Models of Learning in Economics - Nomis Summer School 2022

W Bentley MacLeod
Columbia University and NBER

May 31, 2022

Outline

Introduction

- The purpose these lectures is to discuss two important class of learning models and how they can be used in economics.
- I start with the case of chess, a well understood "hard" problem, to make the point that even for this well defined problem there is not a single, best algorithm.
- Rather, decision making and learning relies upon a variety of algorithms.
- An open question remains on how best to model the heterogeneity in decision processes observed in human populations?

Chess

- We begin with a classic problem in the theory of learning and artificial intelligence, namely how to play chess.
- Chase and Simon (1973a) report the results from an elegant experiment showing that skilled players are better able to recognized and recall legal chess positions.
- However, they are no better than an unskilled player with a random allocation of pieces to the board.
- We use this example to illustrate three distinct approaches to decision making and learning.

Approaches to Decision Making

- Game theory and dynamic programming
- Bayesian decision making
- Rule based decision making

The Game of Chess

- White and black alternative moves, with with while playing at odd dates, $t = 1, 3, 5, ...$ and black playing at even dates, $t = 2, 4, 6, ...$
- Let d_t denote a move at date t and let $\vec{d}_t = \{d_1, d_2, d_3 ... d_t\}$ be a sequence of moves up to date t, where $\vec{d}_0$ represents the start of the game..
- Let D_t be the set consisting of legal moves of length t, and $D\left(\vec{d}_t\right)$ the legal moves at date $t + 1$.
- The sets D_t define the set of all legal moves of length t in chess, while all moves are given by: $D = \cup_{t \geq 1} D_t$.
- The board position after the move at date t is denoted by $x_t = X\left(\vec{d}_t\right)$, where by definition x_0 is the starting position, with white moving at t =1.
- Recently, the governing bodies of chess brought in mandatory 5 repetition and 75 move limits that result in automatic stalemates. Hence, game of chess is a finite game. However, D is a very large set - greater than the number of atoms in

Chess is a Zero Sum Game

- The set of endpoints of the game is given by the set, at which point the outcome is stalemate, white wins or black wins:

$$Z = \left\{ \vec{d}_t | D\left(\vec{d}_t\right) = \emptyset \right\}.$$

- For each $\vec{d} \in Z$ we can define:

$$V\left(\vec{d}\right) = \begin{cases} 1, & \textit{checkmate by white} \\ 0, & \textit{stalemate} \\ -1 & \textit{checkmate by black} \end{cases}$$

The Game Theory Solution

- Game theory, as used in economics, views chess as a "simple" game because it is a finite, zero-sum game of perfect information, and hence has be solved by dynamic programming
  - At date $t+1$ suppose that $V\left(\vec{d}_{t+1}\right) \in \{0, 0.5, 1\}$ (black wins, stalemate, white wins) is defined for each $\vec{d}_{t+1} \in D_{t+1}$.
  - Let:

$$V\left(\vec{d}_t\right) = \begin{cases} \max_{d_{t+1} \in D\left(\vec{d}_t\right)} V\left(\left\{\vec{d}_t, d_{t+1}\right\}\right), & \textit{white moves} - t = 0,2,4 \\ \min_{d_{t+1} \in D\left(\vec{d}_t\right)} V\left(\left\{\vec{d}_t, d_{t+1}\right\}\right), & \textit{black moves} - t = 2,4,6 \end{cases}$$

- The value function is well defined for $\vec{d}_t \in Z$, and thus one can readily verify that it is well defined for moves due to the fact that all plays are of finite play.

Theorem

The value function for chess satisfies $V\left(\vec{d}_t\right) \in \{0, 0.5, 1\}$ for all $\vec{d}_t \in D$.

- Hence, chess is a solved game, and the optimal move, d_t^* at date t for $\vec{d}_t \in D_t \backslash Z$ is given by:

$$d_{t+1}^* = d^*\left(\vec{d}_t\right) = \begin{cases} \arg\max_{d_{t+1} \in D\left(\vec{d}_t\right)} V\left(\left\{\vec{d}_t, d_{t+1}\right\}\right), & t = 1, 3, 5, \ldots \\ \arg\min_{d_{t+1} \in D\left(\vec{d}_t\right)} V\left(\left\{\vec{d}_t, d_{t+1}\right\}\right), & t = 2, 4, 6, \ldots \end{cases}$$

Observations

- Zermelo's Lemma: Suppose at date t (odd) white has a winning position ($V\left(\vec{d}_t\right) = 1$), then white can force a win in a finite number of moves.
- Even though the game is finite, it is not known if $V\left(\vec{d}_0\right)$ is 0, 0.5 or 1!
- This approach suggests two solutions to actual play:
  - If one can assign a value to a future board position, then this can direct choice.
  - If one can assign a stored optimal choice to a current position, this can also direct choice.
- Chess playing programs use a combination of strategies (Maharaj et al. (2022)), though theory models tend to focus on one approach only.

Outline

Savage's Approach to Decision Making

1. Build a model of the uncertainty relevant to the problem at hand.
2. Model the relationship between choices, uncertainty and the outcomes one cars about.
3. Choose a course of action.

Theorem

Under the appropriate conditions a person acts as if they maximize expected utility, where the probability of an event is given by their beliefs.

Savage's Approach Applied to Chess

- Suppose that you are white, the date is t, and the current board position is given by the sequence of moves $\vec{d}_t$.
- Let us suppose that you have an evaluation function, $V\left(\vec{d}_t\right) \in [0,1]$ that is the probability that white will win given the moves $\vec{d}_t$.
- In general, different opponents will have different styles of play. Hence, we can define your payoff from different moves, d_{t+1} given beliefs over your opponent:

$$Q\left(\vec{d}_{t+1}\right) = u\left(\vec{d}_{t+1}\right) + \sum_{d_{t+2} \in D\left(\vec{d}_{t+1}\right)} Pr\left[d_{t+2} | \vec{d}_{t+1}\right] V\left(\vec{d}_{t+1}, d_{t+2}\right)$$

  where $u\left(\vec{d}_{t+1}\right) = 1$ if white checkmates and 0.5 if stalemate.
- See Maharaj et al. (2022) for a discussion of Q learning and how it is applied to chess programs.

Learning about ability

- Consider an infinitely lived agent, with time indexed by $t = 0, 1, 2, ...$
- Each period, the agent chooses effort $a_t \geq 0$ and produces output:

$$y_t = \eta + a_t + \varepsilon_t,$$

  where η denotes the *ability* of the agent, with prior distribution $N\left(m_\eta, \sigma_\eta^2\right)$, while ε_t is an i.i.d. sequence of error terms with distribution $N\left(0, \sigma_\varepsilon^2\right)$.
- In the subsequent development, it is convenient to use *precision* as a parameter. It is defined by the reciprocal of the variance: $\rho_\eta = 1/\sigma_\eta^2$ and $\rho_\varepsilon = 1/\sigma_\varepsilon^2$.

Timing

1. At the beginning of period t, worker is paid $w_t = E\left[y_t | y_0, \ldots, y_{-1}\right]$ - expected productivity as a function of information based upon previous periods' output.
2. The agent selects effort a_t.

Solving for Effort and Wage

- It is assumed that worker chooses effort to maximize expected income:
$$U = E\left\{\sum_{t=0}^{\infty}\left(w_t - g\left(a_t\right)\right)\delta^t\right\},$$
where δ is the discount rate per period, and $g\left(\cdot\right)$ is the dis-utility of effort.
- At the end of the period, the market observes a signal of agent ability given by:
$$x_t = y_t - \hat{a}_t = \eta + \varepsilon_t,$$
where the market correctly anticipates effort $\hat{a}_t$
- The wage offered to the agent will be:
$$w_t = m_t + \hat{a}_t,$$
where $m_t = E\left\{\eta | x_0, x_1, ..., x_{t-1}\right\}$ is the expected ability of the agent in period t.

Bayes Rule 1

- $$E\{\eta|x_0, x_1, ..., x_{t-1}\} = m_t = \frac{\rho_\eta m_\eta + \rho_t \overline{x}_t}{\rho_\eta + \rho_t}, \tag{1}$$
  $$\frac{1}{var(m_t)} = \rho_t = t\rho_\varepsilon, \tag{2}$$

  where $\overline{x}_t = \frac{1}{t}\sum_{s=0}^{t-1} x_s$.

Implications of Strong Beliefs under Bayes Rule

- With addition observations, precision increases and one learns true ability $\lim_{t\to\infty} m_t = \eta$.
- The effect of data has a smaller effect, the more "stongly held" are one's views, as measured by ρ_η:

$$\frac{\partial m_t}{\partial \bar{x}_t} = \frac{\rho_t}{\rho_\eta + \rho_t}.$$

- For example, if one believes that vaccines are not effective, then regardless of the evidence, one will not get a shot.
- More importantly, Bayes rule implies that the same data will have different impact upon different individual's views.

Implications of Weak Beliefs under Bayes Rule

- Notice that we have assumed that $x_t = y_t - \hat{a}_t$, where $\hat{a}_t$ is the expect effort.
- We can use this to compute effort in period $t = 0$.
- Consider a worker who chooses $a_0 > \hat{a}_0$, say $a_\Delta = a_0 - \hat{a}_0$. In this case the change in discounted income is:

$$\Delta W = W(a) - W(\hat{a}_0) = \sum_{t=0}^{\infty} \frac{\rho_t a_\Delta / t}{\rho_\eta + \rho_t} \delta^t$$

$$= \Delta \sum_{t=0}^{\infty} \frac{1}{\rho_\eta / \rho_\varepsilon + t} \delta^t$$

- When the employer has strong beliefs ($\rho_\eta \to \infty$), then $\Delta W \to 0$ - in other words when firm has strong beliefs, then there are *low* incentives to work hard.
- However, when $\delta \to 1$, and firm as weak beliefs, ρ_η is small, then $\Delta W \to \infty$. In other words when a firm has little information about a person, they have high incentives to work hard.

Empirical Applications

- Gibbons and Murphy (1992)
- Farber and Gibbons (1996)
- Altonji and Pierret (2001)
- MacLeod and Urquiola (2015); MacLeod et al. (2017)

The Problem of Drug Choice

- Drug choice is not simply one of choosing the best one conditional upon patient X_i.
  - Physicians must try a drug and then be prepared to change if it does not work - technically this is a correlated multi-armed bandit problem (CMAB).
  - How should they choose the sequence of trials?
  - How can we characterize variation in observed practice styles?

Physician Choice: A 2 Drug 2 Period Example

- Two drugs, $d \in \{0, A, B\}$ with *unknown idiosyncratic* effect $e_d \sim N(\mu_d, 1/\rho_d)$ - $d = 0$ is no drug.
  - Observation error (physician skill ρ): $\varepsilon_t \sim N(0, 1/\rho)$
  - Timing:
    - Period 0: Physician observes patient condition $y_0 = e_{d0} + \varepsilon_0$ and choose drug d_1.
    - Period 1: Physician observes patient condition $y_1 = e_{d1} + \varepsilon_1$ and choose drug d_2.
    - Period 2: Patient condition is $y_2 = e_{d2} + \varepsilon_2$

Physician Decision Rule and Preferences

- Decision rule is given by $\delta = \{d_1, d_2 = \delta_1(y_0, d_1)\}$.
  - Physician cares about outcome today and tomorrow:

$$U(\delta, \zeta) = E\{(1-\zeta)\, y_1 + \zeta y_2 | \delta\}$$

  - When $\zeta = 1$ physician cares only about long term condition, while $\zeta = 0$ implies that physician cares only about short term condition.

Experimentation vs Exploitation.

- Suppose that $\mu_A > \mu_B > 0$ (*ex ante* A is the better treatment) but $\rho_A > \rho_B$ (B has more variance in effects across individuals)
  - If $\zeta = 0$ then optimal choice is $d_1 = A$.
  - If $\zeta = 1$ then optimal choice is $d_1 = B$ - learn about B and then use best drug in period 2.

More Precisely

Proposition

The value of learning about drug $d \in \{A, B\}$ is:

$$V_d = s_d L\left(\frac{\mu_A - \mu_B}{s_d}\right),$$
$$s_d = \sqrt{\frac{\rho}{\rho_d(\rho + \rho_d)}},$$
$$L(x) = (1 - F(x))(f(x) - x), x \geq 0,$$

Where $L(x)$ is unit-normal loss function, and F,f are normal cdf and pdf. Notice $L(0) = f(0)/2$ and $lim_{x \to \infty} L(x) = 0$.

Implications

- From the value of information we can derive the physician's optimal choice. Choose $d = A$ iff:

$$(1-\zeta)(\mu_A - \mu_B) + \zeta(V_A - V_B) \geq 0$$

  - Notice that the value of information:
    - Increases with physician skill
    - Falls with an increase in precision ρ_d
    - Increases when drugs are more similar (a feature of anti-depression drugs is the magnitudes of their effects are quite similar!)
  - Big take away - there are gains from experimenting with choices known to be sub-optimal in expected value.

Outline

The Role of Time

- From cognitive science the well known (in psychology) Hicks-Hyman law and Fitts Laws relate the time required to make a decision the *quality* of a decision.
- The early literature characterizes these decisions using fast thinking (look-up tables - Churchland and Sejnowski (1993), chapter 3) or slow thinking (search Nilsson (1980), MacLeod (2002)) - the dichotomy is popularized in Kahneman (2011)'s *Thinking Fast and Slow (2011).*
- However, the two way classification is a major simplification, and maybe too simple for economics?

The View from Cognitive Science - Newell 1987

State (secs)	Time Units (secs) - t_a	System	World (theory)	Econ
10^7	Months (10^7)			
10^6	weeks (10^7)		Social Band	Institutions
10^5	days (10^7)			
10^4	hours (10^7)	task		Work
10^3	10 min (10^7)	task	Rational Band	and
10^2	minutes (10^7)	task		Consumption
10^1	10 sec (10^7)	unit task		
10^0	1 sec (10^7)	operations	Cognitive Band	Experience
10^{-1}	100 msec (10^7)	deliberate act		and
10^{-2}	10 msec (10^7)	neural circuit		Human Capital
10^{-3}	1msec (10^7)	neuron	Biological Bank	
10^{-4}	100 μs (10^7)	Organelle		

Very Fast, Fast, Slow and Very Slow

- Very fast (millisecond-seconds)- human capital/skill - surgery, trading, flying
- Fast (seconds) - biased decision making/behavioral economics
- Slow (seconds to hours) - Rational choice - pursuit of well defined goals using Savage's model and scientific knowledge
- Very slow (days to years)- RCTs (randomized control trials) and doing science to build better decision making models.

Why is the distinction useful?

- The economic theory of the 1950s can be viewed as developing a *unitary* model based upon extensions of the Arrow-Debreu model.
- Modern economics is characterized by plethora of seemingly conflicting models:
  - Dynamic extensions of general equilibrium theory
  - Structural IO models
  - Behavioral models describing a number of behaviors that are inconsistent with the simple utility maximization model.
- Thinking in terms of the time scale can assist in delineating the *scope* of a model - this is common practice in physics where they *explicitly* use both time and distance scales for determining the scope of a model (e.g. cosmology vs quantum mechanics).

Pattern Recognition in Chess

- In a seminal study, Chase and Simon (1973b) proved that skill at chess is in part memory and simply knowing what is a good move.
- Chess skill is a form of *human capital* in which individuals over years learn more strong positions, and how to make choices in strong position.
- MacLeod (2016) observes that one can associate the skill that is associated with fast think with *human capital.*
- In economics, human capital is the financial cost of acquiring a certain level of skill, and hence can be viewed as explicit measure of *bounded rationality.*
- If individuals had perfect decision making skills then there would be no need for education!

Outline

Learning from Experience (MacLeod (2016))

Consider a situation where each period, $t = 0, 1, 2, ...$ an individual faces the following sequence of events:

1. An event $\omega_t \in \Omega$ is observed, Ω has $N < \infty$ elements.
2. The agent makes a choice $d_t \in D$, using decision rule $\sigma_t(\omega_t)$.
3. The agent evaluates the outcome using criteria $U(d_t|\omega_t)$ and updates the decision rule to $\sigma_{t+1}(.)$ to be used the next period.

The Decision Rules

- Let $D = \{0, 1\}$ and let the decsion rule $\sigma(\omega) \in [0, 1]$ be the probability of choosing $d = 1$.
- The optimal choice is given by $\sigma^*(\omega) = \arg\max_{d \in D} U(d|\omega)$.
- Suppose that at date $t = 0$ the agent has no experience and we set $\Omega_0 = \emptyset$.
- Each period t the agent experiences ω_t and we set $\Omega_t = \Omega_{t-1} \cup \{\omega_t\}$.

Learning by Doing

- Suppose that $\max_{d \in D} U(d|\omega) = u_g > u_b = min_{d \in D} U(d|\omega)$.
- The basic learning by doing begins by guessing, and then following the optimal strategy for experienced events:

$$\sigma_t(\omega_t) = \begin{cases} \sigma^*(\omega_t), & \text{if } \omega_t \in \Omega_t. \\ 1/2 & \text{if } \omega_t \notin \Omega_t. \end{cases}$$

- The idea is simple - after experiencing event ω_t, it becomes salient, and the individual figures out the optimal choice for that event.
- Note: $E(U(\sigma_0(\cdot)|\omega)) = \frac{u_g + u_b}{2}$ and $lim_{t \to \infty} E(U(\sigma_t(\cdot)|\omega)) = u_g$

Outline

Extending to Bayesian Planning

- MacLeod (2016) illustrates an extension of the simply learning by doing to Bayesian learning (see also Jovanovic and Nyarko (1995) for a low dimensional version).
- Suppose that at date t the agent can explore n_t new states at a cost $c(n_t)$.
- The gain from exploring a state is that when it occurs the decision maker gets u_g rather than $\frac{u_g+u_b}{2}$.
- The states she should explore depends upon the probability that they will occur - suppose there are N states, then beliefs are given by $\mu \in \Delta^N$.

Learning

- We can learn about μ if we suppose that μ comes from a distribution over Δ^N give by the Dirichlet distribution:

$$f(\mu|\alpha) = \frac{\Gamma(\alpha_1 + \ldots + \alpha_N)}{\Gamma(\alpha_1)\ldots\Gamma(\alpha_N)} \mu_1^{\alpha_1 - 1} \ldots \mu_N^{\alpha_N - 1}.$$

- If we let $\alpha_i = b$ for all i, then the prior is uniform over states.
- Let $\vec{\alpha}^t$ be parameter at date t, then $Pr[\omega_i] = \frac{\alpha_i^t}{\sum \alpha_j^t}$.
- The cool part is that if we observe state i occurs at date t, then the posterior distribution of μ is Dirichlet with $\alpha_i^{t+1} = \alpha_i^t + 1$ and the other values remain unchanged.
- Notice that with larger b, the effect of learning is smaller - thus b is level of *dogmatism*.

Results

- The optimal number of states explored in period t, n_t, falls with t.
- The optimal number of states explored rises with b.
- If b is sufficiently small, then there is no planning, and person engages in pure learning by doing.

Outline

Extending to Continuous States

- The model can be easily extended to $\Omega \subset \Re^n$, where Ω is compact.
- Given a history of events and corresponding optimal responses at time t, $H_t = \left\{ \left(\omega^1, z^1\right), \left(\omega^2, z^2\right), \ldots, \left(\omega^{t-1}, z^{t-1}\right) \right\}$.
- For a new event ω^t, suppose the individual uses the response for the previously experience event that minimizes the Euclidean distance:

$$\sigma\left(\omega^t | H^t\right) \in \left\{ \begin{array}{c} z^\tau = \sigma^*\left(\omega^\tau\right) | \left\|\omega^\tau - \omega\right\| \leq \left\|\omega^{\tau'} - \omega\right\|, \\ \forall \tau, \tau' \in \{1, \ldots, t-1\}, \left(\omega^\tau, z^\tau\right) \in H^t \end{array} \right\}. \tag{3}$$

Learning Result

Theorem

Suppose that Ω is a compact subset of $\Re^d$, the optimal behavior, $\sigma^*(\omega)$, is Borel measurable, and μ is absolutely continuous with respect to Lebesgue measure. Then

$$\lim_{t\to\infty} E\left\{\left\|\sigma^*\left(\omega^t\right) - \sigma\left(\omega^t | H^t\right)\right\|\right\} = 0, \tag{4}$$

where
$H^t = \{(\omega_1, \sigma(\omega_1, T)), (\omega_2, \sigma(\omega_2, T)), \ldots, (\omega_{t-1}, \sigma(\omega_{t-1}, T))\}$.

Discussion

- This decision rule is a version of the nearest neighbor rule from machine learning - the result shows that it is a *universally consistent* learning rule.
- The main requirement is that there is a stable relationship between events and the optimal choice.
- The history H^t can be viewed as a micro-foundation for human capital.
- It highlights the two ways we learn
  - Learning from others and being taught the right response to an event.
  - Learning by trial and error, and improve with experience.

Long Run Choices: The Unemployment Problem

- In the 1970's to 1980's we witness an increasing gap between the Canadian and US unemployment rates.

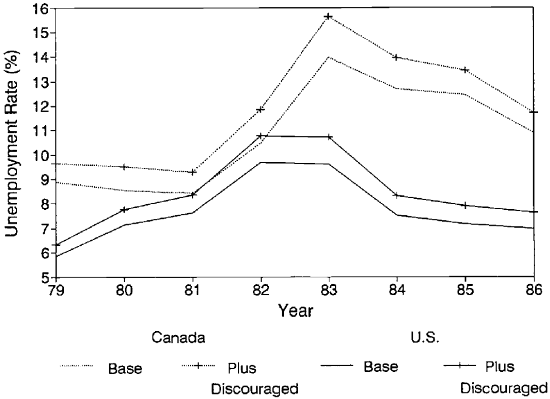

Fig. 5.12 Unemployment rates excluding and including discouraged workers, the United States and Canada

Source: Card and Ridell (1993).

The puzzle - Card and Ridell (1993)

- UI claims. However, an analysis of the regional extended benefit system in Canada suggests that the UI system itself is not the cause of the high level of unemployment at the close of the 1980s. If the same group of workers had entered the pool of unemployment in 1979 as in 1989, our simulations suggest that the average duration of available UI benefits would have been the same.
- supply patterns. These findings point to a significant role of the UI system in accounting for the rise in relative Canadian unemployment, although most of the rise in the Canadian-U.S. unemployment gap remains unexplained.

UI Incentives

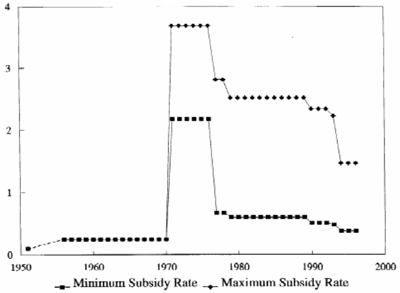

Note: The subsidy rate for minimally qualified claimants is the replacement rate times the maximum number of benefits for a minimally qualified claimant divided by the minimum weeks of employment to qualify to UI (weeks worked by a minimally qualified claimant. The minimum subsidy rate is the subsidy rate in a low unemployment region (under 6 % unemployment rate), while the maximum subsidy rate is the subsidy rate in a high unemployment region (above 15 %).

Fig. 1. UI subsidy rate for minimally qualified claimants.

- Source: Lemieux and MacLeod (2000) - notice that the big bump in unemployment rate differences occurred *after* 1980, at the same time as the big disincentive increase!.

Learning Effects

- Lemieux and MacLeod (2000) use Canadian administrative data to explore a simple idea:
  - The time at which the policy change occurs is not the time at which individuals consider the impact of UI upon labor supply.
  - That occurs when they actually lose their job, at which point they update their labor supply in the light of the current UI parameters.
  - Hence there is a lag between the treatment effect of the policy, and when the recipient observes, and then responds to the treatment.

The Effect of Learning

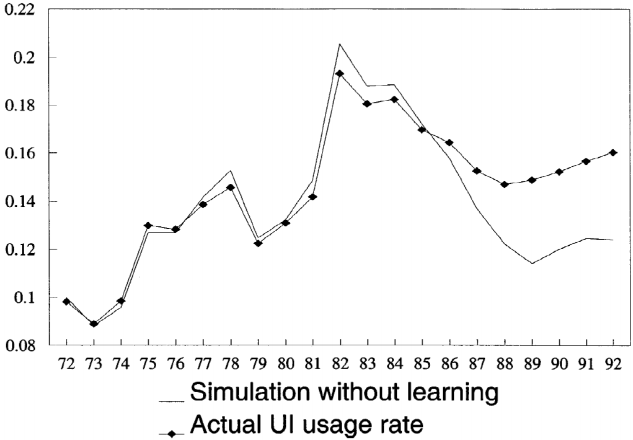

Fig. 4. Simulated impact of learning on UI usage rate, men.

- Source: Lemieux and MacLeod (2000)

Implications

- In this example the long run effects of UI upon behavior are much larger than the short run effects.
- In terms of the *causal impact of UI on unemployment:*
  - *The short run effects for t_y in the order of days/months is small/zero.*
  - The long run effect with $t_w$on the order of years is large and significant.
- However, empirically convincing causal effects typically rely upon short run responds (e.g. regression discontinuity designs), and hence careful work such as Card and Ridell (1993) cannot measure these long run effects.
- A simple mechanism/learning model is needed to figure out what is happening.

Outline

Lessons from Playing Chess

- The fact that a problem is finite and solvable in principle does not imply that it can be solved in practice - uncertainty arises from complexity rather than from risk.
- This implies that there is not a *single* algorithm that is optimal for playing chess.
- Successful algorithms use a variety of approaches - Bayesian learning, random search and experimentation, reinforcement learning and pattern recognition.
- A feature of successful game playing algorithms is experimentation - see Silver et al. (2017).

Implications for Human Decision Making

- These observations are consistent with Marvin Minsky (1986)'s society of theory mind view - since humans must act in complex environments, then they necessarily use a complex set of algorithms.
- Moreover, for the same problem, different individuals may use *different* algorithms.
- An implication of these models is even if two individuals are using the same algorithm to make choice, if they have different life experiences, then they may make different choices with the same data.
- For social scientists the challenge is how to incorporate such heterogeneity into empirically useful models of social behavior and performance

Why the utility maximization model is still the Best for economics

- A feature of all the algorithms we have discussed is that their goal to learn the best strategy for a particular situation.
- In the short run different individuals will make different decisions because of variation in their experiences, and the use of different algorithms to make choices.
- However, if the goals of a population are known, and the algorithms used are "good", then the utility maximization model will provide a good first order model of observed behavior.
- Thus, for many substantive economic questions, the utility maximization model remains the most useful approach.

The Commodification Problem

- The Savage approach to decision making highlights the idea that individuals build their own private models of the world.
- In particular, when commodities are not standardized, then individuals will conceptualize the same commodities in different ways, and hence two individuals with the same preferences (say a desire to be rich) and facing the same decision (which career to choose) may make different decisions.
- This problem is particularly acute today due to the tsunami of new online products - facebook, instagram, twitter and so on for which there does not exist an accepted framework to model these complex commodities.

Questions for Cognitive Scientists

- How do individuals adapt and learn in dynamic environments and deal with new commodities and choices with unforeseeable consequences?
- How do individuals learn from each other and form groups with coherent world views?
- How does one modify mistaken beliefs held by a group?
- Finally, the complexity of human decision making implies that they is always a great deal of variance in observed performance - it would be helpful to find ways to characterize that heterogeneity?

Summary of Learning Algorithms

1. Backwards induction and the Bellman principle.
2. Learning ability from repeated observations of performance.
3. Experimentation and learning from different treatment choices.
4. Pattern recognition and learning by doing in a complex feature space.

Bibliography

Altonji, J. G. and C. R. Pierret (2001). Employer learning and statistical discrimination. *Quarterly Journal of Economics 116*(1), 313–50.

Card, D. and W. C. Ridell (1993). A comparative analysis of unemployment in the United States and Canada. In D. Card and R. H. Freeman (Eds.), *Small Differences that Matter: Labor Markets and Income Maintenance in Canada and the United States*. Chicago, IL: University of Chicago Press.

Chase, W. G. and H. A. Simon (1973a). The mind's eye in chess. In W. G. Chase (Ed.), *Visual Information Processing*. New York, NY: Academic Press.

Chase, W. G. and H. A. Simon (1973b). Perception in chess. *Cognitive psychology 4*(1), 55–81.

Churchland, P. S. and T. J. Sejnowski (1993). *The Computational Brain*. Cambridge, MA: MIT Press.

Farber, H. S. and R. Gibbons (1996, November). Learning and wage dynamics. *Quarterly Journal of Economics 111*(4), 1007–47.