Repeated Interactions

CMPUT 654: Modelling Human Strategic Behaviour

Camerer & Ho (1999)
Chen, Liu, Chen, and Lee (2011)

Lecture Outline

1. Camerer & Ho (1999)

2. Chen, Liu, Chen, and Lee (2011)

Camerer & Ho (1999)

Why:
Extremely influential model of repeated interactions

1. Explain the motivation: Combining **choice reinforcement** with **belief-based** learning
2. Describe the model (including how it **generalizes** choice reinforcement and belief-based learning)
3. Some empirical results (parameter fits, performance)

Belief-Based Learning

Paradigmatic version is **fictitious play** [Robinson 1951]:

1. Estimate strategy of opponent as $\hat{s}_{-i}(a_{-i}) = \dfrac{w(a_{-i})}{\sum_{a'_{-i} \in A_{-i}} w(a'_{-i})}$
2. Best respond to $\hat{s}_{-i}$

- Takes no account of payoffs of other agent
- Requires access to all of own counterfactual payoffs

Choice-Reinforcement Learning

- Each action has an associated **reinforcement** $R(a_i,t)$
- Reinforcements on chosen actions update based on **realized payoffs** (where $0 \leq \phi \leq 1$):

$$R(a_i, t) = \begin{cases} \phi \cdot R(a_i, t-1) + u_i(a_i, a_{-i}(t)) & \text{if } a_i(t) = a_i, \\ \phi \cdot R(a_i, t-1) & \text{otherwise.} \end{cases}$$

Experience-Weighted Attraction

- Action probability is monotonic in **attractions** $A(a_i,t)$:

$$s_i(a_i, t) = \frac{\exp[\lambda A(a_i, t)]}{\sum_{a_i' \in A_i} \exp[\lambda A(a_i', t)]}$$

- Attractions updated according to

$$A(a_i, t) = \frac{\phi \cdot N(t-1) \cdot A(a_i, t-1) + [\delta + (1-\delta) \cdot I[a_i = a_i(t)]] \cdot u_i(a_i, a_{-i}(t)}{N(t)}$$

where

$$N(t) = \rho \cdot N(t-1) + 1$$

Differences from Belief-Based Learning

- Initial attractions can be **arbitrary**
- Attractions can grow outside bounds of payoffs

Empirical Results

- Actually do an out-of-sample check!
- EWA and belief-based perform best on this data
- Stylized behaviour in beauty contests:

    1. Dispersed initial play
    2. Rapid convergence to equilibrium

Game Model	No. of Parameters	Calibration			Validation		
		LL	AIC	BIC	ρ^2	LL	MSD
Median Action (M = 378)							
1-Segment							
Random Choice	0	−677.29	−677.29	−677.29	0.0000	−315.24	0.1217
Choice Reinforcement	8	−341.70	−349.70	−365.44	0.4837	−80.27	0.0301
Belief-based	9	−438.74	−447.74	−465.45	0.3389	−113.90	0.0519
EWA	11	−309.30	−320.30	*−341.94**	0.5271	−41.05	0.0185
2-Segment							
Random	0	−677.29	−677.29	−677.29	0.0000	−315.24	0.1217
Choice Reinforcement	17	−331.25	−348.25	−381.70	0.4858	−66.32	0.0245
Belief-based	19	−379.24	−398.24	−435.62	0.4120	−70.31	0.0250
EWA	23	−290.25	*−313.25**	−358.51	0.5375	*−34.79**	*0.0139**
p-beauty contests (M = 1372)							
1-Segment							
Random	0	−6318.29	−6318.29	−6318.29	0.0000	−2707.84	0.0099
Choice Reinforcement	12	−5910.99	−5922.99	−5954.33	0.0626	−2594.37	0.0101
Belief-based	13	−6083.04	−6096.04	−6129.99	0.0352	−2554.21	0.0097
EWA	15	−5878.20	−5893.20	−5932.38	0.0673	−2381.28	0.0098
2-Segment							
Random	0	−6318.29	−6318.29	−6318.29	0.0000	−2707.84	0.0099
Choice Reinforcement	25	−5910.98	−5935.98	−6001.28	0.0605	−2594.17	0.0101
Belief-based	27	−6083.02	−6110.02	−6180.54	0.0330	−2554.11	*0.0097**
EWA	31	−5771.46	*−5802.46**	*−5883.43**	0.0816	*−2355.00**	0.0098

Chen, Liu, Chen, and Lee (2011)

Why:
Recent, high-performing behavioural model

1. Define the model

2. Empirical results

Market Entry Games

- Binary choice whether to enter a risky market
- Entry payoff is

$$V(t) = 10 - k \times E + G_t$$

  where G_t is randomly $L < 0$ or $H > 0$, with $\mathbb{E}[G_t] = 0$

- Not entering yields G_t / s or $-G_t / s$ with equal probability

Model: I-SAW

- Agents are in one of three **modes**: `explore`, `inertia`, `exploit`
    - `explore` mode: choose from some fixed distribution
    - `inertia` mode: choose the same action as last time
    - `exploit` mode: choose action with highest expected **subjective value**:

$$\text{ESV}(a_i) = (1 - w)(\text{SampleM}(a_i, t)) + w(\text{GrandM}(a_i, t))$$

- Fixed probability ε of entering `explore` mode; else, enter `inertia` mode with high probability when **surprise** is low

BI-SAW

- Exactly the same as I-SAW, except that SampleM is defined differently (only draws from *b* most recent games)
-

Simulation & Estimation

Simulation:

- Parameters chosen for each individual
- Chosen uniformly from a range
    - Lower bound fixed, upper bound learned
- 5000 trajectories sampled from the game, error is averaged

Estimation:

- **Grid search** on upper bounds!