Games and Economic Behavior 99 (2016) 257–274

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

A generalized cognitive hierarchy model of games

Juin-Kuan Chong[a,*,1], Teck-Hua Ho[b], Colin Camerer[c]

[a] National University of Singapore, Singapore
[b] National University of Singapore and University of California, Berkeley, USA
[c] California Institute of Technology, USA

ARTICLE INFO

Article history:
Received 4 June 2014
Available online 5 September 2016

JEL classification:
C72
C91
D03
D83

Keywords:
Cognitive hierarchy
Level-k model
Level-m model
Generalized cognitive hierarchy
Non-equilibrium structural models
Behavioral game theory

ABSTRACT

Subjects in simple games frequently exhibit non-equilibrium behaviors. Cognitive hierarchy (CH) and level k (LK) are two prevailing structural models that capture such behaviors well. This paper proposes a generalized CH (GCH) model that nests a variant of the LK model, called LM. GCH differs from CH in two ways. First, each lower level's actual frequency is exponentially weighted with α to form level-k's belief on relative proportions; α captures *stereotype bias*. CH assumes no stereotype bias ($\alpha = 1$) and LM assumes extreme bias ($\alpha = \infty$). Second, GCH replaces random choice with *minimum aversion* for level 0. Level 0s are more likely to choose strategies that never yield the minimum payoff for any of the opponent's strategies. GCH captures behaviors better than CH and LK in fifty-five $n \times m$ games from four datasets. Robustness tests using three new games further validate GCHs descriptive strength over CH and LK.

1. Introduction

Subjects frequently do not play equilibrium in simple one-shot games. Sometimes these non-equilibrium behaviors persist even when 1) subjects are motivated by substantial financial incentives and/or 2) they play these games repeatedly. These non-equilibrium behaviors pose serious challenges for standard equilibrium models. Recent research has developed structural non-equilibrium models as alternatives and these models have shown promise in explaining and predicting these out-of-equilibrium behaviors in a wide variety of games and markets. See Crawford (2013) for a comprehensive review.

Standard equilibrium models make three fundamental assumptions to generate their predictions: 1) players form subjective beliefs about what their opponents will do, 2) players choose actions to maximize their expected payoffs conditional on their subjective beliefs (i.e., they are subjective utility maximizers), and 3) players' subjective beliefs are always correct (i.e., their beliefs always match their opponents' actual actions). The third assumption is strong because it uses opponents' actual actions to pin down players' *ex ante* beliefs. Since players' beliefs are always correct under standard equilibrium models, no players are ever surprised by their opponents' actions. Non-equilibrium structural models, on the other hand, do not require

* Corresponding author.
E-mail address: bizcjk@nus.edu.sg (J.-K. Chong).
[1] We thank Vince Crawford, the associate editor and three anonymous reviewers. Their constructive comments and suggestions substantially improved the paper. We thank Sadat Reza, Rehan Ali, Zhang Xing and Wu Hang for their research assistance.

http://dx.doi.org/10.1016/j.geb.2016.08.007
0899-8256/

players' subjective beliefs to be always correct. As a result, players can be surprised by their opponents' actions in actual game plays as we observe in many real world situations.

To generate their behavioral predictions, non-equilibrium models posit that players are heterogeneous and have different levels of thinking ability. The proportion of players with thinking level k ($k = 0, 1, \ldots, \infty$) is denoted by $f(k)$ where $\sum_{k=0}^{\infty} f(k) = 1$. These models start with an explicit assumption on how non-strategic level-0 players behave. Higher level players' behaviors are then determined iteratively by assuming they best-respond to lower level players. As a consequence, the aggregate behavioral prediction is simply the sum of each level's behavior weighted by its actual proportion.

Cognitive hierarchy (CH) (Camerer et al., 2004) and level k (LK) models (Nagel, 1995; Stahl and Wilson, 1994, 1995; Costa-Gomes et al., 2001; Costa-Gomes and Crawford, 2006) have been successful in explaining non-equilibrium behaviors, but they differ in their assumptions on how higher level players best respond to lower level players. Both models start with the specification of a level-0 rule. While LK assumes that level k best-responds exclusively to level $k-1$, CH assumes that level k best responds to all lower levels (i.e., $0, 1, \ldots, k-2, k-1$) with the perceived proportion of each lower level equal to its normalized true proportion (i.e., the perceived proportion of level $h = \frac{f(h)}{\sum_{h'=0}^{k-1} f(h')}$, for all $h < k$).

This paper develops a generalization of CH which nests a variant of LK, which we call level m (LM). Like LK, LM assumes each player best-responds to a lower-level player. Unlike LK, however, LM assumes that level k best-responds to the most frequently occurring lower level (i.e., the modal level below k), reflecting stereotype bias. If the modal level happens to be $k-1$, LK and LM make the same prediction for level-k behavior.[2]

Generalized cognitive hierarchy (GCH) assumes that level k best-responds to all lower levels but the relative proportion of each lower level is obtained by first weighting its actual frequency with an exponent parameter α and then normalizing it so that all relative proportions sum to 1. Specifically, level k's belief of the relative proportion of level h is equal to $\frac{f(h)^{\alpha}}{\sum_{h'=0}^{k-1} f(h')^{\alpha}}$ ($\forall h < k$). The relative proportion of GCH reduces to that of CH when $\alpha = 1$ and to that of LM when $\alpha = \infty$. When $\alpha = \infty$, the perceived proportion of level h equals to 1 if $h = \text{argmax}_{h'<k} f(h')$ and equals to 0 otherwise.

The behavioral predictions of non-equilibrium structural models depend critically on level 0's behavioral rule because higher level rules are defined iteratively based on lower level ones. Despite its importance, there is a lack of a plausible model for level 0 that can be applied uniformly across all games. Existing research uses two different models of the level-0 rule. In one model, level 0 is assumed to randomize uniformly across all possible actions (e.g., Camerer et al., 2004). While the random choice rule is parsimonious and general, it appears too simplistic as a plausible rule of behavior. In another model, saliency is used to derive level 0's rule from the payoff structure (e.g., Crawford and Iriberri, 2007). However, saliency is not precisely defined and different people may invoke different saliency principles.

We also develop a behaviorally plausible model of level 0. Our model hypothesizes that level-0 players use *minimum avoidance* as the saliency principle in choosing their actions. Specifically, they prefer strategies that never yield a minimum payoff given any opponent's strategy. Let's illustrate the minimum avoidance principle with a $n \times m$ game matrix. We focus on row player i who has n strategies. Strategy j belongs to the preferred set S_i if strategy j never yields the minimum payoff for any of the opponent's m possible strategies. Our model posits that level-0 players are β times more likely to choose strategies in set S_i than to choose those outside S_i. Formally, level-0 players choose each strategy in S_i with probability $\frac{\beta}{n-|S_i|+|S_i|\cdot\beta}$ and each strategy not in S_i with probability $\frac{1}{n-|S_i|+|S_i|\cdot\beta}$, where $|S_i|$ is the cardinality of the set S_i. Below, we show that this simple model of the level-0 rule captures asymmetric dominance and compromise effects, two empirical regularities well-documented in the individual choice literature. In addition, it predicts that level 0 is less likely to choose dominated strategies, which fixes an obvious weakness of the random choice model. When $\beta = 1$, the model reduces to the random choice rule; when $\beta = \infty$, level 0 never chooses strategies outside the preferred set S_i.

GCH has two parameters (i.e., α and β) more than the CH model. The α parameter allows us to establish an explicit link between CH and a variant of LK (i.e., LM). The β parameter uses minimum avoidance as the saliency principle and allows us to study random choice and saliency as two extreme cases of a well-defined model of level 0.

We fit GCH to 55 $n \times m$ games (with $n, m = 2, 3, 4, 5, 6$) from four distinct datasets. We show that our generalized model describes and predicts behaviors better than the CH and LK models. Both α and β are important and add to the overall fit. The parameter α is always greater than 1, suggesting that higher level players exhibit stereotype bias. The parameter β is always greater than 1, suggesting that level-0 players are averse to minimum payoffs. We also estimate nested special cases of GCH and show that they are rejected in favor of GCH.

To check for robustness, we apply GCH to three new games: a coordination game with an irrelevant alternative (from Amaldoss et al., 2008), a compromise game with private information (from Carrillo and Palfrey, 2009) and a money request game for inducing level-k reasoning (from Arad and Rubinstein, 2012).

In the first game, Amaldoss et al. (2008) add an irrelevant (i.e., dominated) alternative to a coordination game and show that subjects are better able to coordinate their behavior. GCH captures the improved coordination with level 0 choosing the non-minimum strategy with higher probability, and the higher levels showing a more concentrated response because of stereotype bias. This game offers a clean and simple showcase of when and how GCH works. In the second game, we show how GCH and other strategic thinking models can be extended to games with private information. In the fight-or-compromise game, Carrillo and Palfrey (2009) find that people tend to fight less frequently than the Nash

[2] Unlike LK, LM does not predict oscillating behavior in some games such as the market entry and money request games.

equilibrium would predict. While all non-equilibrium models predict more fighting as we move to higher thinking levels, our results show that GCH provides a better balance in calibration and in validation between the compromising lower levels and the fighting higher levels. In the third game, Arad and Rubinstein (2012) propose a game where a player is motivated to request exactly one unit of money less than an opponent. The game was designed to induce and detect level-k reasoning. Eight out of the ten strategies belong to the preferred set S_i. As a consequence, the saliency of minimum aversion becomes less relevant and is correctly captured by GCH and reflected in the insignificant β parameter. All non-equilibrium models correctly estimate mean thinking levels to be between 2 and 3 and seem to capture behaviors equally well.

The rest of the paper is organized as follows. Section 2 describes the model and shows how the GCH model nests CH and LM as special cases. In addition, the similarities and differences between LM and LK models are discussed. Section 3 provides maximum likelihood estimates for GCH, CH, LM, LK and a two-parameter LK model. Section 4 describes the application of GCH to three new games. Section 5 concludes our discussion and suggests future research directions.

2. Model

2.1. Notations

We focus the GCH model on two-player matrix games. Beginning with notation, we index players by i and denote the jth strategy of player i by s_i^j. Player i chooses from m_i possible strategies. Player i's opponent is denoted by $-i$ and the opponent's strategy by $s_{-i}^{j'}$; there are m_{-i} such strategies. Player i's payoff is denoted by $\pi_i(s_i^j, s_{-i}^{j'})$. Players are heterogeneous and choose a decision rule from a rule hierarchy. We index rule levels in the rule hierarchy by k ($k = 0, 1, 2, \ldots, k-1, k, \ldots, \infty$) and denote the proportion of players using the level-k rule by $f(k)$, where $\sum_{k=0}^{\infty} f(k) = 1$. Henceforth, we shall use the terms "level k" and "level-k player" to refer to a player using the level-k rule. Lastly, we assume $f(.)$ follows a one-parameter Poisson distribution with mean and variance τ. Formally, $f(k) = \frac{\tau^k \cdot e^{-k\cdot\tau}}{k!}$.

We denote the probability that player i with rule level k will choose strategy s_i^j by $P_k(s_i^j)$. We posit that level k best responds to a proportionate combination of all lower levels from level 0 to level $k-1$ and will choose the strategy that maximizes the expected payoff. The expected payoff of choosing s_i^j is computed on the basis of level k's belief of what lower levels will do. We denote the expected payoff of strategy s_i^j for level k as $E_k(s_i^j)$. The expected payoff is computed as follows:

$$E_k(s_i^j) = \sum_{j'=1}^{m_{-i}} \pi_i(s_i^j, s_{-i}^{j'})\{\sum_{h=0}^{k-1} g_k(h) \cdot P_h(s_{-i}^{j'})\} \tag{1}$$

where $g_k(h)$ is level k's belief of the proportion of level h where $h < k$; $P_h(s_{-i}^{j'})$ is the choice probability of the corresponding level-h player.

Player i chooses strategy j with probability 1 if j gives the maximum expected payoff. Formally,

$$P_k(s_i^j) = \begin{cases} 1 & \text{if } j = \arg\max_{j''} E_k(s_i^{j''}) \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

and where there is a tie for the maximum, equal probabilities are assigned to all strategies with the maximum.

To fully specify the model, two model components in the expected payoff equation (1) have to be determined: the behavior of level 0 $P_0(s_i^j)$ and the beliefs of opponents $g_k(h)$. Beginning with $P_0(s_i^j)$ for level 0, we can iteratively compute the choice probabilities of level k for all k from the belief $g_k(h)$. In the next two subsections, we specify the two model components and elaborate on the linkage between GCH and CH, and LK.

2.2. Model of level 0

There are two existing models of the level-0 rule. One model assumes that level 0 uniformly randomizes among all possible actions. This model has the advantage of being well-defined and general. However, it may not be adequate for describing behaviors of actual subjects. The second model assumes that level 0 chooses a salient action in the feasible strategy space. This model can be ad-hoc because it is often difficult *ex ante* to determine which action is the most salient in a game.

GCH assumes that level-0 players are averse to receiving minimum payoffs. Specifically, level 0 is more likely to choose from a set of strategies that will never yield the minimum payoff for any of the opponent's strategies. A strategy in this set will always provide a strictly higher payoff than at least one other strategy for any of the opponent's strategies. In short, these strategies will never yield the worst outcome. Level 0 is non-strategic (not best-responding) but is sufficiently motivated to engage in the cognitively simple task (mostly eye-balling the payoff matrix) of constructing a preferred set.

For player i, we define i's "never worst set" S_i as the set of strategies that will never yield the worst payoff given any strategy of the opponent $-i$. Specifically, any strategy j in set S_i will yield a strictly higher payoff than at least one other

Table 1
Game 2 from Cooper and Van Huyck (2003).

	L	R	Data
T	0.4, 0.5	0.6, **0.3**	0.909
B	**0.2**, **0.1**	**0.2**, **0.1**	0.091
Data	0.738	0.262	

Note: payoffs in bold indicate the minimum payoff for the corresponding columns or rows. The never worst set for the row player consists of the dominant strategy T while the never worst set for the column player is empty.

strategy in every possible strategy of $-i$. Formally, if i's strategies are indexed by j as well as j'' and $-i$'s strategies are indexed by j', then $S_i = \{j : \pi_i(s_i^j, s_{-i}^{j'}) > \min_{j''}\{\pi_i(s_i^{j''}, s_{-i}^{j'})\},\ \ \forall j'\}$. GCH posits that level-0 player i is β times more likely to choose strategies in S_i than to choose strategies not in S_i. Hence, the choice probability of level 0 for each strategy s_i^j is:

$$P_0(s_i^j) = \begin{cases} \frac{\beta}{m_i - |S_i| + |S_i| \cdot \beta} & \text{if } j \in S_i \\ \frac{1}{m_i - |S_i| + |S_i| \cdot \beta} & \text{otherwise} \end{cases} \tag{3}$$

where $|S_i|$ is the cardinality of S_i. Note that S_i can be a null set and in this case our model reduces to the uniform randomization rule with level-0 players choosing each strategy with an equal probability of $\frac{1}{m_i}$.

The proposed model of non-strategic level 0 suggests that non-dominated strategies are chosen more frequently than dominated strategies, as one would expect intuitively. Furthermore, the model also captures two well-established empirical regularities in the individual choice literature: the compromise effect (Simonson, 1989) and the asymmetric dominance effect (Huber et al., 1982) where the predicted choices fit the definition of never worst set. We use three games in the datasets to demonstrate how these phenomena, if present, are captured with this model of level 0; the note under each payoff table explains in detail how the never worst sets are derived.

1. *Preference for non-dominated strategies*
   When player i has two strategies and one of them is strictly dominant, a level-0 player chooses this strategy with the probability of $\frac{\beta}{1+\beta}$. Table 1 shows one such game, Game 2 from Cooper and Van Huyck (2003). The dominant strategy is labeled T, which pays 0.4 and 0.6 when the opponent chooses L and R respectively. It is more than what strategy B will pay in both scenarios which is 0.2. GCH allows us to investigate whether a level-0 player who is β times more likely to choose T than to choose B would describe and predict actual behavior better in this kind of game. There are ten such games with a dominant strategy for at least one of the two players in the 55 games we study.
2. *Compromise effect*
   Decision makers have a tendency to select a compromise alternative in individual choice. Specifically, in a choice among three alternatives with two attributes, they often prefer the alternative that yields intermediate payoffs in both attributes to alternatives that yield a high payoff in one attribute and a low payoff in the other attribute. Extending this behavioral phenomenon of the compromise effect to a game setting, a player may prefer strategies that yield intermediate payoffs in all outcomes to strategies that yield minimum payoffs in some outcomes (and possibly higher payoffs in other outcomes). We use Game 10 from Camerer et al. (2004) to illustrate this behavioral tendency. This game is a 3×3 matrix game with the payoff matrix shown in Table 2. Strategy L of a column player is a compromise alternative because it gives intermediate payoffs of 20, 10 and 0 when the row player chooses T, M and B respectively. Strategy L never receives the worst payoff in all possible outcomes. GCH predicts that level 0 chooses L with probability $\frac{\beta}{2+\beta}$ and C or R with probability $\frac{1}{2+\beta}$. In the 55 games we analyze, there are ten games with such a payoff structure.
3. *Asymmetric dominance effect*
   Equilibrium models predict that the choices of players will not be influenced by the presence of a dominated alternative. However, experimental evidence (e.g. Amaldoss et al., 2008) shows that a strategy would appear to be more attractive when it dominates another strategy in the set. We use Game 6 from Costa-Gomes et al. (2001) to illustrate how this phenomenon is captured in the model of level 0. This game is a 3×2 matrix game with the payoff matrix shown in Table 3. Strategy M is a dominated strategy (dominated by B). If row players exhibit an asymmetric dominance effect, they will be more likely to choose strategy B even though this game has a unique pure strategy Nash equilibrium at (T,L). GCH can capture this phenomenon since its level 0 chooses B with probability $\frac{\beta}{2+\beta}$ and strategy T and M with probability $\frac{1}{2+\beta}$. There are seven games with an asymmetric dominance payoff structure for at least of one of the two players in the 55 games we study.

2.3. Model of opponents, $g_k(h)$

In GCH, level k believes that the relative proportion of lower level h is:

Table 2
Game 10 from Camerer et al. (2004).

	L	C	R	Data
T	**−20**, 20	30, **−30**	**−30**, 30	0.00
M	−10, 10	30, 30	0, **0**	0.96
B	0, 0	**−10**, 10	10, **−10**	0.04
Data	0.29	0.58	0.13	

Note: payoffs in bold indicate the minimum payoff for the corresponding columns or rows. The never worst set for the row player consists of strategy M while the never worst set for the column player is strategy L. M is the compromise alternative which gives the intermediate payoff of −10 and 0 when the opponent chooses L and R.

Table 3
Game 6 from Costa-Gomes et al. (2001).

	L	R	Data
T	74, 62	**43**, **40**	0.14
M	**25**, **12**	76, 93	0.11
B	59, 37	94, **16**	0.75
Data	0.70	0.30	

Note: payoffs in bold indicate the minimum payoff for the corresponding columns or rows. The never worst set for the row player consists of strategy B which dominates strategy M. The never worst set for the column player is empty.

$$g_k(h) = \frac{f(h)^\alpha}{\sum_{h'=0}^{k-1} f(h')^\alpha}, \quad \forall h < k. \tag{4}$$

The parameter $\alpha(\geq 1)$ captures the tendency of players exhibiting stereotype bias. When $\alpha = 1$, level k's belief about lower levels equals its normalized true proportion, hence correctly reflecting the relative proportions of the lower levels. When $\alpha > 1$, player k's belief about lower levels is concentrated on the more frequently occurring lower levels. A higher α has two effects. It captures a more concentrated belief and it predicts the concentration to center on the more commonly occurring lower levels. This concentration in belief applies not only to the player's own belief about lower levels, it carries over to the player's belief about the beliefs of lower levels as well. In other words, the player's prediction of lower-level choice probabilities $P_h(s_{-i}^{j'})$ $\forall h$ (in equation (1)) carries the same degree of belief concentration or the same α.

This concentration or simplification of beliefs, known in social psychology literature as the out-group homogeneity effect (Quattrone and Jones, 1980; Park and Rothbart, 1982), leads to stereotype bias. In social psychology, stereotype bias is a simplification of beliefs (where the heterogeneity of out-group members is reduced) when people are faced with a heavy cognitive load or when their cognitive resources are depleted. Hence, this simplification of belief is done to reduce cognitive load and to achieve efficiency in information processing and decision making. The parameter α in GCH captures this behavioral tendency concisely.

While α is a free parameter in GCH, CH assumes $\alpha = 1$ (and $\beta = 1$). When $\alpha = \infty$, level k exhibits extreme stereotype bias and believes that the opponent is only of the most frequently occurring lower rule. We single out the special case $\alpha = \infty$ (and $\beta = 1$) for discussion due to its similarity to the LK model. We call this special case level m (LM). In both LK and LM, a level-k player believes that the opponent is of a singular type. While the singular type is level $k - 1$ in LK, it is the modal lower rule in LM. If the modal rule is level $k - 1$ for all k, then the two models are identical.

For Poisson distributions with non-integral τ, the two models generate an identical prediction for level k up to the mode + 1 (i.e. $\lfloor\tau\rfloor + 1$), but may generate a different prediction for level k above mode + 1.[3] In other words, if there is a difference in prediction between the LK and LM models, the difference comes from the best responses of levels mode + 2 and above.

For a matrix game, the LK model predicts two possible level-k behaviors. For the first type of behavior, LK predicts that best responses above a finite level k' converge to the same strategy. For the second type of behavior, LK predicts that the best responses above a finite level k' cycle through a fixed set of strategies infinitely. If LK predicts the convergent behavior, then the LM model also predicts the same behavior when $\tau \geq k'$. Here is why. For level k above k', best responding to $k - 1$ (in LK) is the same as best responding to level k' (in LM) because both best respond to the same strategy. If LK predicts

[3] For integral τ, the Poisson distribution has two modes, τ and $\tau - 1$. Therefore, all higher levels in LM best-respond to a 50–50 mix of these two modal levels.

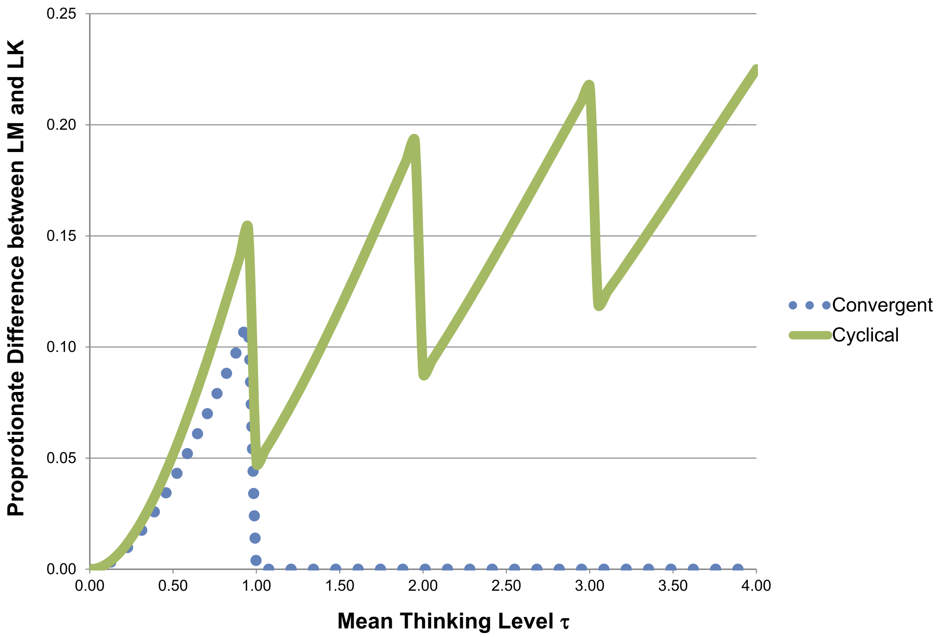

Fig. 1. Average proportion of difference in predicted behaviors between LM and LK models.

the cyclical behavior, the best responses for k' and above in LK (which cycles through a fixed set of strategies infinitely) are different from those in LM (which always predicts the same best response to mode) regardless of τ.

The 2×2 games offer a simple way to empirically demonstrate the point above. With only two strategies to consider, LK and LM producing the same convergent behaviors can be demonstrated using any $\tau > 1$. Since there are only two strategies, LK predicts that convergence starts at no higher than level 2 (who best responds to level 1). And, level 2 in LM will also best respond to the modal level 1 (who chooses the same strategy as level 1 in LK).

To contrast the convergent behaviors and the cyclical behaviors, we separate the 2×2 games in our dataset into two subsets: the first subset consists of 12 games where the LK model makes cyclical predictions and the second subset consists of 9 games where the LK model predicts higher levels' choices converging to the same strategy. For each game at a given τ, we compute the empirical proportion of all level-k players in the LM model who do not pick the same strategy as level k in the LK model. Formally, this proportion of difference is $\sum_k f(k) \cdot I[P_k^{LM}(s_i^j) \neq P_k^{LK}(s_i^j)]$ where $I[\cdot]$ is the indicator function. We average over all games in each subset to produce the average proportion of difference in prediction for the subset. We vary τ over the range of $0 \leq \tau \leq 4$ in steps of 0.05. The result is shown in Fig. 1. The figure shows that LK and LM make the same prediction for the convergent subset when $\tau > 1$; and it shows that the difference in prediction between LK and LM fluctuates in the cyclical subset.

The rest of the paper compares and contrasts the empirical analysis of GCH to other strategic thinking models.

3. Estimation and results

3.1. Likelihood function

Let us first derive the likelihood function. GCH has three parameters to be estimated, namely, τ, α, β. Let $P_k(s_i^j|\tau,\alpha,\beta)$ be GCH's predicted probability of level k choosing strategy s_i^j. We start with level-0 players. $P_0(s_i^j|\tau,\alpha,\beta) = \frac{\beta}{m_i - |S_i| + |S_i| \cdot \beta}$ if s_i^j is in set S_i and $\frac{1}{m_i - |S_i| + |S_i| \cdot \beta}$ if s_i^j is not.

For $k \geq 1$, we use equation (1) to compute the expected payoff for each strategy s_i^j. Players are assumed to choose the strategy that yields the highest expected payoff. Specifically, $P_k(s_i^j|\tau,\alpha,\beta) = 1$ if s_i^j gives the highest expected payoff and 0 otherwise.[4] The predicted choice probability for strategy s_i^j is simply an aggregation of best responses from all thinking levels weighted by their proportion $f(k|\tau)$:

$$R(s_i^j|\tau,\alpha,\beta) = \sum_k f(k|\tau) \cdot P_k(s_i^j|\tau,\alpha,\beta) \tag{5}$$

[4] If more than one strategy give the highest expected payoff, we assume players choose these strategies with equal probabilities.

Table 4
Data descriptions.

Dataset	Total	Number of Games[1] 2x2	3x2	3x3	4x2	4x4	5x5	6x6	Number of Observations
Stahl and Wilson (1995)	12	-	-	12	-	-	-	-	576
Cooper and Van Huyck (2003)	8	8	-	-	-	-	-	-	1496
Costa-Gomes, Crawford and Broseta (2001)	13	3	8	-	2	-	-	-	468
Camerer, Ho and Chong (2004)	22	10	-	7	-	2	1	2	1056
Total	55	21	8	19	2	2	1	2	
Number of Observations	3596	2084	288	912	72	96	48	96	

[1] Dimensions of the payoff matrix, defined as the number of rows multiplied by the number of columns.

where $R(s_i^j|\tau,\alpha,\beta)$ is the aggregate choice probability of strategy s_i^j.

We then aggregate over all subjects i and all strategies j to form the log-likelihood function:

$$LL(\tau,\alpha,\beta)=\sum_i\sum_j I(s_i^j)\cdot\ln R(s_i^j,\tau,\alpha,\beta) \tag{6}$$

where $I(s_i^j)$ is the indicator function; $I(s_i^j)=1$ if player i chooses strategy s_i^j and 0 otherwise.

We maximize the above log-likelihood function to find the optimal values for τ,α and β. This log-likelihood function can be highly non-linear. To avoid being trapped in local optima, we perform an extensive grid search over the parametric space with τ ranges from 0 to 10, α ranges from 1 to 20, and β from 1 to 10.[5] Although computationally demanding, we can run multiple grid searches in parallel. We have also capped the thinking level at 20 because the proportion of thinking levels above 20 is negligible. The step size of grid search for the three parameters is set at 0.01.

The maximum likelihood estimation used four distinct sets of matrix games. These matrix games are two-person, one-shot games. They include the 12 games in Stahl and Wilson (1995), the 8 games in Cooper and Van Huyck (2003), the 13 games in Costa-Gomes et al. (2001), and the 22 games in Camerer et al. (2004).[6] General details of the games are reported in Table 4.[7] The number of strategies varies from two to six. The 55 games differ not only in number of strategies, but also in payoff structures as they were originally designed to test various behavioral hypotheses. Hence, this collection of games represents a rigorous test of the GCH model.

3.2. Results

Using the 55 matrix games, we estimated the GCH model, its two special cases (namely CH and LM), LK and a two-parameter LK model. The two-parameter LK model allows for the possibilities that there might be more level 0 than specified by a Poisson distribution. The two-parameter LK model specifies a bi-modal distribution of types where proportion γ of the population consists of just level 0 and proportion $1-\gamma$ of the population consists of thinking types (including level 0) that are distributed according to a Poisson distribution.[8] When $\gamma=0$, the two-parameter model reduces to the original LK model.

In estimating GCH, we conduct a grid search in the region defined by $0\le\tau\le 10, 1\le\alpha\le 20, 1\le\beta\le 10$. The CH model imposes additional constraints $\alpha=1,\beta=1$; the LM model has $\alpha=\infty,\beta=1$. The left panel of Table 5 reports the fit measures and parameter estimates of the five models. GCH's improvement in fit over the other four models ranges from 1.2% to 4.3%. The improvements are statistically significant using the likelihood ratio test for the nested models CH and LM and the Vuong test for the non-nested models LK and bi-modal LK (Vuong, 1989). These improvements are significant for two reasons. First, GCH adds an average of 100 log-likelihood points over and above the other four models; this difference is large when compared with the differences in log-likelihood among the other four models. Second, it is difficult to have

[5] These upper bounds were never binding in the optimization. Furthermore, for α, more than 90% of the belief is loaded onto the model level(s) at the upper bound of $\alpha=20$ (close to 100% if τ is 3 or less); hence the belief structure is sufficiently close to that of $\alpha=\infty$. We are grateful to one reviewer for highlighting this point.

[6] The 22 games collected by Camerer et al. (2004) replicate games reported in Ochs (1995), Bloomfield (1994), Binmore et al. (2001), Rapoport and Amaldoss (2000), Tang (2001), Goeree et al. (2003), Mookerjee and Sopher (1997), Rapoport and Boebel (1992), Messick (1967), Lieberman (1962) and O'Neill (1987).

[7] The payoff matrices with data are available upon request from the authors.

[8] Hence, the total proportion of level 0 is $\gamma+(1-\gamma)f(0|\tau)$.

Table 5
Calibration Results for GCH, CH, LM, LK, GCH with $\alpha = 1$ and GCH with $\beta = 1$.

Model		GCH		CH		LM		LK		Bi-modal LK		GCH (β = 1)		GCH (α = 1)	
Log-likelihood		**-2637**		-2751		-2742		-2732		-2670		-2696		-2654	
% Improvement of GCH over the model				4.3%		4.0%		3.6%		1.2%		2.2%		0.6%	
Likelihood Ratio Test				227	*	208	*	-		-		117	*	33	*
Vuong Test				-		-		4.68	*	1.94		-		-	
p -value				0.0000		0.0000		0.0000		0.0261		0.0000		0.0000	
Parameter Estimates	τ	1.11	*	1.00	*	1.01	*	1.11	*	1.79	*	1.20	*	0.76	*
	α	2.10	*	-		-		-		-		2.50	*	1.00	f
	β	2.10	*	-		-		-		-		1.00	f	2.99	*
	γ	-		-		-		-		0.34	*	-		-	

[1] Number of observations = 3596.
[2] Bold indicates the best performing model.
[3] * indicates significance at the 1% level.
[4] f indicates a fixed parameter value.
[5] The significance of α and β is tested with respect to the value of 1.

substantial improvements over the existing models in simple games with two or three strategies; 48 of the 55 games are simple games.[9]

The τ estimates of GCH, CH, LM, and LK are remarkably close, between 1.00 and 1.11.[10] This suggests that the average step of thinking is about 1. α is estimated to be significant at 2.10, suggesting that subjects exhibit stereotype bias and give disproportionately more weight to more frequently occurring levels. β is estimated to be significant at 2.10, indicating that level-0 players exhibit minimum aversion and are more than twice as likely to choose strategies without any minimum payoff than to choose strategies with a minimum payoff in at lease one of the outcomes.

To check that each modeling extension is important, we also estimate the special cases of $(\alpha, \beta = 1)$, which provides a new model of opponents, and $(\alpha = 1, \beta)$, which provides a new model of level-0 players. The right panel of Table 5 shows that each modeling extension is crucial in improving GCH's overall fit.

3.3. Why is GCH better?

Let's examine how each modeling extension helps improve the GCH model. Modeling the stereotype bias involves two components: level k's own stereotype bias, and level k's belief that others have stereotype bias. Similarly, modeling minimum aversion also has two components: aversion in level 0's own behavior, and the higher levels' response to level 0's aversion. We examine each of these components below and show how each contributes to the improvement of GCH over CH.

3.3.1. Allowing for stereotype bias: the α parameter

When $\alpha = 1$, there is no stereotype bias because level k's beliefs of proportions of lower levels correspond to their true relative proportions. When $\alpha = \infty$, there is maximum stereotype bias because level k believes that opponents are all of the most frequently occurring lower level. GCH assumes that level-k players exhibit stereotype bias (i.e., $\alpha > 1$) and they project themselves on others and assume that others (i.e., lower level players) exhibit the same bias too. To separate the effects of level-k players' own stereotype bias and their belief that others have stereotype bias, we denote α^* and α^{**} as the parameters for a player's own stereotype bias and the player's belief of others' stereotype bias, respectively. Both parameters affect the expected payoff of level k as follows:

$$E_k(\pi_i(s_i^j)) = \sum_{j'=1}^{m_{-i}} \pi_i(s_i^j, s_{-i}^{j'})\{\sum_{h=0}^{k-1} g_k^*(h) \cdot P_h(s_{-i}^{j'}|g_k^{**}(h))\} \tag{7}$$

where $g_k^*(h)$ is the relative proportion with level k's own stereotype bias α^*, and the lower-level choice probability $P_h(s_{-i}^{j'}|g_k^{**}(h))$ is based on level k's belief about lower levels' stereotype bias α^{**}.[11]

[9] Past estimates of τ for existing strategic thinking models hover around 1 and 2. This suggests that existing models are already capturing most behaviors quite well without the need to go beyond three iterated levels of thinking with at most three different chosen strategies.

[10] The τ estimate for the two-parameter LK model is 1.79 and the bi-modal estimate of level 0 is $\gamma = 0.34$. This proportion of level 0 estimated by the two-parameter LK model is quite close to that estimated by the other models. The range of level 0 for Poisson distribution when $\tau = 1.00$ and 1.11 is between 0.33 and 0.37.

[11] $g_k^{**}(h)$ is the distribution of relative proportions which level k believes level h (with α^{**}) has over lower levels 1 to $h - 1$. Hence, $g_k^{**}(h)$ is distributed as $\frac{f(l)^{\alpha^{**}}}{\sum_{h'=0}^{h-1} f(h')^{\alpha^{**}}}$ $\forall l < h$.

In the GCH model, we have $\alpha^* = \alpha^{**} = \alpha$. We set $\alpha^* = 1$ to suppress level-k players' own stereotype bias, but allow them to believe lower level players have stereotype bias. In other words, the choice probability $P_h(s_{-i}^{j'}|g_k^{**}(h))$ is conditioned on the belief of k on lower levels having stereotype bias $\alpha^{**} = \alpha$ and hence $g_k^{**}(h)$. Comparing this to the full GCH model allows us to gauge the incremental effect of own stereotype bias.

The log-likelihood of the best fitted GCH model is -2637.33 with $\tau = 1.11, \alpha = 2.10$ and $\beta = 2.10$, as reported in Table 5. When we suppress own stereotype bias with $\alpha^* = 1$ only (and $\alpha^{**} = \alpha = 2.10$), the log-likelihood drops to -2675.07. This suggests that own stereotype bias plays a significant role in improving fit in GCH. When we suppress both stereotype biases with $\alpha^* = 1$ and $\alpha^{**} = 1$, the log-likelihood does not change much. This suggests that level k's belief of others' stereotype bias did not contribute to the improvement in fit. In fact, if one were to suppress only level k's belief about lower levels' stereotype bias by having $\alpha^{**} = 1$ (and $\alpha^* = \alpha$), the log-likelihood drops only marginally to -2639.56. Hence, we conclude that capturing level k's own stereotype bias is more crucial than capturing level k's belief that others may have stereotype bias.

3.3.2. Allowing for minimum aversion: the β parameter

The right panel of Table 5 shows that allowing for minimum aversion improves the fit significantly. Does this improvement in fit derive from the behaviors of level 0 or higher level players' beliefs about lower level players? Specifically, does minimum aversion provide a better description of level-0 behavior, or does it capture the beliefs of higher levels (and hence their behaviors) better?

To test and quantify the impact of our proposed model of level 0 on the behavior of higher levels, we suppress the model of level 0 in the behavior of higher levels. In other words, we run a model where the belief of higher levels is based on the behavior of a uniformly random level 0; but the behavior of level 0 is that of minimum aversion. Let us denote level 0's minimum aversion by β^* and higher levels' belief of level-0's minimum aversion by β^{**}. In GCH, $\beta^* = \beta^{**} = \beta$. Formally, the modified expected payoff of level 1 and higher from equation (1) is:

$$E_k(\pi_i(s_i^j)) = \sum_{j'=1}^{m_{-i}} \pi_i(s_i^j, s_{-i}^{j'})\{\sum_{h=0}^{k-1} g_k(h) \cdot P_h(s_{-i}^{j'}|\beta^{**})\} \tag{8}$$

while the behavior of level-0 players is modeled by $P_0(s_i^j|\beta^*)$.

The model with $\beta^* = \beta$ for level-0 behavior and $\beta^{**} = 1$ for higher level players (that is, high level players assume a uniformly randomizing level 0) yields a log-likelihood of -2691.37. This represents a sizeable drop from -2637.33 and suggests that higher levels indeed think that level 0 players are minimum averse. If we further suppress minimum aversion in level 0 with $\beta^* = 1$ (in addition to $\beta^{**} = 1$), there is an additional sizable drop in the log-likelihood to -2716.94. In combination, these results suggest that level 0 players are indeed minimum averse and higher levels believe level 0 are minimum averse and hence adjust their behaviors accordingly.

4. Applications of the GCH model

In this section, we apply the GCH model to three new games with more complex and varied game structures in order to test the robustness of the above findings. The first application involves some coordination games from Amaldoss et al. (2008) where coordination improves with the addition of an irrelevant alternative. This application offers a clean test of whether GCH can capture the underlying asymmetric dominance effect by contrasting the estimation results before and after the addition. The second application is on the so-called "compromise" game studied by Carrillo and Palfrey (2009), where subjects with private signals of their own strengths must decide whether to fight or to compromise. Experimental results show that the fighting rate is substantially lower than the equilibrium prediction of always fighting. While all non-equilibrium models predict that fighting will increase with thinking level, GCH is more adept at capturing the lower fight rate when the private signal of own strength is low. In the third application, GCH is applied to capture the behavior of a money request game studied by Arad and Rubinstein (2012); this game is meant to induce level-k reasoning. GCH is able to capture level-k reasoning as well as level-k models do. In each application, we take the approach of calibrating the non-equilibrium models using data from half of the games and then validating the best fitted models using data from the remaining half which have different design parameters.

4.1. Coordination by an irrelevant alternative

Amaldoss et al. (2008) demonstrated how coordination could be improved by adding a weakly dominated alternative. In a 2×2 non-zero sum game, the row player chooses between A and B and the column player chooses between L and R. There are two pure-strategy equilibria: (A, R) and (B, L). When a weakly dominated (by one of the two original strategies) alternative was added to the choice set of the row player, Amaldoss et al. (2008) showed that the choice proportion of the dominating strategy increased. Specifically, when A′ (weakly dominated by A) was added, the choice proportion of A increased from 47.22% to 59.45%. Similarly, when B′ was added, the proportion of B increased from 52.78% to 83.12% (see Table 5 of Amaldoss et al., 2008). Column players chose R more frequently when A′ was added; the proportion of R increased

Table 6
Model fits and estimates for Amaldoss et al. (2008).

Model		GCH	CH	LM	LK	Bi-modal LK
Amaldoss, Bettman and Payne (2008)						
Log-likelihood						
Calibration (1440 observations)		**-1051**	-1074	-1097	-1109	-1109
% Improvement of GCH over the model			2.2%	4.4%	5.5%	5.5%
Validation (1440 observations)		**-1130**	-1137	-1164	-1168	-1168
% Improvement of GCH over the model			0.6%	3.0%	3.3%	3.3%
Likelihood Ratio Test			46 *	92 *	-	-
Vuong Test			-	-	4.72 *	5.06 *
p -value			0.0000	0.0000	0.0000	0.0000
Parameter Estimates	τ	1.78 *	2.48 *	1.33 *	2.04 *	2.04 *
	α	3.49 *	-	-	-	-
	β	1.49 *	-	-	-	-
	γ	-	-	-	-	0.00
Control session AB in set 6						
Log-likelihood						
Calibration (720 observations)		-489	-489	-490	-489	-489
Parameter Estimates	τ	1.68 *	1.68 *	0.99 *	1.58 *	1.58 *
	α	1.00	-	-	-	-
	β	1.00	-	-	-	-
	γ	-	-	-	-	0.00

[1] Models calibrated using set 6 and validated using set 5, each with two sessions AA'B and ABB'.
[2] Bold indicates the best performing model.
[3] * indicates significance at the 1% level.
[4] The significance of α and β is tested with respect to the value of 1.
[5] There are multiple optima for the bi-modal LK model for the control session AB in set 6.

from 51.81% to 72.92%. Hence, both row and column players had a higher proportion of coordination at (A, R), increasing from an average of 22.22% to 41.67%. Column players also chose L more frequently when B′ was added, increasing from 48.19% to 75.14%. This resulted in better coordination at (B, L), increasing from an average of 23.20% to 63.20% (see Table 7 of Amaldoss et al., 2008).

Similar to Amaldoss et al. (2008), we analyze strategic thinking using the data they generated in study 2, datasets 5 and 6. The two datasets use different payoff matrices. Each dataset has two sessions, one with row strategies (A, B, A′) and another with (A, B, B′).[12] Column players in both sessions have column strategies (L, R). We calibrate using dataset 6 and validate using dataset 5.[13] Table 6 reports the parameters and the calibration and validation results.

GCH provides an improvement in the log-likelihood of 23 points and 58 points over CH and LK respectively. Having a mass point for level 0 in the bi-modal LK model does not help to improve the fit over LK. These improvements are statistically significant even at the 0.001 level using the likelihood ratio test for the nested models and the Vuong test for the non-nested models. GCH also has a significant lead over the CH and LK models in validation. The parameter estimates of the GCH model are $\tau = 1.78, \alpha = 3.49$ and $\beta = 1.49$, compared to $\tau = 2.48$ for the CH model, $\tau = 1.33$ for the LM model and $\tau = 2.04$ for the LK model. The average thinking step is about 2 when we look at the τ estimates from the four models. The stereotype bias parameter is estimated at 3.49. When coupled with the τ estimates, it suggests that level 2 and above respond mostly to level 1. Since $\beta = 1.49$, never worst strategies are 49% more favored by level 0.

To further understand the differences among the three models, we describe the predicted behaviors of level-k players using the (A, B, A′) session in dataset 6 (payoff matrix reproduced in Table 7). As indicated in the table, strategy A in the never worst set is preferred by level-0 row players while level-0 column players randomize between L and R. GCH predicts that the behavior of row players level 3 and above stabilizes into strategy A and the behavior of column players level 2 and above stabilizes into strategy R. Hence, GCH predicts that the game will converge into (A, R), one of the two equilibria, for players level 3 and higher. Similarly, the CH model predicts that row players level 5 and higher will choose A and column players level 5 and higher will choose R. The GCH and CH models predict that higher level players converge into one of

[12] Minimum aversion picks up the asymmetric dominance in the (A, B, A′) sessions where A is favored by the level-0 row player but in the (A, B, B′) session, B is not favored by the level-0 row player. In the (A, B, B′) session the weakly dominant alternative B pays the minimum (same as B′) when the column player picks R, hence no strategy is favored by the level-0 row player. In short, (A, B, B′) is a game with a specific form of asymmetric dominance which GCH does not cover.

[13] Calibrating using dataset 5 yields a substantially better fit for GCH when compared to CH and LK; the log-likelihood for GCH is −1026 versus −1111 for CH and −1160 for LK. One reason that GCH fits better in dataset 5 is due to the higher contribution from capturing minimum aversion.

Table 7
(A, B, A′) Session in Dataset 6 from Amaldoss et al. (2008).

	L	R
A	16,**16**	20, 24
A′	**8**, 12	20, **12**
B	24, 20	**8**, **8**

Note: payoffs in bold indicate the minimum payoff for the corresponding columns or rows. The never worst set for the row player consists of strategy A which weakly dominates strategy A′. The never worst set for the column player is empty. There are two equilibria: (A, R) and (B, L).

the two equilibria. More importantly, the equilibrium converged into is the one coordinated by the irrelevant alternative A′, whereas, the LK model predicts that behaviors of row and column players will infinitely oscillate among the available strategies.

While the aggregate predictions of GCH and CH are similar, the rationales behind them are different. CH has a higher τ at 2.48 (compared to GCH's 1.78), implying that CH requires at least level 4 in order to coordinate at (A, R). Under GCH, level 0 players are minimum averse and higher level thinkers best respond largely to the most common occurring type, so the model requires only level 2 or higher in order to coordinate at (A, R).

To see why GCH provides a more plausible description, let us contrast the estimates for the models before and after the irrelevant alternative A′ was introduced.[14] The lower panel of Table 6 shows the estimates for the various thinking models for the control sessions before the irrelevant alternative A′ was introduced. The τ estimates for all models (except LM) stay in a tight range of (1.58, 1.68). Contrast these with the estimates after A′ was introduced. τ in GCH changes minimally from 1.68 to 1.78, τ in CH changes from 1.68 to 2.48 and τ in LK changes from 1.58 to 2.04. CH and LK attribute more coordination after the introduction of A′ to having more players with higher levels of thinking. The distribution of thinking levels in GCH remains relatively unchanged, so the increased coordination is attributed to the higher levels reacting to minimum aversion of level 0. We find that GCH offers a simpler and more plausible explanation than CH and LK.[15]

4.2. To fight or to compromise

Carrillo and Palfrey (2009) studied a game of two-sided private information. In this two-player game, each player is privately assigned a randomly drawn strength $\sigma \in [0, 1]$ and must decide to either compromise or fight. If at least one player chooses to fight, the player with higher strength is paid 1 and the weaker player is paid 0. If both choose to compromise, they are each paid $M < 1$. The equilibrium prediction is for both to fight regardless of their own strength but fighting is actually observed in only 61.1% of the cases.

Since one player's decision is irrelevant if the other player decides to fight, we only need to consider the best response function conditioned on the other player choosing compromise. Specifically, the best response function depends on each player's own strength and is in the form of threshold: if a player's own strength is σ, the player should fight if $\sigma > \sigma^*$ for some threshold σ^*; otherwise, the player should compromise. The threshold strength is where the expected payoff of fighting equals M, which is the payoff for compromise.

To determine whether level k should fight, let us define $M_k(\sigma)$ as the expected payoff of fighting for level k with strength σ, given that lower levels compromise. In addition, we denote $q^k(\sigma)$ as the posterior strength distribution of level k given that level k compromises.

Beginning with level 0, the expected payoff from fighting is better than that from compromising when $\sigma > M$, and the expected payoff from fighting is worse than that from compromising when $\sigma < M$. With minimum aversion, level 0 will fight with probabilities $P(\mathit{fight}|\sigma > M) = \frac{\beta}{1+\beta}$ and $P(\mathit{fight}|\sigma < M) = \frac{1}{1+\beta}$. In other words, a minimum-averse level 0 player will fight more (less) often when strength is higher (lower) than M, as opposed to a randomizing level 0 who is equally likely to fight and to compromise regardless of strength.

Conditional on level 0 compromising, level 1 and beyond would derive the posterior strength distribution of level 0 as $q^0(\sigma) = \frac{\beta}{M\cdot\beta+(1-M)}$ over the range $[0, M]$ and $q^0(\sigma) = \frac{1}{M\cdot\beta+(1-M)}$ over the range $[M, 1]$. Hence, a minimum averse level 0 would lead higher levels to believe that the strength of level 0 is more likely to be in $[0, M]$.

The expected payoff of fighting for level 1 with strength σ is $M_1(\sigma) = \int_0^\sigma q^0(\sigma')\partial\sigma' \cdot 1 + \int_\sigma^1 q^0(\sigma')\partial\sigma' \cdot 0 = \frac{\beta}{M\beta+(1-M)} \cdot \sigma$ for $\sigma < M$. Level 1 is indifferent between compromise and fight at threshold σ^1 where $M_1(\sigma^1) = M$. Solving for σ^1, the best response of level 1 is to compromise when $\sigma \le \sigma^1 = M \cdot \frac{M\beta+(1-M)}{\beta}$ and fight otherwise. Therefore, level 2 would

[14] We thank the anonymous reviewers for their suggestions.

[15] There is evidence that individual thinking levels may change across games (e.g. Georganas et al., 2013). The above result for the CH and LK models suggest that the average thinking level increases when a dominated alternative is added; presumably more effort is required to evaluate the bigger choice set. However, the GCH model suggests that adding the dominated alternative makes the choice decision simpler for level 0 and higher levels simply respond to the behavior of level 0, with no increased effort.

Table 8
Model fits and estimates for Carrillo and Palfrey (2009).

Model		GCH	CH	LM	LK	Bi-modal LK
Carrilo and Palfrey (2009)						
Log-likelihood						
Calibration (560 observations)		**-215**	-216	-228	-239	-239
% Improvement of GCH over the model			0.6%	6.4%	11.2%	11.2%
Validation (560 observations)		**-173**	-181	-200	-209	-209
% Improvement of GCH over the model			4.9%	15.5%	20.8%	20.8%
Likelihood Ratio Test			2	27 *	-	-
Vuong Test			-	-	3.28 *	3.86 *
p -value			0.2949	0.0000	0.0005	0.0001
Parameter Estimates	τ	1.45 *	1.56 *	1.39 *	1.24 *	1.24 *
	α	1.77 *	-	-	-	-
	β	1.38	-	-	-	-
	γ	-	-	-	-	0.00

[1] Models calibrated using the dataset with M = 0.50 (321 fight) and validated using the dataset with M = 0.39 (368 fight).
[2] Bold indicates the best performing model.
[3] * indicates significance at the 1% level.
[4] The significance of α and β is tested with respect to the value of 1.

assess the posterior strength distribution of level 1 as $q^1(\sigma) = 1$ over the range $[0, \sigma^1]$ and $q^1(\sigma) = 0$ otherwise. Since $\frac{M\beta+(1-M)}{\beta} \leq 1$ (because $\beta \geq 1$), $\sigma^1 \leq M$. Hence, level 1 is more likely to fight against a minimum-averse level 0.

In general, the cumulative strength distribution of lower levels and the expected payoff from level k fighting given that all lower levels compromise can be derived recursively as:

$$M_k(\sigma) = \frac{\sum_{h=0}^{k-1}[g_k(h) \cdot \int_0^{\sigma} q^h(\sigma')\partial\sigma']}{\sum_{h=0}^{k-1} g_k(h) \cdot \sigma^h} \tag{9}$$

where σ^h is the threshold for level h and $g_k(h)$ is the level-k perceived proportion of level h where $h < k$. The threshold of level k, σ^k, is derived from solving $M_k(\sigma^k) = M$. The best response of level k is to compromise if $\sigma \leq \sigma^k$, and fight otherwise.

As delineated in equation (9), the expected payoff from fighting for level k is a weighted combination of the posterior strengths of the lower levels with the weights being the relative proportions of the lower levels (since the payoff of winning the fight is equal to 1). The posterior strength distributions derived by higher level thinkers are skewed towards the lower strengths, making it more profitable for the higher levels to fight at lower thresholds than the lower levels. Put differently, as thinking level k increases, the range of compromise $[0, \sigma^k]$ shrinks and the chance of compromise becomes smaller. More precisely, $\sigma^k \leq \sigma^{k-1}$. In the limit, we will observe the best response of always fighting, which is the equilibrium prediction.

For LK models, since $g_k(k-1) = 1$ and 0 otherwise, one can verify from solving $M_k(\sigma^k) = M$ that the thresholds have a simple structure $\sigma^k = M \cdot \sigma^{k-1}$. For GCH, as $g_k(h) = \frac{f(h)^\alpha}{\sum_{h'=0}^{k-1} f(h')^\alpha}$, $\forall h < k$, the thresholds are influenced by both the distribution of thinking levels and the extent of stereotype bias.

Two game sessions in Carrillo and Palfrey (2009) were run in simultaneous move; one with $M = \frac{7}{18} = 0.39$ and the other with $M = 0.50$.[16] We calibrate using the session with $M = 0.50$ and validate using the session with $M = 0.39$.[17] The estimation results and the parameter estimates are reported in Table 8.

GCH and CH offer the best fit and LK the worst. The bi-modal LK model does not fit better since $\gamma = 0$. The better fit of GCH over LM and LK in calibration is statistically significant at the 0.001 level using the likelihood ratio test and the Vuong test. Although the difference between GCH and CH is not significant in calibration, the performance of GCH in validation is substantially better than the CH and the LK models. The τ estimates from all five models are in the narrow range of [1.24, 1.56]. The stereotype bias parameter α is equal to 1.77, which is slightly lower than the value of 2.10 from the matrix

[16] We only consider simultaneous move. To apply GCH to sequential move requires additional specifications on how players at various levels would react to prior realized and unrealized moves.

[17] Calibrating using the $M = 0.39$ session yields a substantially better fit for GCH when compared to CH and LK where the log-likelihood for GCH is -141 versus -181 for CH and -209 for LK. The main reason that GCH fits better for $M = 0.39$ is the higher contribution from capturing minimum aversion, which accounts for more fighting.

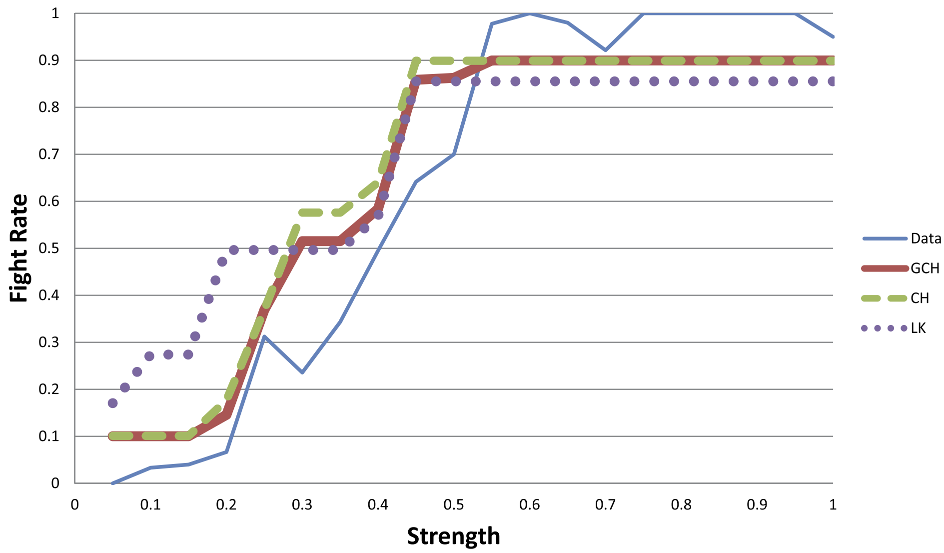

Fig. 2. Empirical and Predicted Fight Rates over Strength [0, 1] with $M = \frac{7}{18}$.

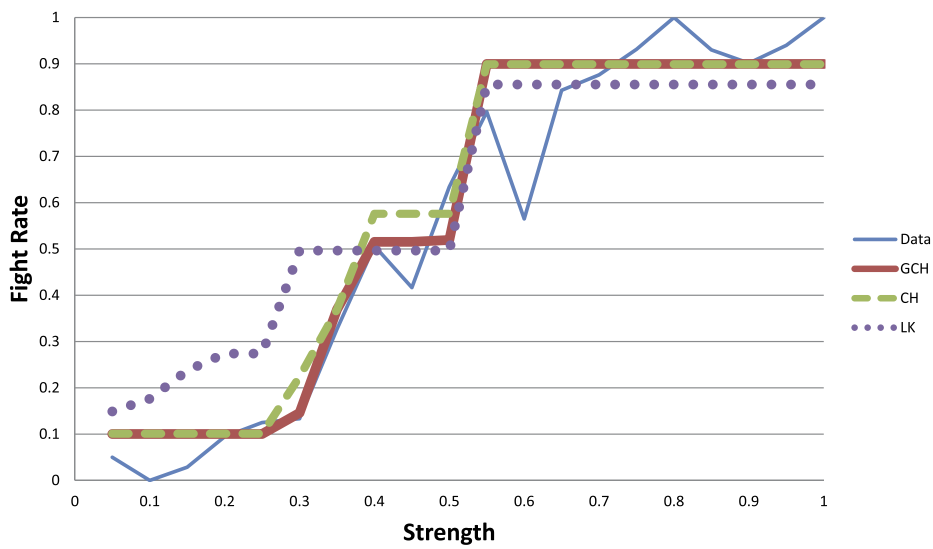

Fig. 3. Empirical and Predicted Fight Rates over Strength [0, 1] with $M = 0.50$.

games. However, the minimum aversion parameter β is lower than that of the matrix game at 1.38 and is not statistically significant. Figs. 2 and 3 better illustrate where GCH improves over LK and CH.[18]

One can observe from both figures that the GCH model (represented by the bold solid step line) fits the data better than the LK model (represented by the dotted line) in two ways. First, the GCH model predicts a wider fight rate range over the entire strength distribution; the predicted fight rate from GCH starts lower and ends higher. This effect is generated by the model of level 0. Second, GCH steps up the fight rate at the right strength level compared to LK. Tracking the actual data closely, GCH's fight rate steps up at the strength level of about 0.2 for $M = 0.39$ and about 0.3 for $M = 0.5$. In other words, the predicted fight rate in LK is almost always higher than that in GCH except at the higher end of the strength signal.

[18] GCH clearly fits better than CH in validation (Fig. 2) but not significantly better in calibration (Fig. 3).

Table 9
Model fits and estimates for Arad and Rubinstein (2012).

Model		GCH		CH		LM		LK		Bi-modal LK	
Arad and Rubinstein (2012)											
Log-likelihood											
Calibration (108 observations)		-213		-214		-216		-215		**-210**	
% Improvement of GCH over the model				0.6%		1.3%		1.0%		-1.4%	
Validation - Cycle (72 observations)		-118		-119		-127		**-117**		-122	
% Improvement of GCH over the model				0.3%		7.6%		-0.9%		3.0%	
Validation - Costless (53 observations)		-202		-202		-211		**-194**		-196	
% Improvement of GCH over the model				0.0%		4.7%		-3.7%		-2.8%	
Likelihood Ratio Test				3		6		-		-	
Vuong Test				-		-		-0.51		-0.73	
p -value				0.2616		0.0585		0.6964		0.7683	
Parameter Estimates	τ	2.33	*	2.35	*	2.00	*	2.11	*	2.30	*
	α	1.06	*	-		-		-		-	
	β	1.21		-		-		-		-	
	γ	-		-		-		-		0.19	*

[1] Models calibrated using the basic version and validated using the cycle and costless versions.
[2] Bold indicates the best performing model.
[3] * indicates significance at 1% level.
[4] The significance of α and β is tested with respect to the value of 1.

4.3. Money request

Arad and Rubinstein (2012) proposed a two-person ten-strategy game aiming to induce level-k reasoning. Each player requests an amount of money ranging from 11 to 20 shekels.[19] Each player will get the amount requested. In addition, a bonus of 20 shekels will be paid to the player whose request is exactly 1 shekel less than the opponent's request. The bonus payment acts as a trigger for level-k reasoning.

Beginning with level 1, 19 is the best response to a wide range of level-0 behaviors, including both uniform random choice and the model of level 0 proposed in this paper. The best response of level 2 to 19 is 18, and so on. Nash equilibrium predicts a 45% share for 17, 18 and 19 (corresponding to levels 3, 2 and 1 respectively in LK), significantly lower than the empirical observation of 75%. In fact, the Nash equilibrium predicts a 50% share for 15 and 16. But the majority of the 20% who chose 11 to 16 described their choices as guesses rather than the result of iterated reasoning, fitting the description of level 0.

We calibrate using the 108 observations from the basic version of the game. Arad and Rubinstein (2012) also run two variants of the basic game called cycle and costless iteration. In the cycle variant, a player who requests 20 shekels will also get a bonus if the opponent's request is 11; this addition to the bonus rule of the basic version completes the cycle where 20 is best response to 11. In the costless iteration variant, any request less than 20 will receive 17 regardless of the requested amount, plus a bonus if the request is one less than the opponent's request. This variant creates the condition that any request below 20 incurs the same cost of 3 shekels. We validate using both variants. Table 9 reports the parameters and the calibration and validation results.

Although the bi-modal LK model provides the best fit in terms of log-likelihood value, its difference with GCH is not statistically significant using the Vuong test. It is however significantly better than LK in calibration. On the other hand, the difference between LK and CH is also not statistically significant.

Several observations on the estimates are worth noting. The γ parameter for level 0 in the bi-modal LK model has a value of 0.19, which roughly matches the observed 20% choosing 11 to 16. The β parameter for minimum aversion is not significant. Minimum aversion does not help to single out "better" strategies as eight out of ten strategies (from 13 to 20) are in the never worst set. The value of τ for all five models varies in a tight range of (2.00, 2.35), capturing the main pattern that majority of the choices fall in the range between 17 and 19.

Level 1 has the same best response of 20 in the cycle and costless sessions. In the cycle session, 20 is the best response because the strategy has a chance to win the bonus not available in the basic version. In the costless session, choosing 19 means giving up 3 shekels relative to choosing 20; hence level 1 finds choosing 19 worse than choosing 20. In summary, the cycle version increases the "benefit" of 20 over 19 and the costless version increases the "cost" of 19 over 20. In both

[19] We thank the anonymous reviewers for their suggestions to include this game in our analysis. While more experimental data have been generated subsequent to Arad and Rubinstein (2012) (e.g., Choo and Kaplan, 2014; Goeree et al., 2014; Lindner and Sutter, 2013), we focus on the data from Arad and Rubinstein (2012).

Table 10
Predicted probabilities of all models for Arad and Rubinstein (2012).

Predicted Probabilities for All Models for the Basic Version of Money Request

Strategy	Data	GCH	CH	LM	LK	Bi-modal LK
11	0.04	0.01	0.01	0.01	0.01	0.03
12	0.00	0.01	0.01	0.01	0.01	0.03
13	0.03	0.01	0.01	0.01	0.02	0.03
14	0.06	0.01	0.01	0.01	0.03	0.04
15	0.01	0.01	0.01	0.01	0.05	0.07
16	0.06	0.07	0.10	0.01	0.11	0.12
17	0.32	0.34	0.34	0.34	0.20	0.19
18	0.30	0.30	0.27	0.28	0.28	0.24
19	0.12	0.24	0.23	0.28	0.27	0.21
20	0.06	0.01	0.01	0.01	0.01	0.03

Predicted Probabilities for All Models for the Cycle Version of Money Request

Strategy	Data	GCH	CH	LM	LK	Bi-modal LK
11	0.01	0.01	0.01	0.01	0.01	0.03
12	0.01	0.01	0.01	0.01	0.01	0.03
13	0.00	0.01	0.01	0.01	0.01	0.03
14	0.01	0.01	0.01	0.01	0.02	0.03
15	0.00	0.01	0.01	0.01	0.03	0.04
16	0.04	0.01	0.01	0.01	0.05	0.07
17	0.10	0.07	0.10	0.01	0.11	0.12
18	0.22	0.34	0.34	0.34	0.20	0.19
19	0.47	0.30	0.27	0.28	0.28	0.24
20	0.13	0.24	0.23	0.28	0.27	0.21

Predicted Probabilities for All Models for the Costless Version of Money Request

Strategy	Data	GCH	CH	LM	LK	Bi-modal LK
11	0.00	0.01	0.01	0.01	0.01	0.03
12	0.04	0.01	0.01	0.01	0.01	0.03
13	0.00	0.01	0.01	0.01	0.01	0.03
14	0.04	0.01	0.01	0.01	0.02	0.03
15	0.04	0.01	0.01	0.01	0.03	0.04
16	0.04	0.01	0.01	0.01	0.05	0.07
17	0.09	0.10	0.10	0.01	0.11	0.12
18	0.21	0.33	0.34	0.34	0.20	0.19
19	0.40	0.27	0.27	0.28	0.28	0.24
20	0.15	0.24	0.23	0.28	0.27	0.21

variants, the best responses of levels 1, 2 and 3 correspond to strategies 20, 19 and 18 respectively. All five models are able to capture this shift in best responses.

Table 10 provides the predicted probabilities of the five models and the data for all three versions of the game. The predicted probabilities in the validation (cycle and costless) sessions are basically the predicted probabilities in the calibration session shifted down by one thinking level, exactly as intended to be induced from the two variants. The bi-modal LK model (and GCH) is slightly worse than LK in validation. Combining results from calibration and validation, we conclude that the three models perform about the same in capturing behavior in the money request games.

The actual behaviors induced by money request games are quite consistent with the intended level-k reasoning, particularly up to level 3. All models (GCH, its special cases, and LK models) are equally adept at capturing the first three levels of thinking (beyond level 0) induced by the games, as evident from the τ estimates and more precisely from the predicted probabilities in Table 10. Given the payoff structure of the money request game, there is little need for GCH to deviate too much from other models in behavioral prediction. First, the payoff structure always induces a best response to the perceived **modal** choice (i.e., choosing one shekel less), making it unnecessary for α to deviate too much from 1; α is estimated to be significant at 1.06. Second, eight out of the ten strategies belong to the never worst set, suggesting that most strategies are no more salient (according to the minimum avoidance principle) than others. Hence, the β parameter is not significantly different from 1. In summary, GCH does not over-fit in situations where simpler models suffice. All models are also capable of capturing the shift in strategy by level 1 and above in the validation sessions.

5. Conclusion

This paper presents a general cognitive hierarchy model which incorporates a more general model of level 0 and a more general model of opponents. The general model of level 0 posits that level 0 is averse to strategies that yield a minimum payoff in at least one scenario. Our empirical results suggest that the never-worst strategies are more than twice preferred over the minimum strategies. We call this effect the minimum aversion effect. It incorporates the asymmetric dominance and compromise effects and captures preference over dominant strategies. This general model of level 0 is not only parsimonious, it is also effective at explaining game behavior resulting in a significant improvement in fit. It not only explains level 0 better, it also explains higher level behaviors better.

Table 11
Game 7 from Stahl and Wilson (1995).

	T	M	B	Data
T	**30**	100	50	0.44
M	40	**0**	90	0.35
B	50	75	**29**	0.21

Note: payoffs in bold indicate the minimum payoff for the corresponding columns or rows. The best response to strategy T is strategy B, the best response to strategy B is strategy M and the best response to strategy M is strategy T, making a full cycle.

The general model of opponents captures the whole spectrum of stereotype bias, from the basic CH model with no stereotype bias to a close cousin of LK, the LM model with maximum stereotype bias. With maximum stereotype bias, a player only focuses on responding to the behavior of the most frequently occurring type among the lower thinking levels. Our empirical results suggest that there exists a mild level of stereotype bias. Incorporating stereotype bias into the model improves the fit significantly. This improvement comes from two sources: the accurate belief of stereotype bias in the lower levels and stereotype bias in the player's own actions.

With careful calibration and validation using a rich repertoire of datasets, the GCH model provides robust and decent performance in one-shot games. More importantly, it provides an intuitive and viable explanation for the observed behaviors. The model has wider applicability. In the last section, we applied the GCH model to three different games with increased complexity.

For future research, we suggest that the GCH model (with some perturbations) be applied to other economic applications to provide a better understanding of the underlying behaviors. In particular, our model of opponents can be parameterized to explain how stereotype bias might vary depending on how a player's initial belief is primed. Our model of level 0 can be parameterized to capture the potential impact of game structure, and more specifically the impact of payoff structure.

Appendix A. Comparison of the LM and LK models

The LK model may generate oscillating behavior as the thinking level increases. To see this, let us consider Game 7 from Stahl and Wilson (1995). This game is a 3×3 symmetric matrix game (see Table 11). To determine LK's prediction in this game, let us begin with the choice of level 1. Believing that the opponent is a level 0 who chooses among the three strategies equally, level 1 best responds by choosing T as it offers the highest expected payoff of $\frac{1}{3} \cdot 30 + \frac{1}{3} \cdot 100 + \frac{1}{3} \cdot 50 = 60$. Level 2 best responds to level 1's choice of T by choosing B which gives a payoff of 50. Level 3 best responds to level 2's choice of B by choosing M yielding a payoff of 90. Level 4 best responds to level 3's choice of M by choosing T paying 100. Level 5 will best respond to T with a choice of B, as a level 2 will do; higher levels will endlessly cycle through the same choice sequence of T, B and M. Given the empirical frequency reported in Table 11, the best fitted τ for LK model is 0.166.

Does the LM model exhibit the same oscillating behavior? The answer is no. Given the empirical data, the best fitted τ for LM model is 0.178. Under LM, level 1's best response is T, which follows from the same calculation as the LK model. Level 2 best responds to the most frequently occurring lower level which in this case is level 0 (since $\tau = 0.178$). Hence, the choice of level 2 is T, the same as level 1. The modal level is level 0 for all higher levels; hence, all higher levels also choose T as the best response. As a consequence, there is no oscillating behavior among higher level thinkers.[20]

The oscillating behavior of the LK model can also occur in games with $N > 2$ players. We illustrate this in a market entry game. In this game, N players independently and simultaneously decide whether or not to enter a market with demand d (where $d < N$). If the number of entrants is greater than d, all entrants earn nothing; otherwise they receive a payoff of 1. Players who stay out receive a payoff of 0.5. Hence, the best response function for a player depends on the number of other players who enter.

Let us derive the best response function of a level-k player. We denote the best response function (or the entry function) of a level k with demand d by $e(k, d)$. In addition, let $E_k(k-1, d)$ be a level-k player's assessment of the total number of entries by competitors who are level $k-1$ or lower. Hence the entry function of a level-k player can be defined as:

$$e(k, d) = \begin{cases} 0 & \text{if } E_k(k-1, d) \geq d \\ 1 & \text{if } E_k(k-1, d) < d \end{cases} \tag{10}$$

Using the above equation, the entry function can be recursively derived starting with $e(0, d)$. In the case of the LK model, a level-k player believes that opponents are only level-$k-1$ players, hence $E_k(k-1, d) = e(k-1, d)$. Using the entry function in equation (10), level k will enter if $k-1$ stays out and vice versa. Therefore, we will see alternating entry and no entry decisions being taken as k goes from level 1 to ∞. In the case of the LM model, the entry function depends on the mode of

[20] In this game, LM fits slightly better than LK. There is a 0.32% improvement in log-likelihood. Note that both models have one parameter.

distribution of thinking level. Given the Poisson distribution, a level-k player, where $k \le mode$, believes that opponents are only level-$k-1$ players, so the entry function will best respond to $E_k(k-1,d) = e(k-1,d)$. For $k > mode$, the level-k player believes that opponents are only modal players and will best respond to $E_k(k-1,d) = e(mode,d)$. In other words, these players will best respond to the modal level by making the same decision; if the modal player enters, all higher levels will stay out and vice versa. Therefore, we see that the prediction of the LM model converges at levels higher than the mode of the distribution.

The next question is whether total entry monotonically increases in demand. Put simply, if we fix the number of players N but increases the demand d, would we observe a corresponding increase in entry from these N players? Assuming a Poisson distribution, we show below that monotonicity in entry is guaranteed when $\tau \le 1.256$, similar to the condition derived for the CH model in Camerer et al. (2004).

We have $e(0,d) = \frac{1}{2}$ for all d, and

$$E_k(k-1,d) = \begin{cases} e(k-1,d) & \text{if } k \le mode \\ e(mode,d) & \text{if } k > mode \end{cases} \tag{11}$$

For $k \ge 1$

$$e(k,d) = \begin{cases} 0 & \text{if } E_k(k-1,d) \ge d \\ 1 & \text{if } E_k(k-1,d) < d \end{cases} \tag{12}$$

We derive the condition of τ for the Poisson distribution when total entries are monotonically increasing in demand. Given that the entry function of level k only depends on a single lower level (whether it is $k-1$ or the mode of the distribution), there is only one single cut-point as d increases. For $\frac{d}{N} < \frac{1}{2}$, level 1 stays out, level 2 enters, and the pattern of odd levels staying out and even levels entering is maintained until the mode. If the mode is even, then all levels above the mode stay out, otherwise, they enter. Hence, if the mode is even, total entry is $(1/2)f(0) + f(2) + f(4) + \cdots + f(mode)$; otherwise, it is $(1/2)f(0) + f(2) + f(4) + \cdots + f(mode-1) + f(mode+1) + f(mode+2) + \cdots$. For $\frac{d}{N} > \frac{1}{2}$, level 1 enters since half of level 0 enter. Level 2 stays out, and the odd–even pattern stays until the mode. If the mode is even, then all levels above the mode enter; otherwise, they stay out. Hence, if the mode is even, total entry is $(1/2)f(0) + f(1) + f(3) + \cdots + f(mode-1) + f(mode+1) + f(mode+2) + \cdots$; otherwise, it is $(1/2)f(0) + f(1) + f(3) + \cdots + f(mode)$.

If the mode is even, the condition of monotonically increasing entries requires total entries in $\frac{d}{N} < 1/2$ to be less than total entries in $\frac{d}{N} > 1/2$. In other words, $1 - f(0) > 2(f(2) + f(4) + \cdots + f(mode))$. This is always true for the uni-modal Poisson distribution given that $f(0)$ is exponentially decreasing in τ. If the mode is odd, we have the condition that $1 - f(0) < 2(f(1) + f(3) + \cdots + f(mode))$, which is satisfied when $\tau \le 1.256$. Hence, total entries are monotonically increasing in demand when $\tau \le 1.256$ in a Poisson distribution of thinking levels.

References

Amaldoss, Wilfred, Bettman, James R., Payne, John W., 2008. Biased but efficient: an investigation of coordination facilitated by asymmetric dominance. Marketing Sci. 27 (5), 903–921.

Arad, Ayala, Rubinstein, Ariel, 2012. The 11–20 money request game: a level-*k* reasoning study. Am. Econ. Rev. 102 (7), 3561–3573.

Binmore, Kenneth, Swierzbinski, Joe, Proulx, Chris, 2001. Does maximin work? An experimental study. Econ. J. 111 (473), 445–464.

Bloomfield, Robert, 1994. Learning a mixed strategy equilibrium in the laboratory. J. Econ. Behav. Organ. 25 (3), 411–436.

Camerer, Colin F., Ho, Teck-Hua, Chong, Juin-Kuan, 2004. A cognitive hierarchy theory of one-shot games. Quart. J. Econ. 119 (3), 861–898.

Carrillo, Juan D., Palfrey, Thomas R., 2009. The compromise game: two-sided adverse selection in the laboratory. Am. Econ. J. Microecon. 1 (1), 151–181.

Choo, Lawrence C.Y., Kaplan, Todd R., 2014. Explaining behavior in the 11–20 game. MPRA Working Paper No. 52808.

Cooper, David, Van Huyck, John, 2003. Evidence on the equivalence of the strategic and extensive form representation of games. J. Econ. Theory 110 (2), 290–308.

Costa-Gomes, Miguel A., Crawford, Vincent P., 2006. Cognition and behavior in two-person guessing games: an experimental study. Amer. Econ. Rev. 96 (5), 1737–1768.

Costa-Gomes, Miguel A., Crawford, Vincent P., Broseta, Bruno, 2001. Cognition and behavior in normal-form games: an experimental study. Econometrica 69 (5), 1193–1235.

Crawford, Vincent P., 2013. Boundedly rational versus optimization-based models of strategic thinking and learning in games. J. Econ. Lit. 51 (2), 512–527.

Crawford, Vincent P., Iriberri, Nagore, 2007. Fatal attraction: salience, naivete, and sophistication in experimental hide-and-seek games. Amer. Econ. Rev. 97 (5), 1731–1750.

Georganas, Sotiris, Healy, Paul J., Weber, Roberton A., 2013. On the persistence of strategic sophistication. Working Paper.

Goeree, Jacob K., Holt, Charles, Palfrey, Thomas, 2003. Risk averse behavior in generalized matching pennies games. Games Econ. Behav. 45 (1), 97–113.

Goeree, Jacob K., Louis, Philippos, Zhang, Jingjing, 2014. Noisy introspection in the "11–20" game. Working Paper.

Huber, Joel, Payne, John W., Puto, Christopher, 1982. Adding asymmetrically dominated alternatives: violations of regularity and the similarity hypothesis. J. Cons. Res. 9 (1), 90–98.

Lieberman, Bernhardt, 1962. Experimental studies of conflict in some two-person and three-person games. In: Criswell, J.H., Solomon, H., Suppes, P. (Eds.), Mathematical Models in Small Group Processes. Stanford University Press, pp. 203–220.

Lindner, Florian, Sutter, Matthias, 2013. Level-*k* reasoning and time pressure in the 11–20 money request game. Econ. Letters 120 (3), 542–545.

Messick, David M., 1967. Interdependent decision strategies in zero-sum games: a computer-controlled study. Behav. Sci. 12 (1), 33–48.

Mookerjee, Dilip, Sopher, Barry, 1997. Learning and decision costs in experimental constant-sum games. Games Econ. Behav. 19 (1), 97–132.

Nagel, Rosemarie, 1995. Unraveling in guessing games: an experimental study. Amer. Econ. Rev. 85 (5), 1313–1326.

Ochs, Jack, 1995. Games with unique, mixed strategy equilibria: an experimental study. Games Econ. Behav. 10 (1), 202–217.

O'Neill, Barry, 1987. Nonmetric test of the minimax theory of two-person zero-sum games. Proc. Natl. Acad. Sci. 84 (7), 2106–2109.

Park, Bernadette, Rothbart, Myron, 1982. Perception of out-group homogeneity and level of social categorization. J. Personal. Soc. Psych. 42 (6), 1051–1068.

Quattrone, George A., Jones, Edward E., 1980. The perception of variability within in-groups and out-groups: implications for the law of small numbers. J. Personal. Soc. Psych. 38 (1), 141–152.

Rapoport, Amnon, Amaldoss, Wilfred, 2000. Mixed strategies and iterative elimination of strongly dominated strategies: an experimental investigation of states of knowledge. J. Econ. Behav. Organ. 42 (4), 483–521.

Rapoport, Amnon, Boebel, Richard B., 1992. Mixed strategies in strictly competitive games: a further test of the minimax hypothesis. Games Econ. Behav. 4 (2), 261–283.

Simonson, Itamar, 1989. Choice based on reasons: the case of attraction and compromise effects. J. Cons. Res. 16 (2), 158–174.

Stahl, Dale O., Wilson, Paul, 1994. Experimental evidence on players' models of other players. J. Econ. Behav. Organ. 25 (3), 309–327.

Stahl, Dale O., Wilson, Paul, 1995. On players models of other players: theory and experimental evidence. Games Econ. Behav. 10 (1), 218–254.

Tang, Fang-Fang, 2001. Anticipatory learning in two-person games: some experimental results. J. Econ. Behav. Organ. 44 (2), 221–232.

Vuong, Quang H., 1989. Likelihood ratio tests for model selection and non-nested hypotheses. Econometrica 57 (2), 307–333.