We can **infer** how **generous** someone is toward us. This may have **evolved** because we don't know **exactly** how much it **costs** them to **help** or **hinder** us. (i.e., there is **NOISE**!)

Evolutionary distribution **without noise**:

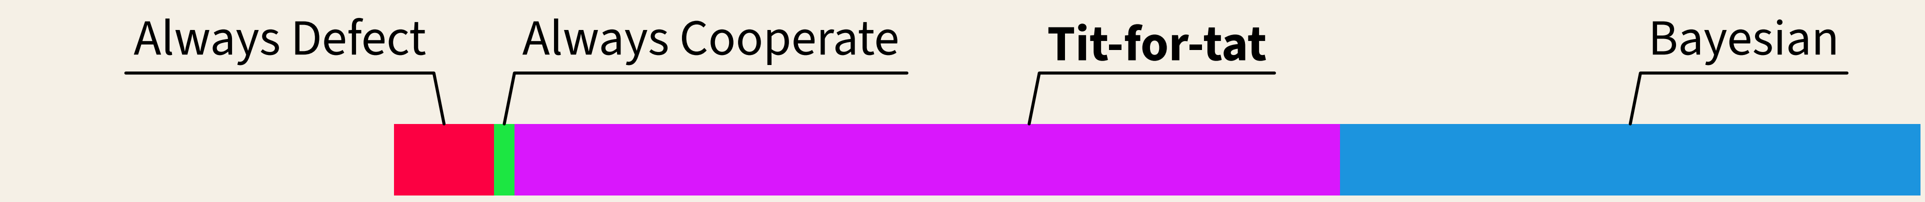

Evolutionary distribution **with noise**:

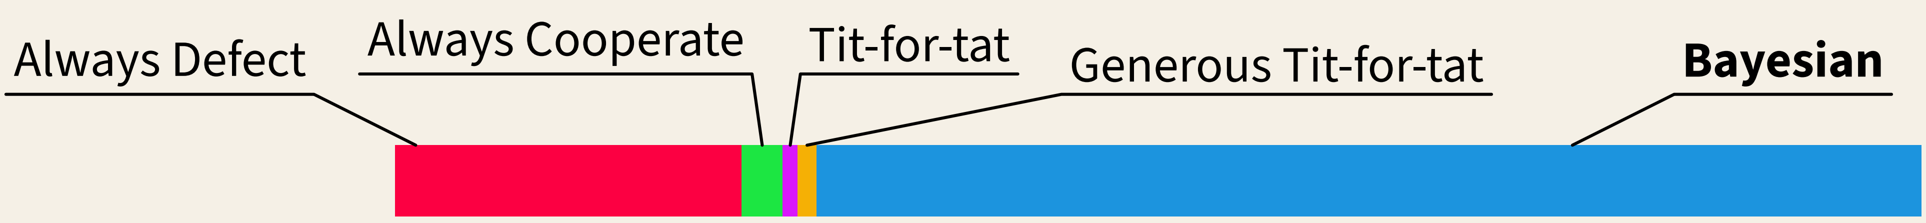

The evolution of reciprocity based on welfare tradeoff ratios in games with asymmetric information

Wenhao Qi, Boyu Wang, Lindsey J. Powell
Department of Psychology, UC San Diego
Presenter: Wenhao (James) Qi

Background

- People are motivated to help or harm another person to varying degrees
- **Welfare tradeoff ratios (WTRs)** quantify the direction and magnitude of such motivations
- **Definition of WTR**: $u_A = w_A + \lambda_{AB} w_B$
  w_A / w_B: Person A/B's (objective) welfare
  λ_{AB}: A's **welfare tradeoff ratio** toward B
  u_A: A's utility, which she tries to maximize
- People can infer another person's **WTR** toward themselves; **How did such abilities evolve?**

Methods

- **Tournament** of different agents playing **repeated alternating games** with each other
- Each game is a **one-player binary** allocation decision; e.g., Option 1 gives \$2 to A and \$8 to B; Option 2 gives \$5 to A and \$2 to B
- The **payoffs change** from round to round
- B's decision tells A whether λ_{BA} is **above or below a threshold**, determined by the payoffs
- **Evolutionary simulation** based on the resulting **pairwise mean payoff matrix**
- Specifically, we look at the **long-run distribution** of a Moran process
- In Experiment 1, players perceive the payoffs **without noise**
- In Experiment 2, **noise** is added to the two players' perceptions **independently**, and neither players knows exactly what the other player sees

Agents (assuming A is self and B is opponent)

- **Always Defect/Cooperate**: λ_{AB} fixed to 0/1
- **Tit-for-tat** (a heuristic strategy): Assumes λ_{BA} = 0 or 1 in each round; Starts with λ_{AB} := 1; When B's decision in the last round distinguishes between λ_{BA} = 0 and 1, sets λ_{AB} := λ_{BA}
- **Generous Tit-for-tat**: Like Tit-for-tat, but unconditionally cooperates with some probability
- **Bayesian**: Does Bayesian inference on λ_{BA} with Hidden Markov Model; Sets λ_{AB} to the median of posterior of λ_{BA}, with slight bias toward 1

Results (see figure ☝)

- **Without noise in payoff perception**: Tit-for-tat performs well; Bayesian can do no better
- **With noise in payoff perception**: Tit-for-tat suffers from misperception and can't cooperate well with itself; Generous Tit-for-tat cooperates better with itself but is less resistant to invasion by Always Defect; Bayesian is robust and successful

Repeated one-player games with variable payoffs:

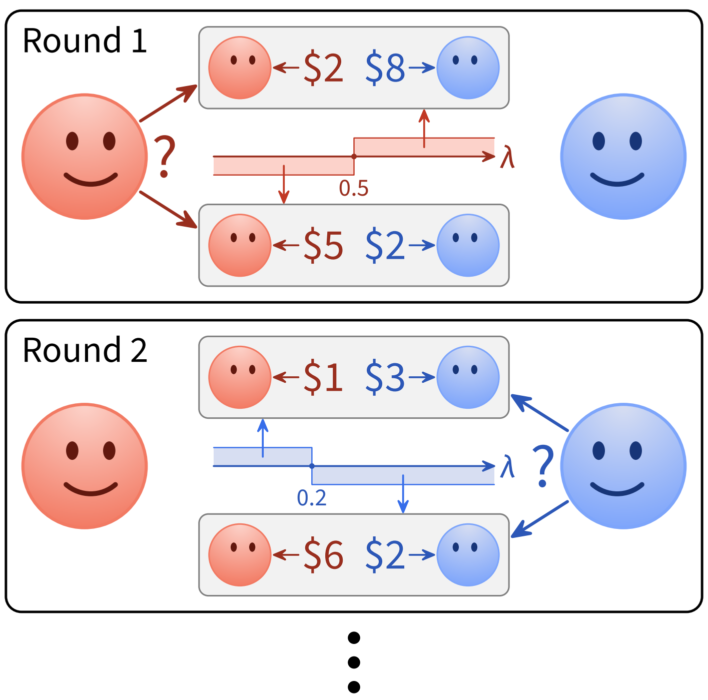

When there is noise in payoff perception: