

Behavioral Game Theoretic Models: A Bayesian Framework For Parameter Analysis

James R. Wright
Department of Computer Science
University of British Columbia
Vancouver, B.C., Canada, V6T 1Z4
jrwright@cs.ubc.ca

Kevin Leyton-Brown
Department of Computer Science
University of British Columbia
Vancouver, B.C., Canada, V6T 1Z4
kevinlb@cs.ubc.ca

ABSTRACT

Studies in experimental economics have consistently demonstrated that Nash equilibrium is a poor description of human players' behavior in unrepeated normal-form games. Behavioral game theory offers alternative models that more accurately describe human behavior in these settings. These models typically depend upon the values of exogenous parameters, which are estimated based on experimental data. We describe methods for deriving and analyzing the posterior distributions over the parameters of such models, and apply these techniques to study two popular models (Poisson-CH and QLk), the latter of which we previously showed to be the best-performing existing model in a comparison of four widely-studied behavioral models [22]. Drawing on a large set of publicly available experimental data, we derive concrete recommendations for the parameters that should be used with Poisson-CH, contradicting previous recommendations in the literature. We also uncover anomalies in QLk that lead us to develop a new, simpler, and better-performing family of models.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Economics, Experimentation, Performance

Keywords

Behavioral game theory, Bounded rationality, Game theory, Cognitive models

1. INTRODUCTION

It is well known that in many settings, the standard game-theoretic assumption that agents will adopt *Nash equilibrium* strategies—where each agent simultaneously responds optimally to all the others—is a poor predictor of actual human behavior [e.g., see 9]. The field of *behavioral game theory* aims to develop models that more accurately describe human behavior, as evaluated using experimental data [2].

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

Copyright © 2012, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

These models typically depend upon parameters such as agents' sensitivity to utility differences and the distribution of cognitive ability in the population. In order for behavioral models to be effective tools for prediction, it is necessary to identify "good" estimates for these parameters. Such parameter values are also interesting for their own sake, because they can offer economic insights into human behavior.

In past work, we conducted an exhaustive meta-study of behavioral game theory models, focusing on the problem of unrepeated, simultaneous-move games [22]. We evaluated four prominent models from the behavioral game theory literature: Quantal Response Equilibrium [QRE; 12], Level- k [Lk; 13, 6], Poisson-Cognitive Hierarchy [Poisson-CH; 3], and Quantal Level- k [QLk; 19]. We evaluated these models according to their out-of-sample prediction performance: we identified parameters for each model that maximized the likelihood of a training data set, and then evaluated each model in terms of the likelihood it assigned to a separate test set. We took data from six publicly available datasets describing lab experiments in which subjects played unrepeated, simultaneous-move games against other human opponents. In the end we observed a result that was statistically significant and quite robust: that the QLk model substantially outperformed all others on the set of all data, and also had the best or nearly the best performance on each individual dataset.

While it effectively identified the best-performing existing model, our previous work offered very little insight about the models' parameter values. For example, because we relied upon a maximum-likelihood approach, we obtained no information about the extent to which parameter values could be changed without a major drop in predictive accuracy, or even about the extent to which individual parameters influence a model's performance. This paper shows how to answer such questions. Specifically, it describes a Bayesian approach for gaining understanding about a behavioral model's entire parameter space. We combine experimental data with explicitly quantified prior beliefs to derive a posterior distribution that assigns probability to parameter settings in proportion to their consistency with the data and the prior [8]. We apply this approach to analyze the posterior distributions for two models: QLk and Poisson-Cognitive Hierarchy. Although Poisson-CH did not demonstrate competitive performance in [22], we analyze it here because it is very simple, and because the literature contains a very specific recommendation for how its parameter should be set: Camerer et al. [3] recommend setting the model's single parameter, which represents agents' mean number of steps of strategic

reasoning, to 1.5. Our own analysis sharply contradicts this recommendation, placing the 99% confidence interval almost a factor of three lower, on the range [0.51, 0.59]. We devote most of our attention to QLk, however, since we previously showed it to achieve strong performance. Our new analysis points out several anomalies in the parameter distributions for QLk, suggesting that a simpler model could be preferable. By exhaustively evaluating a family of variations on QLk, we identify a simpler, more predictive model based in part on the cognitive hierarchy concept.

We pause to contrast our work with a particularly related paper from the economics community. Rogers et al. [17] proposed a unifying framework that generalizes both Poisson-CH and QRE, and compared the fit of several variations within this framework. This is similar to our search of a family of QLk variants, but there are several differences. First, we compare out-of-sample prediction performance, not in-sample fit. Second, Rogers et al. restricted the distributions of types to be grid, uniform, or Poisson distributions, whereas we consider unconstrained discrete distributions. Third, they required different types to have different precisions, while we do not. Finally, we consider level- k beliefs as well as cognitive hierarchy beliefs, whereas they compared only cognitive hierarchy belief models (although their framework in principle allows for both).

There has also been interest by various computer scientists in behavioral solution concepts as an alternative to standard economic models. For example, Wunder et al. [23] extended the cognitive hierarchy model in their analysis of the Lemonade Stand Game, and Gao and Pfeffer [7] used quantal response equilibrium to model bounded rationality in a game where agents respond to payoffs unknown to the modeler.

In the next section, we define the models that we analyze. We then describe the framework that we used for estimating and analyzing parameters and their distributions, and the dataset and specific techniques that we used to derive our estimates. In Section 4 we present the results of our analyses for the Poisson-CH and QLk models. In Section 5, we describe the QLk-like variants that we evaluated, and present the results of our search.

2. BEHAVIORAL MODELS

We begin by formally defining the behavioral models that we analyze. We then relate these models to other widely studied models that we do not consider further, due to their poor predictive performance in our previous study [22].

2.1 Quantal Level- k

Stahl and Wilson [19] proposed a rich model of strategic reasoning that we refer to as the *quantal level- k model* (QLk). The QLk model of human behavior incorporates two primary components: *iterated strategic reasoning*, and *cost-proportional errors*. Iterated strategic reasoning refers to a limit on the number of levels of higher-order belief that agents can maintain.¹ Agents make cost-proportional errors if their rate of making errors increases as errors become less costly. This can be modeled by assuming that agents best respond *quantally*, rather than via strict maximization.

Definition (Quantal best response). Let $u_i(a_i, s_{-i})$ be agent i 's expected utility when playing action a_i against

¹This limit is believed to be quite low; e.g., Arad and Rubinstein [1] found no support for 4th order or higher beliefs.

strategy profile s_{-i} . Then a *quantal best response* $QBR_i(s_{-i}; \lambda)$ by agent i to s_{-i} is a mixed strategy s_i such that

$$s_i(a_i) = \frac{\exp[\lambda \cdot u_i(a_i, s_{-i})]}{\sum_{a'_i} \exp[\lambda \cdot u_i(a'_i, s_{-i})]}, \quad (1)$$

where λ (the *precision*) indicates agents' sensitivity to utility differences. \square

Note that unlike regular best response, which is a set-valued function, quantal best response always returns a single mixed strategy. When $\lambda = 0$, quantal response mixes uniformly over all of the agents' actions; as $\lambda \rightarrow \infty$, quantal best response approaches strict best response.

In the QLk model, agents have one of three levels: level-0, level-1, or level-2. Level-0 agents are nonstrategic, and choose their actions uniformly at random. Level-1 agents quantally best respond to level-0 agents; that is, they believe that the rest of the population consists entirely of level-0 agents. Similarly, level-2 agents quantally respond to level-1 agents. Level-1 and level-2 agents can use different precisions (λ 's), and furthermore level-2 agents' beliefs about level-1 agents' precision can be arbitrarily different from level-1 agents' actual precision.

Definition (QLk model). Let A_i denote player i 's action set. Then the probability distribution $\pi_{i,k}^{QLk} \in \Pi(A_i)$ over actions that QLk predicts for a level- k agent playing as agent i is defined as follows.

$$\begin{aligned} \pi_{i,0}^{QLk}(a_i) &= |A_i|^{-1}, \\ \pi_{i,1}^{QLk} &= QBR_i(\pi_{-i,0}^{QLk}; \lambda_1), \\ \pi_{i,1(2)}^{QLk} &= QBR_i(\pi_{-i,0}^{QLk}; \lambda_{1(2)}), \\ \pi_{i,2}^{QLk} &= QBR_i(\pi_{i,1(2)}^{QLk}; \lambda_2), \end{aligned}$$

where $\pi_{i,1(2)}^{QLk}$ is a mixed-strategy profile representing level-2 agents' (possibly incorrect) beliefs about how level-1 agents play. The overall predicted distribution π^{QLk} of actions is the weighted sum of the distributions for each level: $\pi^{QLk} = \sum_{k=0}^2 \alpha_k \pi_{i,k}^{QLk}$, where $\alpha_0 = 1 - \alpha_1 - \alpha_2$. The QLk model thus has five parameters: $\{\alpha_1, \alpha_2\}$, the proportions of level-1 and level-2 agents ($\alpha_1 + \alpha_2 \leq 1$; any remaining agents are level-0); $\{\lambda_1, \lambda_2\}$, the precisions of level-1 and level-2 agents' responses; and $\lambda_{1(2)}$, level-2 agents' beliefs about the precision of level-1 agents. \square

2.2 Poisson Cognitive Hierarchy

Like QLk, the cognitive hierarchy model [3] aims to model agents with heterogeneous bounds on iterated reasoning. It differs from the QLk model in two key ways. First, agents use standard best response, rather than quantal response. Second, agents best respond to the full distribution of lower-level types, rather than only to the strategy one level below. More formally, every agent has an associated level $m \in \{0, 1, 2, \dots\}$. Let f be a probability mass function describing the distribution of the levels in the population. Level-0 agents play uniformly at random. Level- m agents ($m \geq 1$) best respond to the strategies that would be played in a population described by the truncated probability mass function $f(j \mid j < m)$.

Camerer et al. [3] advocate a single-parameter restriction of the cognitive hierarchy model called *Poisson-CH*, in which f is a Poisson distribution.

Definition (Poisson-CH model). Let $\pi_{i,m}^{PCH} \in \Pi(A_i)$ be the distribution over actions predicted for an agent i with level m by the Poisson-CH model. Let $f(m) = \text{Poisson}(m; \tau)$. Let $BR_i(s_{-i})$ denote the set of i 's best responses to the strategy profile s_{-i} . Let $\pi_i^{PCH} = \sum_{\ell=0}^m f(\ell) \pi_{i,\ell}^{PCH} / \sum_{\ell=0}^m f(\ell)$ be the ‘‘truncated’’ distribution over actions predicted for an agent conditional on that agent’s having level $0 \leq \ell \leq m$. Then π^{PCH} is defined as follows:

$$\pi_{i,0}^{PCH}(a_i) = |A_i|^{-1},$$

$$\pi_{i,m}^{PCH}(a_i) = \begin{cases} |BR_i(\pi_{i,0:m-1}^{PCH})|^{-1} & \text{if } a_i \in BR_i(\pi_{i,0:m-1}^{PCH}), \\ 0 & \text{otherwise.} \end{cases}$$

As with QLk, the overall predicted distribution of actions π^{PCH} is a weighted sum of the distributions for each level: $\pi_i^{PCH} = \sum_{\ell=0}^{\infty} f(\ell) \pi_{i,\ell}^{PCH}$. The mean of the Poisson distribution, τ , is this model’s single parameter. \square

2.3 Other Behavioral Models

We previously evaluated two other behavioral models in addition to those listed above [22]. The first, QRE [12], is a generalization of Nash equilibrium where agents respond quantally instead of best responding. The second, Lk [13, 6], is similar to QLk, except that agents best respond to the next level down rather than quantally responding. We omit these models from our analysis, since they consistently achieved worse predictive performance than QLk in our previous comparison.

3. BAYESIAN PARAMETER ANALYSIS

In this section, we describe this paper’s first contribution, our formal framework for parameter estimation, along with the experimental data and estimation techniques we used.

3.1 Prediction Framework

We begin by noting that any behavioral game theoretic model is a mapping from a game description G and a vector of parameters θ to a distribution over action profiles in G . In other words, a given behavioral model provides us with a way of computing $\Pr(a | G, \theta)$. Our goal is to use this behavioral model in conjunction with our experimental data, \mathcal{D} , to predict future behavior in a given game G . The dataset may or may not contain examples of earlier play in G . The model parameters need to be determined based upon the data. There are two approaches that we can take.

The frequentist approach, taken by Rogers et al. [17] and Wright and Leyton-Brown [22], is to compute a point estimate $\hat{\theta}$ of the model’s parameters based on our dataset, and then use $\hat{\theta}$ to compute the prediction for new games. That is, we use the model

$$\Pr(a | G, \mathcal{D}) = \Pr(a | G, \hat{\theta}). \quad (2)$$

The Bayesian approach is to calculate a posterior distribution over parameter values $\Pr(\theta | \mathcal{D})$, and then use this distribution to integrate out the parameter. This yields the alternate model

$$\Pr(a | G, \mathcal{D}) = \int \Pr(a | G, \theta) \Pr(\theta | \mathcal{D}) d\theta. \quad (3)$$

Observe that Equation (2) can be seen as a special case of (3), where the posterior is a point mass at $\hat{\theta}$.

We use the posterior distribution to analyze model parameters, which is an application of Bayesian estimation. However, on our dataset, we have observed that the frequentist approach achieves better predictive performance than the Bayesian approach. Thus, we use point estimation to compare performance between model variants. We describe both approaches.

3.1.1 Point Estimation

We use the maximum likelihood estimate of the parameters as our point estimate:

$$\hat{\theta} = \arg \max_{\theta} \Pr(\mathcal{D} | \theta).$$

Our dataset consists of a set \mathcal{D} of observations $d_i = (G_i, a_i)$, indicating that an action profile a_i was played in game G_i . The likelihood of a single datapoint is

$$\Pr(d_i | \theta) = \Pr(G_i, a_i | \theta).$$

By the chain rule (of probabilities), this is equivalent to

$$\Pr(d_i | \theta) = \Pr(a_i | G_i, \theta) \Pr(G_i | \theta). \quad (4)$$

We assume that θ and G are independent; i.e., we assume that $\Pr(G_i | \theta) = \Pr(G_i)$. Intuitively, this means that we assume there is a single ‘‘true’’ parameter value θ^* for all games, rather than a separate ‘‘true’’ θ^{G_i} for each game G_i . This allows us to rewrite (4) as

$$\Pr(d_i | \theta) = \Pr(a_i | G_i, \theta) \Pr(G_i). \quad (5)$$

We assume that the datapoints are independent, so the likelihood of the dataset is just the product of the likelihoods of the datapoints:

$$\Pr(\mathcal{D} | \theta) = \prod_{d_i \in \mathcal{D}} \Pr(d_i | \theta). \quad (6)$$

Substituting (5) into (6) gives us

$$\Pr(\mathcal{D} | \theta) = \prod_{d_i \in \mathcal{D}} \Pr(a_i | G_i, \theta) \Pr(G_i). \quad (7)$$

The probabilities $\Pr(G_i)$ are constant with respect to θ , and can therefore be disregarded when maximizing likelihood:

$$\arg \max_{\theta} \Pr(\mathcal{D} | \theta) = \arg \max_{\theta} \prod_{d_i \in \mathcal{D}} \Pr(a_i | G_i, \theta).$$

3.1.2 Bayesian Estimation

We derive an expression for the posterior distribution $\Pr(\theta | \mathcal{D})$ by applying Bayes’ rule, where $p_0(\theta)$ is the prior:

$$\Pr(\theta | \mathcal{D}) = \frac{p_0(\theta) \Pr(\mathcal{D} | \theta)}{\Pr(\mathcal{D})}. \quad (8)$$

Substituting in (7), the posterior distribution is

$$\Pr(\theta | \mathcal{D}) = \frac{p_0(\theta) \prod_{d_i \in \mathcal{D}} \Pr(a_i | G_i, \theta) \Pr(G_i)}{\Pr(\mathcal{D})},$$

and since both $\Pr(G_i)$ and $\Pr(\mathcal{D})$ are constant with respect to θ ,

$$\Pr(\theta | \mathcal{D}) \propto p_0(\theta) \prod_{d_i \in \mathcal{D}} \Pr(a_i | G_i, \theta). \quad (9)$$

Codename	Source	Games	n	Units
SW94	Stahl and Wilson [19]	10	400	\$0.025
SW95	Stahl and Wilson [20]	12	576	\$0.02
CGCB98	Costa-Gomes et al. [5]	18	1566	\$0.022
GH01	Goeree and Holt [9]	10	500	\$0.01
CVH03	Cooper and Huyck [4]	8	2992	\$0.10
RPC09	Rogers et al. [17]	17	1210	\$0.01
HSW01	Haruvy et al. [11]	15	869	\$0.02
HS07	Haruvy and Stahl [10]	20	2940	\$0.02
SH08	Stahl and Haruvy [18]	18	1288	\$0.02
COMBO9	Above datasets	128	3600	\$0.01

Table 1: Names and contents of each dataset. The column headed n indicates the number of observations in the dataset. Units are in expected value.

3.2 Data

We analyzed data from nine experimental studies, summarized in Table 1; this expands beyond the six studies we considered in our previous work [22]. We also constructed a new dataset (COMBO9) containing observations from all nine source datasets. To ensure that each was equally represented, despite their differences in size, we included exactly 400 observations from each dataset (sampled uniformly without replacement).

The precision parameter for quantal response is not *scale invariant*. That is, the correct value of λ can differ depending upon the units in which payoffs are expressed. To ensure consistent estimation of precision parameters, we renormalized all games so that their payoffs were in expected cents.

3.3 Estimation Methods

For all of the models we considered, we used a flat prior for the parameters. Although this prior is improper, it results in a correctly-normalized posterior distribution; the posterior distribution in this case reduces to the likelihood [8]. We computed the posterior distribution for the single-parameter Poisson-CH model by *grid sampling*. That is, we computed the likelihood of the COMBO9 dataset for each value of $\tau \in \{0.01k \mid k \in \mathbb{N}, 0 \leq 0.01k \leq 10\}$, and then normalized by the sum of the likelihoods.

The QLk model has five parameters, and is therefore too high dimensional to grid sample efficiently; approximately 5×10^{12} samples would have been required for a grid of the same granularity as we used for Poisson-CH! Instead, for QLk—and the other high dimensional models introduced in Section 5—we used a sequential Monte Carlo technique called *annealed importance sampling*, or AIS [14]. AIS allows for efficient sampling from high dimensional distributions, similarly to Markov Chain Monte Carlo (MCMC) techniques. However, each sample point generated using AIS is independent, so AIS does not exhibit the random-walk behavior that can plague MCMC samplers. Briefly, the annealed importance sampling procedure is as follows. A sample $\vec{\theta}_0$ is drawn from an easy-to-sample-from distribution P_0 . For each P_j in a sequence of intermediate distributions P_1, \dots, P_{r-1} that become progressively closer to the posterior distribution, a sample $\vec{\theta}_j$ is generated by drawing a sample $\vec{\theta}'$ from a *proposal distribution* $Q(\cdot \mid \vec{\theta}_{j-1})$, and accepted with probability

$$\frac{P_j(\vec{\theta}')Q(\vec{\theta}_{j-1} \mid \vec{\theta}')}{P_j(\vec{\theta}_{j-1})Q(\vec{\theta}' \mid \vec{\theta}_{j-1})}. \quad (10)$$

If the proposal is accepted, $\vec{\theta}_j = \vec{\theta}'$; otherwise, $\vec{\theta}_j = \vec{\theta}_{j-1}$. We repeat this procedure multiple times, and report the distribution of the resulting $\vec{\theta}_r$ values, with each $\vec{\theta}_r$'s contribution weighted according to

$$\frac{P_1(\vec{\theta}_0)P_2(\vec{\theta}_1) \dots P_{r-1}(\vec{\theta}_{r-2})P_r(\vec{\theta}_{r-1})}{P_0(\vec{\theta}_0)P_1(\vec{\theta}_1) \dots P_{r-2}(\vec{\theta}_{r-2})P_{r-1}(\vec{\theta}_{r-1})}. \quad (11)$$

For the initial sampling distribution P_0 , we used a product distribution over the population proportions parameters and the precision parameters. For the population proportion parameter components we used a Dirichlet distribution $\text{Dir}(1, 1, 1)$; this is equivalent to uniformly sampling over the simplex of all possible combinations of population proportions. For the precision parameter components we used the renormalized non-negative half of a univariate Gaussian distribution $\mathcal{N}(0, 2^2)$ for each precision parameter; this gives a distribution that is decreasing in precision (on the assumption that higher precisions are less likely than lower ones), and with a standard deviation of 2, which was large enough to give a non-negligible probability to most previous precision estimates.

The proposal distribution was a product distribution “centered” at the current value, with proportion parameters $\vec{\alpha}'$ sampled from $\text{Dir}(20\vec{\alpha}_{j-1})$, and each precision parameter λ' sampled from $\mathcal{N}(\lambda_{j-1}, 0.2^2)$ (truncated at 0 and renormalized). The “hyperparameters” for the Dirichlet distribution (20) and the precision distributions (0.2^2) were chosen by trial and error on a small subset of the data to make the acceptance rate near to the standard heuristic value of 0.5 [16]. We used 200 intermediate distributions of the form

$$P_j(\vec{\theta}) = \Pr(\vec{\theta} \mid \mathcal{D})^{\gamma_j},$$

with the first 40 γ_j 's spaced uniformly from 0 to 0.01, and the remaining 160 γ_j 's spaced geometrically from 0.01 to 1, as in the original AIS description [14]. We performed 5 Metropolis updates in each distribution before moving to the next distribution in the chain.

For the model performance comparisons in Section 5, we computed the parameter point-estimates using the Nelder-Mead simplex algorithm [15]. To reduce the danger of getting caught in a local minimum, we repeated the optimization from 500 random starting points, and chose the value which gave the highest log likelihood.

4. RESULTS

In this section, we describe the results of our analysis of posterior distributions for the two models defined in Section 2. We start by comparing the posterior distribution for the Poisson-CH model's parameter to a previously published recommendation. We then turn our attention to the QLk model, whose posterior distribution exposes several possible issues with the model.

4.1 Poisson-CH

Camerer et al. [3] recommended setting the τ parameter of the Poisson-CH model to 1.5. Figure 1 gives the cumulative posterior distribution over τ for each of our datasets. Overall, our analysis strongly contradicts Camerer et al.'s recommendation. On COMBO9, the posterior probability of $0.51 \leq \tau \leq 0.59$ is more than 99%. Every other source dataset had a wider high posterior density region than COMBO9 (indicated by the higher slope of COMBO9's cumulative density




Figure 1: Cumulative posterior distributions for the τ parameter of the Poisson-CH model. Bold trace is for the combined dataset; solid trace is for the outlier Stahl and Wilson [19] source dataset; dotted traces are all other source datasets.




Figure 2: Marginal cumulative posterior distribution for the level proportion parameters (α_1, α_2 ; top panel) and precision parameters ($\lambda_1, \lambda_2, \lambda_{1(2)}$; bottom panel) of the QLk model on the combined dataset.

function); this is expected, as smaller datasets lead to less confident predictions. Nevertheless, all but two of the source datasets had median values less than 1.0. Only the Stahl and Wilson [19] dataset (SW94) appears to support Camerer et al.’s recommendation (median 1.43). However, SW94 appears to be an outlier; its high posterior density region is wider than the other distributions, and the distribution is very multimodal, likely due to SW94’s small size.

4.2 QLk

Figure 2 gives the marginal cumulative posterior distributions for each of the parameters of the QLk model. (That is, we computed the five-dimensional posterior distribution, and then extracted from it the five marginal distributions shown here.) We found these distributions surprising for several reasons. First, the models predict many more level-2 agents than level-1 agents. In contrast, it is typically assumed that higher level agents are scarcer, as they perform more complex strategic reasoning. Even more surprisingly, the model predicts that level-1 agents should have much higher precisions than level-2 agents. This is odd if the level-2 agents are to be understood as “more rational”; indeed, precision is sometimes interpreted as a measure of rationality [e.g., see 21, 7]. Third, the distribution of $\lambda_{1(2)}$, the precision that level-2 agents ascribe to level-1 agents, is

very concentrated around very small values ($[0.023, 0.034]$). This differs by two orders of magnitude from the “true” value of λ_1 , which is quite concentrated around its median value of 3.1. Finally, the median value of λ_1 (3.1) is more than 17 times larger than that of λ_2 (0.18). It seems unlikely that level-1 agents would be an order of magnitude more sensitive to utility differences than level-2 agents.

One interpretation is that the QLk model is essentially accurate, and these parameter values simply reflect a surprising reality. For example, the low precision of level-2 agents and the even lower precision that they (incorrectly) ascribe to the level-1 agents may indicate that two-level strategic reasoning causes a high cognitive load, which makes agents more likely to make mistakes, both in their own actions and in their predictions. The main appeal of this explanation is that it allows us to accept the QLk model’s strong performance at face value.

An alternate interpretation is that QLk fails to capture some crucial aspect of experimental subjects’ strategic reasoning. For example, if the higher-level agents reasoned about all lower levels rather than only one level below themselves, then the low value of $\lambda_{1(2)}$ could predict well because it “simulates” a model where level-2 agents respond to a mixture of level-0 and level-1 agents. We investigate this second possibility in the next section.

5. MODEL VARIATIONS

In this section, we investigate the properties of the QLk model by evaluating the predictive power of a family of systematic variations of the model. In the end, we identify a simpler model that dominates QLk on our data, and which also yields much more reasonable marginal distributions over parameter values.

Specifically, we constructed a family of models by extending or restricting the QLk model along four different axes. QLk assumes a maximum level of 2; we varied this by considering maximum levels of 1 and 3 as well. QLk has *inhomogeneous precisions* in that it allows each level to have a different precision; we varied this by also considering *homogeneous precision* models. QLk allows *general precision beliefs* that are not restricted to be accurate; we also constructed models that make the simplifying assumption of *accurate precision beliefs* about lower levels’ precisions. Finally, in addition to *Lk* beliefs, where all other agents are assumed by a level- k agent to be level- $(k-1)$, we also constructed models with *CH* beliefs, where agents believe that the population consists of the true, truncated distribution over the lower levels. We evaluated each combination of axis values; the 17 resulting models² are listed in the first part of Table 2. In addition to the 17 exhaustive axis combinations for models with maximum levels in $\{1, 2, 3\}$, we also evaluated 12 additional axis combinations that have higher maximum levels and 8 parameters or fewer: **ai-QCH4** and **ai-QLk4**; **ah-QCH** and **ah-QLk** variations with maximum levels in $\{4, 5, 6, 7\}$; and **ah-QCH** and **ah-QLk** variations that assume a Poisson distribution over the levels rather than using an explicit tabular distribution. These additional models are listed in the bottom part of Table 2.

5.1 Simplicity Versus Predictive Performance

²When the maximum level is 1, all combinations of the other axes yield identical predictions. Therefore there are only 17 models instead of $3 \cdot 2^3 = 24$.




Figure 3: Model simplicity (number of parameters) versus prediction performance. QLk1, which has far lower performance than the other models, is omitted for scaling reasons.

We evaluated the predictive performance of each model on the COMBO9 dataset using 10-fold cross-validation repeated 10 times, as in [22]. The results are given in the last column of Table 2, and plotted in Figure 3. Performance is measured by how much more likely the test data is according to the given model (with parameters chosen based on a training dataset, separate from the test set) than it is according to assuming that actions are chosen uniformly at random (u.a.r.). E.g., the test data is $10^{18.37}$ times more likely according to the QLk1 model’s prediction than according to a uniform random prediction.

All else being equal, a model with higher performance is more desirable, as is a model with fewer parameters. We can plot an “efficient frontier” of models with the (statistically significant) highest performance for a given number of parameters or fewer; see Figure 3. The original QLk (gi-QLk2) is *not* efficient in this sense; it is dominated by ah-QCH3, which has significantly better predictive performance even though it has fewer parameters (due to restricting agents to homogeneous precisions and accurate beliefs). Our analysis thus argues that the flexibility added by inhomogeneous precisions and general precision beliefs is less important than the number of levels and the choice of population belief.

Conversely, the poor performance of the Poisson variants relative to ah-QCH3 suggests that flexibility in describing the level distribution is more important than the total number of levels modeled. Figure 4 shows the marginal posterior level weight distributions for all four models on the efficient frontier. There is broad agreement among all models on the proportion of level-0 agents. However, in order to get the “right” number of level-0 agents, ah-QCHp (which models the level distribution as a Poisson) must place a great deal of weight on level-1 agents as well; in contrast, the tabular-distribution models all select bimodal distributions that as-

sign relatively little weight to level-1 agents, and more to higher-level agents (level-2 and higher).

In fact, there is a pattern in the models along the efficient frontier: this set consists *exclusively* of models with accurate precision beliefs, homogeneous precisions, and cognitive hierarchy beliefs.³ This suggests that the most parsimonious way to model human behavior in normal-form games is to use a model of this form, with the tradeoff between simplicity (i.e., number of parameters) and predictive power determined solely by the number of levels modeled.

To test the extent of the simplicity/power tradeoff in this family of models, we also evaluated ah-QCH4, ah-QCH5, ah-QCH6, and ah-QCH7, and for comparison purposes also evaluated every other model in our design space having 8 or fewer parameters. For the COMBO9 dataset, adding additional levels yielded small increases in predictive power until level 5, after which it yielded no further, statistically significant improvements. Thus, Figure 3 includes ah-QCH4 and ah-QCH5 as part of the efficient frontier.

Overall, we recommend that practitioners seeking a model for predicting human behavior in simultaneous-move games use a member of the ah-QCH family. How to model the distribution of levels is less clear. Of the models we considered, we recommend ah-QCH3. While it is possible to achieve gains in performance on our (large) dataset by modeling additional levels, these gains are small. Specifically, the gain in

³One might be interested in a weaker definition of the efficient frontier, saying that a model is efficient if it achieves significantly better performance than all *efficient* models with fewer parameters, rather than *all* models with fewer parameters. In this case the efficient frontier consists of all models previously identified as efficient plus ah-QCH7 and gi-QLk3. Our original definition rejected gi-QLk3 because it did not predict significantly better than gh-QLk2, which in turn did not predict significantly better than ah-QCH5.

Name	Max Level	Pop'n Beliefs	Precisions	Prec. Beliefs	#	Log-likelihood vs. u.a.r.
QLk1	1	n/a	n/a	n/a	2	18.37 ± 0.12
gi-QLk2	2	Lk	inhomo.	general	5	29.18 ± 0.03
ai-QLk2	2	Lk	inhomo.	accurate	4	26.75 ± 0.19
gh-QLk2	2	Lk	homo.	general	4	28.64 ± 0.04
ah-QLk2	2	Lk	homo.	accurate	3	26.18 ± 0.03
gi-QCH2	2	CH	inhomo.	general	5	28.17 ± 0.16
ai-QCH2	2	CH	inhomo.	accurate	4	27.39 ± 0.18
gh-QCH2	2	CH	homo.	general	4	27.90 ± 0.03
ah-QCH2	2	CH	homo.	accurate	3	27.44 ± 0.02
gi-QLk3	3	Lk	inhomo.	general	9	30.57 ± 0.17
ai-QLk3	3	Lk	inhomo.	accurate	6	29.54 ± 0.27
gh-QLk3	3	Lk	homo.	general	7	30.35 ± 0.20
ah-QLk3	3	Lk	homo.	accurate	4	27.27 ± 0.03
gi-QCH3	3	CH	inhomo.	general	10	30.35 ± 0.24
ai-QCH3	3	CH	inhomo.	accurate	6	29.96 ± 0.11
gh-QCH3	3	CH	homo.	general	8	30.29 ± 0.12
ah-QCH3	3	CH	homo.	accurate	4	29.47 ± 0.02
ai-QLk4	4	Lk	inhomo.	accurate	8	30.05 ± 0.26
ah-QLk4	4	Lk	homo.	accurate	5	27.30 ± 0.03
ah-QLk5	5	Lk	homo.	accurate	6	27.11 ± 0.11
ah-QLk6	6	Lk	homo.	accurate	7	27.02 ± 0.10
ah-QLk7	7	Lk	homo.	accurate	8	26.99 ± 0.12
ah-QLkp	*	Lk	homo.	accurate	2	27.34 ± 0.02
ai-QCH4	4	CH	inhomo.	accurate	8	29.86 ± 0.20
ah-QCH4	4	CH	homo.	accurate	5	29.82 ± 0.05
ah-QCH5	5	CH	homo.	accurate	6	30.20 ± 0.04
ah-QCH6	6	CH	homo.	accurate	7	30.25 ± 0.03
ah-QCH7	7	CH	homo.	accurate	8	30.33 ± 0.02
ah-QCHp	*	CH	homo.	accurate	2	27.70 ± 0.02

Table 2: Model variations, evaluated on the Combo9 dataset. The column headed # indicates the number of parameters in the model; the models with max level of * used a Poisson distribution.

prediction from adding two parameters to **ah-QCHp** to yield **ah-QCH3** was more than twice as great as was the gain in prediction from adding two parameters to **ah-QCH3** to yield **ah-QCH5**. Since more complex models are more prone to over-fitting—particularly with smaller amounts of data—we believe that **ah-QCH3** offers the best tradeoff between robustness and experimental performance. A practitioner working with a large dataset and interested in maximal prediction quality might also investigate **ah-QCH5**. Another possibility, which we have begun to explore in our current work, is to use a parametric distribution of levels whose shape is a better match to the experimental data than the Poisson distribution.

5.2 Parameter Analysis for **ah-QCH3**

We are now in a position to answer some of the questions from Section 4.2 by examining marginal posterior distributions from a member of our new model family, plotted in Figure 5.

We first note that, in contrast to QLk’s multimodal, jagged parameter CDFs, the parameter CDFs for **ah-QCH3** are smooth and (nearly) unimodal. This suggests that **ah-QCH3** is a much more robust model; its prediction quality is less likely to change drastically as a result of small changes in parameter values.

Second, the posterior distribution for the precision parameter λ is concentrated around 0.20, which is very close to the QLk model’s estimate for λ_2 . This suggests that QLk’s much lower estimate for $\lambda_{1(2)}$ may have been the closest that the

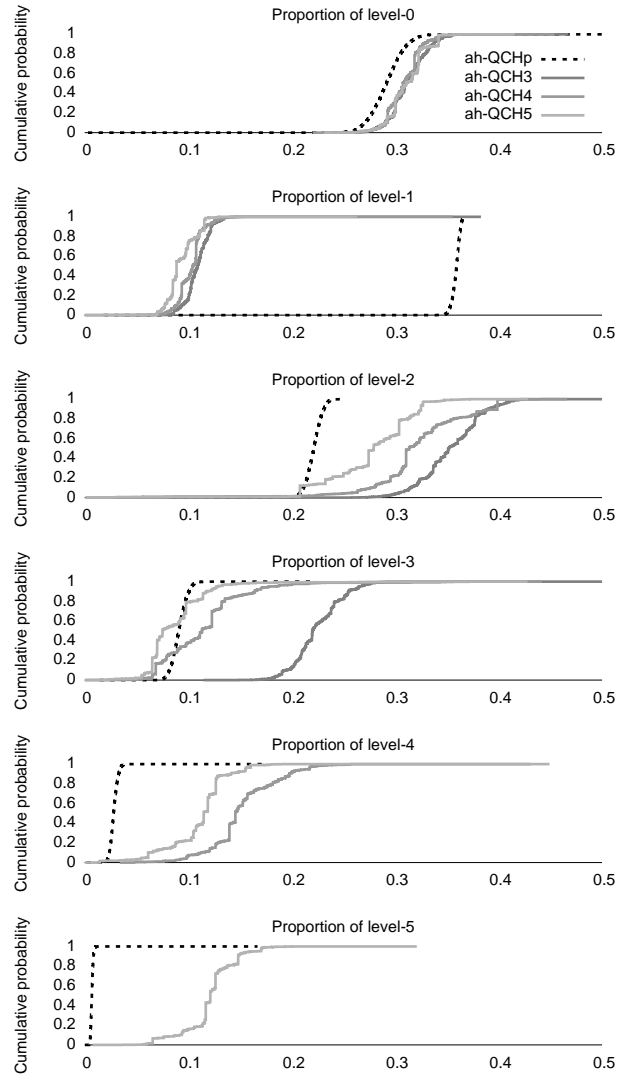


Figure 4: Marginal cumulative posterior distributions of levels of reasoning for efficient frontier models.

model could get to having the level-2 agents best respond to a mixture of level-0 and level-1 agents (as in cognitive hierarchy). It is unclear whether the order-of-magnitude differences and counterintuitive ordering of λ_1 and λ_2 are similar effects where QLk’s parameters are set in a way that “simulates” the assumptions of a more accurate model. Interestingly, like QLk, the **ah-QCH3** model predicts more level-2 agents than level-1. In fact, the **ah-QCH3** model predicts even fewer level-1 agents than QLk. This provides some support for QLk’s seemingly counterintuitive prediction that level-1 agents are less common than more sophisticated types.

6. CONCLUSIONS

We showed how Bayesian parameter analysis can be used to gain understanding about the sensitivity of behavioral game theoretic models to their parameters. We derived concrete recommendations for the use of an existing model, Poisson-CH, which differed substantially from advice in the

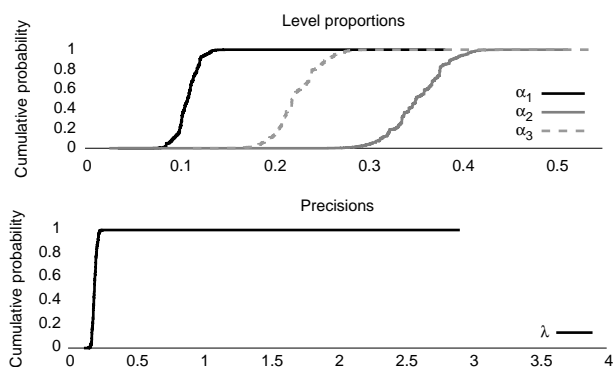


Figure 5: Marginal cumulative posterior distributions for the level proportion parameters ($\alpha_1, \alpha_2, \alpha_3$; top panel) and precision parameter (λ ; bottom panel) of the ah-QCH3 model on the combined dataset.

literature. We also uncovered anomalies in the best-performing existing model (QLk) that led us to a new, simpler, better-performing model (ah-QCH3). More broadly, the family of accurate, homogeneous-precision QCH models allows the modeler to trade off complexity against performance along an efficient frontier of models simply by adjusting the model of the distribution of levels.

In our ongoing work, we are now investigating parametric distributions of levels that match the observed distribution of levels. This would permit a model that uses higher levels, thus yielding higher performance, without requiring more parameters. Other promising directions include exploring applications of ah-QCH models for modeling behavior in practical settings such as markets or bargaining, and applying Bayesian analysis to additional behavioral models.

References

- [1] Ayala Arad and Ariel Rubinstein. The 11-20 money request game: A level- k reasoning study. <http://www.tau.ac.il/~aradaya1/moneyrequest.pdf>, May 2011.
- [2] Colin F. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, 2003.
- [3] Colin F. Camerer, Teck-Hua Ho, and Juin-Kuan Chong. A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 119(3):861–898, 2004.
- [4] David J. Cooper and John B. Van Huyck. Evidence on the equivalence of the strategic and extensive form representation of games. *Journal of Economic Theory*, 110(2):290–308, 2003.
- [5] Miguel Costa-Gomes, Vincent P. Crawford, and Bruno Broseta. Cognition and behavior in normal-form games: an experimental study. Discussion paper 98-22, UCSD, 1998.
- [6] Miguel Costa-Gomes, Vincent P. Crawford, and Bruno Broseta. Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5):1193–1235, 2001.
- [7] Xi Alice Gao and Avi Pfeffer. Learning game representations from data using rationality constraints. In *UAI*, pages 185–192, 2010.
- [8] Jeff Gill. *Bayesian methods: A social and behavioral sciences approach*. CRC press, 2002.
- [9] Jacob K. Goeree and Charles A. Holt. Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, 91(5):1402–1422, 2001.
- [10] Ernan Haruvy and Dale O. Stahl. Equilibrium selection and bounded rationality in symmetric normal-form games. *Journal of Economic Behavior and Organization*, 62(1):98–119, 2007.
- [11] Ernan Haruvy, Dale O. Stahl, and Paul W. Wilson. Modeling and testing for heterogeneity in observed strategic behavior. *Review of Economics and Statistics*, 83(1):146–157, 2001.
- [12] Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38, 1995.
- [13] Rosemarie Nagel. Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5):1313–1326, 1995.
- [14] Radford M. Neal. Annealed importance sampling. *Statistics and Computing*, 11(2):125–139, 2001.
- [15] J. A. Nelder and R. Mead. A simplex method for function minimization. *Computer Journal*, 7(4):308–313, 1965.
- [16] Christian P. Robert and George Casella. *Monte Carlo statistical methods*. Springer Verlag, 2004.
- [17] Brian W. Rogers, Thomas R. Palfrey, and Colin F. Camerer. Heterogeneous quantal response equilibrium and cognitive hierarchies. *Journal of Economic Theory*, 144(4):1440–1467, July 2009.
- [18] Dale O. Stahl and Ernan Haruvy. Level- n bounded rationality and dominated strategies in normal-form games. *Journal of Economic Behavior and Organization*, 66(2):226–232, 2008.
- [19] Dale O. Stahl and Paul W. Wilson. Experimental evidence on players’ models of other players. *Journal of Economic Behavior and Organization*, 25(3):309–327, 1994.
- [20] Dale O. Stahl and Paul W. Wilson. On players’ models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254, 1995.
- [21] Georg Weizsäcker. Ignoring the rationality of others: evidence from experimental normal-form games. *Games and Economic Behavior*, 44(1):145–171, 2003.
- [22] James R. Wright and Kevin Leyton-Brown. Beyond equilibrium: Predicting human behavior in normal-form games. In *AAAI*, pages 901–907, 2010.
- [23] Michael Wunder, Michael Kaisers, John Robert Yaros, and Michael Littman. Using iterated reasoning to predict opponent strategies. In *AAMAS*, pages 593–600, May 2011.