(\alpha+\beta)$ so that when precisions are equal, as in our setting, $\mu = 1/2$. Given our definition of μ and the fact that $y = 0$ player i 's updated estimate of the state is $\mathbb{E}_i[\theta] = \mu x_i$.

Player i chooses a_i to maximize (1) and from the first-order condition 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.

L0: $L0$ agent is non-strategic and randomly chooses a number between the two signals.² Then the average action of $L0$ players with signals $(x_i, 0)$ will be $\mathbb{E}_i[\theta] = \frac{1}{2}x_i$, which is equal to μx_i given our notations.

L1: $L1$ expects that other players are $L0$ players. It means that $L1$ player believes that the average action of other players will be equal to *their own* estimated state: $\bar{a}_{-i} = \mathbb{E}_{-i}[\theta] = \mu x_{-i}$. In

²Type $L0$ is often viewed as the starting point of player's analysis of others' actions rather than the type that is actually present in the game. From this perspective it is natural to assume that type $L0$ randomly chooses an action between the two signals. This is also the approach that is commonly used in the literature. An alternative specification, related to truthful $L0$ in Crawford and Iriberri (2007a), is to assume that $L0$ completely ignores all strategic aspects of the game and the information available to other subjects. In our setting, the average of $L0$ choices, which is what determines the $L1$ -behavior, coincides under these two approaches.

the setting when x_i is private $\mathbb{E}_i [x_{-i}] = \mathbb{E}_i [\theta]$ from which it follows that

$$\mathbb{E}_i [\bar{a}_{-i}] = \mathbb{E}_i [\mathbb{E}_{-i} [\theta]] = \mu (\mu x_i) = \mu^2 x_i.$$

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

$$a_{L1}^{pr} = (1 - r) \mu x_i + r \lambda \mu^2 x_i.$$

When x is public $\mathbb{E}_i [x_{-i}] = x_i$ and so

$$a_{L1}^{pu} = (1 - r) \mu x_i + r \lambda \mu x_i.$$

L2: $L2$ expects that other players are $L1$ players. In the setting with private signals $L1$ players choose

$$a_{-i} = (1 - r) \mu x_{-i} + r \lambda \mu^2 x_{-i}.$$

Again using $\mathbb{E}_i [x_{-i}] = \mathbb{E}_i [\theta]$, we obtain

$$\mathbb{E}_i [\bar{a}_{-i}] = (1 - r) \mu^2 x_i + r \lambda \mu^3 x_i,$$

so an $L2$ player will choose the action equal to

$$a_{L2}^{pr} = (1 - r) \mu x_i + r \lambda ((1 - r) \mu^2 + r \lambda \mu^3) x_i.$$

By the same logic in the setting with both signals being public

$$a_{L2}^{pu} = (1 - r) \mu x_i + r \lambda ((1 - r) \mu + r \lambda \mu) x_i.$$

L3: Following the same calculations one can show that when x_i is private an $L3$ player would play

$$a_{L3}^{pr} = [(1 - r) \mu + (1 - r) r \lambda \mu^2 + (1 - r) r^2 \lambda^2 \mu^3 + r^3 \lambda^3 \mu^4] x_i,$$

and when x_i is public they will play

$$a_{L3}^{pu} = [(1 - r) + (1 - r) r \lambda + (1 - r) r^2 \lambda^2 + r^3 \lambda^3] \mu x_i,$$

L ∞ : By induction we can obtain the action choice of players with infinite levels of reasoning. When x_i is private it is equal to

$$a_{NE}^{pr} = \frac{(1 - r)}{1 - \mu \lambda r} \mu x_i,$$

and when x_i is public it is

$$a_{NE}^{pu} = \frac{(1 - r)}{1 - \lambda r} \mu x_i.$$

In both cases it coincides with NE prediction.

Above we derived level- k predictions under the assumption that subjects are capable of correctly estimating signals received by others. For the setting with private signals we also consider an alternative level- k model where players can not perform appropriate Bayesian updating in estimating x_{-i} . We call it a naive level- k model and we assume that subjects are naive in the sense that they simply think that the other players' private signal is exactly the same as their own. Mathematically, this is equivalent to the case when subjects receive two public signals. Notice also that in this case if $\lambda = 1$ then all level- k players will play action μx_i regardless of k . We summarize our findings in Table 1.

	private x_i	public x_i ; or private x_i with naive update
$L0$	μx_i	μx_i
$L1$	$(1 - r + r\lambda\mu)\mu x_i$	$(1 - r + r\lambda)\mu x_i$
$L2$	$\left[(1 - r) \sum_{i=0}^1 (r\lambda\mu)^i + r^2\lambda^2\mu^2 \right] \mu x_i$	$\left[(1 - r) \sum_{i=0}^1 (r\lambda)^i + r^2\lambda^2 \right] \mu x_i$
$L3$	$\left[(1 - r) \sum_{i=0}^2 (r\lambda\mu)^i + r^3\lambda^3\mu^3 \right] \mu x_i$	$\left[(1 - r) \sum_{i=0}^2 (r\lambda)^i + r^3\lambda^3 \right] \mu x_i$
NE	$\frac{(1 - r)}{1 - \mu\lambda r} \mu x_i$	$\frac{(1 - r)}{1 - \lambda r} \mu x_i$

Table 1: Level- k actions in settings with public and private signals. Variable x_i is the non-zero signal.

Proposition 1 *In the setting with private signals level- k players underweight the public signals as compared to the NE.*

Proof. The weight that subjects put on the public signal is 1 minus the weight on the private signal so it is enough to show that level- k players overweight the private signal. If a_{L_n} is the action taken by an L_n player then $a_{L_{n+1}} = a_{L_n} - [r(r\lambda\mu)^n - (r\lambda\mu)^{n+1}]\mu x_i < a_{L_n}$, where the latter follows from the fact that λ and μ are between 0 and 1. Therefore, $\{a_{L_n}\}$ is a decreasing sequence and since it converges to the NE $a_{L_n} > a_{NE}$ for each n . This implies that the weight level- n players put on the private signals is greater than the NE predicts. The proof for the case of naive update is similar. ■

This proposition has an important implication with regards to the MS model. One of the main results of Morris and Shin (2002) is that the coordination motive forces players to place too much weight on the public signal relative to the weight that would be used by the social planner. As a result, information is not aggregated efficiently and public disclosure of more information could be detrimental to the social welfare. However, as it follows from Proposition 1 the detrimental effect of public disclosure may be less than predicted by theory if agents are not fully rational. Specifically, level- k players, whether naive or not, put higher weight on the private signal — and consequently lower weight on the public signals — than NE predicts.

4 Experimental Design

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

4.1 Payoff Function and Signals

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

$$u_i(a_i, a_{-i}) = \max\{2000 - (1 - r)(a_i - \theta)^2 - r(a_i - \lambda\bar{a}_{-i})^2, 0\}, \quad (4)$$

where a_i is the action of subject i , θ is the true state of the world, \bar{a}_{-i} is the average of all other subjects' actions, $\lambda \in [0, 1]$ is the weight put on \bar{a}_{-i} , and $r \in [0, 1]$ is the relative importance of matching the weighted average of other investors' actions.

Since the payoff function given by (4) is more complicated than those in typical laboratory experiments, every effort was made to ensure that subjects understood the payoff structure. First, expression (4) was presented in a simpler way

$$u_i(a_i, a_{-i}) = 2000 - (1 - r)(a_i - \theta)^2 - r(a_i - \lambda\bar{a}_{-i})^2. \quad (5)$$

The fact that payoffs could not be negative was explained verbally in the instructions. Second, we took advantage of the fact that each term had a very simple and intuitive interpretation. We started by verbally explaining that there are three factors that will determine the payoff: mismatching the underlying state, mismatching $\lambda\bar{a}_{-i}$, and their relative importance r . After this was understood, we presented the actual mathematical form, explained the meaning of each term, and went through several numerical examples. Finally, during the actual experiment at the end of each period the second and third terms in (5) were calculated and displayed together with a_i , θ , and $\lambda\bar{a}_{-i}$. This proved to be very helpful for participants since it highlighted how each term in (5) affects the payoff.

Notice that we bound the period profit away from 0. Otherwise, subjects may incur a large loss in a single period of the experiment that would be impossible to recover even if subjects receive the maximum of 2000 each period afterwards.³

All treatments had a similar structure and differed only in two aspects: parameter values and whether signal x_i was public or private. State θ and signals, whether public or private, were generated prior to the experiment. For each round t , state θ is generated randomly according to a uniform distribution on $[400, 700]$. Given θ , the signals are independently drawn from a normal distribution $N(\theta, 3600)$. Signal y is public and is the same for all subjects. Signal x_i can be public or private. In treatments when it is private different subjects in a group observe different signals.

³This can potentially affect the equilibrium prediction. When the maximum of (5) is negative the agent would be indifferent between all actions. One can show that this happens only when the two signals are very far apart. In our experiment this happened exactly 4 times which is approximately 0.1% of all observations.

When it is public all subjects observe the same signal. Signals and the state were generated in such a way so that each period *all* groups of subjects received the same signals and the underlying state was 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 received signal 105, a (different) member with signal 72 and so on.

Subjects were not informed about the distributions used for state and signal generation. This was done for two reasons. First, we did not want to overwhelm participants with mathematical details. Second, knowing the support of θ would make the Bayesian update more complicated, especially, when signals are outside of the support. Instead, we verbally explained how to interpret each signal and how to use them to infer the value of the state. The precise wording is given in the instructions in the Appendix.

After the state and signals are generated, we normalize them by subtracting y from each of them so that triple (θ, x_i, y) becomes $(\theta - y, x_i - y, 0)$ and normalized signal y , therefore, is always 0. Both normalized signals are then displayed on the computer screens and the payoffs are calculated using normalized state value, $\theta - y$. Note that normalized x -signals and the normalized state could be negative. The reason for using the normalization is three-fold: first, it substantially simplifies the environment. It is easier to make a decision with signals 0 and 43 rather than with signals 529 and 572. Second, this guarantees that subjects know that y was indeed a public signal. Third, it makes our setting similar to the standard beauty contest setting.

4.2 Treatment and Session Description

Four treatments were designed for this study. In the first treatment signal x_i was private and $\lambda = 1$. We label this treatment Pr-A as the non-zero signal was private and the participants must match the average action of other investors. In the second treatment we set $\lambda = 1/2$ so that subjects need to match θ and *one-half* of the average action of their opponents. The latter consideration makes the game related to the p -beauty contest with $p = 1/2$. In this treatment each participant still receives a private signal. We label the treatment Pr-H where the H represents that individuals must now match one-half of the average action. In our third treatment $\lambda = 1/2$ 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 public signal and subjects need to match one-half of the average action. Our final treatment is identical to the Pu-H treatment except that participants are *required* to choose an action between the two signals. This treatment is labeled PuR-H as it is identical to Pu-H except that the domain of actions is restricted.

Among the treatments we conduct PuR-H is the closest to the beauty contest setting. First, subjects' choices are restricted to $[0, x_t]$ which makes the game dominance solvable.⁴ Second, all

⁴To see this, recall that the best response is given by $a_i = (1 - r)x_i + r\lambda\bar{a}_{-i}$. 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 $[(1 + r\lambda)(1 - r)x_i/2, (1 + r\lambda)(1 - r)x_i/2 + r^2\lambda^2 x_i]$ and so on. By repeating this procedure we will get a

information is common knowledge and, finally, when r gets closer to 1 the subjects' goal becomes to match $\frac{1}{2}\bar{a}_{-i}$. One notable difference from the BC model is that here *all* subjects, not just the player who is closest to one-half of the average, are paid. However, the tournament aspect of the BC is still retained in that subjects with actions closer to $\lambda\bar{a}_{-i}$ in PuR-H receive higher payoffs than those farther away. Table 2 provides a summary of the treatments, their mnemonic names, and the number of subjects in each treatment.

Treatment	y	x_i	λ	Restricted Domain?	# of subjects
Pr-A	0	private	1	No	19
Pr-H	0	private	$\frac{1}{2}$	No	13
Pu-H	0	public	$\frac{1}{2}$	No	17
PuR-H	0	public	$\frac{1}{2}$	Yes	11

Table 2: Description of experimental sessions.

Sessions are based on one of the four treatments described above, and each consists of 6 phases with 10 rounds in each phase, for a total of 60 rounds. Within each phase the value of r is fixed but r is different 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 finally 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. Note that 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 the following two directions. First, by changing r we vary the strength of the coordination motive. This 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 becomes weaker. If this is correct it would suggest 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. Indeed, level- k reasonings have been applied to classical settings with private information, such

sequence of intervals with length $r^k \lambda^k x_i$, and this sequence will shrink to a point, which is NE.

as the winner’s curse in common value auctions and overbidding in private value auctions (see Crawford and Iriberry, 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.

4.3 Procedures

Sessions were conducted at UNC Charlotte between March 2008-March 2009. 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 was 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 $-r(a_i - \lambda\bar{a}_{-i})^2$). 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 signals or one was a public signal and the other was a private signal only observable to that specific subject. When all decisions were submitted, \bar{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, \bar{a}_{-i} and their payoff for the current round.

Subjects’ 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.

5 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 , that is the session phase, are fixed. To make results comparable across periods and treatments we normalize the non-zero signal to be 100 and adjust subjects' actions as well as 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.

5.1 NE and Subjects' Behavior




Figure 1: Subjects' behavior and NE in all four treatments. On the y -axis is the average weight that subjects put on the zero signal. Solid line is NE; dotted line, a_n , is the average weight over all actions; dash-dotted line, a_{15_n} , is the average over actions that were taken when the distance between the signals is at least 15.

First, we compare subjects' behavior with NE predictions. For each treatment and each r we calculate the average normalized action a_n and plot it on Figure 1 together with the normalized NE prediction. One caveat with normalization is that when x is close to zero it becomes too sensitive to small changes in behavior. For example, if $x = 4$ then two close actions 1 and 2 will be normalized

as 25 and 50. To address this problem we calculate and plot an additional variable $a15_n$ which is the average of all normalized actions that were taken by subjects observing $|x| \geq 15$.

As we see from Figure 1, in all four treatments subjects' actions tend to be higher than NE predicts. In other words, subjects tend to overweight the non-zero signal which is private in Pr-A and Pr-H and public in Pu-H and PuR-H. In a setting similar to our Pr-A treatment Cornand and Heinemann (2009) also observed the overweighting of private signals by subjects.

r	0.15	0.30	0.50	0.65	0.80	0.95
Pr-A	41.21	31.77	40.02	27.05	32.66	39.13
Pr-H	49.12	35.74	45.12	21.68	38.00	34.02
Pu-H	23.03	22.68	25.91	9.45	14.99	12.86
PuR-H	21.54	14.54	20.80	8.96	17.14	17.02

Table 3: 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.

Another interesting observations is that for intermediate values of r NE is matched very closely. It happens in all treatments when $r = 0.65$ and in treatments with private signals when $r = 0.5$. To test whether this was a result of aggregation or if subjects did play NE in those phases we calculated the average absolute deviation of each action from NE (in normalized units). The results are displayed in Table 3.⁵ It can be immediately seen that in each treatment the phase $r = 0.5$ had the second-largest or largest average deviation from NE. In contrast, the $r = 0.65$ phase has consistently the best performance as compared with other phases within the same treatment. Given that $r = 0.65$ was the last phase this might indicate that with experience subjects begin to play closer to the NE prediction.

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

5.2 Level- k Models and Subjects' Behavior

As we showed in Section 3 level- k players will overweight non-zero signal regardless of k . Therefore, in aggregate subjects behavior is consistent with level- k predictions. In this section we study how well level- k reasoning performs on individual level.

A few things need to be taken into account. First, level- k behavior is often considered as a way to describe subjects' initial responses in which case it should be more pronounced in the beginning of the phase. Second, subjects' behavior can change as the experiment proceeds. They can switch from one level to another one, presumably a more sophisticated one, or they could experiment in the beginning of the phase and then converge to a particular level. Finally, even those subjects

⁵As before we calculated the average absolute deviations for the entire sample and the average deviations excluding observations with $|x| < 15$, but only the latter is reported in the paper. Results are qualitatively similar. The only major difference is an outlier in treatment Pr-H with $r = 0.3$, where the average deviation for the entire sample was 125.22. The presence of the outlier can be also seen on Figure 1.

who did go through some levels of reasoning are likely to be biased towards integers, particularly those ending with 0 and 5. All this suggests that a simple search for subjects whose actions are equal to Lk predictions every period is likely to overlook many instances of level- k behavior.

r	0.15	0.30	0.50	0.65	0.80	0.95
Pr-A	5.26	2.63	13.16	21.05	7.89	21.05
Pr-H	7.69	0.00	11.54	34.62	11.54	19.23
Pu-H	5.88	17.65	17.65	52.94	64.71	91.18
PuR-H	4.55	36.36	18.18	63.64	63.64	86.36

Table 4: Percentage of subjects who closely followed one of the levels of reasoning. We say that a subject’s behavior closely followed one of the levels of reasoning during a given half of the phase if two conditions hold. First, the average deviation of subjects’ choices from this level was the smallest when compared to other levels. Second, the average deviation was less than 10 (in normalized units).

To address these concerns we used the following approach to elicit level- k behavior. We divided the ten periods of each phase into two halves: the first five periods and second five periods. Within each group of five periods for each subject we calculated average absolute deviations of subjects’ normalized actions from normalized level- k and NE predictions. We say that *a subject’s behavior closely followed one of the levels of reasoning* during a given half of the phase if two conditions hold. First, the average deviation of subjects’ choices from this level was the smallest as compared to other levels. Second, the average deviation was less than 10 (in normalized units). The main advantage of this approach over other alternatives⁶ is that it allows us to study separately subjects’ behavior in the beginning and in the end of each phase. In particular, it enables us to evaluate how well level- k models capture subjects’ initial responses and how subjects’ behavior evolves with time.

Table 4 shows the percentage of subjects who closely followed one of the levels of reasoning during either half of a particular phase. Several things can be noticed. First, for any given r in treatments with two public signals the success rate of level- k models is higher. The only exception is $r = 0.15$ but in either case the success rate for $r = 0.15$ is extremely low in all four treatments. Second, the highest success rate occurs in treatments with two public signals *and* $r = 0.95$ where more than 85% of subjects followed some level of reasoning. This suggests that level- k models perform best when there is no private information and the coordination motive is the strongest. The third result, while in some sense being a corollary of the previous two, is worth mentioning separately. In treatments with private signals even when $r = 0.95$ the success rate of level- k models

⁶In addition to the criterion used in calculating Table 4 we tried the following alternatives. We performed the calculations using only the first five periods of each phase (as in Crawford and Iriberry, 2007a) as well as pooling the data from all *all* ten periods of each phase. In addition to having the threshold of 10 normalized units we considered thresholds of 5 and 15 normalized units. We also used actual values instead of normalized ones and tried thresholds of 5, 10 and 15. In all these cases, the qualitative picture does not change. Level- k models perform better in treatments with public signals and in phases with high r . Quantitatively, numbers change as compared to Table 4 depending on whether the criterion is more or less favorable to level- k reasoning. If it is more favorable, say because of a higher threshold, then all frequencies are higher. If it is less favorable, say, because of a lower threshold or because we consider all 10 periods of the phase instead of two five-period intervals, then all frequencies are lower.




Figure 2: Individual behavior of first 10 subjects in treatment Pr-A (the first two rows) and in treatment Pr-H (the last two rows). The value of r is 0.95. The ten phase periods are on the x -axis. Actions as well as level- k predictions are on the y -axis. Dots are actions; the lowest (dotted) line is NE; solid line is $L1n$; dashed line on the top is $L1$; dashed-dotted line in the middle is $L2$. To increase the scale of images we use absolute values so that all levels and actions are positive.

is relatively low and, in particular, is much lower than in treatments with public signals.

To visualize the difference in subject behavior between treatments with private and public signals, we take phase $r = 0.95$ of all four treatments and plot individual choices of the first 10 subjects together with level- k predictions. Figure 2 plots decisions of the first 10 subjects from the two treatments with private signals and Figure 3 shows decisions of the first 10 subjects from the two treatments with public signals.

Looking at subjects' choices in treatments with private signals we see that the majority of subjects do not consistently follow a particular level of reasoning. For example, in treatment Pr-H subjects 1 and 2 were typically playing higher than any level of reasoning, subjects 3, 4, 5, 6, 7 and 8 were not playing as high, however, their choices did not follow any particular level. Only subject 10 more or less consistently played $L1n$, which is $L1$ with naive update, throughout the entire phase.

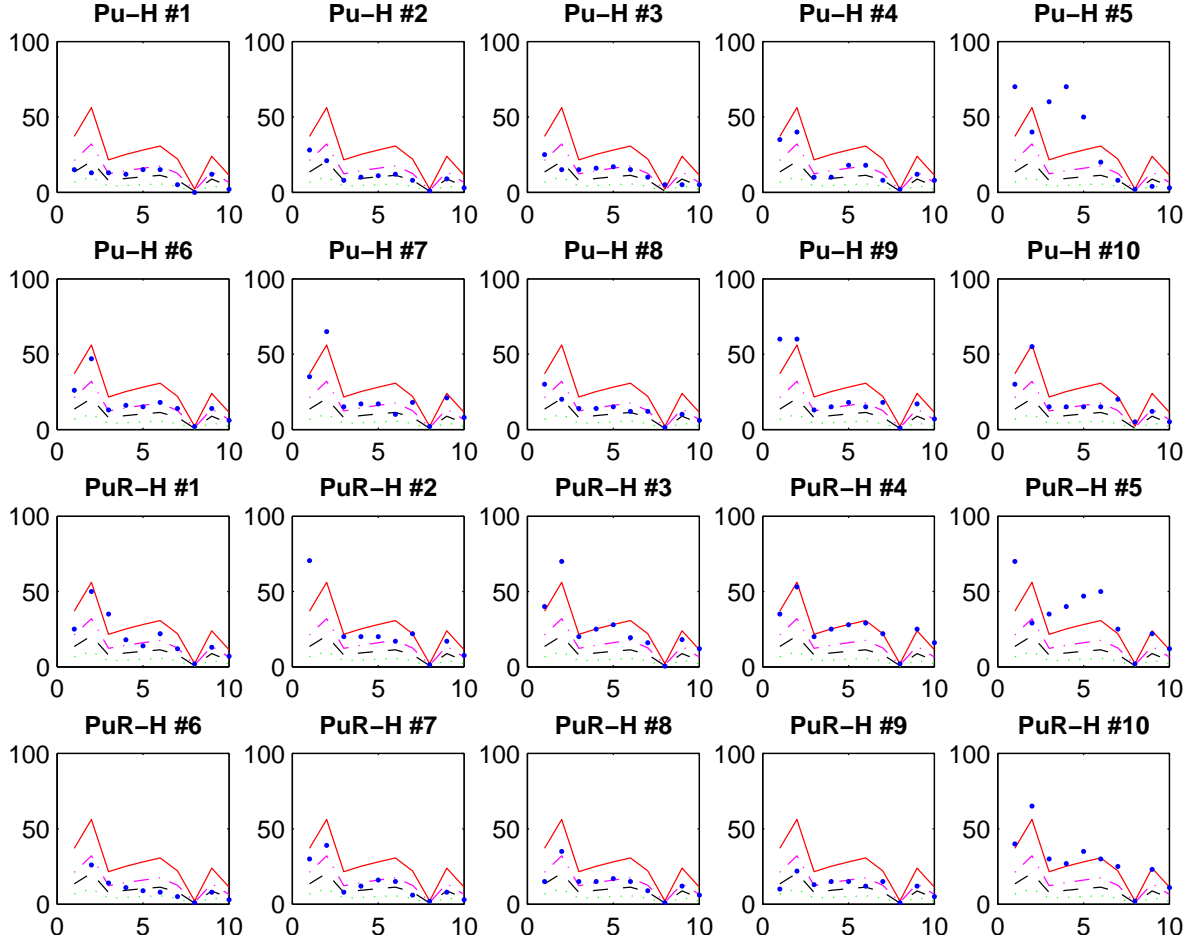


Figure 3: Individual behavior of first 10 subjects in treatment Pu-H (the first two rows) and in treatment PuR-H (the last two rows). The value of r is 0.95. The ten phase periods are on the x -axis. Actions as well as level- k predictions are on the y -axis. Dots are actions. Lines for level- k predictions are ordered from top to bottom. The solid line on the top is $L1$. The dash-dotted line below it is $L2$, the dashed line below it is $L3$ and the dashed line on the bottom is NE . To increase the scale of images we use absolute values so that all levels and actions are positive.

In contrast, in treatments with public signals many subjects closely and consistently followed some level. For example, in treatment Pu-H subject 2 followed $L3$, subjects 3, 6, 8, 9 and 10 followed $L2$. In treatment PuR-H subjects 3 and 4 followed $L1$, subject 6 uses $L2$ in the beginning but switches to $L3$ in the end, subjects 8 and 9 follow $L2$. Overall, these two figures visualize what we already observed in Table 4, i.e. in treatments with public signals the performance of level- k reasoning is much better than in treatments with private signals.

Result 2: *In treatments with private signals and in phases with low r only a few subjects followed levels of reasoning. In treatments with public signals and high values of r level- k models did the best with the majority of subjects following some level of reasoning.*

The next question is how subjects' behavior evolves over time. First, we study how the predictive

power of the level- k model changes within the phase. Given that level- k reasoning is usually viewed as the way to describe subjects' initial responses we measure the performance of level- k models separately in the beginning and in the end of each phase. In Table 5 we calculate percentages of level- k subjects in the first five periods of each phase as well as percentages of level- k subjects in the last five periods of each phase. The criterion for attributing subjects' behavior to a particular level is the same as before. The average absolute deviation from this level prediction should be the smallest compared to other levels and should not be more than 10 normalized units.

r	First 5 rounds						Second 5 rounds					
	0.15	0.30	0.50	0.65	0.80	0.95	0.15	0.30	0.50	0.65	0.80	0.95
Pr-A	5.26	5.26	10.53	31.58	0.00	26.32	5.26	0.00	15.79	10.53	15.79	15.79
Pr-H	0.00	0.00	7.69	46.15	15.38	7.69	15.38	0.00	15.38	23.08	7.69	30.77
Pu-H	11.76	17.65	11.76	58.82	41.18	94.12	0.00	17.65	23.53	47.06	88.24	88.24
PuR-H	9.09	45.45	9.09	63.64	54.55	72.73	0.00	27.27	27.27	63.64	72.73	100.00

Table 5: Percentage of subjects who followed closely one of the levels of reasoning in the beginning of the phase (the left table) and in the end of the phase (the right table). Frequencies are computed similarly to Table 4.

Comparing the left and right parts of Table 5 we see that there is no clear pattern with regards to whether subjects are likely to exhibit level- k behavior in the beginning or in the end of the phase. In 10 phases out of 24 level- k behavior was more common in the beginning, in 4 phases frequencies were equal and in the remaining 10 phases it was more common in the end. Keeping the information structure fixed does not reveal any clear pattern either. In treatments with public information as well as in treatments with private information level- k behavior was as common in the beginning as it was in the end.

Subject #	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
$r=0.8$																	
1st half						L1		L2					L1	L1	L3	L1	L1
2nd half	L2	L2		L1	NE	L1	L1	L2	L1	L2		L3	L1	L1	NE	L1	L3
$r=0.95$																	
1st half	L2	L3	L2	L2		L2	L2	L2	L2	L2	L1	L3	L1	L1	L3	L1	L3
2nd half	NE	L3		L2	L3	L2	L1	L2	L2		L1	NE	L2	L1	L3	L1	L3

Table 6: Subjects' behavior in Pu-H treatment in phases with $r = 0.8$ and $r = 0.95$.

Next, we turn our attention to the $r = 0.8$ and $r = 0.95$ phases of the two treatments with public signals which is where level- k behavior was the most prominent. From Table 5 we see that in these phases level- k reasoning was more pronounced during the second half rather than during the initial five periods. Looking at the evolution of subjects' choices over time we have the following results. Among those subjects who followed some level of reasoning in *both* halves of a phase there were 22 cases (out of 34) when subjects stayed with the same level in both halves, 11 cases when subjects switched to a higher level, and 1 case when a subject switched to lower level (subject 7, phase $r = 0.95$, treatment Pu-H). This is consistent with Nagel (1995) who also found that subjects

tend to adhere to the same level of reasoning throughout the entire study.⁷ For those subjects who followed some level in one half of a phase, only 3 subjects did so in the first half while 15 subjects did so in the second half. In other words, it was more likely for subjects to switch to level- k behavior rather than to abandon it. Similarly, looking at subjects' behavior between phases we see that those subjects who followed some level in the $r = 0.8$ phase would typically continue doing so, though perhaps using a different level, in the next ($r = 0.95$) phase. Table 6 summarizes the behavior of subjects in Pu-H treatment. For brevity sake we omit the similar table for the PuR-H treatment.

Subject #	Pr-A							Pr-H					
	2	6	8	12	14	17	18	24	28	29	30	32	
r=0.8													
1st half								L2n		L1			
2nd half	NE	NE			Ln							L2n	
r=0.95													
1st half	L2		L1	L1		Ln	Ln						L2
2nd half				Ln	Ln	Ln		L3n		L1n	L1n	L3n	

Table 7: Subjects' behavior in treatments with private signals. Only those subjects whose behavior could be attributed to any level are shown.

In treatments with private signals, the picture is considerably less clear. Table 7 shows all participants in treatments with private signals whose actions could be attributed to a particular level in at least one half-phase. The difference between Tables 6 and 7 is immediate. First, as we should expect given our previous results, only a few subjects followed some level of reasoning. Second, among those who did only three subjects followed some level during the entire phase. For the remaining subjects, their actions can often be attributed to some level in the second half (7 instances) rather than in the first half (3 instances). Third, there is much less consistency in the level- k behavior between phases. In particular, only 2 subjects (14 and 32) followed level- k reasoning in both phases.

Result 3: *We do not observe that subjects are more likely to follow level- k reasoning in the beginning of the phase as compared to the end, or vice versa. However, in phases where level- k behavior was most prominent, level- k reasoning was more common in the end.*

Result 4: *In treatments with public signals and high r subjects were more likely to switch to Lk -behavior and within Lk behavior they were more likely to switch to higher levels. In treatments*

⁷Nagel (1995) finds that subjects' choices decreased over time which one may interpret as the evidence that subjects learn to play with higher levels of reasoning. Nagel argues that this interpretation is false. The declining pattern of choices is not because subjects learned to play with higher levels of reasoning but because subjects adjusted downwards their beliefs about the average action given the outcome of the previous play. Our experiment is different in that every period new signal – or the support of the beauty contest game – are randomly drawn. Therefore, subjects cannot directly use the information of the average action of the previous play to guess the average action of the current play. This makes learning much slower and more difficult in our setting. As a result we can observe many subjects staying with the same level of reasoning without adjusting subjects' beliefs as in Nagel (1995).

with private signals such a consistent pattern is either not observed or is considerably weaker.

The last question we want to address in this section is how frequently different levels of reasoning were followed by subjects. This information is given in Table 8 where we count how many times subjects followed a particular level of reasoning during a half of the corresponding phase.

	Pr-A				Pr-H				Pu-H and PuR-H			
	L1	L2	Ln	NE	L1(n)	L2(n)	L3n	NE	L1	L2	L3	NE
0.15	-	-	2	-	-(1)	-(-)	-	1	2	-	-	1
0.30	-	-	1	-	-(-)	-(-)	-	-	11	-	-	3
0.50	1	-	4	-	-(1)	-(-)	1	1	4	2	-	4
0.65	2	-	4	2	5(-)	-(-)	-	4	7	6	-	19
0.80	-	-	1	2	1(-)	-(2)	-	-	18	11	3	4
0.95	2	1	5	-	-(2)	1(-)	2	-	15	23	10	2
Total:	5	1	17	4	6(4)	1(2)	3	6	57	42	13	33
Total %:	2.2%	0.4%	7.5%	1.8%	3.8(2.6)%	0.6(1.3)%	1.9%	3.8%	17%	12.5%	3.9%	9.8%

Table 8: Each entry shows how many times subjects followed a given level within one half of a corresponding phase. Data from treatments Pu-H and PuR-H are combined. In treatment Pr-H columns $L1(n)$ and $L2(n)$ show both sophisticated and naive levels. Numbers for naive levels are in parenthesis. Numbers in the last row are calculated as a percentage of *all* phase-halves played by *all* subjects within a given treatment.

Several things can be noticed upon inspecting Table 8. First, in treatments with private signals naive levels were used more often than the “*sophisticated*” ones: 17 versus 6 in Pr-A and 9 versus 7 in Pr-H. Second, in the treatments with public signals as well as in the Pr-H treatment, $L1$, regardless of naivety, was followed most frequently. In Pr-A we cannot separate $L1n$ from higher naive levels but we still observe that $L1$ was followed more often than $L2$. Third, in the $r = 0.8$ and $r = 0.95$ phases of Pu-H and PuR-H treatments higher levels of reasoning were used more often than in phases with lower r . In particular, in the $r = 0.95$ phase level $L2$ was followed more frequently than $L1$. Finally, we see the effect of learning. In the last phase of each treatment, the one with $r = 0.65$, NE was followed most often.

Result 5: *In all four treatments subjects were more likely to follow the first level of reasoning independent of naivety. One notable exception is the $r = 0.95$ phase in treatments with public signals where $L2$ was most common. In treatments with private signals naive levels of reasoning were considerably more common.*

5.3 Maximum Likelihood Estimation of Level- k and CH Models

In the analysis above, we have shown that the predictive power of level- k reasoning is weaker as we introduce asymmetric information or reduce the weight of coordination component. In this subsection, we compare the results of maximum likelihood estimation across different treatments to further substantiate this finding. We will perform the ML estimations of the following two models: the standard level- k model where Lk type plays best response to the population that

consists entirely of $Lk - 1$ type, and the cognitive hierarchy model (CH) where Lk plays the best response to the population that consists of the mixture of lower types $L0, \dots, Lk - 1$.

Pr-A (1140 obs.)	0.15	0.30	0.50	0.65	0.80	0.95
L1	0.013	0.008	0.040	0.138	0.061	0.038
L2	0.000	0.000	0.040	0.020	0.000	0.027
Ln	0.113	0.109	0.155	0.175	0.097	0.111
NE	0.000	0.000	0.000	0.000	0.025	0.043
LL	-862.412	-867.572	-864.334	-844.267	-853.605	-865.363
Pr-H (780 obs.)						
L1	0.003	0.020	0.010	0.043	0.039	0.031
L2	0.000	0.000	0.018	0.070	0.012	0.030
L1n	0.054	0.000	0.019	0.000	0.011	0.112
L2n	0.048	0.004	0.000	0.035	0.003	0.000
NE	0.000	0.000	0.018	0.081	0.000	0.067
LL	-593.383	-598.368	-599.411	-581.163	-584.006	-591.093
Pu-H (1020 obs.)						
L1	0.020	0.076	0.048	0.096	0.265	0.226
L2	0.010	0.000	0.000	0.000	0.209	0.403
NE	0.010	0.041	0.027	0.204	0.071	0.220
LL	-781.777	-774.376	-779.115	-738.627	-711.186	-645.063
PuR-H (660 obs.)						
L1	0.000	0.010	0.031	0.112	0.180	0.288
L2	0.017	0.000	0.030	0.000	0.223	0.425
NE	0.017	0.125	0.060	0.320	0.021	0.079
LL	-506.461	-499.027	-501.407	-468.224	-474.126	-420.138

Table 9: Results of the ML estimation of level- k models are shown. Estimates are in bold when they are significantly different from 0 at the 5% level.

Results of ML estimation of the standard level- k model are given in Table 9. In the estimation we assume that $L0$ players choose their action at random and Lk players play the best response to the $Lk - 1$ strategy plus an error that is uniformly distributed around the best response. The estimation is performed using normalized values and so the support of the error is set equal to 10 normalized units, or in other words to 10% of the distance between two signals. In the literature the error is usually modeled as having a logistic distribution. We opted for the uniform distribution because subjects are often biased towards “nice” numbers such as integer numbers or those ending with 0 and 5 — something that we also observe in our data. The uniform distribution seems to be a more natural way to capture the error generated by such a bias. In all treatments we directly estimate the shares of $L1$, $L2$ and NE types.⁸ In addition to that, in treatments with private signals we estimate shares of naive types. The share of type $L0$ is then calculated as one minus the sum of the shares of other types.

Comparing the estimates within treatments we see that the log-likelihood roughly increases as r

⁸Adding type $L3$ does not significantly change the results. It changes estimates slightly for the phase with $r = 0.95$ only.

increases from 0.15 to 0.95. In other words, the predictive power of the level- k model increases with higher r , which is consistent with our discussion in previous sections. This feature is prominent for treatments Pu-H and PuR-H, while less so for treatments with private signals. In the latter the phase with $r = 0.65$ tends to perform better than the phase with $r = 0.8$ which could be due to aforementioned learning effect. Another result that follows from Table 9 is that larger values of r generate larger and, most importantly, significantly positive estimates of shares of different types. For example, in treatments with public signals the shares of both $L1$ and $L2$ are significantly positive when $r \geq 0.8$. Furthermore, when $r = 0.95$ types $L1$ and $L2$ are estimated to cover more than 70% of the sample in PuR-H and more than 60% in Pu-H. Again as before this effect is less pronounced for treatments with private signals, however, in Pr-A we still observe more significance as r gets larger.

To compare the performance of the level- k model across treatments we use two criteria: log-likelihood and the share of $L0$. Given that any choice that cannot be attributed to strategic levels of reasoning is automatically labeled as $L0$, the size of the share of $L0$ is a good indicator of the predictive power of the level- k models. In particular, one can interpret the estimate of the share of $L0$ as the frequency of choices that cannot be explained by level- k models.

LL	0.15	0.30	0.50	0.65	0.80	0.95
Pr-A	-0.757	-0.761	-0.758	-0.741	-0.749	-0.759
Pr-H	-0.761	-0.767	-0.768	-0.745	-0.749	-0.758
Pu-H	-0.766	-0.759	-0.764	-0.724	-0.697	-0.632
PuR-H	-0.767	-0.756	-0.760	-0.709	-0.718	-0.637
L0	0.15	0.30	0.50	0.65	0.80	0.95
Pr-A	0.87	0.88	0.76	0.67	0.82	0.78
Pr-H	0.89	0.98	0.93	0.77	0.94	0.76
Pu-H	0.96	0.88	0.92	0.70	0.46	0.15
PuR-H	0.97	0.87	0.88	0.57	0.58	0.21

Table 10: The top part of the table shows log-likelihood across treatments divided by the number of observations and the bottom part shows shares of $L0$.

Table 10 displays both the log-likelihood and shares of $L0$ in different treatments. To make results comparable across different treatments we normalized log-likelihoods by the number of observations in a given treatment. Across all four treatments, the log-likelihood is highest when $r = 0.95$ and the non-zero signal is public (treatments Pu-H and PuR-H). Within treatments with private signals the highest log-likelihood is achieved in the last phase with $r = 0.65$. As for the share of $L0$, it never goes below 75% in treatments with private signals, with the only exception being the phase $r = 0.65$ in Pr-A. Although the performances across all four treatment are comparable for low values of r , the share of $L0$ drops considerably in treatment Pu-H and PuR-H for high values of r . In the phase with $r = 0.95$ the share of $L0$ is 21% in PuR-H and 15% in Pu-H. These results provide statistical support to the earlier finding that the standard level- k model performs very well when there is no private information and the coordination motive is strong. In other cases its performance is quite poor: it fails to explain more than three-quarters of observed choices.

Next, we estimate the cognitive hierarchy model (CH) that was introduced in Camerer et al. (2004). The idea behind the CH model is that higher types believe that opponents' population is a mixture of lower types. For example, type $L2$ believes that some opponents are $L1$ and others are $L0$ and plays accordingly. Camerer et al. assume that there is a correct distribution of different types given by 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 $f(1)/(f(0) + f(1))$.

The estimation of treatments with public signals is straightforward. As before we assume that each type plays the best response plus the error that is uniformly distributed around the best response. The support of the error is 10 normalized units. In estimation we assumed that the highest type in the population is $L3$. For treatments with private signals, we adjust the CH estimation to account for naive types. First, we assume that both naive and sophisticated types are unaware of higher types. Second, naive types are unaware of the sophisticated types. For sophisticated types we considered two alternatives: one when sophisticated types are unaware of naive types and another one when they are aware. We report only the former but the results for the latter were very similar. Finally, we assume that the entire population is divided into two groups: naive players and sophisticated players. The size of each group is an estimation parameter. Within each group types are distributed according to Poisson distribution with parameter τ .

Pr-A (1140 obs.)	0.15	0.30	0.50	0.65	0.80	0.95
τ	0.052	0.084	0.118	0.338	0.148	0.238
Share of Naive	0.229	0.216	0.272	0.310	0.182	0.249
LL	-774.983	-780.276	-751.247	-663.875	-759.356	-765.729
Pr-H (780 obs.)						
τ	0.217	0.151	0.150	0.344	0.252	0.644
Share of Naive	0.969	0.406	0.420	0.128	0.386	0.627
LL	-542.811	-568.764	-570.062	-520.169	-544.556	-496.538
Pu-H (1020 obs.)						
τ	0.135	0.223	0.177	0.241	0.971	4.960
LL	-742.858	-705.694	-727.776	-699.512	-534.554	-371.490
PuR-H (660 obs.)						
τ	0.105	0.213	0.205	0.583	0.899	5.268
LL	-488.229	-459.415	-464.687	-426.619	-395.477	-236.843

Table 11: The ML estimation of the CH model.

Table 11 shows the results of estimation. A higher τ implies that there are larger fractions of subjects who can do higher levels of reasoning. As we can see, the estimates of τ increase as r increases, especially when $r = 0.95$. In general, estimates of τ are higher for treatments with public signals and, as before, the log likelihood increases as r increases. Thus, just as the standard level- k model, the CH model predicts subjects behavior better when r is higher and the information is public. Another interesting observation that follows from Table 11 is that the CH model explains

behavior much better than the standard level- k model. In particular, log-likelihood is higher for the CH model despite the fact that the CH model has less parameters (one in Pu-H, PuR-H and two in Pr-A and Pr-H).

Result 6: *ML estimates confirm that level- k models perform the best when the coordination motive is the strongest and information is public. In the remaining cases level- k models perform poorly and put most of the weight on the non-strategic L0 type.*

Result 7: *In all treatments the CH model fits subjects' behavior better than the standard level- k model, despite the fact that the CH model has fewer parameters than the standard level- k model.*

We conclude the section by comparing the effectiveness of estimated level- k models in predicting the aggregated subject behavior relative to the NE prediction. In Table 12 we calculate the difference (in normalized units) between observed and predicted behavior for NE and CH models which we call prediction errors. In Table 13 we compare the performance of the three models by calculating the ratios of their prediction errors. The left part compares prediction errors of the standard level- k model versus NE and the right part compares CH versus NE. A number higher than 1 means that NE outperforms the level- k alternative in predicting the average action of the subjects, and a number lower than 1 means that the level- k alternative outperforms NE.

	NE						CH					
	0.15	0.30	0.50	0.65	0.80	0.95	0.15	0.30	0.50	0.65	0.80	0.95
Pr-A	25.06	32.50	2.13	2.05	14.52	26.11	21.16	24.18	17.71	18.42	16.47	15.07
Pr-H	15.63	63.08	1.31	4.69	32.72	18.12	10.54	52.29	20.49	17.57	1.07	14.71
Pu-H	24.90	12.33	20.78	2.30	18.93	12.51	21.32	5.01	6.16	22.87	1.08	0.91
PuR-H	12.69	7.12	18.62	1.19	16.82	16.65	9.01	0.26	4.28	17.89	3.87	5.48

Table 12: Average deviation (in normalized units) of the observed behavior from the predicted behavior in NE and the CH model.

We see from both tables that in treatments with public signals both level- k models considerably outperform NE with the usual exception of the $r = 0.65$ phase. In treatments with private signals level- k models perform somewhat worse. NE does considerably better in two phases ($r = 0.65$ and $r = 0.5$) and slightly better in the $r = 0.8$ phase of Pr-A treatment. In phases with $r \geq 0.8$ level- k models tend to outperform NE and more so in treatments with public signals where the prediction errors can differ by a factor of 20. Looking at absolute numbers we also see that in treatments with high r and public signals the average prediction error of NE was between 12.5 and 18.9 whereas the CH prediction error was much lower and varied between 0.9 and 5.5.

Finally, we can use Table 13 to compare the relative performance of CH and standard level- k models. In general, the two models have similar performance and depending on the phase and the treatment either model may perform better than the other. There are few instances when the CH model considerably — i.e. by a factor of at least 2 — outperforms the standard level- k model such as the $r = 0.8$ phase of Pr-H or the $r = 0.95$ phase of Pu-H treatment. However, in these instances both models predict the aggregated behavior reasonably closely which is the main reason for having such large ratios. In both cases the prediction error of both models was less than four normalized

	Level-k/NE						CH/NE					
	0.15	0.30	0.50	0.65	0.80	0.95	0.15	0.30	0.50	0.65	0.80	0.95
Pr-A	0.84	0.73	8.29	9.44	1.15	0.59	0.84	0.74	8.31	8.99	1.13	0.58
Pr-H	0.65	0.81	16.07	3.54	0.09	1.11	0.67	0.83	15.65	3.75	0.03	0.81
Pu-H	0.84	0.36	0.25	8.66	0.05	0.26	0.86	0.41	0.30	9.94	0.06	0.07
PuR-H	0.69	0.07	0.20	13.22	0.35	0.20	0.71	0.04	0.23	15.02	0.23	0.33

Table 13: The left part of the table shows the ratio of average deviations of level- k and NE predictions. The right part shows the ratio of average deviations of CH and NE predictions. Numbers greater than 1 mean that NE performs better and numbers less than 1 mean that the level- k (CH) model performs better.

units. Overall, based on models' ability to predict subjects' aggregated behavior we cannot say that one level- k model is unambiguously better than another. However, it still should be remembered that the CH model has a smaller number of parameters and leads to a higher log-likelihood.

Keeping in mind the usual exception of the $r = 0.65$ phase we conclude with the following result:

Result 8: *In general, both level- k models predict the aggregated behavior better than NE. Their advantage is particularly pronounced in treatments with public signals and strong coordination motives.*

6 Concluding Remarks

The goal of this paper is to determine the setting in which level- k thinking most appropriately describes subjects' behavior. To do that we generalize the classical beauty contest setting by using the Morris and Shin (2002) framework that allows us to introduce private information and vary the strength of the coordination motive. Having the experimental design based on the the MS model generates an environment that is more complex than the one typically used in the level- k literature. Despite this complexity we confirm the finding in the existing beauty contest literature that level- k models are indeed successful in predicting subject behavior when the game setting is close to the classical beauty contest, when information is symmetric and the coordination motive is strong. Moreover, most subjects choose their levels of reasoning consistently in the sense that they either adhere to one particular level or switch to higher levels. However, as we move away from the classical setting, the predictive power of the level- k models diminishes quickly. Only a handful of subjects play according to level- k reasoning and those who do tend to use it in a rather inconsistent manner.

We conjecture that the reason for these results is as follows. When the coordination motive weakens, the behavior of other players becomes less important; as such, subjects are less likely to try to predict it. This is true regardless of whether information is symmetric or not. The introduction of private information into the model weakens level- k behavior 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 since 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. Our results suggest that some subjects substituted naive level- k thinking for standard level- k thinking in settings with private 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-A

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 ;
- the average investments made by others: $a_{average} = \frac{a_1 + a_2 + a_3}{3}$;

Example: Assume that the other three investors in your group invested 150, 200 and 250. The average amount invested by the others is $a_{average} = 200$.

At the time when you make decisions you will **NOT** know either of these two factors. You will not know the average amount invested by others, $a_{average}$, 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	$(0+s)/2$	$(0+s)/2 \pm 40$	$(0 + s)/2 \pm 80$

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 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 know 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 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 the average investments made by others. Your final profit will be calculated by the following formula:

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

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 the average investments made by others.

It is possible that the project quality and 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 it is *more important to match the quality*. Let quality, q , be 10, and a_{average} be 60. 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 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 it is *more important to match the investments of others*. As before assume that $q = 10$ and $a_{\text{average}} = 60$. 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 \cdot 0.001 + 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 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] Allen, Franklin, Stephen Morris and Hyun Song Shin (2006). “Beauty Contests and Iterated Expectations”, *Review of Financial Studies*, 19, 161-177.
- [2] Angeletos, George-Marios, Guido Lorenzoni and Alessandro Pavan (2007). “Wall Street and Silicon Valley: A Delicate Interaction.” Working paper, MIT.
- [3] Angeletos, George-Marios, and Alessandro Pavan (2007). “Efficient Use of Information and Social Value of Information.” *Econometrica*, 75(4), 1103–42.
- [4] Bacchetta, Philippe, and Eric van Wincoop (2005). “A Theory of the Currency Denomination of International Trade”, *Journal of International Economics* 67, 295-319.
- [5] Bosch-Domenech, Antoni, Jose Garcia Montalvao, Rosemarie Nagel, and Albert Satorra (2002). “One, Two, (Three), Infinity, . . . : Newspaper and Lab Beauty-Contest Experiments.” *American Economic Review*, 92, 1687–1701.
- [6] Camerer, Colin, Teck-Hua Ho, and Juin-Kuan Chong (2004). “A Cognitive Hierarchy Model of Games.” *Quarterly Journal of Economics*, 119, 861–898.
- [7] Cornand, Camille, and Frank Heinemann (2009). “Measuring agents’ Overreaction to Public Information in Games with Strategic Complementarities.” Work in progress, Department of Economics and Management, Technical University Berlin.
- [8] Costa-Gomes Miguel, Vincent Crawford and Bruno Broseta (2001) “Cognition and behavior in normal-form games : An experimental study”, *Econometrica* vol. 69, no 5, pp. 1193-1235.
- [9] Costa-Gomes, Miguel, and Vincent Crawford, (2006). “Cognition and Behavior in Two-Person Guessing Games: An Experimental Study”, *American Economic Review*, vol. 96, pp. 1737-1768.

- [10] Costa-Gomes, Miguel, Vincent Crawford, and Nagore Iriberry (2009), “Comparing Models of Strategic Thinking in Van Huyck, Battalio, and Beil’s Coordination Games,” *Journal of the European Economic Association*, forthcoming.
- [11] Crawford, Vincent, and Nagore Iriberry (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, 1721–1770.
- [12] Crawford, Vincent, and Nagore Iriberry (2007b). “Fatal Attraction: Saliency, Naivete, and Sophistication in Experimental Hide-and-Seek Games.” *American Economic Review*, 97, 1731–1750.
- [13] Dewan, Torun, and David Myatt (2007). “Leading the Party: Coordination, Direction and Communication.” *American Political Science Review*, 101(4), 825–843.
- [14] Dewan, Torun, and David Myatt (2008). “The Qualities of Leadership: Communication, Direction and Obfuscation.” *American Political Science Review*, 102(3), 351–368.
- [15] Gneezy, Uri (2005). “Step-Level Reasoning and Bidding in Auctions.” *Management Science*, 51, 1633–1642.
- [16] Ho, Teck-Hua, Colin Camerer, and Keith Weigelt (1998). “Iterated Dominance and Iterated Best Response in Experimental ‘*p*-Beauty Contests’.” *American Economic Review*, 88, 947–969.
- [17] Morris, Stephen and Hyun Song Shin (2002). “The Social Value of Public Information.” *American Economic Review* 92 (2002), pp. 1521–1534.
- [18] Nagel Rosemarie, 1995. “Unraveling in Guessing Games: An Experimental Study.” *American Economic Review*, Vol. 85, No. 5, pp. 1313–1326.
- [19] Stahl, Dale and Paul Wilson (1995). “On Players’ Models of Other Players: Theory and Experimental Evidence.” *Games and Economic Behavior*, 10, 218–254.
- [20] Stahl, Dale and Paul Wilson (1994). “Experimental Evidence on Players’ Models of Other Players.” *Journal of Economic Behavior and Organization*, 25, 309–327.