Predicting Human Behavior In Games

James Wright

March 18, 2014

Behavioral Game Theory

- Many of game theory's recommendations are very counter-intuitive.
- Do people actually follow them?

Behavioral Game Theory

- Many of game theory's recommendations are very counter-intuitive.
- Do people actually follow them?
- No. A large body of experiments demonstrates otherwise.
- Behavioral game theory: Aims to model actual human behavior in games.

Fun Game: Traveler's Dilemma

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

Fun Game: Traveler's Dilemma

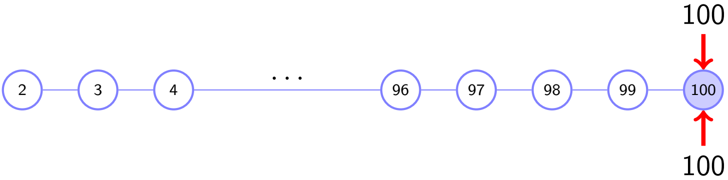

- Two players pick a number (2-100) simultaneously.
- If they pick the same number, that is their payoff.

Fun Game: Traveler's Dilemma

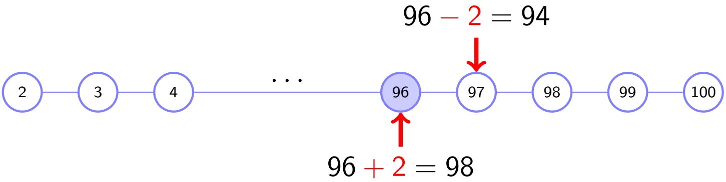

- Two players pick a number (2-100) simultaneously.
- If they pick the same number, that is their payoff.
- If they pick different numbers:
  - Lower player gets lower number, plus bonus of 2.
  - Higher player gets lower number, minus penalty of 2.

Fun Game: 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.
- Give this game a try. Play any opponent only once.

Fun Game: 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 30.
  - Higher player gets lower number, minus penalty of 30.
- Now play a different opponent with a larger penalty.

Fun Game: Traveler's Dilemma

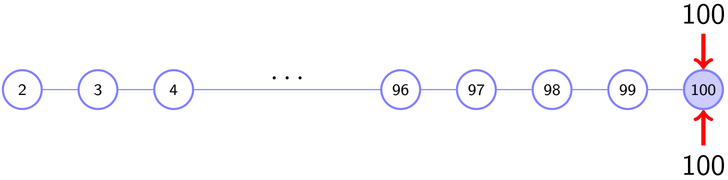

- Two players pick a number (2-100) simultaneously.
- If they pick the same number, that is their payoff.
- If they pick different numbers:
  - Lower player gets lower number, plus bonus of 30.
  - Higher player gets lower number, minus penalty of 30.
- Traveler's Dilemma has a unique Nash equilibrium.

Fun Game: Traveler's Dilemma

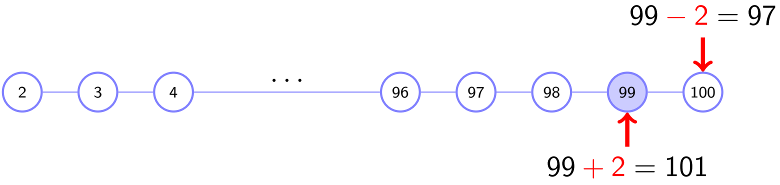

- Two players pick a number (2-100) simultaneously.
- If they pick the same number, that is their payoff.
- If they pick different numbers:
  - Lower player gets lower number, plus bonus of 30.
  - Higher player gets lower number, minus penalty of 30.
- Traveler's Dilemma has a unique Nash equilibrium.

Fun Game: Traveler's Dilemma

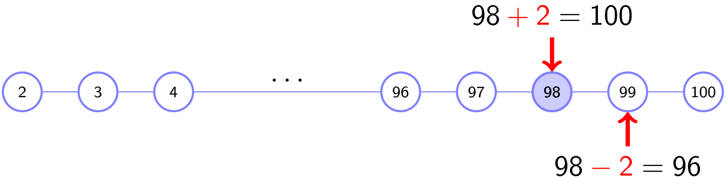

- Two players pick a number (2-100) simultaneously.
- If they pick the same number, that is their payoff.
- If they pick different numbers:
  - Lower player gets lower number, plus bonus of 30.
  - Higher player gets lower number, minus penalty of 30.
- Traveler's Dilemma has a unique Nash equilibrium.

Fun Game: Traveler's Dilemma

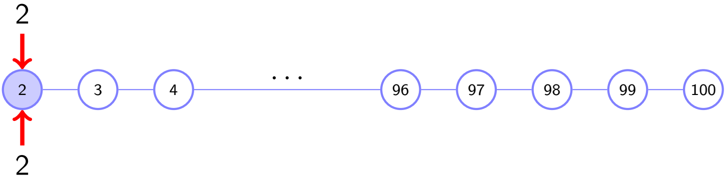

- Two players pick a number (2-100) simultaneously.
- If they pick the same number, that is their payoff.
- If they pick different numbers:
  - Lower player gets lower number, plus bonus of 30.
  - Higher player gets lower number, minus penalty of 30.
- Traveler's Dilemma has a unique Nash equilibrium.

Comparing Behavioral Models

[Wright & Leyton-Brown 2010]

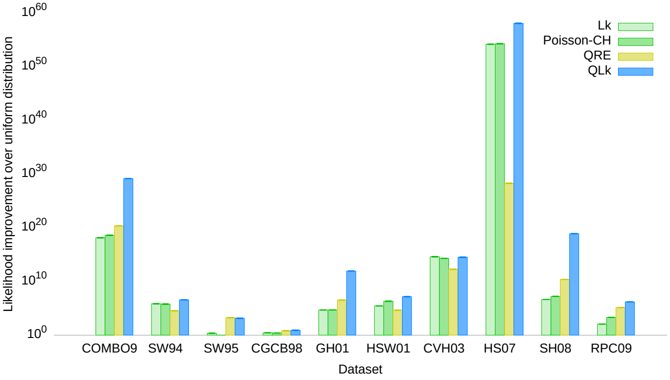

- Many behavioral models have been proposed.
- First study to compare prediction performance of several at once.

Comparing Behavioral Models

[Wright & Leyton-Brown 2010]

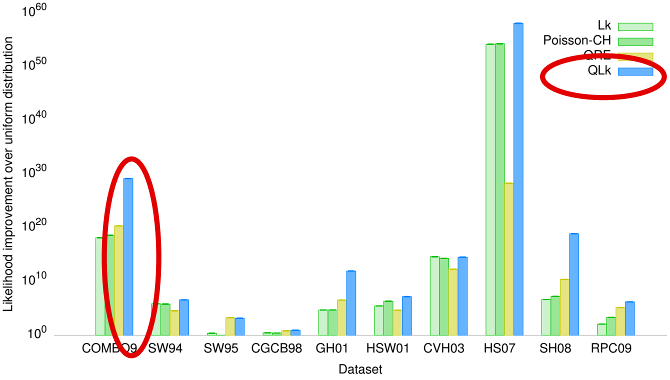

- Many behavioral models have been proposed.
- First study to compare prediction performance of several at once.
- One model performed clearly better than the others.

Two main ideas

1. Quantal utility maximization instead of utility maximization.
2. Iterative reasoning instead of equilibrium.

Two main ideas

1. Quantal utility maximization instead of utility maximization.
2. Iterative reasoning instead of equilibrium.

Iterative reasoning

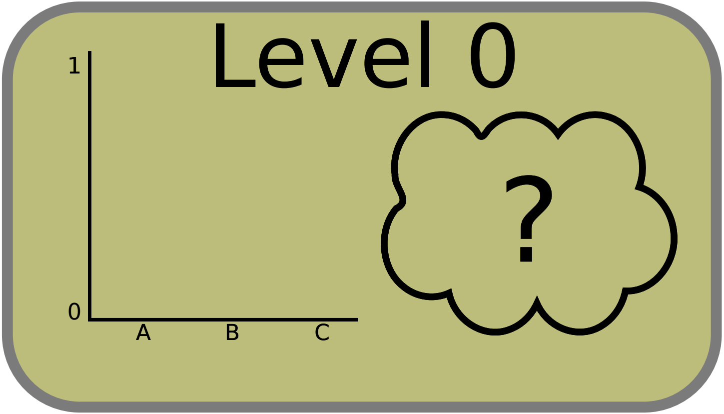

Iterative reasoning

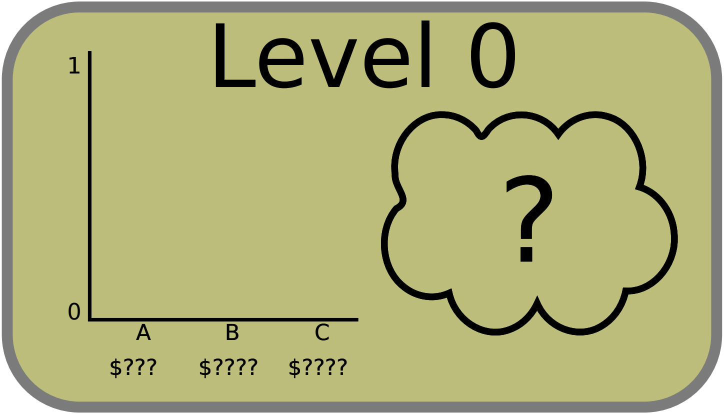

Iterative reasoning

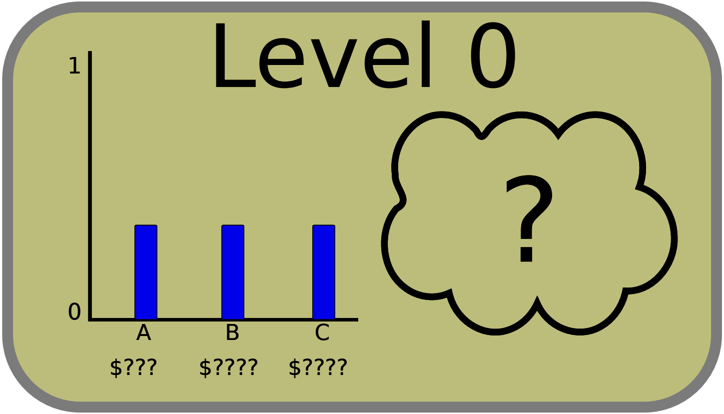

Iterative reasoning

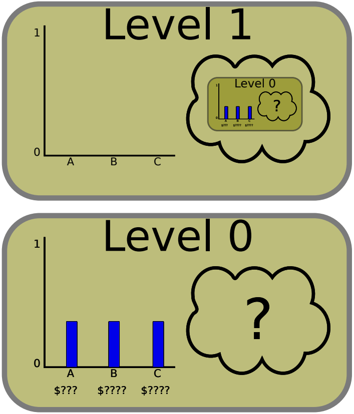

Iterative reasoning

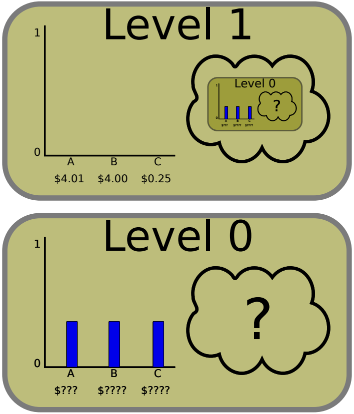

Iterative reasoning

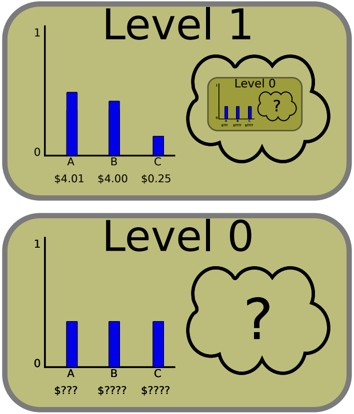

Iterative reasoning

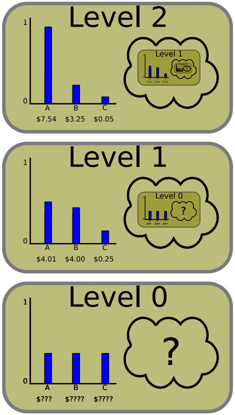

Level-0

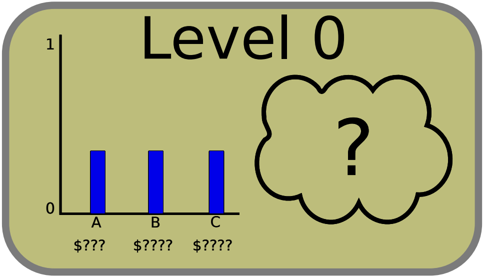

- Level-0 agents' actions influence the behavior of every other level.
- Predictions of iterative models can change dramatically if level-0 predictions change.

Level-0

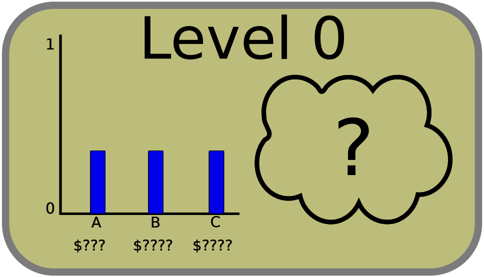

- Level-0 agents' actions influence the behavior of every other level.
- Predictions of iterative models can change dramatically if level-0 predictions change.
- It is unlikely that anyone actually picks actions uniformly.
  - Not knowing expected value is different from knowing *nothing*.
  - Level-0 agents could use all sorts of heuristics.
- Can we do a better job of predicting level-0 actions?

Level-0 meta-model

[Wright & Leyton-Brown, 2014 (submitted)]

- Define a "meta-model" that predicts a distribution of level-0 actions.

Level-0 meta-model

[Wright & Leyton-Brown, 2014 (submitted)]

- Define a "meta-model" that predicts a distribution of level-0 actions.
  - Based on features of the actions that don't require beliefs about the other agents' actions.
- Use an existing iterative model (quantal cognitive hierarchy) on top of the improved level-0 prediction to make predictions.

Features

Five binary features:

1. Minmin Unfairness
2. Maxmax payoff ("Optimistic")
3. Maxmin payoff ("Pessimistic")
4. Minimax regret
5. Efficiency (Total payoffs)

Linear model

For each action, compute weighted sum of informative features, plus a noise weight:

$$w_0 + \sum_{f \in F} w_f I(f) f(a_i)$$

Linear model

For each action, compute weighted sum of informative features, plus a noise weight:

$$w_0 + \sum_{f \in F} w_f I(f) f(a_i)$$

(An action is informative if it can distinguish at least one pair of actions.)

Linear model

For each action, compute weighted sum of informative features, plus a noise weight:

$$w_0 + \sum_{f \in F} w_f I(f) f(a_i)$$

(An action is informative if it can distinguish at least one pair of actions.)
Predict each action w.p. proportional to its weighted sum.

Example

	A	B	C
X	$100, 20$	$10, 67$	$30, 40$
Y	$40, 35$	$50, 49$	$90, 70$
Z	$41, 21$	$42, 22$	$40, 23$

Example

	A	B	C
X	$100, 20$	$10, 67$	$30, 40$
Y	$40, 35$	$50, 49$	$90, 70$
Z	$41, 21$	$42, 22$	$40, 23$

- Minimax regret is not informative (all have max-regret 60)

Example

	A	B	C
X	$100, 20$	$10, 67$	$30, 40$
Y	$40, 35$	$50, 49$	$90, 70$
Z	$41, 21$	$42, 22$	$40, 23$

- Minimax regret is not informative (all have max-regret 60)
- $50, 49$ is the fairest outcome, so Y is minmin unfairness.

Example

	A	B	C
X	$100, 20$	$10, 67$	$30, 40$
Y	$40, 35$	$50, 49$	$90, 70$
Z	$41, 21$	$42, 22$	$40, 23$

- Minimax regret is not informative (all have max-regret 60)
- $50, 49$ is the fairest outcome, so Y is minmin unfairness.
- Y and Z have min payoff 40 (vs. 10 for X)

Example

	A	B	C
X	$100, 20$	$10, 67$	$30, 40$
Y	$40, 35$	$50, 49$	$90, 70$
Z	$41, 21$	$42, 22$	$40, 23$

- Minimax regret is not informative (all have max-regret 60)
- $50, 49$ is the fairest outcome, so Y is minmin unfairness.
- Y and Z have min payoff 40 (vs. 10 for X)
- Y leads to the best total utility $(90 + 70 = 160)$

Example

	A	B	C
X	$100, 20$	$10, 67$	$30, 40$
Y	$40, 35$	$50, 49$	$90, 70$
Z	$41, 21$	$42, 22$	$40, 23$

- Minimax regret is not informative (all have max-regret 60)
- $50, 49$ is the fairest outcome, so Y is minmin unfairness.
- Y and Z have min payoff 40 (vs. 10 for X)
- Y leads to the best total utility $(90 + 70 = 160)$
- X has the highest best-case utility (100)

Example

	A	B	C
X	$100, 20$	$10, 67$	$30, 40$
Y	$40, 35$	$50, 49$	$90, 70$
Z	$41, 21$	$42, 22$	$40, 23$

- Minimax regret is not informative (all have max-regret 60)
- $50, 49$ is the fairest outcome, so Y is minmin unfairness.
- Y and Z have min payoff 40 (vs. 10 for X)
- Y leads to the best total utility ($90 + 70 = 160$)
- X has the highest best-case utility (100)

Action X's weight: $w_0 + w_{\text{maxmax}}$
Action Y's weight: $w_0 + w_{\text{minmin}} + w_{\text{total}} + w_{\text{fairness}}$
Action Z's weight: $w_0 + w_{\text{minmin}}$

Data & Parameters

Name	Source	Games	*n*
SW94	[Stahl and Wilson, 1994]	10	4005
SW95	[Stahl and Wilson, 1995]	12	576
CGCB98	[Costa-Gomes et al., 1998]	18	15662
GH01	[Goeree and Holt, 2001]	10	500
CVH03	[Cooper and Van Huyck, 2003]	8	2992
HSW01	[Haruvy et al., 2001]	15	869
HS07	[Haruvy and Stahl, 2007]	20	2940
SH08	[Stahl and Haruvy, 2008]	18	1288
COMBO8	400 samples from each	111	3200

- Set parameters (weights, level frequencies, etc.) and evaluated performance using cross validation on combined dataset:
  1. Divide data into 10 equal-sized random *folds*
  2. At step t: Choose maximum-likelihood parameters for dataset minus fold t (*training folds*) and compute likelihood of fitted model on fold t (*test folds*).
- Report sum of likelihoods of test folds.

Performance results

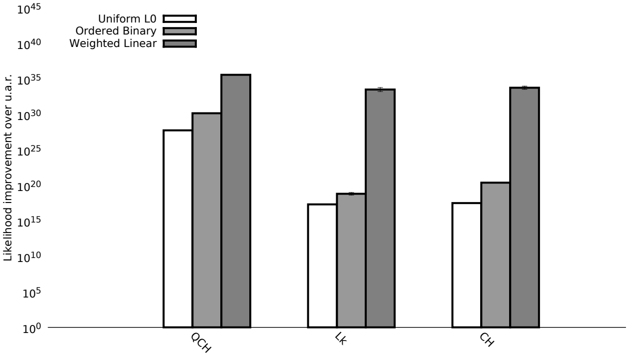

Three iterative models:

1. Quantal Cognitive Hierarchy
2. Level-k
3. Cognitive Hierarchy

Three level-0 meta-models:

1. Uniform L0
2. Ordered Binary
3. Weighted Linear

Parameter analysis

- Maximum likelihood fits do not tell us how important or identified each feature is.

Parameter analysis

- Maximum likelihood fits do not tell us how important or identified each feature is.
- The models produce probabilistic predictions.
- So we can compute a posterior distribution over parameters:

$$\Pr(\ldots, w_0, w_{\text{fairness}}, w_{\text{maxmax}}, \ldots \mid \mathcal{D})$$

- Distribution tells us how important and/or identified parameters are.

Parameter analysis: Weights

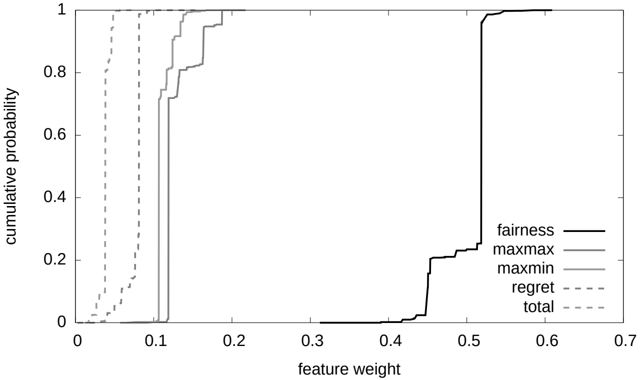

- Fairness is by far the highest weighted feature.
- All the features seem reasonably well identified.

Parameter analysis: Levels

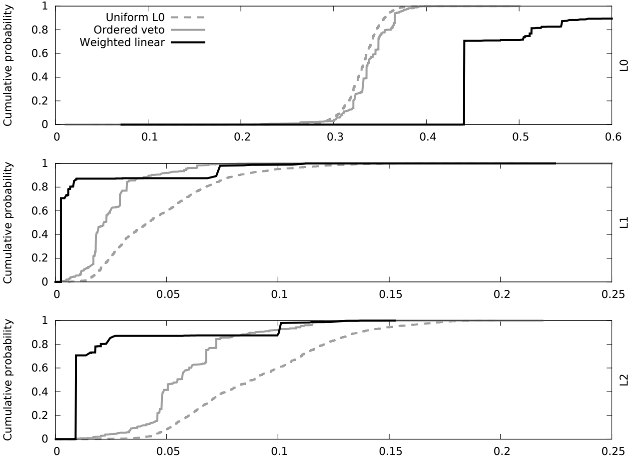

- Weighted linear $\implies$ lower variance estimates
- $\sim$Half the population is level-0!

Conclusions

- Weighted linear meta-model for level-0 agents dramatically improved the performance of all three iterative models.
  - Almost erases the difference between the models themselves.
- Strong evidence for the existence of level-0 agents.
  - For any meta-model, including uniform!
  - Contrary to conventional wisdom.

Thanks!

- Weighted linear meta-model for level-0 agents dramatically improved the performance of all three iterative models.
  - Almost erases the difference between the models themselves.
- Strong evidence for the existence of level-0 agents.
  - For any meta-model, including uniform!
  - Contrary to conventional wisdom.

Bibliography

Cooper, D. and Van Huyck, J. (2003).
Evidence on the equivalence of the strategic and extensive form representation of games.
JET, 110(2):290–308.

Costa-Gomes, M., Crawford, V., and Broseta, B. (1998).
Cognition and behavior in normal-form games: an experimental study.
Discussion paper 98-22, UCSD.

Goeree, J. K. and Holt, C. A. (2001).
Ten little treasures of game theory and ten intuitive contradictions.
AER, 91(5):1402–1422.

Haruvy, E. and Stahl, D. (2007).
Equilibrium selection and bounded rationality in symmetric normal-form games.
JEBO, 62(1):98–119.

Haruvy, E., Stahl, D., and Wilson, P. (2001).
Modeling and testing for heterogeneity in observed strategic behavior.
Review of Economics and Statistics, 83(1):146–157.

Rogers, B. W., Palfrey, T. R., and Camerer, C. F. (2009).
Heterogeneous quantal response equilibrium and cognitive hierarchies.
JET, 144(4):1440–1467.

Stahl, D. and Haruvy, E. (2008).
Level-*n* bounded rationality and dominated strategies in normal-form games.
JEBO, 66(2):226–232.

Stahl, D. and Wilson, P. (1994).
Experimental evidence on players' models of other players.
JEBO, 25(3):309–327.

Stahl, D. and Wilson, P. (1995).

On players' models of other players: Theory and experimental evidence.
GEB, 10(1):218–254.