Stochastic Choice

Tomasz Strzalecki

Summer School in Economic Theory, IIAS Jerusalem, 2023

Notation

X ································ set of alternatives

$x, y, z \in X$ ······················ typical alternatives

$A, B, C \subseteq X$ ····················· finite choice problems (menus)

$\rho(x, A)$ ····················· probability of x being chosen from A

ρ ························· stochastic choice function

Main Model: Bayesian Expected Utility

- A general model that nests as special cases things Mike talked about
  - choice between lotteries
  - perception of numerosity
  - many other applications!
- **The** model in economics
- Useful benchmark to orient yourself during this summer school

Main Model: Bayesian Expected Utility

- The agent makes choices by maximizing **utility**
- The utility has a specific form: **expected** utility
- The expectation is formed using **Bayes** rule

Main Model: Bayesian Expected Utility

$\rightarrow$ The agent makes choices by maximizing **utility**

- The utility has a specific form: **expected** utility
- The expectation is formed using **Bayes** rule

Utility Maximization

- *Utility* of item x is $U(x)$
- *Utility function* is $U : X \to \mathbb{R}$
- Collection of observed menus $\mathcal{M}$
- *Choice function* is $\chi : \mathcal{M} \to X$ such that $\chi(A) \in A$ for all $A \in \mathcal{M}$
- *U represents* χ if $\forall_{A \in \mathcal{M}}\ \chi(A) = x$ iff $U(x) = \max_{y \in A} U(y)$
- Key assumption: U does not depend on A

Revealed Preference Theory: Given a choice function χ, does there $\exists\ U : X \to \mathbb{R}$ that represents it?

Utility Maximization

Axiom (Sen's α): If $x \in A \subset B$ and $\chi(B) = x$, then $\chi(A) = x$.

Proposition: Suppose X is finite and that $\mathcal{M}$ contains all pairs and triples. A choice function χ satisfies Sen's α if and only if $\exists\ U : X \to \mathbb{R}$ that represents it.

Proof (Necessity): Suppose U represents χ and $x \in A \subseteq B$ such that $\chi(B) = x$. Then $U(x) \geq U(y)$ for all $y \in B$. This implies that $U(x) \geq U(y)$ for all $y \in A$. So $\chi(A) = x$. $\square$

Proof (Sufficiency): Also easy but we will skip (even though this is the more interesting part).

Utility Maximization

Comments:

- χ is our *primitive* (what the analyst observes)
- U is our *representation*
- Representation is *as-if* (we don't claim the agent actually maximizes U, just that they behave as-if they do)
  - analogy: physical objects fall down to minimize their distance to earth, but they don't know how to solve any minimization problems
  - if you don't buy as-if-ness, that's OK. It might be interesting to think which parts of the brain do the maximization. But I won't do it here (maybe others will)

Stochastic Choice

- So far we assumed that observed choices are *deterministic*: given a menu the agent always chooses the same element

- What if observed choices are stochastic?

- Well-documented for perceptual choices

  – given two objects of similar weight, the same agent will sometimes pick x as heavier and sometimes pick y

- Also well-documented for economic choices

  – given two lotteries, the same agent will sometimes pick x and sometimes pick y

Stochastic Choice

- We need to change our *primitive*

- Observed probability of choosing x from menu A is $\rho(x, A)$

- Collection of all probability distributions over X is $\Delta(X)$

- *Stochastic choice function* is $\rho : \mathcal{M} \to \Delta(X)$ such that $\sum_{x \in A} \rho(x, A) = 1$

- Two interpretations:
  - within-subject experiments
  - between-subject experiments

Random Utility

- We need to change our *representation*

- Random utility (RU) function $\widetilde{U}$

  – formally, $(\Omega, \mathbb{P})$ is a probability space and $\widetilde{U} : \Omega \to \mathbb{R}^X$

- $\widetilde{U}$ *represents* ρ iff

$$\begin{aligned} \rho(x, A) &= \mathbb{P}(\widetilde{U}(x) = \max_{y \in A} \widetilde{U}(y)) \\ &= \mathbb{P}(\omega \in \Omega : \widetilde{U}(\omega, x) \geq \max_{y \in A} \widetilde{U}(\omega, y)) \end{aligned}$$

- Key assumption: $\mathbb{P}$ does not depend on A

Question: What are the axioms on ρ such that a RU representation exists?

Regularity

Axiom (Regularity): If $x \in A \subseteq B$, then $\rho(x, A) \geq \rho(x, B)$

This is an extension of Sen's α.

Proposition: If ρ has a RU representation, then it satisfies Regularity.

Notation: $N(x, A) := \{\widetilde{U}(x) = \max_{y \in A} \widetilde{U}(y)\}$.

Proof: As we already established, if x maximizes U on B then x maximizes U on A. So the event $N(x, A)$ is a superset of the event $N(x, B)$, so it must have a larger probability. $\square$

Violations of Regularity

1. **Choice Overload**: tasting booth in a supermarket
   - 6 varieties of jam — 70% people purchased no jam
   - 24 varieties of jam — 97% people purchased no jam

2. **Asymmetric dominance effect**: adding a "decoy" option raises demand for the targeted option

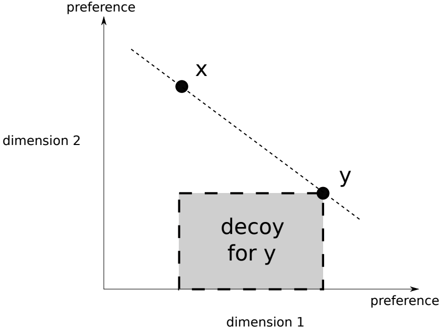

Regularity

Proposition: If X has 3 elements, then Regularity is equivalent to RU representation.

Proof Idea:

- For each A, the sets $N(x, A)$ form a partition of Ω as x ranges over A
- ρ defines a probability distribution over the cells of each partition
- We have as many partitions as there are menus
- Regularity ensures that they are consistent with a single $\mathbb{P}$

Beyond $|X| = 3$

Comments:

- Unfortunately, when $|X| > 3$, Regularity alone is not enough
- More axioms are needed, but hard to find economic interpretation
- More elegant axioms if X consists of lotteries $\rightsquigarrow$ later in this lecture

Beyond $|X| = 3$

Axiom (Block and Marschak, 1960): For all $x \in A$

$$\sum_{B \supseteq A} (-1)^{|B \setminus A|} \rho(x, B) \geq 0.$$

Theorem (Falmagne, 1978): If X is finite and $\mathcal{M} = 2^X \setminus \{\emptyset\}$, TFAE:

(i) ρ has a random utility representation

(ii) ρ satisfies the Block–Marschak axiom

Comments:

- Necessity of this axiom follows from the inclusion-exclusion formula (Möbus transform)
- There are other axioms in the literature (due to McFadden and Richter) but they are even worse

Additive Random Utility (ARU)

- Let $v \in \mathbb{R}^X$ be a deterministic utility function
- Let $\tilde{\epsilon} : \Omega \to \mathbb{R}^X$ be a random *unobserved utility shock* or *error*
  - the distribution of $\tilde{\epsilon}$ has a density and full support

Definition: ρ has an *ARU* representation if it has a RU representation with

$$\tilde{U}(x) = v(x) + \tilde{\epsilon}(x)$$

Special Cases: it is often assumed that $\tilde{\epsilon}(x)$ are i.i.d. across $x \in X$

- *Logit*, where $\tilde{\epsilon}(x)$ has an "Type I Extreme Value" (TIEV) distribution
- *Probit*, where $\tilde{\epsilon}(x)$ has a Normal distribution

Positivity

Full support of $\tilde{\epsilon}$ ensures that all items are chosen with positive probability

Axiom (Positivity): $\rho(x, A) > 0$ for all $x \in A$

Comments:

- This leads to a non-degenerate likelihood function—good for estimation
- Positivity cannot be rejected by any finite data set

Proposition: If X is finite and ρ satisfies Positivity, TFAE:

(i) ρ has a RU representation

(ii) ρ has a ARU representation

The Luce Model

Definition: ρ has a *Luce representation* iff there exists $w : X \to \mathbb{R}_{++}$ s.t.

$$\rho(x, A) = \frac{w(x)}{\sum_{y \in A} w(y)}$$

Intuition 1: $w(x)$ is the "response strength" associated with x. Choice probability is proportional to the response strength.

Intuition 2: The Luce representation is like a conditional probability: the probability distribution on A, is the conditional of the probability distribution on the grand set X.

Equivalent Model: You can also rewrite this as "softmax"

$$\rho(x, A) = \frac{e^{v(x)}}{\sum_{y \in A} e^{v(y)}}$$

for some deterministic utility function $v : X \to \mathbb{R}$

Axioms for Luce/Logit

Axiom (Luce's IIA). For all $x, y \in A \cap B$

$$\frac{\rho(x, A)}{\rho(y, A)} = \frac{\rho(x, B)}{\rho(y, B)},$$

whenever the probabilities are positive.

Proposition: TFAE:

(i) ρ satisfies Positivity and Luce's IIA

(ii) ρ has a Luce representation

(iii) ρ has a logit representation (i.e., ARU with i.i.d. TIEV shocks)

Summary so Far

$$
\begin{aligned}
RU &= \text{Regularity plus other Axioms} \\
RU + \text{ Positivity} &= ARU \\
ARU + i.i.d. + TIEV &= \text{ logit} \\
&= \text{ Luce} \\
&= \text{ Positivity + Luce's IIA Axiom}
\end{aligned}
$$

Main Model: Bayesian Expected Utility

- The agent makes choices by maximizing **utility**

$\rightarrow$ The utility has a specific form: **expected** utility

- The expectation is formed using **Bayes** rule

Expected Utility

- Now $X = \Delta(Z)$, where Z is the set of prizes
- Typical items are now $p, q, r \in X$, called *lotteries*

Definition: U has an *EU form* if for some function $u \in \mathbb{R}^Z$

$$U(p) := \mathbb{E}_p u := \sum_{z \in Z} u(z)p(z)$$

- The function u is called the Bernoulli utility function.
- When Z is money, then concavity of u corresponds to risk aversion

Expected Utility

- Key property of EU is *linearity in probabilities*
- For any $p, q \in \Delta(Z)$ and $\alpha \in (0,1)$ define a new lottery $\alpha p + (1-\alpha)q$ that attaches probability $\alpha p(z) + (1-\alpha)q(z)$ to each prize $z \in Z$

Proposition: For finite Z, U has an EU form iff

$$U(\alpha p + (1-\alpha)q) = \alpha U(p) + (1-\alpha)U(q)$$

Random Expected Utility (REU)

Definition: ρ has a *REU representation* if has a RU representation where with probability one $\tilde{U}$ has an EU form:

$$\tilde{U}(p) := \mathbb{E}_p \tilde{u}$$

for some random Bernoulli utility function $\tilde{u} \in \mathbb{R}^Z$

REU—Linearity

Definition: $\alpha A + (1-\alpha)q := \{\alpha p' + (1-\alpha)q : p' \in A\}$

Axiom (Linearity). For any $\alpha \in (0,1)$ and $p \in A$ and $q \in X$

$$\rho(p, A) = \rho(\alpha p + (1-\alpha)q, \alpha A + (1-\alpha)q)$$

Idea: Linearity of U applied utility by utility

$$\tilde{u}_\omega \in N(p, A) \Longleftrightarrow \tilde{u}_\omega \in N\left(\alpha p + (1-\alpha)q, \alpha A + (1-\alpha)q\right)$$

Violation of Linearity: Allais Paradox

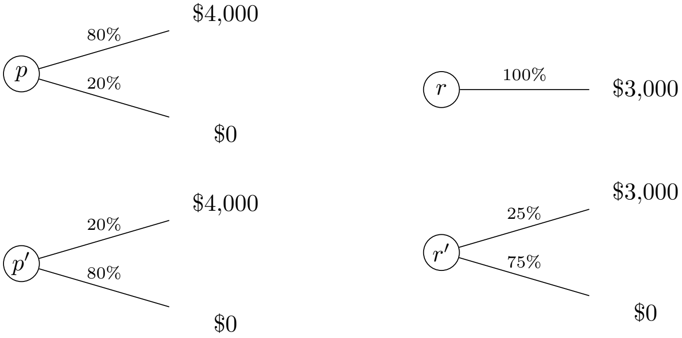

- Note that $p' = .25p + .75\delta_0$ and $r' = .25r + .75\delta_0$
- Kahneman and Tversky (1979) show that $\rho(r, \{p, r\}) = 0.84$ but $\rho(r', \{p', r'\}) = 0.37$

REU—Axioms

Notation: $Ext(A)$ is the set of extreme points of A

Axiom (Extremeness). $\rho(Ext(A), A) = 1$

Idea: The indifference curves are linear, so maximized at an extreme point of the choice set (modulo ties)

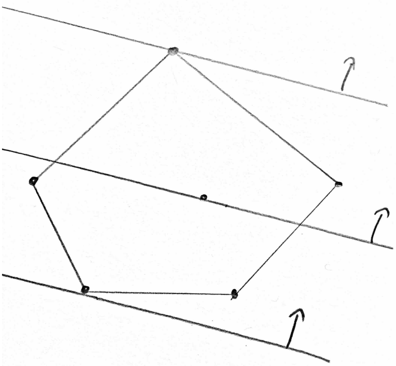

REU—Axiomatization

Theorem† (Gul and Pesendorfer, 2001). ρ has a REU representation iff it satisfies

- Regularity
- Extremeness
- Linearity
- Continuity†

A different model

- Let U_θ be a family of vNM forms with CARA or CRRA indexes (allow for risk-aversion and risk-loving)
- Higher θ is more risk-aversion

Model 1 (à la REU): There is a probability distribution $\mathbb{P}$ over error shocks $\tilde{\epsilon}$ to the preference parameter θ

$$\rho_\theta^{REU}(p, A) = \mathbb{P}\{U_{\theta+\tilde{\epsilon}}(p) \geq U_{\theta+\tilde{\epsilon}}(q) \text{ for all } q \in A\}$$

Model 2 (à la ARU): There is a probability distribution $\mathbb{P}$ over error shocks $\tilde{\epsilon}$ to the expected value, $\tilde{\epsilon}$ i.i.d. over lotteries

$$\rho_\theta^{ARU}(p, A) = \mathbb{P}\{U_\theta(p) + \tilde{\epsilon}(p) \geq U_\theta(q) + \tilde{\epsilon}(q) \text{ for all } q \in A\}$$

Comment: In Model 2, preferences over lotteries are not vNM!

Comparing the two models

Observation 1: Model 1 has intuitive properties:

- If p FOSD q, then $\rho_\theta^{REU}(p, \{p, q\}) = 1$
- If p SOSD q, then $\rho_\theta^{REU}(p, \{p, q\})$ is increasing in θ

Observation 2: Model 2 not so much:

- If p FOSD q, then $\rho_\theta^{ARU}(p, \{p, q\}) < 1$
- If p SOSD q, then $\rho_\theta^{ARU}(p, \{p, q\})$ is not monotone in θ

Theorem If p SOSD q, then $\rho_\theta^{ARU}(p, \{p, q\})$ is strictly decreasing for large enough θ.

Main Model: Bayesian Expected Utility

- The agent makes choices by maximizing **utility**
- The utility has a specific form: **expected** utility

$\rightarrow$ The expectation is formed using **Bayes** rule

Timing of Beliefs

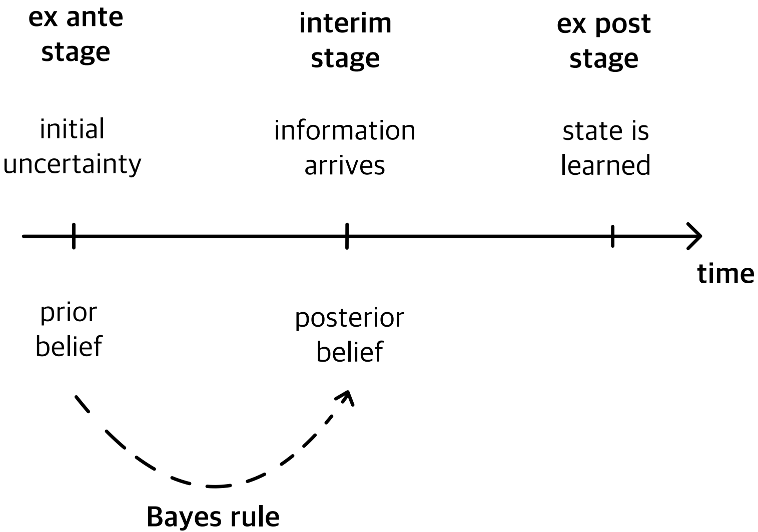

Ex Ante Stage

- S is set of *states* of the world
- $p \in \Delta(S)$ is *prior* of the agent (initial belief)
- $v : X \times S \to \mathbb{R}$ (deterministic) utility function of the agent
- For any belief $q \in \Delta(S)$ the expected utility of x is denoted by

$$\mathbb{E}_q v(x) := \sum_{s \in S} q(s) v(x, s)$$

- Agent faced ex ante with menu $A \subseteq X$ solves $\max_{x \in A} \mathbb{E}_p v(x)$
- Observed choices of agent are deterministic

Interim Stage

- Agent receives a message $m \in M$ (a "noisy mental representation" or a privately observed signal)

- $\beta : S \to \Delta(M)$ is the signal structure (a.k.a. *Blackwell experiment*)

- For each message m there is a *posterior belief* $q(\cdot|m) \in \Delta(S)$

- Posterior is given by the *Bayes rule*

$$q_m(s) = q(s|m) = \frac{\beta(m|s)p(s)}{\sum_{s'} \beta(m|s')p(s')}$$

- Given message m agent solves $\max_{x \in A} \mathbb{E}_{q_m} v(x)$

Interim Choice Probabilities

- Agent: does not know s learns m

- Analyst: knows s, does not learn m

- Observed choices are stochastic. Choice probability in state s is

$$\rho^s(x, A) = \beta\Big(\Big\{m \in M : \mathbb{E}_{q_m} v(x) = \max_{y \in A} \mathbb{E}_{q_m} v(y)\Big\} \,\Big|\, s\Big)$$

- So now instead of ρ we have a collection $(\rho^s)_{s \in S}$

Observation: If β does not depend on the menu, then each ρ^s has a RU representation

Example 1: character recognition

- In each trial the subject is briefly shown a character, say c or e
- $X = \{c, e\}$, $S = \{s^c, s^e\}$
- $v(c, s^c) = v(e, s^e) = 1$, $v(c, s^e) = v(e, s^c) = 0$
- $M = \mathbb{R}$ random perception
- $\beta : S \to \Delta(M)$ is the signal with density $b(m|s)$
- Bayes rule says:

$$\frac{q(s^c|m)}{q(s^e|m)} = \frac{b(m|s^c)}{b(m|s^e)} \frac{p(s^c)}{p(s^e)}$$

- Optimal to choose c if $\frac{q(s^c|m)}{q(s^e|m)} > 1$

Example 1: character recognition

- Bayes rule says:

$$\frac{q(s^c|m)}{q(s^e|m)} = \frac{b(m|s^c)}{b(m|s^e)} \frac{p(s^c)}{p(s^e)}$$

- Optimal to choose c if $\frac{q(s^c|m)}{q(s^e|m)} > 1$

- Let $p := p(s^c)$ and $\ell(m) := \frac{b(m|s^c)}{b(m|s^e)}$

- Optimal to choose c if $\ell(m) > \frac{p}{1-p}$

- Let $L(k) := \{m \in M : \ell(m) > k\}$. Notice $k > k'$ implies $L(k) \subseteq L(k')$

- So we have $\rho^{s,p}(c) = \beta(L(\frac{1-p}{p})|s) \rightsquigarrow$ increasing function of p

Example 1: character recognition

- Imagine you run this experiment in batches
- In each batch of trials the frequency of c is different
- BEU model predicts that ρ depends on frequency
  - assuming that agent somehow adapts to the frequency in each batch
  - perhaps you throw out initial trials in each batch (adaptation phase)

Example 1: character recognition

- BEU model predicts that ρ depends on frequency

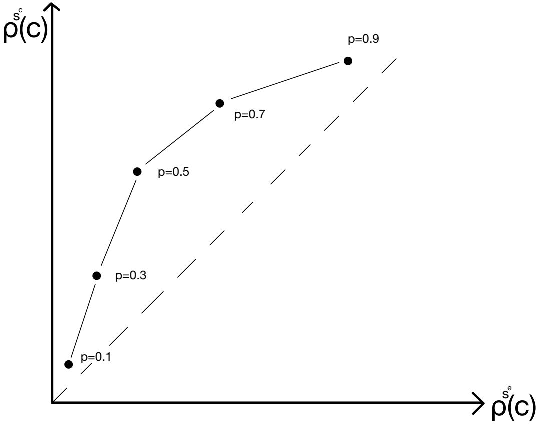

- Another prediction of BEU: making the task harder shifts the curve toward the diagonal

Example 2: weight discrimination

- $X = \{\ell, r\}$ - physical objects
- $s \in \mathbb{R}^2_+$ - true weight of each object
- $p \in \Delta(S)$ prior is such that weight is i.i.d. over objects
- $m \in \mathbb{R}^2_+$ - perception of weight of each object
- $m_x \sim N(s_x, \nu^2)$ - signal structure
- If menu is $A = \{\ell, r\}$ then observed choice probability is

$$\rho^s(\ell, A) = \Phi\Big(\frac{s_\ell - s_r}{\nu\sqrt{2}}\Big) \rightsquigarrow \text{ psychometric function}$$

Example 3: Weber's law

- Instead of $m_x \sim N(s_x, \nu^2)$
- Define $m_x \sim N(\log s_x, \nu^2)$
- If menu is $A = \{\ell, r\}$ then observed choice probability is

$$\rho^s(\ell, A) = \Phi\Big(\frac{\log(s_\ell / s_r)}{\nu\sqrt{2}}\Big) \rightsquigarrow \text{ Weber's law}$$

Example 4: choice between lotteries

- This is (a version of) Mike's model
- For each lottery, agent has a noisy perception of payoffs and probabilities
- Maximizes EU given their posterior of what the lottery is
- The simple version where only payoffs are imperfectly perceived is a special case of REU
  - can interpret random perception/posterior as random Bernoulli utility

Example 5: economic example

- Agent is a HR recruiter who is hiring an applicant
- $S := \{0, 1\}$ is the qualification of the applicant (low or high)
- Interview can either be a flop or go well: $M = \{m_0, m_1\}$
- Signal is symmetric with precision $b := \beta(m_1|s = 1) = \beta(m_0|s = 0)$
- $A := \{0, 1\}$ is menu of choices (either pass or make a hire).
- Utility of hiring a qualified applicant equals 1 and an unqualified applicant, -1. The utility of not hiring is zero.
- The analyst who observes s but not m sees high-skilled applicants hired b percent of the time and low-skilled applicants being hired $1 - b$ percent of the time.

Bayes Rule

- BEU assumes Bayesian updating
- Even though there is massive evidence against it:
  - base-rate neglect
  - confirmation bias
  - gambler's fallacy
  - hot-hand fallacy
- So what? BEU also assumes EU even though evidence against
  - yet BEU can produce behavior similar to Prospect Theory, etc!
- Difference between assuming these things at the level of representation and at the level of behavior

Summary

- RU (random tastes): Regularity plus other messy axioms
- REU (random risk aversion): Regularity plus Linearity plus Extremeness
- BEU (random perception): Axioms?
- In all of these there was an invariance assumption
  - the distribution of utilities independent of menu
  - the signal structure independent of menu

Going forward: active learning

- So far, learning was *passive* (β was fixed)

- In models of *active learning* the agent can choose β at a cost

  $\rightsquigarrow$ pay attention

- What is the appropriate cost function?

  – mutual information $\rightsquigarrow$ *rational inattention* (Sims and his followers)

    – tractable but has lots of problems

    – for example the psychometric function in the weight discrimination task is a step function, instead of a smooth S-shaped function

  – many other costs have been proposed

- Next lecture: dynamic model of active learning

Example of active learning

	s_1	s_2
x	0	2
y	1	1
z	2	0

- Prior is $(\frac{1}{2}, \frac{1}{2})$
- Cost of learning the state perfectly is 0.75
- No other learning possible (cost infinity)
- $\rho(x, \{x, y\}) = 0$, $\rho(x, \{x, y, z\}) = \frac{1}{2}$
- Violation of Regularity because adding z adds incentive to learn about the state