Computationally rational decision makers

Behavioral economists point out that decision makers do not seem to act like the “rational men” assumed by classical decision theory and game theory:

- they cooperate in prisoner’s dilemma
- they give money in the dictator game
- they exhibit confirmation bias and first-impression bias

They attempt to explain these deviations from rationality using insights from psychology.

- The explanations seems somewhat *ad hoc*
- Different approaches are used to explain different behaviors

Our goal:

> *To construct an overarching theory, by taking computational limitations and communication constraints seriously*

A motivating example

Wilson [2002/2014] considers a decision problem where a decision maker (DM) needs to make a single decision.

- Nature is in one of two states: 0, 1
- The DM wants to "match" nature's state
- Nature's state is static: it doesn't change
- The DM gets one of k independent signals, which are correlated with nature's state, at each time step
  - For each signal σ, the agent knows the probability of seeing σ conditional on the true state of nature being $0/1$.
- The game ends at each step with some small constant probability. At that point the agent must make a decision.

Computational limitations

How can we model computationally bounded agents?

- There isn't a "right" answer here, but one standard choice is to think of agents as *probabilistic* finite automata.
  - Just like a deterministic automaton, but each transition has a probability.
  - E.g., in state s_1, if I see a 0, I go to s_2 with probability $1/3$ and stay in s_1 with probability $2/3$.

Computational limitations

How can we model computationally bounded agents?

- There isn't a "right" answer here, but one standard choice is to think of agents as *probabilistic* finite automata.
  - Just like a deterministic automaton, but each transition has a probability.
  - E.g., in state s_1, if I see a 0, I go to s_2 with probability $1/3$ and stay in s_1 with probability $2/3$.

- Suppose we have a prior probality on 0 or 1.
- Every time we see a signal, we can use Bayes' Rule to update the probability.
  - You should know how to do this!
- But an automaton can't implement Bayes' rule.
  - It can't keep track of all the possible posterior probabilities.

Computational limitations

How can we model computationally bounded agents?

- There isn't a "right" answer here, but one standard choice is to think of agents as *probabilistic* finite automata.
  - Just like a deterministic automaton, but each transition has a probability.
  - E.g., in state s_1, if I see a 0, I go to s_2 with probability $1/3$ and stay in s_1 with probability $2/3$.

- Suppose we have a prior probality on 0 or 1.
- Every time we see a signal, we can use Bayes' Rule to update the probability.
  - You should know how to do this!
- But an automaton can't implement Bayes' rule.
  - It can't keep track of all the possible posterior probabilities.

So what is an optimal automaton for Wilson's decision problem?

The optimal automaton

Wilson proves that the optimal automaton has the following structure:

- The DM ignores all but the strongest signals for 0 and 1
- The states can be laid out "linearly": $-n, \ldots, 0, \ldots, n$
  - Intuitively, state 0 represents "indifference"
  - more positive/negative means more likely to be 1/0
- The automaton moves right/left if it gets a strong signal for 1/0

The optimal automaton

Wilson proves that the optimal automaton has the following structure:

- The DM ignores all but the strongest signals for 0 and 1
- The states can be laid out "linearly": $-n, \ldots, 0, \ldots, n$
  - Intuitively, state 0 represents "indifference"
  - more positive/negative means more likely to be 1/0
- The automaton moves right/left if it gets a strong signal for 1/0
- Key point: The probability of moving left/right decreases the further out to the right/left the agent is.
  - "Don't bother me; I've made up my mind!"

The punch line

The optimal automaton with $2n+1$ states has this structure:

- independent of n;
- transition probability depends on n and signal strength.

The optimal automaton exhibits “human-like” behavior:

- It ignore evidence
- It exhibits confirmation bias
- The order that evidence is received matters!
  - First-impression bias
- Belief polarization:
  - Two people that initially have have only slightly different beliefs can end up with very different beliefs

Supposedly irrational behavior may be quite rational!

The punch line

The optimal automaton with $2n+1$ states has this structure:

- independent of n;
- transition probability depends on n and signal strength.

The optimal automaton exhibits “human-like” behavior:

- It ignore evidence
- It exhibits confirmation bias
- The order that evidence is received matters!
    - First-impression bias
- Belief polarization:
    - Two people that initially have have only slightly different beliefs can end up with very different beliefs

Supposedly irrational behavior may be quite rational!

What I've Been Doing

I've been trying to apply this paradigm to more settings:

- What happens if nature's state changes
    - Think of the stock market!

What I've Been Doing

I've been trying to apply this paradigm to more settings:

- What happens if nature's state changes
  - Think of the stock market!
- Multi-armed bandits
- A ranger-poacher problem
- ...