Beyond Equilibrium:
Predicting Human Behavior in Normal-Form Games

Kevin Leyton-Brown, University of British Columbia

Based on joint work with James R. Wright

OpLog Seminar, September 10, 2012

Overview

Context

- Motivation: Predict human behavior in strategic settings.
- Our focus: Unrepeated "initial play" in simultaneous-move, 2-player games.
- Game theory: Studies idealized rational agents, not human agents.
- Behavioral game theory: Aims to extend game theory to modeling human agents.
  - There are a wide range of BGT models in the literature.
  - Historically, BGT has been most concerned with explaining behavior, often on particular games, rather than predicting it.
  - No study compares a wide range of models, considers predictive performance, or looks at such a large, heterogeneous set of games.

Contribution

Our contributions:

- Compared predictive performance of the most prominent solution concepts for our setting:
  - Nash equilibrium, plus
  - Four models from behavioral game theory

  . . . using nine experimental datasets from the literature
- Bayesian sensitivity analysis:
  - Yields new insight into existing model (Poisson-CH)
  - Argues for a novel simplification of an existing model (Quantal level-k)

Overview

Example: Traveler's Dilemma

- Two players pick a number (2-100) simultaneously.

Example: Traveler's Dilemma

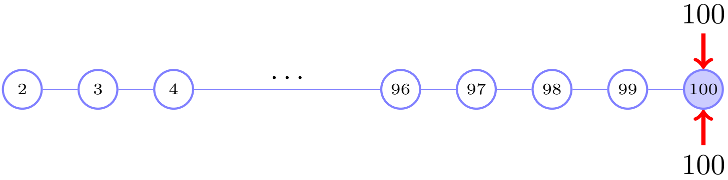

- Two players pick a number (2-100) simultaneously.
- If they pick the same number, that is their payoff.

Example: Traveler's Dilemma

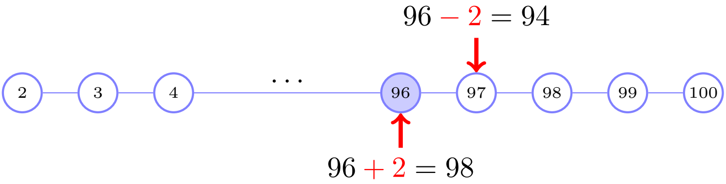

- Two players pick a number (2-100) simultaneously.
- If they pick the same number, that is their payoff.
- If they pick different numbers:
  - Lower player gets lower number, plus bonus of 2.
  - Higher player gets lower number, minus penalty of 2.

Example: Traveler's Dilemma

- Two players pick a number (2-100) simultaneously.
- If they pick the same number, that is their payoff.
- If they pick different numbers:
  - Lower player gets lower number, plus bonus of 2.
  - Higher player gets lower number, minus penalty of 2.

Example: Traveler's Dilemma

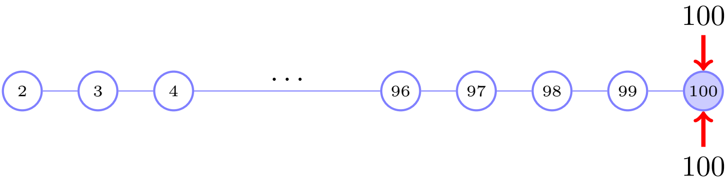

- Two players pick a number (2-100) simultaneously.
- If they pick the same number, that is their payoff.
- If they pick different numbers:
  - Lower player gets lower number, plus bonus of 2.
  - Higher player gets lower number, minus penalty of 2.
- Traveler's Dilemma has a unique Nash equilibrium.

Example: Traveler's Dilemma

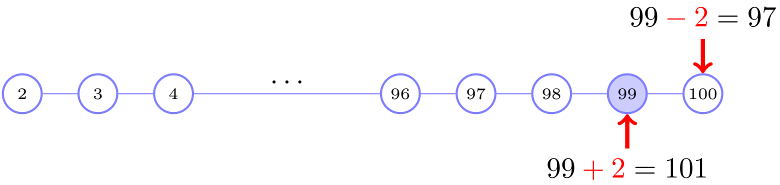

- Two players pick a number (2-100) simultaneously.
- If they pick the same number, that is their payoff.
- If they pick different numbers:
  - Lower player gets lower number, plus bonus of 2.
  - Higher player gets lower number, minus penalty of 2.
- Traveler's Dilemma has a unique Nash equilibrium.

Example: Traveler's Dilemma

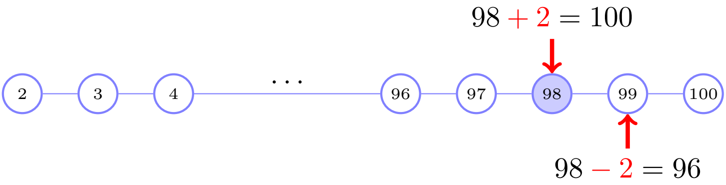

- Two players pick a number (2-100) simultaneously.
- If they pick the same number, that is their payoff.
- If they pick different numbers:
  - Lower player gets lower number, plus bonus of 2.
  - Higher player gets lower number, minus penalty of 2.
- Traveler's Dilemma has a unique Nash equilibrium.

Example: Traveler's Dilemma

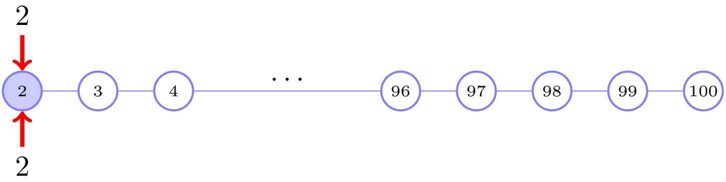

- Two players pick a number (2-100) simultaneously.
- If they pick the same number, that is their payoff.
- If they pick different numbers:
  - Lower player gets lower number, plus bonus of 2.
  - Higher player gets lower number, minus penalty of 2.
- Traveler's Dilemma has a unique Nash equilibrium.

Nash equilibrium and human subjects

- Nash equilibrium often makes counterintuitive predictions.
    - In Traveler's Dilemma: The vast majority of human players choose 97–100. The Nash equilibrium is 2.
- Modifications to a game that don't change Nash equilibrium predictions at all can cause large changes in how human subjects play the game [Goeree & Holt 2001].
    - In Traveler's Dilemma: When the penalty is large, people play much closer to Nash equilibrium.
    - But the size of the penalty does not affect equilibrium.

Nash equilibrium and human subjects

- Nash equilibrium often makes counterintuitive predictions.
  - In Traveler's Dilemma: The vast majority of human players choose 97–100. The Nash equilibrium is 2.
- Modifications to a game that don't change Nash equilibrium predictions at all can cause large changes in how human subjects play the game [Goeree & Holt 2001].
  - In Traveler's Dilemma: When the penalty is large, people play much closer to Nash equilibrium.
  - But the size of the penalty does not affect equilibrium.
- Clearly Nash equilibrium is not the whole story.
- Behavioral game theory proposes a number of models to better explain human behavior.

BGT model: Quantal response equilibrium (QRE)

Cost-proportional errors: Agents are less likely to make high-cost mistakes than low-cost mistakes.

QRE model [McKelvey & Palfrey 1995] parameter: (λ)

- Agents quantally best respond to each other.

$$QBR_i(s_{-i}, \lambda)(a_i) = \frac{e^{\lambda u_i(a_i, s_{-i})}}{\sum_{a_i' \in A_i} e^{\lambda u_i(a_i', s_{-i})}}$$

- Precision parameter $\lambda \in [0, \infty)$ indicates how sensitive agents are to utility differences.
  - $\lambda = 0$ means agents choose actions uniformly at random.
  - As $\lambda \to \infty$, QBR approaches best response.

Nice story—but is QRE a good model?

Let's say we pay a bunch of people to play games against each other, and gather some data. Now we'd like to know how good a job our QRE model does. How would we do that?

Two issues:

- have to set the model's parameter (λ) to use it at all;
- must ensure that we do this in a way that generalizes to new play by the same people.

Scoring a model's performance

- We randomly partition our data into different sets: $\mathcal{D} = \mathcal{D}_{train} \cup \mathcal{D}_{test}$
- We choose parameter value(s) that maximize the likelihood of the training data:

$$\vec{\theta}^* = \arg\max_{\vec{\theta}} \Pr(\mathcal{D}_{train} \,|\, \mathcal{M}, \vec{\theta}).$$

  - a tricky non-convex optimization problem
- We score the performance of a model by the likelihood of the test data:

$$\Pr(\mathcal{D}_{test} \,|\, \mathcal{M}, \vec{\theta}^*).$$

- To reduce variance, we repeat this process multiple times with different random partitions and average the results

BGT models: Iterative strategic reasoning

- Level-0 agents choose actions non-strategically.
    - In this work (and most others), uniformly at random
- Level-1 agents reason about level-0 agents.
- Level-2 agents reason about level-1 agents.
- There's a probability distribution over levels.
    - Higher-level agents are "smarter"; scarcer

- Predicting the distribution of play: weighted sum of the distributions for each level.

BGT model: Lk

Lk model [Costa-Gomes et al. 2001] parameters: $(\alpha_1, \alpha_2, \epsilon_1, \epsilon_2)$

- Each agent has one of 3 levels: level-0, level-1, or level-2.
- Distribution of level $[2, 1, 0]$ agents is $[\alpha_2, \alpha_1, (1 - \alpha_1 - \alpha_2)]$
- Each level-k agent makes a “mistake” with prob ϵ_k, or best responds to level-$(k-1)$ opponent with prob $1 - \epsilon_k$.
  - Level-k agents believe all opponents are level-$(k-1)$.
  - Level-k agents aren't aware that level-$(k-1)$ agents will make “mistakes”.

$$
\begin{aligned}
IBR_{i,0} &= A_i,\\
IBR_{i,k} &= BR_i(IBR_{-i,k-1}),\\
\pi^{Lk}_{i,0}(a_i) &= |A_i|^{-1},\\
\pi^{Lk}_{i,k}(a_i) &= \begin{cases} (1-\epsilon_k)/|IBR_{i,k}| & \text{if } a_i \in IBR_{i,k},\\ \epsilon_k/(|A_i| - |IBR_{i,k}|) & \text{otherwise.} \end{cases}
\end{aligned}
$$

BGT model: Cognitive hierarchy

Cognitive hierarchy model [Camerer et al. 2004] parameter: (τ)

- An agent of level m best responds to the truncated, true distribution of levels from 0 to $m-1$.
- Poisson-CH: Levels are assumed to have a Poisson distribution with mean τ.

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

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

$$\pi^{PCH}_{i,0:m-1} = \frac{\sum_{\ell=0}^{m-1} \pi^{PCH}_{i,\ell} \text{Pr}(\text{Poisson}(\tau) = \ell)}{\sum_{\ell=0}^{m-1} \text{Pr}(\text{Poisson}(\tau) = \ell)}$$

BGT model: QLk

QLk model [Stahl & Wilson 1994] parameters: $(\alpha_1, \alpha_2, \lambda_1, \lambda_2, \lambda_{1(2)})$

- Distribution of level $[2, 1, 0]$ agents is $[\alpha_2, \alpha_1, (1 - \alpha_1 - \alpha_2)]$
- Each agent quantally responds to next-lower level.
- Each QLk agent level has its own precision (λ_k), and its own beliefs about lower-level agents' precisions $(\lambda_{\ell(k)})$.

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

Overview

Model comparisons: Nash equilibrium vs. BGT

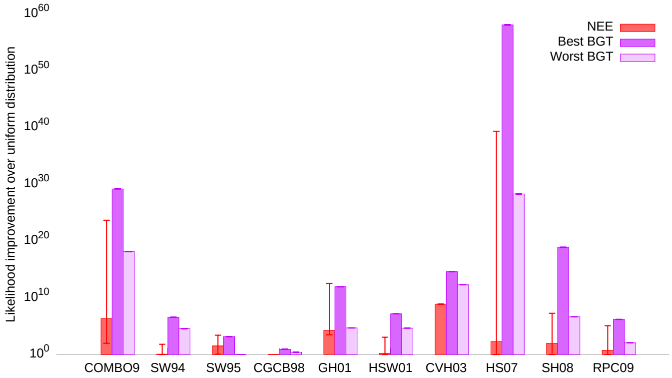

Model comparisons: Nash equilibrium vs. BGT

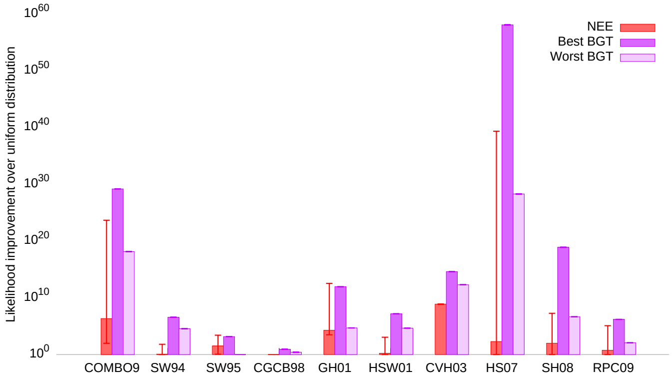

- Average NEE virtually always worse than every BGT model (only exception: `SW95`).
- All NEE significantly worse than best BGT model in most datasets.

Model comparisons: Lk and CH vs. QRE

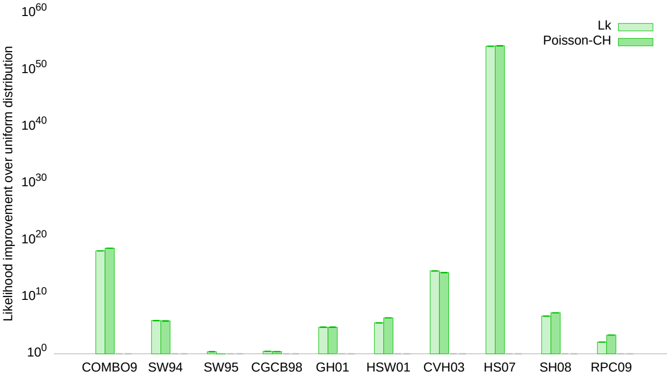

- Lk and Poisson-CH performance was strikingly similar.

Model comparisons: Lk and CH vs. QRE

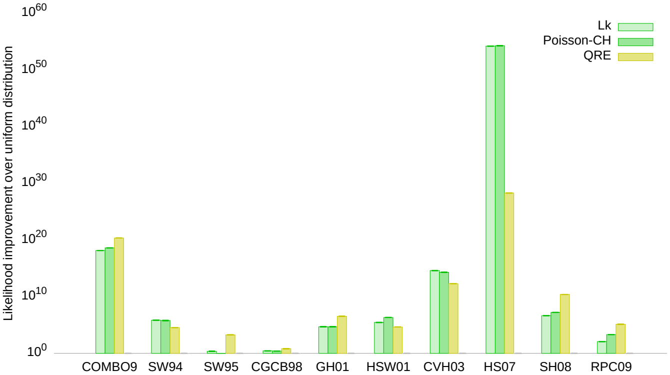

- Lk and Poisson-CH performance was strikingly similar.
- No consistent ordering between Lk/Poisson-CH and QRE.
  - Iterative strategic reasoning and quantal response appear to capture distinct phenomena.

Model comparisons: QLk

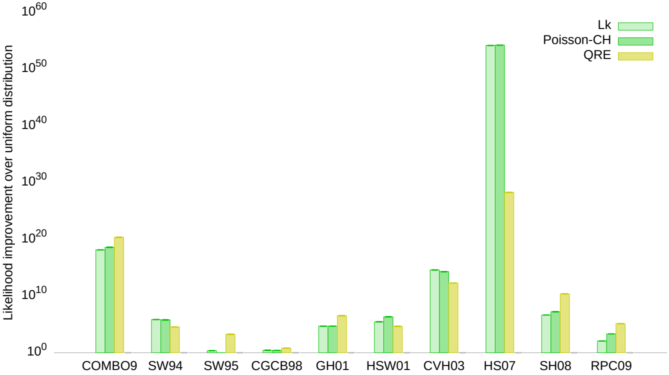

- So perhaps a model with both iterative and quantal response components would perform best?

Model comparisons: QLk

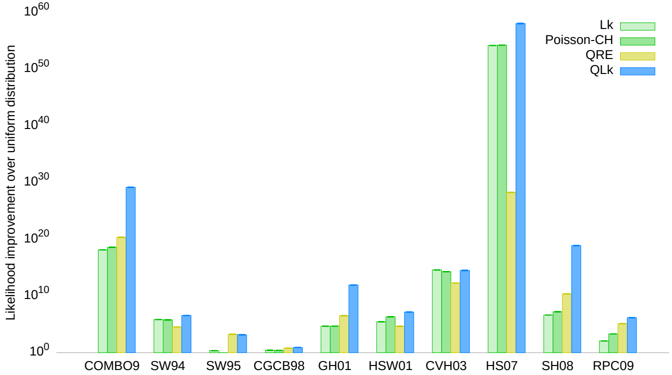

- So perhaps a model with both iterative and quantal response components would perform best?
- In fact, on every dataset, QLk is either the best predictive model or very similar to the best.

Overview

Taking Stock of What We Have Done

Take-home message so far

QLk is the best of the models for prediction.

Question

How strongly does the data argue for particular parameter values?

Posterior distributions

A posterior distribution gives the probability of each possible combination of parameter values, given the data, e.g.:

$$\Pr(\alpha_1 = 0.1, \alpha_2 = 0.3, \lambda = 0.1 \,|\, \mathcal{D})$$

- Maximum likelihood only tells us the most likely parameter setting, given the data.
- The posterior distribution over parameter settings describes the relative probability of all possible parameter settings.
- Individual parameters can be analyzed by inspecting the marginal posterior distribution.

$$\Pr(\alpha_1 = 0.1 \,|\, \mathcal{D}) = \iint \Pr(\alpha_1 = 0.1, \alpha_2 = \alpha_2', \lambda = \lambda' \,|\, \mathcal{D}) d\alpha_2' d\lambda'$$

  - Flat distributions indicate less important parameter values.
  - Sharp distributions indicate a high degree of certainty.

Warm-up: Poisson-CH

Regarding the single parameter (τ) for the Poisson-CH model:

> *"Indeed, values of τ between 1 and 2 explain empirical results for nearly 100 games, suggesting that a τ value of 1.5 could give reliable predictions for many other games as well."* [Camerer et al. 2004]

Warm-up: Poisson-CH's Posterior Distribution

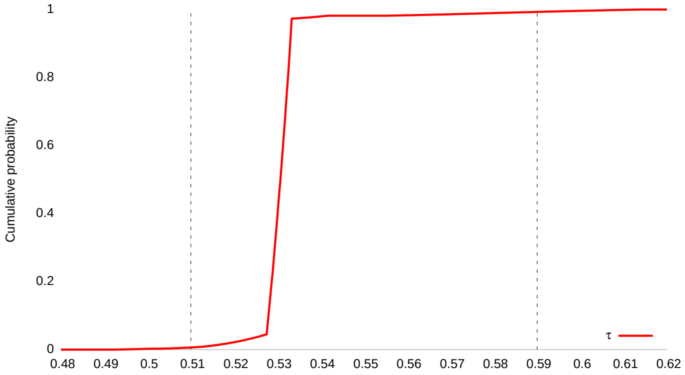

Our analysis gives 99% posterior probability that the best value of τ is 0.59 or less.

Refresher: QLk's Parameters

QLk has 5 different parameters:

- α_1: Proportion of level-1 agents.
- α_2: Proportion of level-2 agents.
- λ_1: Precision of level-1 agents.
- λ_2: Precision of level-2 agents.
- $\lambda_{1(2)}$: Level-2 agents' belief about level-1 agents' precision.

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

Posterior distributions: QLk

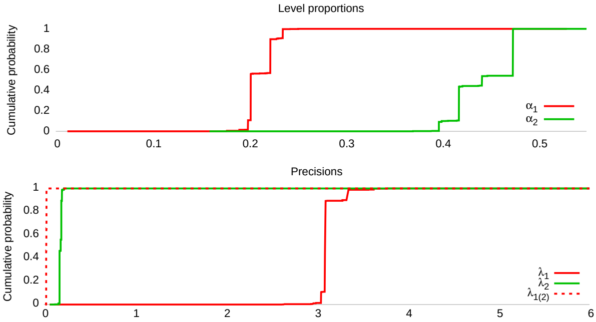

Posterior distributions: QLk

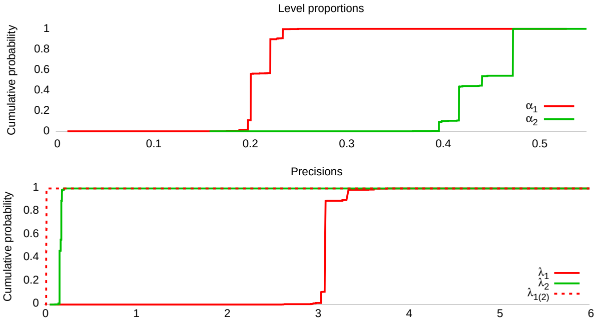

Some surprises:

1. α_1, α_2: Best fits predict more level-2 agents than level-1.

Posterior distributions: QLk

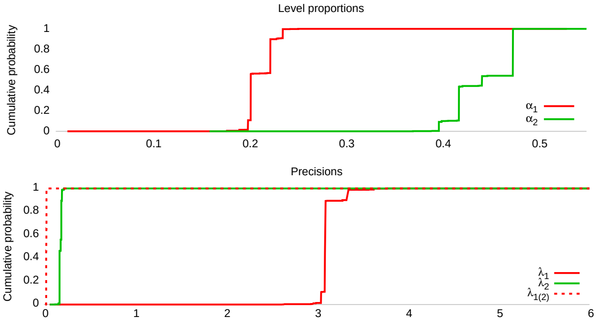

Some surprises:

1. α_1, α_2: Best fits predict more level-2 agents than level-1.
2. λ_1, λ_2: Level-2 agents have lower precision than level-1 agents.
3. $\lambda_1,\ \lambda_{1(2)}$: Level-2 agents' beliefs are very wrong.

Maybe QLk isn't quite the right model

We constructed a family of models by systematically varying QLk:

1. Top level:
   - 1, **2**, 3, 4, 5, 6, 7, Poisson
2. Precisions: Homogeneous or **inhomogeneous**.
3. Precision beliefs: Accurate or **general**.
4. Population beliefs: **Lk** or CH.

We evaluated all variations leading to ≤ 8 parameters.

Model variations: Efficient frontier

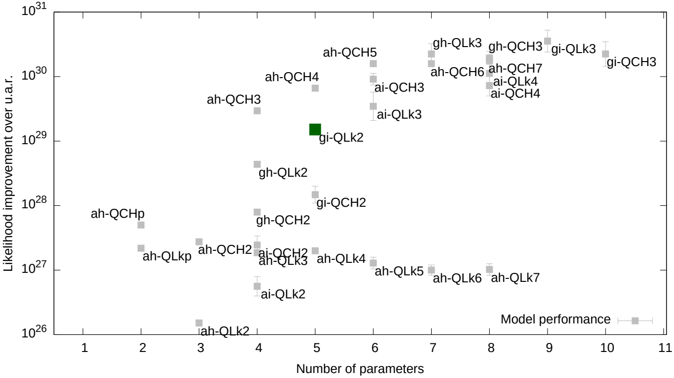

Model variations: Efficient frontier

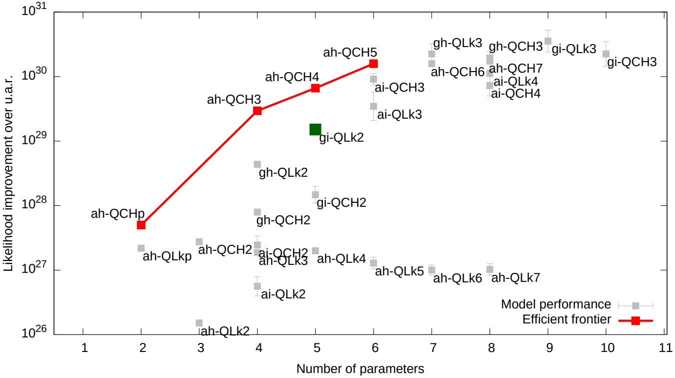

- Efficient frontier: best performance for # of parameters.
- QLk (gi-QLk2) is not on the efficient frontier.
- Best models all have accurate precision beliefs, homogeneous precision, cognitive hierarchy population beliefs.

Thinking back to QLk

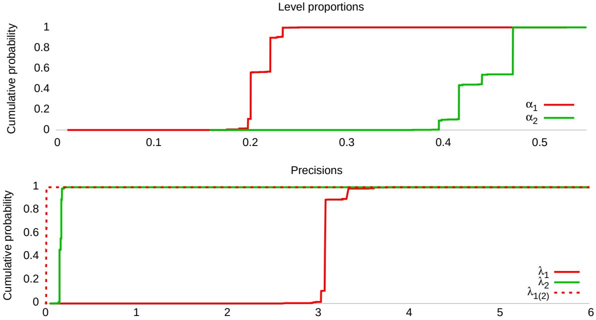

Recall...

- α_1, α_2: Best fits predict more level-2 agents than level-1.
- λ_1, λ_2: Level-2 agents have lower precision than level-1 agents.
- λ_1, $\lambda_{1(2)}$: Level-2 agents' beliefs are very wrong.

ah-QCH3: Posterior distribution

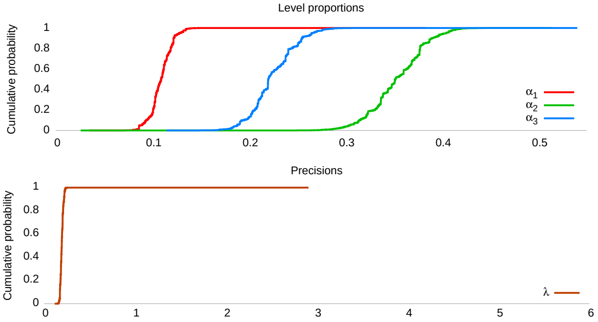

- More robust model: small parameter changes less likely to change prediction quality.
  - Smooth, unimodal distributions for level proportions.
- Distribution for λ is unimodal, with narrow confidence region
- Still more agents of type 2 than 1.

Marginal distributions comparison

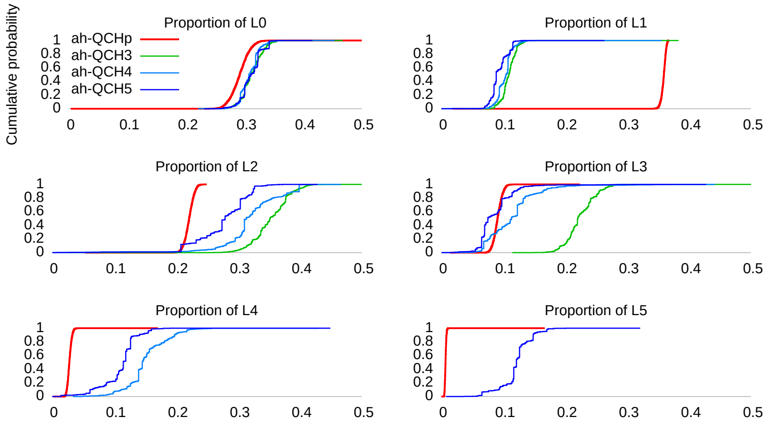

- Poisson QCH matches tabular L0 proportions very closely.
- To do so, forced to match most other proportions poorly.

Marginal distributions comparison

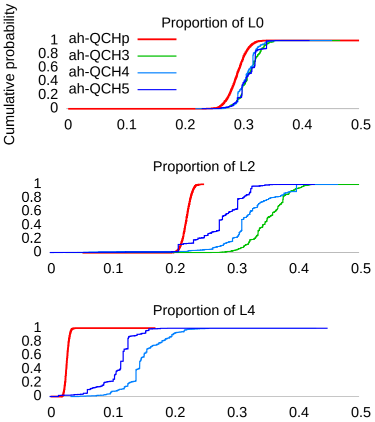

- Poisson QCH matches tabular L0 proportions very closely.
- To do so, forced to match most other proportions poorly.
- If L0 were treated specially, could Poisson match others?

Spike-Poisson model

Spike-Poisson QCH model parameters: $(\tau, \epsilon, \lambda)$

- An ah-QCH model with precision λ.
- Proportion distribution f is a mixture of Poisson distribution and a "spike" distribution of L0 agents:

$$f(m) = \begin{cases} \epsilon + (1-\epsilon)\text{Poisson}(m;\tau) & \text{if } m = 0, \\ (1-\epsilon)\text{Poisson}(m;\tau) & \text{otherwise.} \end{cases}$$

Spike-Poisson performance

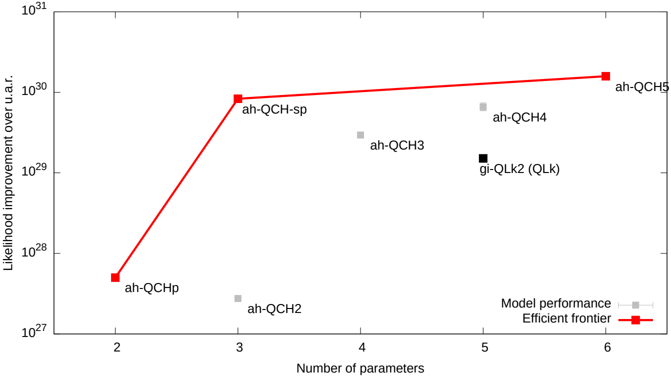

- Spike-Poisson QCH outperforms all other ah-QCH models except for ah-QCH5.
- Only three parameters, fewer even than ah-QCH3.

Summary

- Compared predictive performance of four BGT models.
  - BGT models typically predict human behavior better than Nash equilibrium-based model.
  - QLk has best performance of the four.
- Bayesian sensitivity analysis of parameters.
  - Parameters for QLk are counterintuitive, hard to identify.
  - Using CH beliefs and a single precision for all agents yields more identifiable parameter values, superior predictive performance.
    - Even with fewer parameters!

Thank you!

- Compared predictive performance of four BGT models.
  - BGT models typically predict human behavior better than Nash equilibrium-based model.
  - QLk has best performance of the four.
- Bayesian sensitivity analysis of parameters.
  - Parameters for QLk are counterintuitive, hard to identify.
  - Using CH beliefs and a single precision for all agents yields more identifiable parameter values, superior predictive performance.
    - Even with fewer parameters!