Behavioral Game Theory

Based on joint work with James Wright

Talk Overview

Computational Problems in Domination

- Identifying strategies dominated by a pure strategy
  - polynomial, straightforward algorithm
- Identifying strategies dominated by a mixed strategy
  - polynomial, somewhat tricky LP
- Identifying strategies that survive iterated elimination
  - repeated calls to the above LP
- Asking whether a strategy survives iterated elimination under all elimination orderings
  - polynomial for strict domination (elimination doesn't matter)
  - NP-complete otherwise

Rationalizability

- Rather than ask what is irrational, ask what is a best response to some beliefs about the opponent
  - assumes opponent is rational
  - assumes opponent knows that you and the others are rational
  - ...
- Equilibrium strategies are always rationalizable; so are lots of other strategies (but not everything).
- In two-player games, rationalizable $\Leftrightarrow$ survives iterated removal of strictly dominated strategies.

Formal definition

Definition (Correlated equilibrium)

Given an n-agent game $G = (N, A, u)$, a correlated equilibrium is a tuple (v, π, σ), where v is a tuple of random variables $v = (v_1, \ldots, v_n)$ with respective domains $D = (D_1, \ldots, D_n)$, π is a joint distribution over v, $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a vector of mappings $\sigma_i : D_i \mapsto A_i$, and for each agent i and every mapping $\sigma_i' : D_i \mapsto A_i$ it is the case that

$$\sum_{d \in D} \pi(d) u_i \left(\sigma_1(d_1), \ldots, \sigma_i(d_i), \ldots, \sigma_n(d_n)\right) \geq \sum_{d \in D} \pi(d) u_i \left(\sigma_1(d_1), \ldots, \sigma_i'(d_i), \ldots, \sigma_n(d_n)\right).$$

Existence

Theorem

For every Nash equilibrium σ^* there exists a corresponding correlated equilibrium σ.

- This is easy to show:
  - let $D_i = A_i$
  - let $\pi(d) = \prod_{i \in N} \sigma_i^*(d_i)$
  - σ_i maps each d_i to the corresponding a_i.
- Thus, correlated equilibria always exist

Remarks

- Not every correlated equilibrium is equivalent to a Nash equilibrium
  - thus, correlated equilibrium is a weaker notion than Nash
- Any convex combination of the payoffs achievable under correlated equilibria is itself realizable under a correlated equilibrium
  - start with the Nash equilibria (each of which is a CE)
  - introduce a second randomizing device that selects which CE the agents will play
  - regardless of the probabilities, no agent has incentive to deviate
  - the probabilities can be adjusted to achieve any convex combination of the equilibrium payoffs
  - the randomizing devices can be combined

Talk Overview

Computing CE

$$\sum_{a \in A | a_i \in a} p(a) u_i(a) \geq \sum_{a \in A | a_i \in a} p(a) u_i(a'_i, a_{-i}) \quad \forall i \in N, \forall a_i, a'_i \in A_i$$

$$p(a) \geq 0 \quad \forall a \in A$$

$$\sum_{a \in A} p(a) = 1$$

- variables: $p(a)$; constants: $u_i(a)$

Computing CE

$$\sum_{a \in A | a_i \in a} p(a) u_i(a) \geq \sum_{a \in A | a_i \in a} p(a) u_i(a_i', a_{-i}) \quad \forall i \in N, \forall a_i, a_i' \in A_i$$

$$p(a) \geq 0 \quad \forall a \in A$$

$$\sum_{a \in A} p(a) = 1$$

- variables: $p(a)$; constants: $u_i(a)$
- we could find the social-welfare maximizing CE by adding an objective function

$$\text{maximize:} \quad \sum_{a \in A} p(a) \sum_{i \in N} u_i(a).$$

Why are CE easier to compute than NE?

$$\sum_{a \in A | a_i \in a} p(a) u_i(a) \geq \sum_{a \in A | a'_i \in a} p(a) u_i(a'_i, a_{-i}) \quad \forall i \in N, \forall a_i, a'_i \in A_i$$

$$p(a) \geq 0 \quad \forall a \in A$$

$$\sum_{a \in A} p(a) = 1$$

- intuitively, correlated equilibrium has only a single randomization over outcomes, whereas in NE this is constructed as a product of independent probabilities.
- To change this program so that it finds NE, the first constraint would be

$$\sum_{a \in A} u_i(a) \prod_{j \in N} p_j(a_j) \geq \sum_{a \in A} u_i(a'_i, a_{-i}) \prod_{j \in N \setminus \{i\}} p_j(a_j) \quad \forall i \in N, \forall a'_i \in A_i.$$

- This is a nonlinear constraint!

Talk Overview

Fun Game

- Guess 95% of the average

Fun Game

- Guess 95% of the average
- Guess 40% of the average

Behavioral Game Theory

- 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.

This Talk

We:

- Compared predictive performance of:
  - Nash equilibrium, plus
  - Four prominent models from behavioral game theory
  - Using six 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)

Talk Overview

Nash equilibrium and human subjects

- Nash equilibrium often makes counterintuitive predictions.
  - In Traveler's Dilemma: The vast majority of human players choose 97–100.
- 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.
- 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.

Behavioral game theory models

Themes:[1]

1. Quantal response: Agents best-respond with high probability rather than deterministically best responding.
2. Iterative strategic reasoning: Agents can only perform limited steps of strategic "look-ahead".

One model (QRE) is based on quantal response, two models (Lk, CH) are based on iterative strategic reasoning, and one model (QLk) incorporates both.

[1] Recall: we restrict attention to unrepeated, simultaneous-move games.

BGT model: Quantal response equilibrium (QRE)

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.

BGT models: Iterative strategic reasoning

- Level-0 agents choose 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".

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

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.

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

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

$$TBR_{i,m} = BR_i\left(\sum_{\ell=0}^{m-1} \Pr(\text{Poisson}(\tau) = \ell)\pi_{-i,\ell}^{PCH}\right)$$

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 $(\lambda_{\ell(k)})$.

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

Prediction using Nash equilibrium

- We would like to compare BGT models' prediction performance to Nash equilibrium.
- Unmodified Nash equilibrium is not suitable for predictions:
  1. Games often have multiple Nash equilibria.
  2. A Nash equilibrium will often assign probability 0 to some actions.

Prediction using Nash equilibrium

- We would like to compare BGT models' prediction performance to Nash equilibrium.
- Unmodified Nash equilibrium is not suitable for predictions:
  1. Games often have multiple Nash equilibria.
  2. A Nash equilibrium will often assign probability 0 to some actions.
- We constructed two different Nash-based models to deal with multiple equilibria:
  - UNEE (Uniform Nash Equilibrium with Error):
    - Predict the average of all Nash equilibria, + error.
  - NNEE (Nondeterministic Nash Equilibrium with Error):
    - Predict the post-hoc "best" Nash equilibrium, + error.
- Both models avoid probability 0 predictions via a tunable error probability.

Talk Overview

Experimental setup: Overview

What do we need to compare predictive models?

1. Metric to measure performance
2. Test to evaluate generalization performance
3. Experimental data describing real human play

Key issue: must set models' free parameters using data (a tricky optimization problem), then test generalization performance.

1. Performance Metric

- We score the performance of a model by the likelihood of the test data:
$$\mathbf{P}(\mathcal{D}_{test} \mid \mathcal{M}, \vec{\theta}^*).$$
- To evaluate this score (and make meaningful model predictions) we must choose the parameters to maximize the likelihood of the training data:
$$\vec{\theta}^* = \arg\max_{\vec{\theta}} \mathbf{P}(\mathcal{D}_{train} \mid \mathcal{M}, \vec{\theta}).$$

2. Test to evaluate Generalization Performance

- We estimate generalization performance using 10-fold cross-validation.
- Problem: this estimate may depend upon the particular partition into folds.
- We average over multiple (again, 10) cross-validation runs.
- We can then compute 95% confidence interval by assuming a t-distribution of these averages [Witten & Frank 2000].
  - This assumption is validated by the law of large numbers as the number of cross-validation repetitions grows.

3. Experimental data

- Subjects played 2-player normal form games once each.
- Each action by an individual player is a single observation.
- Data from six experimental studies, plus a combined dataset:
  - SW94: 400 observations from [Stahl & Wilson 1994]
  - SW95: 576 observations from [Stahl & Wilson 1995]
  - CGCB98: 1296 observations from [Costa-Gomes et al. 1998]
  - GH01: 500 observations from [Goeree & Holt 2001]
  - CVH03: 2992 observations from [Cooper & Van Huyck 2003]
  - RPC09: 1210 observations from [Rogers et al. 2009]
  - ALL6: All 6974 observations

Talk Overview

Model comparisons: Nash equilibrium vs. BGT

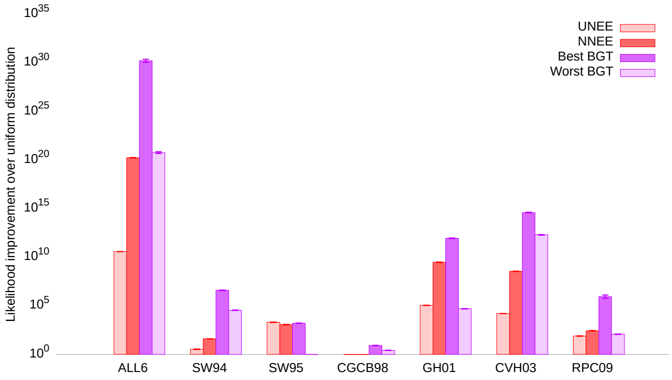

- UNEE almost always worse than every BGT model (exceptions: GH01, SW95).
- Even NNEE worse than QLk and QRE in most datasets.

Model comparisons: Lk and CH vs. QRE

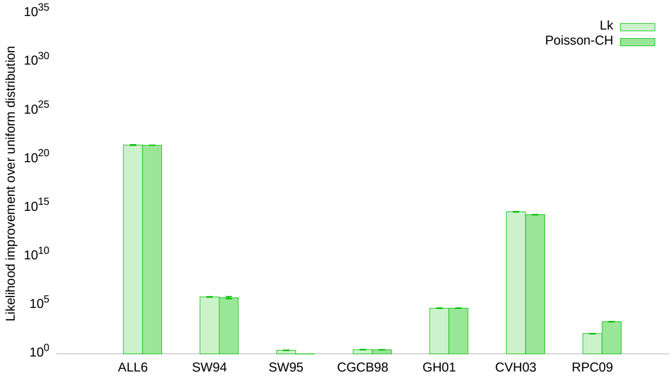

- Lk and Poisson-CH performance was roughly similar.
- No consistent ordering between Lk/Poisson-CH and QRE.
  - Iterative strategic reasoning and quantal response appear to capture distinct phenomena.

Model comparisons: Lk and CH vs. QRE

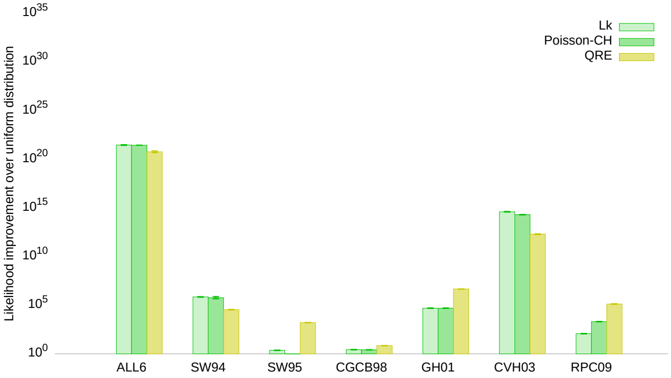

- Lk and Poisson-CH performance was roughly similar.
- No consistent ordering between Lk/Poisson-CH and QRE.
  - Iterative strategic reasoning and quantal response appear to capture distinct phenomena.

Model comparisons: QLk

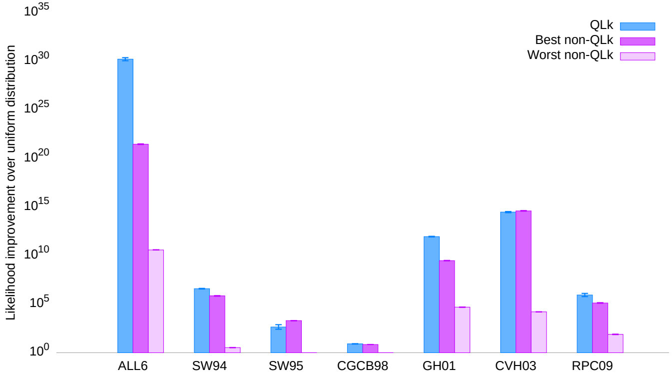

- We would expect a model with both iterative and quantal response components to perform best.
- That is the case: QLk is the best predictive model on almost every dataset.

Talk Overview

Taking Stock of What We Have Done

Take-home message so far:

QLk is the best of the models for prediction.

For the rest of the talk, we will concentrate on gaining a deeper understanding of QLk. We will:

- use Bayesian methods to understand QLk's parameter space
  - determining which ranges of values parameters can take
  - identifying the most important parameters
- indulge ourselves by digressing to consider Poisson-CH ☺

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.

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

Bayesian Sensitivity Analysis

Two questions:

1. How sure are we about the parameter values we fit?
   - That is, how strongly does the data argue for particular parameter values?

Bayesian Sensitivity Analysis

Two questions:

1. How sure are we about the parameter values we fit?
    - That is, how strongly does the data argue for particular parameter values?
2. How important are the different parameters?
    - We say that a parameter is important if the model's predictions are substantially degraded if we change its value.

Posterior distributions

- Maximum likelihood only tells us the most likely parameter setting, given the data.
- The posterior distribution over parameter settings describes the relative probability (normalized likelihoods) of all possible parameter settings.
- Individual parameters can be analyzed by inspecting the marginal posterior distribution.
  - Flat distributions indicate uncertainty about parameter values.
  - Sharp distributions indicate a high degree of certainty.

Posterior distributions: Monte Carlo sampling

- The posterior distribution rarely has an analytic representation.
- We use Monte Carlo sampling to draw an approximate sample of values from the joint posterior distribution.
- Expectations taken over these approximate samples are unbiased estimators of the true expectations.

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

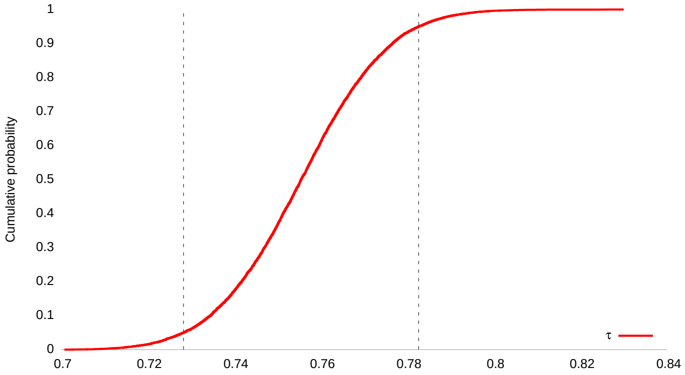

Our analysis gives 99% posterior probability that the best value of τ is 0.8 or less.

Posterior distributions: QLk

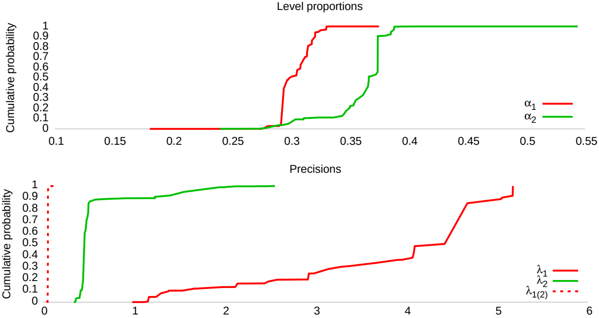

Posterior distributions: QLk

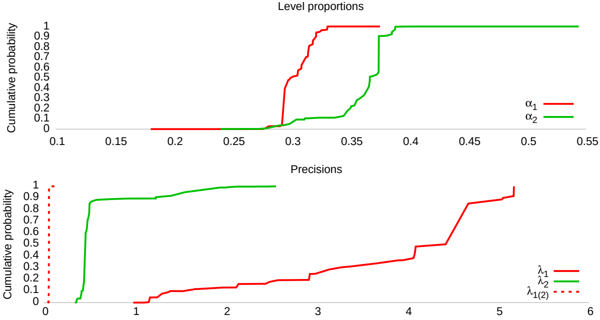

Some surprises:

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

Posterior distributions: QLk

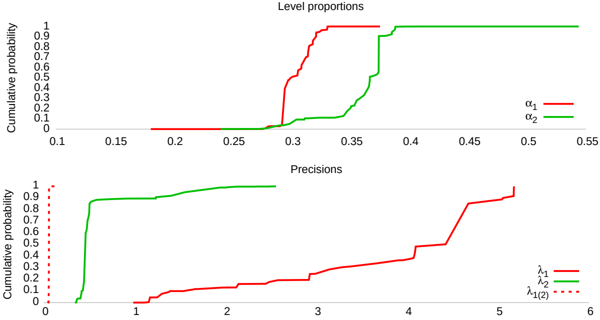

Some surprises:

1. α_1, α_2: Best fits predict more level-2 agents than level-1.
2. λ_1 is not very identifiable from data; multimodal.

Posterior distributions: QLk

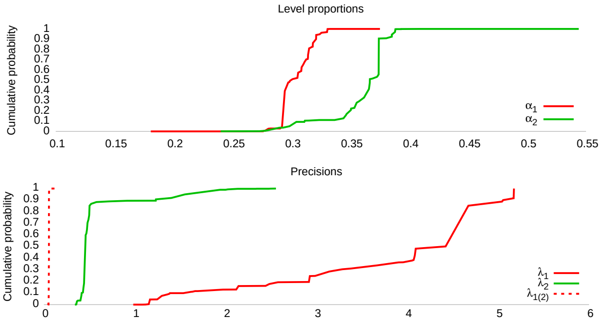

Some surprises:

1. α_1, α_2: Best fits predict more level-2 agents than level-1.
2. λ_1 is not very identifiable from data; multimodal.
3. λ_1, λ_2: Level-2 agents have lower precision than level-1 agents.
4. λ_1, $\lambda_{1(2)}$: Level-2 agents' beliefs are very wrong.

Sensitivity measures: Main and total effect

Definition

The main effect of a parameter θ_j is the percentage of variance in the prediction that is reduced when θ_j is fixed to its true value.

Problem: If a parameter has most of its influence from interactions with other parameters, it will have a low main effect.

- Could compute second-order effects, third-order, ...
- There are exponentially many of these!

Definition

The total effect of a parameter θ_j is the sum of all main and higher-order effects that θ_j participates in.

We estimate both of these quantities using Monte Carlo sampling.

Parameter importance: QLk

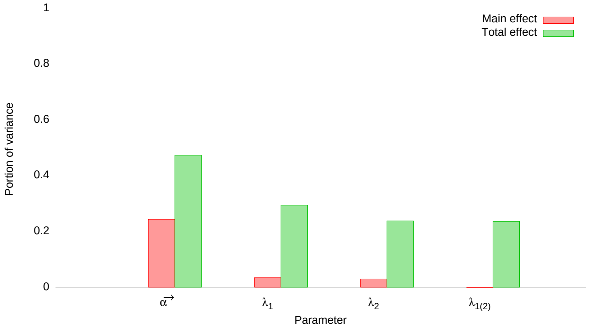

- High interaction effects for all parameters.
- Precision parameters influence mostly through interactions.
- Proportion parameters ($\vec{\alpha}$) about twice as important as precision parameters ($\lambda_1, \lambda_2, \lambda_{1(2)}$).

Homogeneous QLk

- QLk is not very sensitive to its individual precision parameters.
- The precision parameters are also hard to identify.
- Would a single precision parameter λ serve just as well?

Definition (Homogeneous QLk model)

$$
\begin{aligned}
\pi_{i,0}^{HQLk}(a_i) &= |A_i|^{-1}, \\
\pi_{i,1}^{HQLk} &= QBR_i(\pi_{-i,0}^{HQLk}, \lambda), \\
\pi_{i,2}^{HQLk} &= QBR_i(\pi_{-i,1}^{HQLk}, \lambda).
\end{aligned}
$$

Homogeneous QLk: Parameter importance

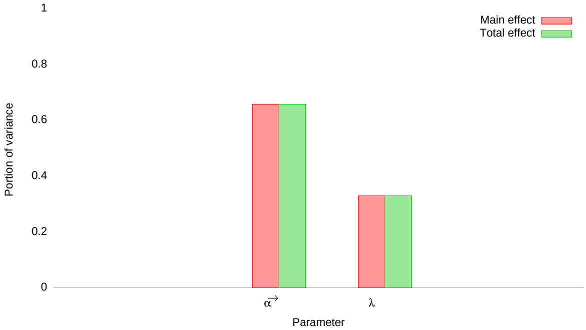

- Virtually no interaction effects.
- Proportion parameters ($\vec{\alpha}$) still about twice as important as precision (λ).

Thinking back to QLk

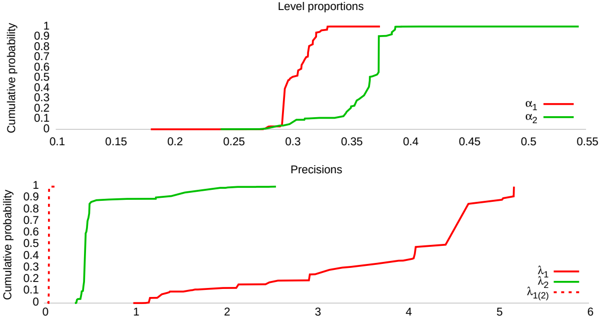

Recall...

- α_1, α_2: Best fits predict more level-2 agents than level-1.
- λ_1 is not very identifiable from data; multimodal.
- 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.

Homogeneous QLk: Posterior distribution

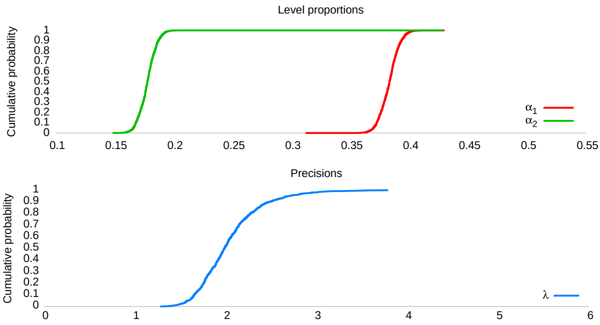

- More agents of type 1 than 2.
- More confident identification of precision (λ):
  - Unimodal;
  - Narrower confidence region.

Homogeneous QLk: Performance

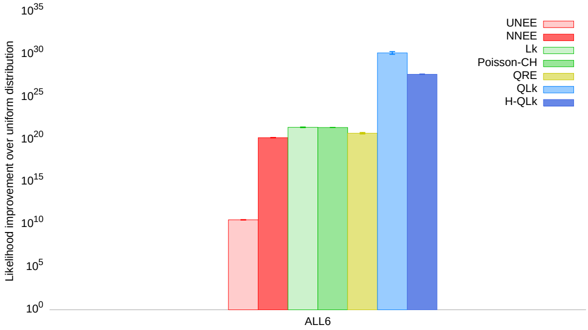

Performance of H-QLk and QLk are similar; H-QLk still outperforms all other models.

Summary

- Compared predictive performance of four BGT models.
  - BGT models typically predict human behavior better than Nash equilibrium-based models.
  - QLk has best performance of the four.
- Bayesian sensitivity analysis of parameters.
  - Best parameter for Poisson-CH is likely much lower than previously thought.
  - Parameters for QLk are counterintuitive, hard to identify, interaction-laden.
  - Using a single precision for all agents yields better more intuitive parameter values, similar predictive performance.