amse
école d'économie d'aix-marseille
aix-marseille school of economics

Working Papers / Documents de travail

Inclusive Cognitive Hierarchy in Collective Decisions

Yukio Koriyama
Ali Ozkes

WP 2018 - Nr 19

Inclusive Cognitive Hierarchy in Collective Decisions*

Yukio Koriyama[†] Ali Ozkes[‡]

June 22, 2018

Abstract

We study the implications of structural models of non-equilibrium thinking, in which players best respond while holding heterogeneous beliefs on the cognitive levels of others. We introduce an *inclusive* cognitive hierarchy model, in which players are capable of projecting the self to others in regard to their cognitive level. The model is tested in a laboratory experiment of collective decision-making, which supports inclusiveness. Our theoretical results show that inclusiveness is a key factor for asymptotic properties of deviations from equilibrium behavior. Asymptotic behavior can be categorized into three distinct types: naïve, Savage rational with inconsistent beliefs, and sophisticated.

JEL classification: C92, D72, D91

Keywords: Collective decision-making, Bounded rationality, Cognitive hierarchy, Information aggregation

We are thankful to Colin Camerer, Nobuyuki Hanaki, Guillaume Hollard, Keigo Inukai, Nagore Iriberri, Jean-François Laslier, Uğur Özdemir, Thomas Palfrey, Eduardo Perez-Richet, Phil Reny, Ariel Rubinstein, seminar participants at Ecole Polytechnique, Paris School of Economics, Toulouse School of Economics, GATE at Saint-Etienne, University of Chicago, Caltech, University of Bayreuth, Hitotsubashi University, Keio University, Kyoto University, Osaka University, 5th Murat Sertel Workshop in Istanbul, SAET Conference 2014, Transatlantic Theory Workshop 2014, French Symposium on Games 2015, ASFEE 2015, GAM-SOC Workshop 2015, Social Choice and Welfare Conference 2016, Game Theory Congress 2016, and BEAM Conference for valuable comments. Financial support by Investissements d'Avenir, ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047 and IDEX ANR-10-IDEX-0001-02 PSL MIFID is gratefully acknowledged.

[†]CREST, Ecole Polytechnique, Université Paris-Saclay. Email: yukio.koriyama@polytechnique.edu.

[‡]Aix-Marseille Univ., CNRS, EHESS, Centrale Marseille, AMSE. E-mail: ali-ihsan.ozkes@univ-amu.fr.

1 Introduction

In her presidential speech to the American Political Science Society, Ostrom (1998) called for a behavioral theory of collective action based on models of boundedly rational individuals, pointing to the fact that behavior in social dilemmas is affected by many structural variables, such as group size. The objective is put forward with an understanding of the "thin" model of rationality as a limiting case of bounded or incomplete rationality, as in Selten (1975).

Collective decision-making processes entail most of the time, if not always, asymmetric information. And studying strategic thinking in private information games is crucial, as they are common tools for modeling interactions in bargaining, matching, contracts, financial markets, and political systems, *inter alia*. If strategic naïveté is prevalent, peculiarities due to private information can be more crucial than what is predicted by equilibrium analysis, and thus policy responses may be astray, as noted by Brocas et al. (2014).

Models of non-equilibrium strategic thinking have been proposed to explain structural deviations from equilibrium in a variety of games. A sizable part of bounded rationality literature is devoted to the models of cognitive hierarchy, starting with the *level-k* model due to Nagel (1995) and Stahl and Wilson (1995), which allow for heterogeneity among individuals in levels of strategic thinking. In the level-k model, a foundational level (level-0) represents a strategically naïve initial approach to a game, and a level-k player (hereafter Lk), where $k \geq 1$, is assumed to best respond to others with level $k-1$. The construction of levels resonates with *rationalizability*, as in Bernheim (1984), due to the fact that the decisions made by a level-k player survive k rounds of iterated elimination of strictly dominated strategies in two-person games.

Closely related, the *cognitive hierarchy* (CH) model introduced by Camerer et al. (2004) allows for heterogeneity in the beliefs on others' levels.[1] For each $k \geq 1$, a level-k player best responds to a mixture of strictly lower levels, induced by the truncation up to level $k-1$ from the underlying level distribution, which is either obtained from maximum likelihood estimations applied to data, or calibrated from previous estimates. The set of level-1 strategies in the CH model (hereafter $CH1$) is exactly the same as that of $L1$, and Lk and CHk differ for $k \geq 2$. Experimental studies provide various sorts of evidence that the CH model delivers a better fit for explaining the actual behavior of players in certain games. Common to these models is the assumption that level-k players do not assign any probability to levels higher than or equal to k, and this emerges from the idea that the cognitive limits among players have in fact a hierarchical structure.

In both of the aforementioned models, *self-awareness* as described in Camerer et al. (2004) is precluded,[2] that is, the possibility for players to assign positive probability to the events in which other players have the same cognitive level as themselves is ruled out.

[1] Crawford et al. (2013) provide a fine review of these models and applications.

[2] Self-awareness is probably not the best term here, since what is meant is the capacity for projecting the self to others, which has little to do with the capacity for introspection often meant by self-awareness in, for example, psychology.

Here, we propose, instead, an *inclusive cognitive hierarchy* (ICH) model, in which players are allowed to project themselves to others in regard to their cognitive level. The ICH model thus allows for inclusiveness while maintaining the hierarchical structure of cognitive levels and the partial consistency induced by truncation of the underlying level distribution as in other cognitive hierarchy models. In this paper, we study the role of inclusiveness both empirically and theoretically. On the empirical side, we present the results of our laboratory experiment on binary collective decision making. On the theoretical side, we provide theorems which describe asymptotic properties of the deviation from rational behavior.

Since Condorcet's classical work in 1785, mathematical support has been provided for the idea of increasing the accuracy of collective decisions by including more individuals in the process. In his seminal *Essai*, Condorcet considered non-strategic individuals voting to make a decision on a binary issue where each alternative is commonly preferred to the other in one of the two states of the world. Each individual independently receives an imperfectly informative private signal about the true state of the world and votes accordingly. Under majority rule, the probability of reaching a correct decision monotonically increases with the size of the electorate and converges to certainty in the limit.

Although enabling strategic behavior may imperil the validity of the basic assumptions in the original model, the asymptotic property survives in various circumstances of collective decision-making (e.g. Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1997)). Studies have documented conditions under which the probability of making a right decision increases and converges to certainty as the group size increases, even when strategic players may vote against their signals.

However, assuming full rationality, especially complete mutual consistency of beliefs, may be too demanding when the strategy space is big and a large number of players are involved. In large games, the collective behavior of strategic players may differ qualitatively from that in small games. This intuition finds support from experimental evidence and information aggregation by a large group is not, to say the least, exempt from it, especially when a major presumption of the strategic models is that voters take into account their probability of being pivotal, as in Downs (1957). As shown by Esponda and Vespa (2014), perfectly accurate hypothetical thinking to extract information from others' strategies that is required in most strategic voting models might be too strong as an assumption, even untenable, when it is assumed for all individuals.[3] Battaglini et al. (2008) document in another experimental study an increase in irrational, non-equilibrium play as the size of the electorate increases.[4] Collective performances are correlated across challenges, as demonstrated in Woolley et al. (2010), hence a good knowledge about the behavioral basis in collective decision-making

[3]In their experiment, 50 to 80% of the participants behaved non-strategically when voting was simultaneous, and they find that non-optimal behavior is typically due to difficulty in extracting information from hypothetical events.

[4]As Camerer (2003) stresses (Chapter 7), the effect of group size on behavior in strategic interactions is a persistent phenomenon, especially towards coordination.

procedures is essential to understanding the more general phenomena of our societies. We argue in this paper that the cognitive hierarchy approach can be fruitful in that regard, once inclusiveness is incorporated.

There are at least three reasons why we believe that the consequences implied by inclusiveness should be studied. First, the observed limitations of existing models in the extant literature on strategic thinking call for studies with an explicit focus on inclusiveness,[5] for which our experimental analysis provides a clear support. Second, our post-experimental questionnaire also suggests inclusiveness.[6] Figure 1(a) shows that a fairly large proportion (96%) of our experimental subjects exhibit a positive degree of inclusiveness: 'sometimes', 'most of the time', or 'always'. They are further asked to provide a subjective estimation of the percentage among other members who used the 'same reasoning' (Figure 1(b)). The answers vary and a vast majority of them are far from 0%, which suggests that ruling out inclusiveness would be too extreme as an assumption.[7] Third, our analytical results imply that the presence of inclusiveness simply matters (Theorem 1) in a certain class of games (Theorem 2). A novel finding of this paper is that a lack of the inclusiveness condition may imply that the deviation from rational behavior asymptotically diverges away without a bound, which we argue is not coherent with the spirit of the cognitive hierarchy models.[8] One of the objectives of this paper is to shed light on the role of the behavioral assumptions that account for the asymptotic properties of the deviation from the equilibrium thinking.

Related Literature

Gerling et al. (2005) provide an extensive survey on the studies of collective decision-making in committees and the Condorcet Jury Theorem. Palfrey (2016) provides a comprehensive survey on experiments in political economy and particularly in strategic voting. Costinot and Kartik (2007) study voting rules and show that the optimal voting rule is the same when players are sincere, playing according to Nash equilibrium, to level-k, or a mixture of these. Bhattacharya et al. (2017) test experimentally the theoretical predictions about individual behavior and group decisions under costly information acquisition. They find poor support for the comparative statics predictions which are delivered theoretically. Alaoui and Penta (2016) introduce a model of strategic thinking that endogenizes individuals' cognitive bounds

[5] Colman et al. (2014) point to the observed weak performance of the CH model in common interest games. Georganas et al. (2015) conclude that level-k models have a high performance in some games, but not in others. "Equilibrium plus noise" models such as quantal response equilibrium (cf. Goeree et al. (2016)) often miss systematic patterns in participants' deviations from equilibrium, which is a main feature observed in our experimental data.

[6] As the questionnaire is not incentivized, the interpretation should be performed with care.

[7] Heterogeneity of the answers implies that the participants' degree of inclusiveness was also heterogeneous. This observation is coherent with the assumption of hierarchical belief structure that we describe in detail in Section 2.

[8] Evidence from experimental studies on public good games suggest that as group size increases, individual behavior bears convergent and stabilizing tendencies. See Isaac et al. (1994).

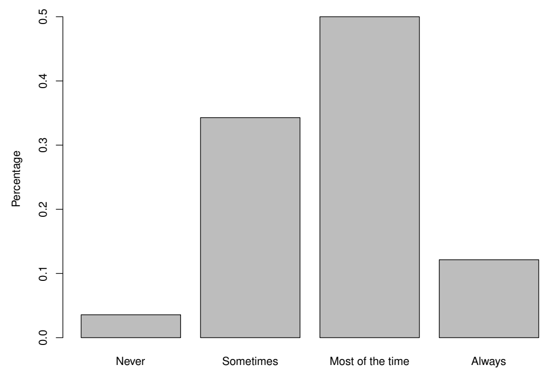

(a) "When you made decisions, did you think that the other participants in your group used exactly the same reasoning as you did?"

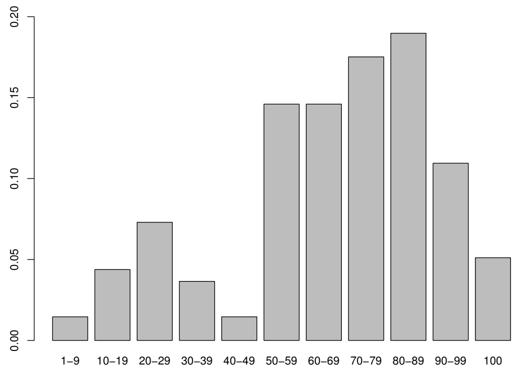

(b) "What is the percentage of the other participants using the same reasoning, according to your estimation?"

Figure 1: Responses in the post-experimental questionnaire.

as a result of a cost-benefit analysis. Their framework allows individuals to reason about opponents whom they regard as more sophisticated as well. Hanaki et al. (2018) study how the strategic environmental effects depend on the group size in beauty contest games, finding support for the presence of the effects in large groups.

The paper proceeds as follows. We introduce the inclusive cognitive hierarchy model formally in the following section. We furthermore provide a numerical comparison of the individual behaviors and the performance of collective decisions under different specifications of the cognitive hierarchy in a game of information aggregation. In Section 3 we introduce our experimental design that features novelties due to our modeling concerns and signal setup. Section 4 provides the results of the experiment, and the models are compared in terms of the fit to the data. Section 5 provides theoretical results focusing on linear quadratic games. We present our main theorems, which provide a characterization of games according to the asymptotic properties of the strategic thinking. We conclude by summarizing our findings and presenting further research questions in Section 6. The proofs of the theorems are relegated to Appendix.

2 The Model

Let (N, X, u) be a symmetric normal-form game where $N = \{1, \ldots, n\}$ is the set of players, $X \subset \mathbb{R}$ is a convex set of pure strategies, and $u : X^n \to \mathbb{R}^n$ is the payoff function. Each player forms a belief on the cognitive levels of the other players. Let $g_k(h)$ denote the probability that a k^{th}-level player assigns independently for each of the other players to belong to the h^{th}-level.

In the standard *level-k model*, a naïve, non-strategic behavior is specified as the initial level (*level*-0, or $L0$). For $k \geq 1$, a level-k (Lk) player holds the belief that all of the other players belong to exactly one level below herself:

$$g_k(h) = \begin{cases} 1 & \text{if } h = k-1, \\ 0 & \text{otherwise.} \end{cases} \tag{L}$$

In the cognitive hierarchy model, each k^{th}-level (CHk) player best responds to a mixture of lower levels. Let $f = (f_0, f_1, \ldots)$ be a distribution over $\mathbb{N}$ which represents the composition of levels. Each k^{th}-level player holds a belief on the distribution of the other players' levels that is a truncation up to one level below herself:

$$g_k(h) = \frac{f_h}{\sum_{m=0}^{k-1} f_m}, \text{ for } 0 \leq h \leq k-1 \text{ and } k \geq 1. \tag{CH}$$

Thus, these two models share the following assumption:

Assumption 1 (Strong overconfidence) $g_k(h) = 0$ *for all* $h \geq k$*.*

Assumption 1 enables us to say that what we call "levels" here indeed has a hierarchical structure. To see this, consider $g_k(h)$ in the format of a k-h matrix. Assumption 1 implies that the upper-diagonal entries are all zeros, and thus the remaining non-zero elements have a pyramid structure with strictly lower-diagonal entries. Each level-k player assigns non-zero probabilities only to the levels strictly lower than herself. In that sense, players are assumed to be *overconfident.*[9] Alternatively, the following assumption can be considered:

Assumption 2 (Weak overconfidence) $g_k(h) = 0$ *for all* $h > k$*.*

Assumption 2 is weaker than Assumption 1. As in Assumption 1, zero probability is assigned for all strictly upper-diagonal entries, and thus a hierarchical structure among levels is still preserved. However, the diagonal entries are not restricted to be zero. A level-k player is allowed to assign a non-zero possibility for the other players to have the same level as herself.

In what follows, we formally introduce the *inclusive cognitive hierarchy (ICH) model*, in which the strong overconfidence condition is weakened to allow for inclusiveness.

[9] Camerer and Lovallo (1999) report on experimental evidence for overconfident behavior in the case of the market entry game. When ability is a payoff-relevant variable in a strategic interaction, evidence shows that players tend to be overconfident. See Benoît and Dubra (2011). On the other hand, Azmat et al. (2018) find an underestimation of students' grades in the absence of feedback.

2.1 Inclusive Cognitive Hierarchy

Fix an integer $K > 0$ that prescribes the highest level considered in the model.[10] In the ICH model, we consider a sequence of mixed strategies $\sigma = (\sigma_0, \ldots, \sigma_K)$, in which for each $k \in \{1, ..., K\}$, $\sigma_k \in \Delta(X)$ is a best reply, assuming that the other players' levels are drawn from the truncation of the underlying distribution f *up to level* k. Note that a sequence of truncated distributions $g = (g_1, \ldots, g_K)$ is uniquely defined from f. As in the other models, we focus on level-symmetric profiles in which all players of the same level use the same mixed strategy.

Definition 1 *A sequence of level-symmetric strategies* $\sigma = (\sigma_0, \ldots, \sigma_K)$ *is called* ***inclusive cognitive hierarchy strategies*** *when there exists a distribution* f *over* $\mathbb{N}$ *under which*

$$supp\left(\sigma_k\right) \subset \arg\max_{x_i \in X} \mathbb{E}_{x_{-i}}\left[u\left(x_i, x_{-i}\right)|g_k, \sigma\right], \ \forall k \in \{1, \cdots, K\},$$

where g_k *is the truncated distribution induced by* f *such that*

$$g_k\left(h\right) = \frac{f_h}{\sum_{m=0}^{k} f_m} \text{ for } h \in \{0, \cdots, k\}, \tag{ICH}$$

and the expectation over x_{-i} *is drawn, for each player* $j \neq i$*, from a distribution*

$$\gamma_k\left(\sigma\right) := \sum_{m=0}^{k} g_k\left(m\right)\sigma_m.$$

We note that Definition 1 is analogous to the definitions used in the standard level-k model and the cognitive hierarchy model. It simply replaces the assumptions on beliefs, (L) and (CH), with (ICH). Building on previous studies (as developed by Camerer et al., 2004), we maintain the assumption that the underlying distribution f of levels follows a Poisson distribution with coefficient τ:

$$f_k = \frac{\tau^k}{k!}e^{-\tau}.$$

Note that the expectation of the distribution is τ, which thus represents the overall expected level among the players.

2.2 Condorcet Jury Theorem

A group of n individuals makes a binary collective decision $d \in \{-1, 1\}$. The true state of the world is also binary, $\omega \in \{-1, 1\}$, with a common prior of equal probabilities. The payoff is a function of the realized state and the collective decision as follows:

$$\tilde{u}(\omega, d) = \begin{cases} 0 & \text{if } \omega \neq d, \\ q & \text{if } \omega = d = 1, \\ 1 - q & \text{if } \omega = d = -1, \end{cases}$$

[10]We assume $f_i > 0$ for all $i \leq K$. For the truncated distribution to be well-defined, it is sufficient to assume $f_0 > 0$, but we restrict ourselves to the cases where all levels are present with a positive probability.

with $q \in (0,1)$ for all individuals.[11] Each individual $i \in \{1, \cdots, n\}$ receives a private signal $s_i \in S$, distributed independently conditional on the true state ω. A collective decision is made by the majority rule. Upon receiving signal s_i, individual i casts a vote $v_i \in \{0,1\}$, and the collective decision is determined by the sign of $(\sum_i v_i - n/2)$.

We do not restrict the signal space S to be binary.[12] We assume $S \subset \mathbb{R}$ so that S is an ordered set, and we assume that the commonly known distribution satisfies the monotone likelihood ratio property, that is, the posterior distribution $\Pr[\omega = 1|s_i]$ is increasing in s_i.

A strategic Condorcet Jury Theorem claims that asymptotic efficiency is obtained among the rational individuals with homogeneous preferences and costless information acquisition, as described above. More precisely, it claims that, under the Nash equilibrium behavior, the probability of making a right decision converges to one as n goes to infinity. In the following subsection we ask whether the asymptotic, collective efficiency would be obtained under the cognitive hierarchy models in which individuals may show systematic deviations from the Nash behavior.

2.3 Asymptotic efficiency: a numerical example

In order to underline the differences implied by different behavioral assumptions, we provide numerical computation results using the game described above. The model parameters are chosen so that the game coincides exactly with the one with asymmetric payoffs in our experiments (Section 3).[13] Our aim here is to highlight numerically the behavioral consequences of the inclusiveness condition, by comparing the probabilities of making a correct group decision and the individual strategies predicted by each model. Four different behavioral specifications are compared: Nash equilibrium (NE), the standard level-k model (L), the Poisson cognitive hierarchy model (CH), and the Poisson inclusive cognitive hierarchy model (ICH).

Figure 2(a) shows the probability of making a correct decision as a function of the group size. Convergence to one in Nash equilibrium reflects the strategic Condorcet Jury Theorem, as a benchmark for the cognitive hierarchy models under our consideration. The ICH model also exhibits asymptotic efficiency. The quality of group decisions shows a stark contrast

[11] The assumption of symmetric prior is without loss of generality, since we allow the payoffs of the two types of right decisions to be heterogeneous. Although the preferences are often represented by a loss function for wrong decisions in the standard CJT models, we equivalently use a gain function for right decisions in accordance with our experiment, which awards positive points to right decisions, rather than subtracting points for wrong decisions.

[12] Even though a number of CJT models assume a binary signal space, we believe that this is not the right assumption for information aggregation problems. Even under the binary state space, there are uncountably many ways to update the prior belief, and thus the set of possible beliefs spans a continuous space. Assuming the signal space to be binary implies that there are only two ways of Bayesian update, which is far from innocuous.

[13] Both the error term and the Poisson coefficient that are estimated for the experimental data in Section 4 are calibrated here as the average of the estimated values.

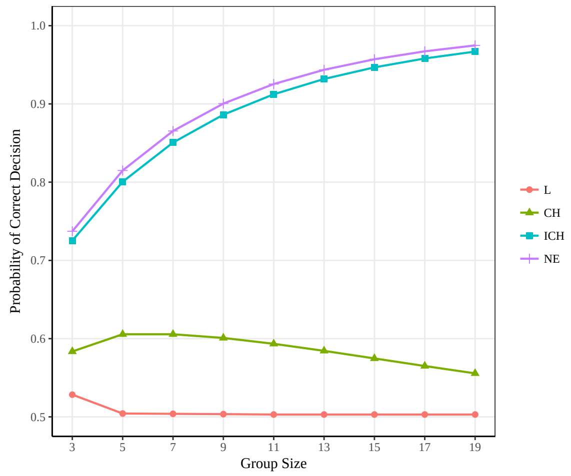

(a) Probability of correct group decision as a function of group size.

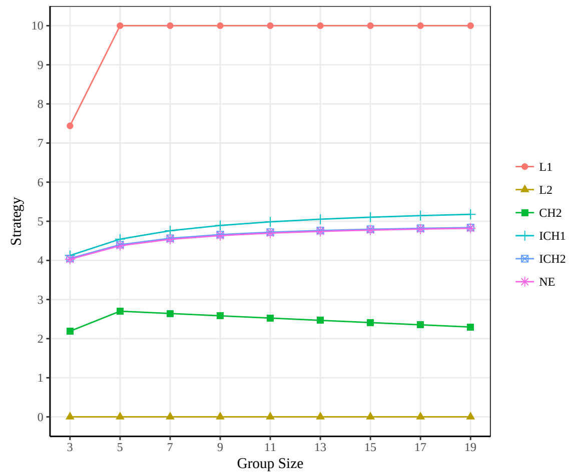

(b) Level-1 and level-2 strategies. Note that $CH1 = L1$.

Figure 2: A numerical example.

both in the L model and in the CH model with the ICH model. The probability approaches 0.5 in large groups, which is as bad as pure noise, implying that the Condorcet Jury Theorem fails to the worst extent without considering the inclusiveness of the agents.

The key difference in asymptotic properties between CH and ICH stems from the inclusiveness condition. With inclusiveness, each player in the ICH model chooses the optimal behavior that maximizes the expected utility given that the other players may have the same level. As is shown analytically in Section 5, the distance of the ICH strategies from the Nash equilibrium does not diverge away even in a large group, since the strategy at each level is obtained as a best reply to the other players' strategies, which may match each other with a positive probability. A highest-level player in the ICH model is thus capable of choosing his strategy fully rationally, correcting the biases caused by the lower levels. Without inclusive players, the deviations from the rational behavior accumulate in the group and thus deteriorate the quality of the collective decision.

Figure 2(b) depicts level-1 and level-2 strategies under the L, CH, and ICH models, and the Nash equilibrium, as a function of the group size n.[14] In the L model, strategies hit the boundary for both level-1 (except for $n = 3$) and level-2 values. In the CH model, level-1 is the same as in the L model, while the level-2 strategy is decreasing, meaning a divergence from the Nash behavior. In the ICH model, strategies are increasing in both level-1 and level-2, in accordance with the Nash behavior. We observe that the convergence of the ICH2 strategy is remarkably quick, so that it visibly coincides with the Nash equilibrium.

These figures demonstrate a stark contrast among the behavioral assumptions under our consideration. In particular, the figures suggest that the inclusiveness condition plays a key

[14]Strategy is defined as a cutoff value in the interval $[0, 10]$. See Section 4.1 for a precise definition.

role in describing the asymptotic behavior. In the following sections, we show the results from our stylized laboratory experiment which provides statistical evidence for our scrutiny of the models.

3 Experimental Design

All of our computerized experimental sessions were held at the Ecole Polytechnique Experimental Economics Laboratory.[15] In total we had 140 actual participants in 7 sessions, in addition to the pilot sessions with more than 60 participants. In each session, 20 participants took part in 4 phases (together with a short trial phase) which lasted about one hour in total. Earnings were expressed in experimental currency units (ECUs) and exchanged for cash, to be paid immediately following the session. Participants earned an average of about 21 Euros, including a default 5 Euros for participation. Complete instructions and details can be found in our online appendix.[16] The instructions pertaining to the entire experiment were read aloud at the beginning of each session. Before each phase, the changes from the previous phase were read aloud, and an information sheet providing the relevant details of the game was distributed. These sheets were exchanged with new ones before each phase.

We employed a within-subject design where each participant played all 4 phases consecutively in a session. Each phase contained 15 periods of play, and thus each participant played for a total of 60 periods under a direct-response method. Since the question of our research relates to the strategic aspects of group decisions, our experiment was presented to participants as an abstract group decision-making task where neutral language was used to avoid any reference to voting or elections of any sort.

In the beginning of each period, the computer randomly formed groups of participants, of a size that was commonly known and predetermined for each phase (either $n = 5$, 9, or 19).[17] Then, a box was shown to each participant with one hundred squares (to be referred as *cards* from now on), all colorless (gray in *z-Tree*). At the same time, the unknown true color of the box for each group was determined randomly by the computer. The participants were informed that the color of the box would be either blue or yellow, with equal probability. It was announced that the blue box contained 60 blue and 40 yellow cards, whereas the yellow box contained 60 yellow and 40 blue cards.

After confirming to proceed to the next screen, 10 cards drawn by the computer with random locations in the box were shown to the participants, this time with the true colors. These draws were independent among all participants but were drawn from the same box

[15] Both the z-Tree program (Fischbacher (2007)) and the website for participant registration were developed by Sri Srikandan, to whom we are very much grateful.

[16] The online appendix can be found at http://sites.google.com/site/ozkesali.

[17] In the phase with $n = 9$, two groups of 9 randomly-chosen members were formed at each period. Having 20 participants in total in each session, 2 randomly-chosen participants were 'on hold' during the period. The same method is applied in the phase with $n = 19$. A group of 19 was formed and thus one randomly-chosen participant was on hold.

in the same group. Having observed the 10 randomly-drawn cards, the participants were required to choose either blue or yellow by clicking on the corresponding button. Then, the decision for the group was reached by majority rule, which was resolute every time, since we only admitted odd numbers for group sizes and abstention was not allowed. Once all participants in a group had made their choices, the true color of the box, the number of members who chose blue, the number of members who chose yellow, and the earnings for that period were revealed on the following screen. A new period started after everyone confirmed.

In one of the four phases, the group size was set at $n = 5$ and the payoffs were symmetric. Each participant earned 500 ECUs for any correct group decision (i.e., a blue decision when true color of the box was blue, or a yellow decision when true color of the box was yellow). In the case of an incorrect decision, no award was earned. In the other three phases, each treatment differed only in the size of the groups (5, 9, or 19) where asymmetric payoffs were fixed. The correct group decision when the true color of the box was blue awarded each participant in the group with 900 ECUs, whereas the correct group decision when the true color of the box was yellow awarded them with 200 ECUs.[18] Lastly, we implemented a random-lottery incentive system where the final payoffs at each phase were determined by the payoffs from a randomly-drawn period.[19]

Let us underline that asymmetry in remuneration is introduced in our experiment in order to observe the effect of a prior bias on the participants' behavior. It is not surprising that an informative strategy (i.e. voting for the choice favored by the signal), or one close to it, is employed by a large majority of the participants under symmetric awards.[20] When it is commonly known that one of the alternatives may provide a larger award, in addition to the change of the symmetric Nash equilibrium shifting towards the ex ante preferable alternative, each individual's behavior may shift, and furthermore, such shifts may be heterogeneous across individuals. Consequentially, each individual may hold heterogeneous beliefs over the strategies employed by the other individuals in the group. The accumulated effects of such heterogeneous belief formation may hamper the performance of group decision-making, which is one of our main concerns in this paper.

At the beginning of each session, as part of the instructions, participants played through two mandatory trial periods. Each session concluded after a short questionnaire. According to the anonymously-recorded questionnaire, 44% of the participants were female. The age distribution was as follows: 31% between the age of 19 and 22, 26% between 23 and 29,

[18]We also conducted pilot sessions with the rewards of 800 ECUs and 300 ECUs. As there was less of a marked contrast in the observed deviations from Nash behavior, we decided to run the rest of the sessions with the rewards 900 ECUs and 200 ECUs.

[19]Participants were told both verbally and through info sheets that in the case where the lottery picked a period for remuneration in which a participant had been on hold, the payoff in that phase for this participant was set at 500 ECUs, which is about the average of the winning points.

[20]Behaviors close to the informative strategy are indeed observed in our experiments with unbiased payoffs. See histogram in Figure 3(a).

14% between 30 and 39, and 29% between 40 and 67. Heterogeneity in their professions was relatively high: 46% administrative staff, 37% undergraduate students ("Polytechniciens"), 12% Ph.D. students, 1% master students, and 3% researchers. 6% of the participants had previously taken an advanced course in game theory, while 14% had taken an introductory course. 39% said that they had some notions about game theory, while 41% claimed to have no knowledge of game theory.

4 Experimental Results

In this section we present our experimental results at both the individual and group levels.

4.1 Cutoff Strategies

Under our experimental design, a pure strategy of an individual is a function from the realized signal to a binary vote. It is straightforward to show that the best reply of an individual, given any belief on the strategies of other group members, is a cutoff strategy. There exists a threshold for each individual such that she votes for blue if and only if her signal induces a higher posterior probability of a blue state than the threshold. Since the posterior belief over the two states varies monotonically as a function of the number of blue cards among the 10 revealed ones, a cutoff strategy in our experiment is that each individual votes for blue when the number of observed blue cards is higher than the cutoff value, and for yellow otherwise. Special cases include voting for one of the colors regardless of the signal. The cutoff value is considered as an extreme value (either 0 or 10) for such a behavior.

We have observed in our data that a majority of participants used a cutoff strategy with randomization. Randomization occurs with two or more realized values of the signal, with the degree of randomization varying monotonically in the right direction (i.e. a higher probability of voting for blue given more blue cards in the signal). We regard such a behavior as a consequence of decision-making with an error or other uncertainties which are not explicitly formalized in the model.[21]

More precisely, let $x_i \in [0, 10]$ be the (continuous) cutoff strategy of voter i, and let $s_i^t \in \{0, 1, \cdots, 10\}$ be the (discrete) signal realized in period t, i.e. the number of blue cards out of the ten revealed ones. We assume that the probability that i casts a vote $v_i^t \in \{0(\text{yellow}), 1(\text{blue})\}$ depends on the distance between the signal and the cutoff value, following the logistic distribution:

$$\varphi\left(v_i^t | s_i^t, x_i, \varepsilon_i\right) = \frac{\left(\exp \varepsilon_i \left(s_i^t - x_i\right)\right)^{v_i^t}}{1 + \exp \varepsilon_i \left(s_i^t - x_i\right)} \tag{1}$$

[21]In the post-experiment questionnaire, a few participants expressed reasonings which seemed to have no clear connection with any Bayesian update, such as "I chose yellow when I saw three or more yellow cards aligned in a row, since I thought it was a strong sign that the box is yellow." We assume that such deviations from rationality are accounted for by the error term.

where ε_i is the error coefficient. As φ pins down a distribution over voting profiles $(v_i)_{i=1}^n$, the social decision d made by the simple majority rule is a function of the realized signals $s = (s_i)_{i=1}^n$, cutoff values $x = (x_i)_{i=1}^n$, and the error coefficients $\varepsilon = (\varepsilon_i)_{i=1}^n$. Hence, for fixed ε, we can write the expected payoff as a function of the cutoff strategy profile $x \in [0, 10]^n$:[22]

$$u(x) = \mathbb{E}_{\omega,s}\left[\tilde{u}\left(\omega, d\left(s, x, \varepsilon\right)\right)\right].$$

Let T be the set of periods under consideration. Provided the observation $(s_i^t, v_i^t)_{t\in T}$, our estimation of the cutoff strategies is obtained by maximizing the likelihood:

$$\mathcal{L}\left(\left(s_i^t, v_i^t\right)_{t\in T}\Big| x_i, \varepsilon_i\right) = \prod_{t\in T} \varphi\left(v_i^t s_i^t, x_i, \varepsilon_i\right).$$

Some robustness checks are in order. During each period t, what we observe for each individual i is only one realized signal s_i^t and the casted vote v_i^t. Therefore, a cutoff estimation requires a pool of observations over periods. We checked estimations over different intervals of periods,[23] suggesting that no change in strategy occurred over time, including the possibility of learning. We thus concluded that the estimation is robust including observations over all 15 periods.[24]

We also tested different assumptions on ε_i: (i) common across phases for each individual, (ii) common across individuals for each phase and session, (iii) common across individuals and sessions for each phase, and (iv) common for individuals, sessions and phases. Comparing the estimated cutoff values between any pair of different assumptions on ε_i, the R^2 values are all close to one, implying that the estimation is robust.[25] In the following, we use the estimated cutoff values under assumption (iii), as our main concern here is the change in subjects' behavior as a function of the group size.

In Figure 3, histograms of the estimated cutoff values are shown for each phase. Several remarks are in order. First, we see a clear shift in the distribution from the symmetric payoffs to the asymmetric ones. Most notably, for each of the group sizes of 5, 9 and 19 with asymmetric payoffs, a peak in the frequencies is clearly visible on the intervals $[0, 1)$, representing 6%, 9% and 9% of all cutoff values, respectively. As the cutoff value 0 corresponds to the behavior of voting for blue regardless of the obtained signal, the presence

[22]Technically, we can extend the strategy space beyond the interval $[0, 10]$, given our specification of the errors in (1). However, we chose not to do so, as the main intuition of the model is unchanged, and we prefer to avoid any possible confusion caused by a cutoff value defined outside of the signal space.

[23] Our trials include $T = \{1, \cdots, 15\}$, $\{1, \cdots, 10\}$, $\{6, \cdots, 15\}$, $\{1, \cdots, 7\}$, $\{9, \cdots, 15\}$, $\{1, \cdots, 5\}$, $\{6, \cdots, 10\}$, and $\{11, \cdots, 15\}$. The t-tests did not reject the null hypothesis for any pairs of intervals at the $p = 0.10$ level,

[24]We also counted the number of votes consistent with the cutoff estimation obtained from $T = \{1, \cdots, 15\}$. Overall, 90.3% of actions are consistent. Inconsistent actions are spread across periods, and the t-statistics of the comparison between the first and last 7 periods are 1.06, 1.32 and 1.39 for the number of inconsistent blue actions, yellow actions and the sum, respectively, none statistically significant at the $p = 0.10$ level.

[25]The R^2 values for the 6 pairs are respectively, 0.90, 0.89, 0.89, 0.99, 0.99 and 0.99.

of the peaks suggests that a certain amount of participants used the signal-independent voting strategy, or at least one close to it. Second, about half of the estimated cutoff values are included in the interval $[4, 5)$ with asymmetric payoffs. The percentages in this interval for group sizes of 5, 9, and 19 are 54%, 51%, and 66%, respectively. Note that the unbiased strategy is represented by the cutoff value of 5. A cutoff value lower than 5 corresponds to a strategy biased in favor of voting for blue, the ex ante optimal choice. According to our estimation, roughly between one half and two thirds of participants used a cutoff strategy slightly biased towards the ex ante optimal choice. Third, no single player used a cutoff strategy higher than 8 in any phase with asymmetric payoffs. It is worth underlining that no signal-independent voting behavior to the other extreme direction (i.e. a cutoff value of 10, which corresponds to voting regardless of the signal for yellow, the ex ante suboptimal alternative) is observed with asymmetric payoffs. Fourth, a non-negligible amount of voting behaviors in favor of yellow are observed, even though they are rather a minority. The frequencies of cutoff values higher than 5 are 15%, 17% and 9%, respectively, in the three phases with asymmetric payoffs.

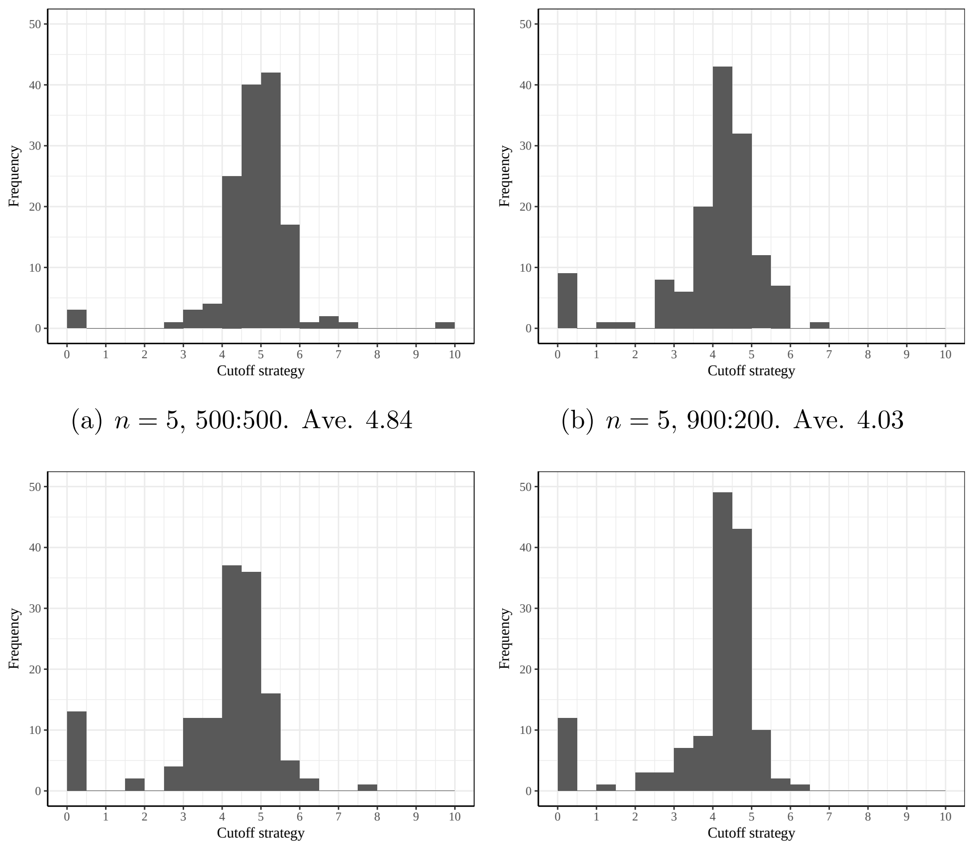

(a) $n = 5$, 500:500. Ave. 4.84

(b) $n = 5$, 900:200. Ave. 4.03

(c) $n = 9$, 900:200. Ave. 4.01

(d) $n = 19$, 900:200. Ave. 3.98

Figure 3: Histogram of the estimated cutoff strategies, 140 observations.

4.2 The model fit

In what follows, we evaluate three models: level-k (L), cognitive hierarchy (CH), and inclusive cognitive hierarchy (ICH), estimating the parameters which fit best to our experimental data.

Our aim is to find out the sequence of level strategies $\sigma = (\sigma_1, \cdots, \sigma_K)$ and the level distribution $f = (f_0, \cdots, f_K)$, which best fits the observed cutoff values $(x_i)_{i \in N}$, provided an exogenously fixed level-0 strategy σ_0.[26] Throughout the paper, we maintain the assumption that f follows the truncated Poisson distribution with coefficient τ. Estimation of the level distribution f thus boils down to finding the best-fit τ. Level strategy σ_k ($k \geq 1$) is the best response of a player holding a belief g_k. Remember that the distinction between the models stems uniquely from the way in which g_k is defined (see (L), (CH) and (ICH) in Section 2).

We assume that the observed cutoff values $(x_i)_{i \in N}$ are drawn from the distribution $\sum_{k=0}^{K} f_k \sigma_k$ with a logistic error. More precisely, the probability that the realized cutoff strategy is x_i is:

$$\phi(x_i | f, \sigma, \rho) = \sum_{k=0}^{K} f_k \mathbb{E}_y \left[\ell(x_i | y, \rho) | \, \sigma_k \right]$$

where expectation over y is drawn from the distribution σ_k,[27] and

$$\ell(x_i | y, \rho) = \frac{\exp\left(\frac{x_i - y}{\rho}\right)}{\rho \left(1 + \exp\left(\frac{x_i - y}{\rho}\right)\right)^2}$$

is the density function of the logistic distribution with mean y and the error coefficient ρ. Our estimation maximizes the likelihood, choosing the Poisson distribution f and the error coefficient ρ:

$$\prod_{i \in N} \phi(x_i | f, \sigma, \rho).$$

Given σ_0, the variables to be estimated are thus τ and ρ.

4.3 Level-0 Strategy

Before proceeding to the comparison of the models, we briefly discuss the choice of the level-0 strategy, which can be supported by the idea of *saliency*. As discussed in Crawford and Iriberri (2007), *inter alia*, some naturally occurring landscapes that are focal across the strategy space may constitute salient non-strategic features of a game and attract naïve assessments. For instance, a strategy space represented by a real interval, say $[m, M]$, may have its minimal point m, its maximal point M, and its midpoint $\frac{m+M}{2}$ as salient locations. In our game, two of these deserve close attention. First, if we expect that a non-strategic,

[26] Remember that K is the highest level under consideration. Our estimation sets $K = 2$ due to a heavy computational burden for $K \geq 3$ and large n.

[27] σ_k could be a mixed strategy, as is often assumed for $k = 0$.

level-0 player would evaluate her choices while disregarding others' strategic incentives, such a behavior corresponds to a strategy of the ex ante favored choice, i.e. always voting for blue. A salient location would then be 0. Second, if we expect that a non-strategic player would choose a strategy which maximizes the probability of making a right decision regardless of the winning point (and thus is not payoff-maximizing), then a salient strategy is the midpoint 5. Furthermore, a uniform randomization over all available pure strategies is often chosen as the level-0 strategy in the literature (see discussion in Camerer et al. (2004)).

Table 1 provides a comparison based on the maximum likelihood estimation described in the previous subsection, under the ICH model with the group size $n = 5$. We used a grid search, with both τ and ρ varied from 0.1 to 10 with an increment of 0.1. Among the four salient strategies of σ_0 we specified above, the highest log-likelihood is attained at $\sigma_0 = 0$, followed by the uniform distribution, with the lowest at $\sigma_0 = 5$ and 10. Note that the estimated τ^* hits the upper bound grid value 10 for $\sigma_0 = 5$ and 10. Provided the Poisson assumption imposed on f, a high value of τ^* means that the model can best fit the data by assigning the largest possible probability to the highest level-K (and the smallest to level-0). Therefore, one can expect at best the model to provide as a good fit as the symmetric Nash equilibrium. We thus conclude that setting such values of the level-0 strategy is of limited interest for describing the deviation from rational behavior, since these values of σ_0 can provide only a limited capacity of explaining our data beyond the Nash model.[28] Between the uniform σ_0 and $\sigma_0 = 0$, the latter fits better, which is coherent with the observations from our experiments with asymmetric payoffs that (i) there is a peak at 0 for all n, and (ii) no single player used a cutoff strategy higher than 8 in any phase. In what follows, we continue our analysis by setting the level-0 strategy to be $\sigma_0 = 0$.

	τ^*	ρ^*	$ICH1$ (σ_1)	$ICH2$ (σ_2)	LL
$\sigma_0 = 0$	3.9	2.5	4.45	4.28	-190.7
$\sigma_0 = 5$	10	1.6	4.17	4.23	-222.9
$\sigma_0 = 10$	10	1.6	3.84	4.21	-225.6
$\sigma_0 =$unif$[0, 10]$	4.3	2.3	4.00	4.20	-213.3

Table 1: Comparison of level-0 specifications. Maximum log-likelihood under the ICH model with $n = 5$.

4.4 Results

Table 2 summarizes the comparison of the models by maximum likelihood estimation.

[28]Similarly, we observed that τ^* hits the boundary value at $\sigma_0 = 5$ and $\sigma_0 = 10$ for $n = 9$ and $n = 19$, as well as under the L and the CH models, which implies that the prediction power of the models is quite limited under the assumption of $\sigma_0 = 5$ or $\sigma_0 = 10$.

		τ^*	ρ^*	σ_1	σ_2	LL
	L	0.1	0.4	10	0	-415.3
$n = 5$	CH	8.0	0.6	10	1.76	-366.4
	ICH	3.9	2.5	4.45	4.28	-190.7
	L	0.1	0.4	10	0	-416.3
$n = 9$	CH	10.0	0.7	10	2.52	-334.3
	ICH	5.0	2.3	4.91	4.67	-210.6
	L	0.1	0.4	10	0	-414.8
$n = 19$	CH	9.1	0.7	10	2.40	-338.0
	ICH	6.7	2.4	5.11	4.84	-210.7

Table 2: Comparison of the models by maximum log-likelihood. Pooled across sessions, $N = 140$. $\sigma_0 = 0$. Grid search for both the Poisson coefficient τ and the error coefficient ρ, from 0.1 to 10 with increment 0.1.

4.4.1 Level-k Model

An Lk player (level-k player in the L model) maximizes her payoff holding a belief that all other individuals play the $L(k-1)$ strategy. Since each individual has an incentive of correcting biases caused by all other players, the best reply of a Lk player is biased toward the opposite direction with respect to the Nash equilibrium as compared to the $L(k-1)$ strategy, and the degree of amplification increases as n increases.[29] We see in Table 2 that the cutoff strategy of the $L1$ player hits $\sigma_1 = 10$ as a response to the $L0$ strategy $\sigma_0 = 0$ for all values of $n = 5$, 9, and 19. A similar argument applies to $L2$ in the opposite direction, implying $\sigma_2 = 0$. Such an oscillation continues in the L model, and a bang-bang solution is obtained perpetually as k increases.

For all n, we observe that the most likely value τ^* hits the lower bound 0.1. This is not surprising, given the bang-bang strategies $\sigma_1 = 10$ and $\sigma_0 = \sigma_2 = 0$. As the truncated Poisson density of level 1, $g_2(1) = \frac{\tau}{1+\tau+\tau^2/2}$, is minimized at $\tau = 0$, the model fits best to the data observed in Figure 3, at the lowest value of τ. We underline that Battaglini et al. (2010) are the first to observe a limitation of the level-k model in binary-state, binary-action committee decision making in a laboratory experiment.

4.4.2 CH Model

The CH model stipulates that a CHk player (a level-k player in the CH model) maximizes her expected payoff holding a belief that other $n-1$ players have levels up to $k-1$. In particular, a $CH1$ player holds the belief that all other players have level 0, which is exactly

[29] Note that the best reply is always well-defined, since the probability of being pivotal is always non-zero, given the probabilistic voting action of each player specified by the logistic error function (1).

the same as the belief of an $L1$ player. In our game, the $CH1$ strategy is $\sigma_1 = 10$ for all $n = 5$, 9 and 19.

In Table 2, we observe that the values of τ^* are high for all n. This means that the best-fitting Poisson distribution assigns the highest density to level 2. Since the CH model attributes the corner values $\sigma_0 = 0$ and $\sigma_1 = 10$ for levels 0 and 1, the model fits to the data by assigning minimal densities to these corner strategies. We do not observe either a particular trend in τ^*, or a convergence of the $CH2$ strategies to Nash equilibrium, as n increases.[30] A key observation is an increasing sensitivity of the $CH2$ strategy for large n. Not only is the best-reply function in our game decreasing, but the *slope* of the best-reply function becomes steeper as n increases. Thus, the sensitivity of the best reply to the belief over the other players' strategies also increases as n increases. As we show later in Theorem 1, such an increasing sensitivity leads to a divergence of the CH strategy in asymptotically expanding games, as n increases. We provide further discussions in Section 5.

4.4.3 ICH Model

We observe that the estimated strategies of $ICH1$ and $ICH2$ are both increasing in n.[31] Comparing these values with the Nash equilibrium, the increase of both $ICH1$ and $ICH2$ is in line with the increase of the Nash equilibrium strategy with respect to n (Figure 2(b)). The intuition is that the Nash strategy monotonically converges to the unbiased strategy (i.e. 5), since all individuals equally share the prior bias caused by the asymmetric payoffs, and such an individual share converges to zero as n increases. As we discuss later in Section 5, $ICH1$ would converge to a value opposite to the prior bias with respect to the Nash equilibrium, and $ICH2$ would approach to the Nash equilibrium (Theorem 1). Our ICH estimations from the data are consistent with these theoretical predictions.

The estimated values of τ^* are 3.9, 5.0 and 6.7, for $n = 5$, 9 and 19, respectively. For these values of τ, it is interesting to compute the probability assigned to their own level by an inclusive player. For a level-2 player, it is $g_2(2) = \frac{\tau^2/2}{1+\tau+\tau^2/2}$, which is equal to 61%, 68% and 74%, respectively. For a level-1 player, it is $g_1(1) = \frac{\tau}{1+\tau}$, which is equal to 80%, 83% and 87%, respectively. These values are consistent with the responses we observed in Figure 1(b).

Another observation is that the best-fitting τ values are increasing as the group size increases. Given that τ is the expectation of the level drawn from the Poisson distribution,

[30] As a robustness check, the session-wise estimation shows that the difference of σ_2 is not statistically significant between $n = 5$ and $n = 9$. We observe that an increase from $n = 9$ to $n = 19$ is significant at the $p = 0.10$ level by the t-test, although such an increase is not observed in the histogram of all estimated strategies in Figure 3.

[31] This observation is robust with the session-wise estimations. The differences are statistically significant at the $p = 0.01$ level under the Wilcoxon test. The t-statistics for $ICH1$ are 13.9 and 6.94, and for $ICH2$ are 9.65 and 9.74, from $n = 5$ to $n = 9$ and from $n = 9$ to $n = 19$, respectively, implying that all differences are statistically significant at the $p = 0.01$ level.

a larger τ corresponds, ceteris paribus, to a higher value in the expected cognitive levels. Therefore, an increase in the estimated values of τ may be interpreted as evidence that the average cognitive level increases as the group size grows larger. This observation is not consistent with the findings of Guarnaschelli et al. (2000), in which evidence of lower expected levels with larger groups is reported, reflecting a larger cognitive load in large groups.

4.4.4 Comparison

Our maximum likelihood estimations show that the ICH model best fits the data for all $n = 5$, 9 and 19. The result is robust for the estimations done by each session separately: the log-likelihood is improved under the ICH model and the differences are significant at the $p = 0.01$ level, with the Wilcoxon test both against L and CH. From the observations above, we conclude that our laboratory experiments provide clear evidence that the inclusive cognitive hierarchy model (ICH) fits better to the data as compared with the standard level-k model (L) and the cognitive hierarchy model (CH).

Our data also provides us with statistics which allow a comparison of the fits between the above models and Nash equilibrium. The log-likelihood of the best fitting symmetric Nash equilibrium to our data is -222.9, -246.2 and -244.2 with $\rho^* = 1.6$, 1.3 and 1.3, for $n = 5$, 9 and 19, respectively. The Bayesian information criterion thus suggests that the ICH model fits better than Nash equilibrium, followed by the L and CH models. Our observation that Nash equilibrium fits better than the L and CH models is consistent with the findings of Battaglini et al. (2010) in an experiment on the swing voter's curse.

5 Asymptotic Properties

The main findings from our experimental results in the previous section are that (i) the maximum likelihood estimation exhibits a better fit of the ICH model to the data than L and CH models, and (ii) clear deviation from Nash behavior is observed, in an information aggregation problem with asymmetric payoffs.

Our objective in this section is to provide theoretical explanations according to the asymptotic property of the games. We show below that the distance of the L and CH strategies from the Nash equilibrium diverges, while that of the ICH strategy is bounded (Theorem 1) for the games in which the best-reply functions are asymptotically expanding. On the other hand, the strategies in all of those models are bounded (Theorem 2) for the games in which the best-reply functions are not asymptotically expanding. These analytical results suggest that whether or not the inclusiveness condition matters in describing the asymptotic behavior depends on the asymptotic property of the best-reply functions.

For the sake of tractability, we focus on linear quadratic games (Currarini and Feri (2015)) which have desirable features for our analysis. First, they are fully aggregative games (Cornes and Hartley (2012)), in which the payoff of each player is affected by the action profile of the players through the aggregate of the strategies of all players and her own strategy. This fits well to our current objective, as our goal here is to understand analytically how the optimal strategy of a player would be affected by the belief over the type of the other players. The fact that the strategy of the other players appears explicitly in an aggregative form allows us to obtain straightforward insights on the relationship between the shape of the best-reply functions and the players' beliefs over the strategies of the other players. Second, in a more technical convenience, linear quadratic games have a property such that, when a player holds a stochastic belief over the strategies of the other players, the maximizer of her expected payoff coincides with the best reply against the pure strategy which takes the expected value of the aggregate. This is because the linearity of the derivative allows us to switch the order of the partial derivative and the expectation. Then, facing heterogeneous beliefs over the other players' strategies, our analysis can simply focus on the best reply against the expectation of the beliefs, which provides us with a high tractability of the models.[32]

Consider n individuals, each of whom takes an action $x_i \in \mathbb{R}$. The payoff of player i in a linear quadratic game is a function of her own action x_i and the aggregate of the other players' actions $X_{-i} = \sum_{j \neq i} x_j$ in the following form:

$$u_i(x_i, X_{-i}) = \lambda^t x + x^t \Gamma x \tag{2}$$

where $x = (x_i \quad X_{-i})^t$, $\lambda = (\lambda_x \quad \lambda_X)^t$ and

$$\Gamma = \begin{pmatrix} \gamma_{11} & \gamma_{12} \\ \gamma_{12} & \gamma_{22} \end{pmatrix}.$$

There are several games of interest which fall into the class of linear quadratic games.

Example **1** *(A simple quadratic game) Suppose* $u_i(x_i, x_{-i}) = -\left(\sum_j x_j\right)^2$. *Then,* $\lambda^t = (0 \quad 0)$ *and*

$$\Gamma = \begin{pmatrix} -1 & -1 \\ -1 & -1 \end{pmatrix}.$$

Example **2** *(Cournot competition) Consider a Cournot competition. Suppose that the inverse demand function is linear* $P(Q) = a - bQ$, *and each firm has a constant marginal cost* c_i. *Let* q_i *be the quantity produced by firm* i *and* $Q_{-i} := \sum_{j \neq i} q_j$. *The profit of firm* i *is:*

$$\Pi_i = q_i (a - b(q_i + Q_{-i}) - c_i).$$

Then, $\lambda^t = (a - c_i \quad 0)$ *and*

$$\Gamma = \begin{pmatrix} -b & -\frac{b}{2} \\ -\frac{b}{2} & 0 \end{pmatrix}.$$

[32] Obtained insights could be extended to a game with more general payoff functions, to the extent that the second-degree Taylor expansion of the payoff function with respect to the aggregate strategy provides an approximation.

Example 3 *(Keynesian beauty contest games) Suppose that each of n players chooses a number x_i simultaneously, and each player's payoff is quadratic with respect to the distance between her own choice and the average of all players' choices multiplied by a constant $p \in (0,1)$. Then,*

$$u_i(x_i, X_{-i}) = -\left(x_i - p\left(\frac{x_i + X_{-i}}{n}\right)\right)^2.$$

Then, $\lambda^t = (0 \quad 0)$ and

$$\Gamma = \begin{pmatrix} -\left(1-\frac{p}{n}\right)^2 & \left(1-\frac{p}{n}\right)\frac{p}{n} \\ \left(1-\frac{p}{n}\right)\frac{p}{n} & -\left(\frac{p}{n}\right)^2 \end{pmatrix}.$$

Example 4 *(Public good provision game) Suppose that each agent contributes x_i to a public good and the cost is quadratic:*

$$u_i(x_i, X_{-i}) = \theta_i(x_i + X_{-i}) - c_i x_i^2.$$

Then, $\lambda^t = (\theta_i \quad \theta_i)$ and

$$\Gamma = \begin{pmatrix} -c_i & 0 \\ 0 & 0 \end{pmatrix}.$$

We impose some regularity conditions on the linear quadratic game in the form (2). First, we assume $\gamma_{11} < 0$. This implies that u_i has a unique maximizer for any X_{-i} and thus the best-reply function is well-defined. It is straightforward to show that the game defined by (2) has a unique symmetric Nash equilibrium:

$$x^* := -\frac{\lambda_x}{2(\gamma_{11} + (n-1)\gamma_{12})}.$$

We assume that the denominator is non-zero so that the symmetric Nash equilibrium is well-defined. By applying a parallel transformation $y_i := x_i - x^*$, (2) becomes:

$$u_i = \lambda^t x + x^t \Gamma x = \lambda_y^t y + y^t \Gamma y + c$$

where $Y_{-i} = \sum_{j \neq i} y_j$, $y = (y_i \quad Y_{-i})^t$, $\lambda_y = (0 \quad \lambda_Y)^t$, and λ_Y and c are independent of y. As the terms $\lambda_y^t y = \lambda_Y Y_{-i}$ and c have no strategic consequence on player i's behavior (i.e. the best-reply function of player i is unaffected), we can assume $\lambda_Y = 0$ and $c = 0$ without loss of generality. Therefore, in the following, we focus our attention on the games with the payoff function:

$$u_i = y^t \Gamma y, \tag{3}$$

with $\gamma_{11} < 0$ and $\gamma_{11} + (n-1)\gamma_{12} \neq 0$ (as in Angeletos and Pavan (2007)). Notice that there is a unique symmetric Nash equilibrium $y_i^* = 0$ for all i.

The first-order condition of player i is:

$$\frac{\partial u_i}{\partial y_i} = 2\gamma_{11} y_i + 2\gamma_{12} Y_{-i}.$$

When player i holds a stochastic belief over the strategies of the other players, the aggregate of the other players' strategies is a random variable $\tilde{Y}_{-i}$. Since the first-order condition is linear in Y_{-i} in quadratic games, the best reply against a mixed-strategy profile coincides with the best reply against the aggregate strategy which takes deterministically the expected value of the random variable:

$$BR_i\left(\tilde{Y}_{-i}\right) = -\frac{\gamma_{12}}{\gamma_{11}}\mathbb{E}\left[\tilde{Y}_{-i}\right]. \tag{4}$$

In order to describe asymptotic properties, we consider a sequence of linear quadratic games in which the number of players increases. More precisely, let $G(n) = \langle n, \mathbb{R}, (u_i^n)_{i=1}^n \rangle$ be a normal-form game with n players where the set of pure strategies is fixed as the set of real numbers $\mathbb{R}$,[33] and u_i^n is the payoff function of player i which satisfies (3). We analyze asymptotic properties of the strategies under the sequence of games $\{G(n)\}_{n=2}^{\infty}$.

Remember that the three models under our scrutiny here, L, CH and ICH, differ only in the assumption imposed on players' beliefs on the types of the other players. For each model, the strategy in each level is defined in the same way as in Definition 1. The only difference is that the frequency $g_k(h)$, assigned in the belief of a level-k player to the event in which each of the other players should be level-h, is specified by the equation (ICH) in Section 2 in the ICH model, but it is replaced by (L) (resp. (CH)) in the L (resp. CH) model. It is worth emphasizing that our results in this section do not hinge on the Poisson assumption concerning the underlying distribution f_k. We consider a sequence of level-symmetric strategies $\sigma = (\sigma_k)_{k\geq 0}$ where for each $k \geq 1$, σ_k maximizes the expected payoff under the belief $g_k(h)$ (Definition 1).

For each game $G(n)$ and in each of the three models, the level-0 strategy σ_0 is exogenously given, allowing the possibility of a mixed strategy. In order to make the comparison explicit across the models for $k \geq 1$, we add a superscript which represents the model, such as σ_k^L, σ_k^{CH}, and σ_k^{ICH}. Note that, by (4), $\sigma_k^M(n)$ are all pure strategies for $k \geq 1$ for each model $M \in \{L, CH, ICH\}$.

We assume that the following limit exists, allowing infinity:

$$A := \lim_{n\to\infty}\left|\frac{\gamma_{12}}{\gamma_{11}}n\right| \in \mathbb{R}_{\geq 0} \cup \{\infty\}.$$

Remember that $-\frac{\gamma_{12}}{\gamma_{11}}$ is the slope of the best-reply function (4). Since $\tilde{Y}_{-i}$ is the sum of the strategies of the other players, A is the limit of the slope of player i's best-reply function, as a function of the *average* of the other players.

First, we consider the case $A = \infty$. In such games, we say that the sequence of the games is **asymptotically expanding**, denoting the property that the sensitivity of one's

[33]The assumption of the one-dimensional, unbounded strategy space allows us to obtain clear insights on the convergence and/or divergence of the strategies. In the games with a compact, one-dimensional strategy space, these insights could be inherited with some adjustments, e.g. divergence corresponds to a bang-bang corner solution.

strategy to the average strategy of the other players increases without a bound. We show that the strategies diverge from the Nash equilibrium in the L and the CH models, while it is bounded in the ICH model.

THEOREM **1** *Consider a sequence of games* $\{G(n)\}_{n=2}^{\infty}$ *in which the payoff functions satisfy (3) for each* n*. Consider any* σ_0 *and let* $\mu := \mathbb{E}[\sigma_0]$*. Suppose* $A = \infty$*. For any* $\mu \neq 0$*,* $\lim_{n\to\infty} \left|\sigma_k^L(n)\right| = \infty$ *and* $\lim_{n\to\infty} \left|\sigma_k^{CH}(n)\right| = \infty$*, while* $\lim_{n\to\infty} \left|\sigma_1^{ICH}(n)\right| < \infty$ *and* $\lim_{n\to\infty} \left|\sigma_k^{ICH}(n)\right| = 0$ *for all* $k \geq 2$*.*

Among the examples described above, the sequence $\{G(n)\}_{n=2}^{\infty}$ is asymptotically expanding ($A = \infty$) in the simple quadratic game (Example 1) and in the linear Cournot competition (Example 2). A common feature of these games is that the aggregate of all players' strategies enters into each player's payoff in a way that the aggregate term does not dissipate for large n. When $A = \infty$, we show that the behaviors in the ICH model show a stark contrast with those in the L or in the CH model. The presence of the inclusiveness condition thus leads to an intrinsic difference in the prediction. Moreover, we show that the ICH strategy converges to the Nash equilibrium for any level $k \geq 2$.

We interpret these properties of the ICH model as accurately capturing the phenomenon which is frequently observed in the behavioral data. In addition to a naïve strategy (level-0), we often observe sophisticated behaviors, with possibly heterogeneous degrees of sophistication. What is implied by Theorem 1 for asymptotically expanding games is that there are fundamentally three degrees of strategic sophistication: naïve (level-0), partially sophisticated (level-1), and highly sophisticated (level-2 or more). Since the strategies of level-2 or higher all converge to the Nash equilibrium, behaviors in this class of games fall into one of the following three classes asymptotically: (i) naïve strategy which does not maximize the expected payoff, (ii) level-1 strategy which maximizes the payoff but under an inconsistent belief, and (iii) fully sophisticated strategy, which maximizes the payoff under the consistent belief.

Now, consider a sequence of games which satisfies the same conditions assumed in Theorem 1, except for that in A.

THEOREM **2** *Suppose* $A < \infty$*. For any* μ*,* $\left|\sigma_k^L(n)\right|$*,* $\left|\sigma_k^{CH}(n)\right|$ *and* $\left|\sigma_k^{ICH}(n)\right|$ *are all bounded as* $n \to \infty$*, for all* $k \geq 1$*.*

In the standard Keynesian beauty contest games (Example 3), we have:

$$A = \lim_{n\to\infty} \left| \frac{\left(1 - \frac{p}{n}\right) \frac{p}{n} n}{-\left(1 - \frac{p}{n}\right)^2} \right| = p < \infty.$$

In the games with $A < \infty$, the slope of the best-reply function is bounded as n goes to infinity. Hence, even in a game with a large number of players, the optimal strategy of a

		Data		L		CH		ICH	
Name	N	mean	s.d.	τ	s.d.	τ	s.d.	τ	s.d.
Lab	86	35.1	19.6	1.06	12.4	1.31	9.6	0.90	12.9
Classroom	138	26.8	17.7	1.90	12.7	2.53	8.6	1.51	14.1
Take-home	119	25.2	17.0	2.11	12.5	2.85	8.3	1.65	14.2
Theorists	54	17.8	24.3	3.46	10.4	5.15	6.0	2.42	14.3
Conference	92	16.8	20.1	3.72	9.9	5.64	5.6	2.54	14.3
E-mail	150	22.2	20.7	2.56	11.9	3.58	7.5	1.93	14.3
Newspapers	7893	23.1	20.2	2.41	12.1	3.34	7.7	1.85	14.3

Table 3: MoM estimation of beauty contest games. The first moment (mean) is matched.

player does not diverge. In the beauty contest games, we see that the aggregate term is relevant in each player's payoff to the degree of the *average* of all players.

Most of the remarkable results in the literature of strategic thinking have considered the games with $A < \infty$, such as the beauty contest games. Other examples include finite games, such as market entry games, coordination games, or centipede games, to name a few. Even for the finite games, as long as the game is dominance solvable, we can consider that the game falls into the class of $A < \infty$, in the sense that the infinite iteration of applying the best-reply function leads to a convergence to the Nash equilibrium.

It is rather straightforward to see why dominance-solvable games have been the most frequently chosen object of research in the literature. When the iterative application of the best-reply function leads to the unique Nash equilibrium, it corresponds to the high-level strategies converging to the Nash equilibrium. Such a property is often considered to fit well to the idea that the deviation from the Nash equilibrium caused by the limited cognitive ability of the low-level players dissipates as the cognitive limit goes to infinity.

Implied by Theorem 2 is that the inclusiveness condition has little relevance asymptotically in the games with $A < \infty$. It does not follow, however, that the inclusiveness condition makes no difference in the prediction power of the models.[34]

Table 3 shows the comparison of the L, CH, and ICH models fitting to the data in beauty contest games. The data sets are exactly the ones studied in Bosch-Domènech et al. (2002).[35] Following Camerer et al. (2004), the Poisson coefficient τ is estimated by a method of moment, matching the mean. By comparing the second moment, we observe that the ICH model fits slightly better to the data than L and CH.

[34] Papers in the literature provide comparative studies over different models, e.g. Breitmoser (2012).

[35] We express our deepest gratitude to Rosemarie Nagel for providing us with the data.

6 Conclusion

We have introduced a cognitive hierarchy model which allows inclusiveness in the belief over the cognitive levels of other players. We find that asymptotic properties of the group decision-making, especially asymptotic efficiency, exhibit a stark contrast depending on whether the inclusiveness condition is admitted in the model or not. Results from our laboratory experiment provide evidence for (i) systematic deviations from Nash equilibrium behavior, and (ii) a better fit to the data under the model with inclusiveness.

Our theoretical analysis implies that the asymptotic property of the slope of the best-reply function is the key ingredient to determining whether the asymptotic properties differ between the models with and without inclusiveness. Even with an increasing sensitivity of the best-reply function to the beliefs, inclusiveness prevents divergence of the strategies from fully rational behavior. Since the same property is shared by the best-reply function of the collective decision-making studied in this paper, the presence of the inclusiveness prevents strategies from diverging away from the symmetric Nash equilibrium, and hence provides asymptotic efficiency as the group size increases.

Most of the games studied with cognitive hierarchy models share the property that the best-reply function is asymptotically non-expansive (i.e. $A < \infty$). In such games, our analysis implies that the presence of inclusiveness has little relevance, at least asymptotically. Therefore, inclusiveness does not imperil the insights obtained in the existing cognitive hierarchy models, even though its explanatory power may vary as a function of model settings and parameters. A major example is the classical Keynesian beauty contest game. The data provides us with evidence that the model with inclusiveness fits the data toward the same direction as those without inclusiveness. Similar insights are inherited in the games with best-reply functions that are 'contractive' in a broader sense, such as dominance-solvable games, coordination games, and market entry games, among others. We hence do not insist on an intrinsic improvement of the predictive power of the cognitive hierarchy model with the presence of the inclusiveness condition in this class of games. The main message of this paper is that there are games in the other class in which the presence of inclusiveness matters. We think there are interesting games in this class, e.g. Cournot competitions, that are worth pursuing in further analysis.

Our interests go beyond the analytical results obtained in this paper. A crucial difference induced by the presence of inclusiveness is the existence of 'sophisticated' players. A highest-level player in the ICH model best replies holding the correct belief concerning the distribution of the levels of other players. This is not the case in the cognitive hierarchy models without inclusiveness. Players are supposed to be Savage rational, but full consistency of their beliefs is not postulated even for the highest level. In that sense, simply the existence of fully-sophisticated players may suffice to convey our message. However, our model consists of players who are naïve (level-0), best-replying but with inconsistent beliefs (level-1), and sophisticated (level-2). Rubinstein (2016) proposes a typology of players with two types, e.g. "instinctive" and "contemplative". The three types emerging from our approach based on cognitive hierarchy theory can be seen as a richer typology in the same vein that allows for an intermediary type. The beauty of the cognitive hierarchy models lies, we believe, in the heterogeneous degrees of belief inconsistency that can be explicitly accommodated. We would like to further understand the role of heterogeneous degrees of inconsistent beliefs under the existence of fully-sophisticated players. We leave this for future research.

References

Alaoui, L. and A. Penta (2016): "Endogenous depth of reasoning," *The Review of Economic Studies*, 83, 1297–1333.

Angeletos, G.-M. and A. Pavan (2007): "Efficient use of information and social value of information," *Econometrica*, 75, 1103–1142.

Austen-Smith, D. and J. S. Banks (1996): "Information aggregation, rationality, and the Condorcet jury theorem," *American Political Science Review*, 34–45.

Azmat, G., M. Bagues, A. Cabrales, and N. Iriberri (2018): "What you don't know... Can't hurt you? A field experiment on relative performance feedback in higher education," *mimeo.*

Battaglini, M., R. B. Morton, and T. R. Palfrey (2008): "Information aggregation and strategic abstention in large laboratory elections," *The American Economic Review*, 194–200.

——— (2010): "The swing voter's curse in the laboratory," *The Review of Economic Studies*, 77, 61–89.

Benoît, J.-P. and J. Dubra (2011): "Apparent overconfidence," *Econometrica*, 79, 1591–1625.

Bernheim, B. D. (1984): "Rationalizable strategic behavior," *Econometrica*, 1007–1028.

Bhattacharya, S., J. Duffy, and S. Kim (2017): "Voting with endogenous information acquisition: Experimental evidence," *Games and Economic Behavior*, 102, 316–338.

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

Breitmoser, Y. (2012): "Strategic reasoning in p-beauty contests," *Games and Economic Behavior*, 75, 555–569.

Brocas, I., J. D. Carrillo, S. W. Wang, and C. F. Camerer (2014): "Imperfect choice or imperfect attention? Understanding strategic thinking in private information games," *The Review of Economic Studies*.

Camerer, C. (2003): *Behavioral game theory: Experiments in strategic interaction*, Princeton University Press.

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

Camerer, C. and D. Lovallo (1999): "Overconfidence and Excess Entry: An Experimental Approach," *The American Economic Review*, 306–318.

Colman, A. M., B. D. Pulford, and C. L. Lawrence (2014): "Explaining strategic coordination: Cognitive hierarchy theory, strong Stackelberg reasoning, and team reasoning." *Decision*, 1, 35.

Cornes, R. and R. Hartley (2012): "Fully aggregative games," *Economics Letters*, 116, 631–633.

Costinot, A. and N. Kartik (2007): "On Optimal Voting Rules under Homogeneous Preferences," Working paper.

Crawford, V. P., M. A. Costa-Gomes, and N. Iriberri (2013): "Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications," *Journal of Economic Literature*, 51, 5–62.

Crawford, V. P. and N. Iriberri (2007): "Fatal attraction: Salience, naivete, and sophistication in experimental Hide-and-Seek games," *The American Economic Review*, 1731–1750.

Currarini, S. and F. Feri (2015): "Information sharing networks in linear quadratic games," *International Journal of Game Theory*, 44, 701–732.

Downs, A. (1957): *An economic theory of democracy*, New York: Harper Collins Publishers.

Esponda, I. and E. Vespa (2014): "Hypothetical thinking and information extraction in the laboratory," *American Economic Journal: Microeconomics*, 6, 180–202.

Feddersen, T. and W. Pesendorfer (1997): "Voting behavior and information aggregation in elections with private information," *Econometrica*, 1029–1058.

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

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

Gerling, K., H. P. Grüner, A. Kiel, and E. Schulte (2005): "Information acquisition and decision making in committees: A survey," *European Journal of Political Economy*, 21, 563–597.

Goeree, J. K., C. A. Holt, and T. R. Palfrey (2016): *Quantal Response Equilibrium: A Stochastic Theory of Games*, Princeton University Press.

Guarnaschelli, S., R. D. McKelvey, and T. R. Palfrey (2000): "An experimental study of jury decision rules," *American Political Science Review*, 407–423.

Hanaki, N., Y. Koriyama, A. Sutan, and M. Willinger (2018): "The strategic environment effect in beauty contest games," *mimeo*.

Isaac, R. M., J. M. Walker, and A. W. Williams (1994): "Group size and the voluntary provision of public goods: experimental evidence utilizing large groups," *Journal of Public Economics*, 54, 1–36.

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

Ostrom, E. (1998): "A behavioral approach to the rational choice theory of collective action: Presidential address, American Political Science Association, 1997," *American Political Science Review*, 92, 1–22.

Palfrey, T. R. (2016): "Experiments in Political Economy," in *The Handbook of Experimental Economics*, Princeton University Press, vol. 2, chap. 6, 347–434.

Rubinstein, A. (2016): "A typology of players: Between instinctive and contemplative," *The Quarterly Journal of Economics*, 131, 859–890.

Selten, R. (1975): "Reexamination of the perfectness concept for equilibrium points in extensive games," *International Journal of Game Theory*, 4, 25–55.

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

Woolley, A. W., C. F. Chabris, A. Pentland, N. Hashmi, and T. W. Malone (2010): "Evidence for a collective intelligence factor in the performance of human groups," *Science*, 330, 686–688.

A Mathematical Appendix

A.1 Proof of Theorem 1

Proof: Let
$$\alpha_n := -\frac{\gamma_{12}}{\gamma_{11}}(n-1).$$
By (4), α_n is the slope of the best-reply function with respect to the *average* of the other players' strategies. Using α_n, we can explicitly write the level-k strategy under each of the three models, L, CH, and ICH. Note that these models differ only in the belief held by each player, specified in equations (L), (CH) and (ICH) in Section 2.[36] By definition, $\lim_{n\to\infty}|\alpha_n| = A$.

In the L model, the strategy of the level-$(k+1)$ player is defined as the best reply to the level-k player. By (4) and (L),

$$\sigma_{k+1}^{L}(n) = \alpha_n \sigma_k^{L}(n) \quad \text{for } k \geq 0.$$

Hence,
$$\sigma_k^{L}(n) = (\alpha_n)^k \mu \quad \text{for } k \geq 1.$$
Therefore, for any $\mu \neq 0$ and any $k \geq 1$, we have $\lim_{n\to\infty}\left|\sigma_k^L(n)\right| = \infty$ if $A = \infty$, and bounded if $A < \infty$.

In the CH model, by (4) and (CH),

$$\sigma_k^{CH}(n) = \alpha_n \left(\sum_{h=0}^{k-1} g_k^{CH}(h)\, \sigma_h^{CH}(n)\right). \tag{5}$$

Especially, $\sigma_1^{CH}(n) = \alpha_n \mu$. For the sake of induction, assume that $\sigma_h^{CH}(n)$ is a polynomial of degree h with respect to α_n for $h \leq k-1$. Then, by (5), $\sigma_k^{CH}(n)$ is a polynomial of degree k with respect to α_n. Therefore, we have:

$$\sigma_k^{CH}(n) = \varphi_k(\alpha_n)\, \mu$$

where $\varphi_k(\cdot)$ is a polynomial of degree k. Therefore, for any $\mu \neq 0$ and any $k \geq 1$, we have $\lim_{n\to\infty}\left|\sigma_k^{CH}(n)\right| = \infty$ if $A = \infty$, and bounded if $A < \infty$.

In the ICH model, by (4) and (ICH),

$$\sigma_k^{ICH}(n) = \alpha_n \left(\sum_{h=0}^{k} g_k^{ICH}(h)\, \sigma_h^{ICH}(n)\right).$$

[36] In the proof, we write the (possibly mixed) strategy of a level-0 player as $\sigma_0 = \mu$, identifying it with its expected value, since expectation is the only relevant term which determines the best reply in the linear quadratic games.

Hence,

$$\sigma_k^{ICH}(n) = \frac{\alpha_n \sum_{h=0}^{k-1} g_k^{ICH}(h)\, \sigma_h^{ICH}(n)}{1 - \alpha_n g_k^{ICH}(k)}. \tag{6}$$

Now, suppose $A = \infty$. For $k = 1$,

$$\sigma_1^{ICH}(n) = \frac{\alpha_n g_1^{ICH}(0)\, \sigma_0}{1 - \alpha_n g_1^{ICH}(1)}.$$

As $\lim_{n\to\infty} |\alpha_n| = \infty$, we have $\lim_{n\to\infty} \sigma_1^{ICH} = -\frac{g_1^{ICH}(0)}{g_1^{ICH}(1)}\mu = -\frac{f_0}{f_1}\mu$.[37]

For $k = 2$, by (6),

$$\sigma_2^{ICH}(n) = \frac{\alpha_n \left(g_2^{ICH}(0)\, \sigma_0 + g_2^{ICH}(1)\, \sigma_1^{ICH}(n)\right)}{1 - \alpha_n g_2^{ICH}(2)}.$$

As $\lim_{n\to\infty} |\alpha_n| = \infty$, we have:

$$\lim_{n\to\infty} \sigma_2^{ICH}(n) = -\frac{g_2^{ICH}(0)\, \mu + g_2^{ICH}(1)\left(-\frac{f_0}{f_1}\mu\right)}{g_2^{ICH}(2)}.$$

Since $\frac{g_2^{ICH}(0)}{g_2^{ICH}(1)} = \frac{f_0}{f_1}$, we have $\lim_{n\to\infty} \sigma_2^{ICH}(n) = 0$. For $k > 2$,

$$\begin{aligned}
&\lim_{n\to\infty} \sigma_k^{ICH}(n) \\
= &\lim_{n\to\infty} \left(\frac{\sum_{h=0}^{k-1} g_k^{ICH}(h)\, \sigma_h^{ICH}(n)}{\frac{1}{\alpha_n} - g_k^{ICH}(k)}\right) \\
= &-\frac{1}{g_k^{ICH}(k)} \left(g_k^{ICH}(0)\, \mu + g_k^{ICH}(1)\left(-\frac{f_0}{f_1}\mu\right) + \sum_{h=2}^{k-1} g_k^{ICH}(h) \lim_{n\to\infty} \sigma_h^{ICH}(n)\right).
\end{aligned}$$

The first two terms in the bracket cancel out, since $\frac{g_k^{ICH}(0)}{g_k^{ICH}(1)} = \frac{f_0}{f_1}$. For the sake of induction, assume $\lim_{n\to\infty} \sigma_h^{ICH}(n) = 0$ for $2 \leq h \leq k-1$. Then, $\lim_{n\to\infty} \sigma_k^{ICH}(n) = 0$. ■

A.2 Proof of Theorem 2

Proof: Suppose $A < \infty$. Then by (6), for $k \geq 1$,

$$\lim_{n\to\infty} \sigma_k^{ICH}(n) = \frac{A \sum_{h=0}^{k-1} g_k^{ICH}(h)\, \sigma_h^{ICH}(n)}{1 - A g_k^{ICH}(k)}.$$

Especially, for $k = 1$,

$$\lim_{n\to\infty} \sigma_1^{ICH}(n) = \frac{A g_1^{ICH}(0)\, \mu}{1 - A g_1^{ICH}(1)} < \infty.$$

[37] Remember that g_k is the truncated distribution induced by f, the underlying distribution over levels defined in Definition 1.

For the sake of induction, assume $\lim_{n\to\infty}\left|\sigma_h^{ICH}(n)\right| =: s_h < \infty$ for $1 \leq h \leq k-1$. Then, for $k \geq 2$,

$$\lim_{n\to\infty}\left|\sigma_k^{ICH}(n)\right| \leq \left|\frac{A\sum_{h=0}^{k-1} g_k^{ICH}(h)\, s_h}{1 - A g_k^{ICH}(k)}\right| < \infty.$$

■