Stochastic Choice

Tomasz Strzalecki

Hotelling Lectures in Economic Theory
Econometric Society European Meeting, Lisbon, August 25, 2017

Table of Contents

Lecture 1

Introduction
Random Utility/Discrete Choice
- Properties of RU
- Special Cases
- Random Expected Utility (REU)

Learning
Attention
- Optimal Attention
- Random Attention

Deliberate Randomization

Lecture 2

Introduction
Dynamic Random Utility
Dynamic Discrete Choice
Decision Times

Plan

Purpose: Overview where the field is and where it seems to be going

Lecture 1: Static Choice

- Random Utility (and Discrete Choice)
- Learning, Attention, Deliberate Randomization

Lecture 2: Dynamic Choice

- Dynamic Random Utility
- Dynamic Discrete Choice
- Drift-Diffusion Models

Disclaimer

- I won't get too deeply into any one area
- The ES monograph (in preparation) fills in more details
  - Theorem† means there are some terms I did not define
  - Theorem‡ means that additional technical conditions are needed
- I cover mostly work in decision theory. I am not an expert on neighboring fields, such as discrete choice econometrics, structural IO and labor, experimental economics, psychology and economics, cognitive science. Happy to talk if you are one.
- All comments welcome at tomasz_strzalecki@harvard.edu

Notation

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

$x \in X$ ·················· typical alternative

$A \subseteq X$ ·················· finite choice problem (menu)

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

ρ ························ stochastic choice function (rule)

Stochastic Choice

- **Idea**: The analyst/econometrician observes an agent/group of agents

- **Examples**:

  – Population-level field data: McFadden (1973)

  – Individual-level field data: Rust (1987)

  – Between-subjects experiments: Kahneman and Tversky (1979)

  – Within-subject experiments: Tversky (1969)

Is individual choice random? Why?

Stylized Fact: Choice can change, even if repeated shortly after

- Tversky (1969), Hey (1995), Ballinger and Wilcox (1997), Hey (2001), Agranov and Ortoleva (2017)

Possible reasons:

- Randomly fluctuating tastes
- Noisy signals
- Attention is random
- People just like to randomize
- Trembling hands
- Experimentation (experience goods)

Questions

1. What are the properties of ρ (axioms)?

   - Example: "*Adding an item to a menu reduces the choice probability of all other items*"

2. How can we "explain" ρ (representation)?

   - Example: "*The agent is maximizing utility, which is privately known*"

Goals

1. Better understand the properties of a model. What kind of predictions does it make? What axioms does it satisfy?
   - Ideally, prove a *representation theorem* (ρ satisfies Axioms A and B if and only if it has a representation R)

2. Identification: Are the parameters pinned down uniquely?

3. Determine whether these axioms are reasonable, either normatively, or descriptively (testing the axioms)

4. Compare properties of different models (axioms can be helpful here, even without testing them on data). Outline the modeling tradeoffs

5. Estimate the model, make a counterfactual prediction, evaluate a policy (I won't talk about those here)

Testing the axioms

- Axioms expressed in terms of ρ, which is the limiting frequency
- How to test such axioms when observed data is finite?
- Hausman and McFadden (1984) developed a test of Luce's IIA axiom that characterizes the logit model
- Kitamura and Stoye (2016) develop tests of the static random utility model based on axioms of McFadden and Richter (1990)
- I will mention many other axioms here, without corresponding "tests"

Richness

- The work in decision theory often assumes a "rich" menu structure
  - Menu variation can be generated in experiments
  - But harder in field data
  - But don't need a full domain to *reject* the axioms

- The work in discrete choice econometrics often assumes richness in "observable attributes"
  - I will abstract from this here

- The two approaches lead to somewhat different identification results
  - Comparison?

Introduction

Random Utility/Discrete Choice
- Properties of RU
- Special Cases
- Random Expected Utility (REU)

Learning

Attention
- Optimal Attention
- Random Attention

Deliberate Randomization

Random Utility

Idea: Choice is random because:

- There is a population of heterogenous individuals
- Or there is one individual with varying preferences

Models:

- Random Utility
- Discrete Choice

Notation:

$(\Omega, \mathcal{F}, \mathbb{P})$ $\cdots\cdots\cdots$ probability space that carries all random variables

Random Utility (RU)

- Let $\tilde{U} : \Omega \to \mathbb{R}^X$ be a random utility function on X

- $C(x, A)$ is the event in which the agent chooses x from A

$$C(x, A) := \{\omega \in \Omega : \tilde{U}_\omega(x) \geq \tilde{U}_\omega(y) \text{ for all } y \in A\}$$

- T is the event in which there is a tie

$$T := \{\omega \in \Omega : \tilde{U}_\omega(x) = \tilde{U}_\omega(y) \text{ for some } x \neq y\}$$

Definition: ρ has a *random utility representation* if there exists $(\Omega, \mathcal{F}, \mathbb{P})$ and $\tilde{U} : \Omega \to \mathbb{R}^X$ such that $\mathbb{P}(T) = 0$ and

$$\rho(x, A) = \mathbb{P}\big(C(x, A)\big)$$

Key assumption:

- $\mathbb{P}$ is independent of the menu; it's the structural invariant of the model
- Menu-dependent $\mathbb{P}$ can trivially explain any ρ

Discrete Choice (DC)

- 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 a *discrete choice* representation if it has a RU representation with $\tilde{U}(x) = v(x) + \tilde{\epsilon}(x)$

This is sometimes called the additive random utility model

Discrete Choice (DC)

- The fact that $\tilde{\epsilon}$ has a density rules out ties
- The full support assumption on $\tilde{\epsilon}$ ensures that all items are chosen with positive probability

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

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

Ways to deal with ties

- Prohibit them outright by assuming
  - $\mathbb{P}(T) = 0$
  - density on $\tilde{\epsilon}$
- But sometimes more convenient to allow ties
  - Use a tiebreaker (Gul and Pesendorfer, 2006)
  - Change the primitive (Barberá and Pattanaik, 1986; Lu, 2016; Gul and Pesendorfer, 2013)
- I will skip over the details in this talk

Equivalence

Theorem: If X is finite and ρ satisfies Positivity, then the following are equivalent:

(i) ρ has a random utility representation

(ii) ρ has a discrete choice representation

Questions:

- What do these models assume about ρ?
- Are their parameters identified?
- Are there any differences between the two approaches?

Introduction

Random Utility/Discrete Choice
Properties of RU
Special Cases
Random Expected Utility (REU)

Learning

Attention
Optimal Attention
Random Attention

Deliberate Randomization

Axiomatic Characterizations

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

Intuition When we add an item to a menu, existing items have to "make room" for it.

Examples of violation:

1. Iyengar and Lepper (2000): tasting booth in a supermarket
   - 6 varieties of jam — 70% people purchased no jam
   - 24 varieties of jam — 97% people purchased no jam
2. Huber, Payne, and Puto (1982): adding a "decoy" option raises demand for the targeted option

Decoy Effect

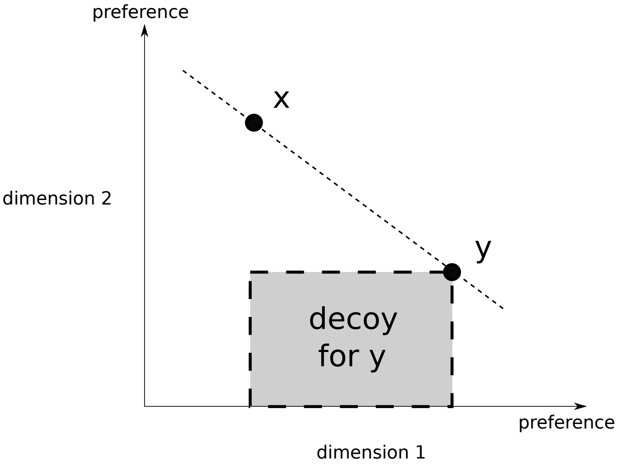

Axiomatic Characterizations

Theorem (Block and Marschak, 1960). If ρ has a random utility representation, then it satisfies Regularity. Moreover, Regularity is sufficient if $|X| \leq 3$.

Comments:

- Unfortunately, when $|X| > 3$, Regularity alone is not enough
- More axioms are needed, but they are hard to interpret
- More elegant axioms if X consists of lotteries (Gul and Pesendorfer, 2006) $\rightsquigarrow$ later in this lecture

Identification of Utilities/Preferences

- Since utility is ordinal, we cannot identify its distribution—at best we can hope to pin down the distribution of ordinal preferences
- But it turns out we can't even do that

Example (Fishburn, 1998). Suppose that $X = \{x, y, z, w\}$. The following two distributions over preferences lead to the same ρ.

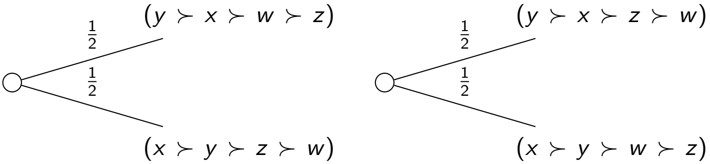

Note that these two distributions have disjoint supports!

Identification of "Marginal" Preferences

Theorem (Falmagne, 1978). If $\mathbb{P}_1$ and $\mathbb{P}_2$ are random preference representations of the same ρ, then for any $x \in X$

$$\mathbb{P}_1(x \text{ is } k\text{-th best in } X) = \mathbb{P}_2(x \text{ is } k\text{-th best in } X)$$

for all k

Identification in DC

Theorem: If $(v_1, \tilde{\epsilon}_1)$ is a DC representation of ρ, then for any $v_2 \in \mathbb{R}^X$ there exists $\tilde{\epsilon}_2$ such that $(v_2, \tilde{\epsilon}_2)$ is another representation of ρ

Comments:

- So can't identify v (even ordinally) unless make assumptions on unobservables
- If assume a given distribution of $\tilde{\epsilon}$, then can pin down more
- Also, stronger identification results are obtained in the presence of "observable attributes" (see, e.g. Matzkin, 1992)

Introduction

Random Utility/Discrete Choice

- Properties of RU
- Special Cases
- Random Expected Utility (REU)

Learning

Attention

- Optimal Attention
- Random Attention

Deliberate Randomization

i.i.d. DC

- It is often assumed that $\tilde{\epsilon}_x$ are i.i.d. across $x \in X$
  - logit
  - probit

- In i.i.d. DC the binary choice probabilities are given by $\rho(x, \{x, y\}) = F(v(x) - v(y))$ where F is the cdf of $\tilde{\epsilon}_x - \tilde{\epsilon}_y$
  - such models are called Fechnerian

The Luce Model

- In the logit model the choice probabilities are given by the closed-form

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

- This is known as the Luce representation

Axiom (Luce's IIA). For all $x, y \in A \cap B$ whenever the probabilities are positive

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

Theorem (Luce, 1959; McFadden, 1973): The following are equivalent

(i) ρ satisfies Positivity and Luce's IIA

(ii) ρ has a Luce representation

(iii) ρ has a logit representation

Evidence

- Luce's IIA axiom is routinely violated
  - Blue bus/red bus problem (Debreu, 1960)
  - Actually, blue bus/red bus is a problem for all i.i.d. DC models

- Fix: relax the i.i.d. assumption
  - nested logit
  - GEV (generalized extreme value)
  - multivariate probit
  - mixed logit

- Another axiom that i.i.d. DC satisfies: Strong Stochastic Transitivity
  - often violated in experiments (Rieskamp, Busemeyer, and Mellers, 2006)

Generalizations of Luce

- Elimination by aspects (Tversky, 1972)
- Random Attention (Manzini and Mariotti, 2014)
- Attribute rule (Gul, Natenzon, and Pesendorfer, 2014)
- Additive Perturbed Utility (Fudenberg, Iijima, and Strzalecki, 2015)
- Perception adjusted Luce (Echenique, Saito, and Tserenjigmid, 2013)

Introduction

Random Utility/Discrete Choice
- Properties of RU
- Special Cases
- Random Expected Utility (REU)

Learning

Attention
- Optimal Attention
- Random Attention

Deliberate Randomization

Random Expected Utility (REU)

- Gul and Pesendorfer (2006) study choice between lotteries
- Specify the RU model to $X = \Delta(Z)$, where Z is a finite set of prizes

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

$$\tilde{U}(x) := \sum_{z \in Z} \tilde{u}(z) x(z)$$

for some $\tilde{u} \in \mathbb{R}^Z$

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

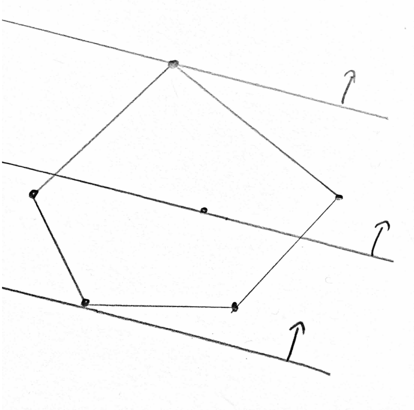

REU—Axioms

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

$$\rho(x, A) = \rho(\alpha x + (1-\alpha)y, \{\alpha x' + (1-\alpha)y : x' \in A\})$$

Idea: Just like the vNM Independence axiom

REU—Gul and Pesendorfer (2006) Results

Theorem† (Characterization). ρ has a REU representation if and only if it satisfies

- Regularity
- Extremeness
- Linearity
- Continuity†

Theorem† (Uniqueness). In a REU representation the distribution over ordinal preferences is essentially† identified.

REU—Comments

- Simple axioms
- Better identification results
- Stronger assumptions: Allais (1953) paradox is a rejection of Linearity
  - We'll see soon what happens if vNM is relaxed
- Gul and Pesendorfer (2006) also developed a version with tie-breakers, need to weaken continuity
- Model used as a building block for a lot to come
- This is only one possible specification of risk preferences . . .

Measuring Risk Preferences

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

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

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

Model 2 (a la DC): There is a probability distribution $\mathbb{P}$ over error shocks $\tilde{\epsilon}$ to the expected value

$$\rho_\theta^{DC}(x, A) = \mathbb{P}\{U_\theta(x) + \tilde{\epsilon}(x) \geq U_\theta(y) + \tilde{\epsilon}(y) \text{ for all } y \in A\}$$

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

Measuring Risk Preferences

Notation:

- FOSD—First Order Stochastic Dominance
- SOSD—Second Order Stochastic Dominance

Observation 1: Model 1 has intuitive properties:

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

Observation 2: Model 2 not so much:

- If x FOSD y, then $\rho_\theta^{DC}(x, \{x, y\}) < 1$
- If x SOSD y, then $\rho_\theta^{DC}(x, \{x, y\})$ is not monotone in θ

Measuring Risk Preferences

Theorem: (Wilcox, 2008, 2011; Apesteguia and Ballester, 2017) There exists $\bar{\theta}$ such that $\rho_\theta^{DC}(x, \{x, y\})$ is strictly decreasing for $\theta > \bar{\theta}$.

Comments:

- This biases parameter estimates
- Subjects may well violate FOSD and SOSD. Better to model these violations explicitly rather than as artifacts of the error specification?
- A similar lack of monotonicity for discounted utility time-preferences
- Apesteguia, Ballester, and Lu (2017) study a general notion of single-crossing for random utility models

Recap

- RU and DC are equivalent as far as ρ is concerned
- But have different parameters:
  - distribution over preferences
  - deterministic v and random $\tilde{\epsilon}$
- Suggestive of different parametric specifications

Introduction

Random Utility/Discrete Choice
- Properties of RU
- Special Cases
- Random Expected Utility (REU)

Learning

Attention
- Optimal Attention
- Random Attention

Deliberate Randomization

Learning

- In RU choice is stochastic because preferences are fluctuating

- Another possible reason: choices are driven by agent's noisy signals

- This is a special case of RU
  - with "preferences" equal to "expected utility conditional on the signal"

Question: Can any RU ρ be represented this way?

Answer: Depends if the model is rich enough to permit a separation of tastes and beliefs

Learning—probabilistic model

- Fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a random utility $\tilde{U} : \Omega \to \mathbb{R}^X$
- Let $\mathcal{G}$ represent the information the agent is learning
- Conditional on the signal the agent maximizes $\mathbb{E}[\tilde{U}(x)|\mathcal{G}]$

Comments:

- Choices are random because they depend on the signal realization
  - No information ($\mathcal{G}$ trivial) $\Rightarrow$ choices are deterministic
  - Full information ($\mathcal{G} = \mathcal{F}$) $\Rightarrow$ this is just a RU model
  - In general, the finer the $\mathcal{G}$, the more random the choices, keeping $(\Omega, \mathcal{F}, \mathbb{P})$ constant
- ρ has a (probabilistic) learning representation iff it has a RU representation
- Strictly special case of RU in a dynamic setting (Frick, Iijima, and Strzalecki, 2017) $\rightsquigarrow$ Lecture 2

Learning—statistical model

S ························ set of unknown states

$p \in \Delta(S)$ ················ prior belief

$v : S \to \mathbb{R}^X$ ·············· state-dependent utility function

$\mathbb{E}_p v(x)$ ··················· (ex ante) expected utility of x

- Signal structure: in each state s there is a distribution over signals
- For each signal realization, posterior beliefs are given by the Bayes rule
- The prior p and the signal structure $\Rightarrow$ distribution μ over posteriors
  - Often convenient to work with μ directly
  - For each posterior $\tilde{q}$ the agent maximizes $\max_{x \in A} \mathbb{E}_{\tilde{q}} v(x)$

Learning—statistical model

For each s, the model generates a choice distribution $\rho^s(x, A)$

– In some lab experiments the analyst can control/observe s

An average of ρ^s according to the prior p generates $\rho(x, A)$

Comments:

- The class of ρ generated this way equals the RU class
- For each s conditional choices ρ^s also belong to the RU class
  - Consistency conditions of ρ^s across s?
- The (statistical) learning model becomes a strictly special case of RU when specified to Anscombe–Aumann acts (Lu, 2016)

Learning—the Lu (2016) model

- Random Utility model of choice between Anscombe–Aumann acts
- This means $X = \Delta(Z)^S$
  - In each state the agent gets a lottery over prizes in a finite set Z
- Random Utility $\tilde{U}(x) = \sum_{s \in S} v(x(s))\tilde{q}(s)$, where
  - v is a (deterministic) vNM form over $\Delta(Z)$
  - $\tilde{q}$ is the (random) posterior over S
- The distribution over $\tilde{q}$ is given by μ

Learning—the Lu (2016) model

Theorem‡ (Characterization). ρ has a (statistical) learning representation iff it satisfies the Gul and Pesendorfer (2006) axioms *plus* S-independence†, Non-degeneracy†, and C-determinisim†.

- Ties dealt with by changing the primitive (3rd kind)

Theorem‡ (Uniqueness). The prior p is unique, the information structure μ is unique and the utility function v is cardinally-unique.

- In fact, the parameters can be identified on binary menus
- *Test functions*: calibration through constant acts

Theorem‡ (Comparative Statics). Fix v and p and consider two information structures μ and μ'. ρ is "more random" than ρ' if and only if μ is Blackwell-more informative than μ'.

More about learning

- Models of learning so far:
  - the probabilistic model (information is $\mathcal{G}$)
  - the statistical model (information is μ)
    - the Lu (2016) model

- In all of them information is independent of the menu

- But it could reasonably depend on the menu:
  - if new items provide more information
  - or if there is limited attention $\rightarrow$ next section

Example

	$\tilde{U}_\omega(steak\ tartare)$	$\tilde{U}_\omega(chicken)$	$\tilde{U}_\omega(fish)$
$\omega = good\ chef$	10	7	3
$\omega = bad\ chef$	0	5	0

- *fish* provides an informative signal about the quality of the chef
  - $\mathcal{G}^{\{s,c,f\}}$ gives full information:
    - if the whole restaurant smells like fish $\rightarrow$ chef is bad
    - if the whole restaurant doesn't smell like fish $\rightarrow$ chef is good
  - $\rho(s, \{s, c, f\}) = \rho(c, \{s, c, f\}) = \frac{1}{2}$ and $\rho(f, \{s, c, f\}) = 0$

- in absence of f get no signal
  - $\mathcal{G}^{\{s,c\}}$ gives no information
  - $\rho(s, \{s, c\}) = 0$, $\rho(c, \{s, c\}) = 1$ (if prior uniform)

- violation of the Regularity axiom!
  - menu-dependent information is like menu-dependent (expected) utility

Bayesian Probit

- Natenzon (2016) develops a Bayesian Probit model of this, where the agent observes noisy signal of the utility of each item in the menu
  - signals are jointly normal and correlated
  - model explains decoy effect, compromise effect, and similarity effects
  - correlation $\Rightarrow$ new items shed light on relative utilities of existing items
- Note: adding an item gives Blackwell-more information about the state, the state is *uncorrelated* with the menu

Example (Luce and Raiffa, 1957)

	$\tilde{U}_\omega(\textit{steak tartare})$	$\tilde{U}_\omega(\textit{chicken})$	$\tilde{U}_\omega(\textit{frog legs})$
$\omega = \textit{good chef}$	10	7	3
$\omega = \textit{bad chef}$	0	5	0

- *frog legs* provides an informative signal about the quality of the chef
  - only good chefs will attempt to make *frog legs*
  - so $\{s, c, f\}$ signals $\omega = \textit{good chef}$
  - so $\{s, c\}$ signals $\omega = \textit{bad chef}$

- this implies
  - $\rho(s, \{s, c, f\}) = 1$, $\rho(c, \{s, c, f\}) = \rho(f, \{s, c, f\}) = 0$
  - $\rho(s, \{s, c\}) = 0$, $\rho(c, \{s, c\}) = 1$ (if prior uniform)

- so here the menu is directly *correlated* with the state
  - unlike in the *fish* example where there is no correlation
  - Kamenica (2008)–model where consumers make inferences from menus (model explains choice overload and compromise effect)

Learning—recap

- Information independent of menu
  - Special case of RU (or equivalent to RU depending on the formulation)
  - More informative signals $\Rightarrow$ more randomness in choice

- Information depends on the menu
  - More general than RU (can violate Regularity)
  - Two flavors of the model:
    - more items $\Rightarrow$ more information (Natenzon, 2016)
    - correlation between menu and state (Kamenica, 2008)
  - General analysis? Axioms?

Introduction

Random Utility/Discrete Choice
- Properties of RU
- Special Cases
- Random Expected Utility (REU)

Learning

Attention
- Optimal Attention
- Random Attention

Deliberate Randomization

Optimal Attention

- Imagine now that the signal structure is chosen by the agent
  - instead of being fixed
- The agent may want to choose to focus on some aspect
  - depending on the menu
- One way to model this margin of choice is to let the agent choose attention optimally:
  - *Costly Information Acquisition* (Raiffa and Schlaifer, 1961)
  - *Rational Inattention* (Sims, 2003)
  - *Costly Contemplation* (Ergin, 2003; Ergin and Sarver, 2010)

Value of Information

For each information structure μ its value to the agent is

$$V(\mu) = \sum_{\tilde{q} \in \Delta(S)} [\max_{x \in A} \mathbb{E}_{\tilde{q}} v(x)] \mu(q)$$

Comment: Blackwell's theorem says the value of information is always positive: more information is better

Optimal Attention

- For every menu A, the agent chooses μ to maximize:

$$\max_{\mu} \; V(\mu) - C(\mu)$$

- where $C(\mu)$ is the cost of choosing the signal structure μ
  - could be a physical cost
  - or mental/cognitive
- this is another case where information depends on the menu A
  - this time endogenously

Optimal Attention

Special cases of the cost function:

- Mutual information: $C(\mu) = \sum_{q \in \Delta(S)} \phi^{KL}(q)\mu(q)$

  where $\phi^{KL}(q)$ is the relative entropy of q with respect to the prior

- Separable cost functions $C(\mu) = \sum_{q \in \Delta(S)} \phi(q)\mu(q)$

  for some general ϕ

- Neighborhood-based cost functions (Hébert and Woodford, 2017)

- General cost functions: C is just Blackwell-monotone and convex

Questions:

- Is it harder to distinguish "nearby" states than "far away" states?
  - Caplin and Dean (2013), Morris and Yang (2016), Hébert and Woodford (2017)

Optimal Attention

- Matejka and McKay (2014) analyze the mutual information cost function used in Sims (2003)
  - they show the optimal solution leads to weighted-Luce choice probabilities ρ^s
  - can be characterized by two Luce IIA-like axioms on ρ^s
  - demonstrate a violation of Regularity

Example (Matejka and McKay, 2014): $\rho(x, \{x, y, z\}) > \rho(x, \{x, y\})$ because adding z adds incentive to learn about the state

	s_1	s_2
x	0	1
y	0.5	0.5
z	k	$-k$

Optimal Attention

- Caplin and Dean (2015) characterize general cost C
  - assume choice is between Savage acts
  - assume the analyst knows the agent's utility function and the prior
  - can be characterized by two acyclicity-like axioms on ρ^s
  - partial uniqueness: bounds on the cost function
- Lin (2017) also characterizes general cost C
  - building on Lu (2016) and De Oliveira, Denti, Mihm, and Ozbek (2016)
  - the utility and prior are recovered from the data
  - can be characterized by a relaxation of REU axioms plus the De Oliveira, Denti, Mihm, and Ozbek (2016) axioms
  - essential uniqueness of parameters: minimal cost function unique

Introduction

Random Utility/Discrete Choice
Properties of RU
Special Cases
Random Expected Utility (REU)

Learning

Attention
Optimal Attention
Random Attention

Deliberate Randomization

Random Attention

- In the Optimal Attention model, paying attention meant optimally choosing an informative signal about its utility (at a cost)

- In the Random Attention model, attention is exogenous (and random)
  - $\tilde{\Gamma}(A) \subseteq A$ is a random *Consideration Set*
  - $v \in \mathbb{R}^X$ is a deterministic utility function
  - for each possible realization $\tilde{\Gamma}(A)$ the agent maximizes v on $\tilde{\Gamma}(A)$
  - so for each menu we get a probability distribution over choices

Random Attention

- Manzini and Mariotti (2014)
  - each $x \in A$ belongs to $\tilde{\Gamma}(A)$ with prob $\gamma(x)$, independently over x
  - if $\tilde{\Gamma}(A) = \emptyset$, the agent chooses a default option
  - axiomatic characterization, uniqueness result
    - turns out this is a special case of RU
- Brady and Rehbeck (2016)
  - allow for correlation
  - axiomatic characterization, uniqueness result
    - now can violate Regularity
- Cattaneo and Masatlioglu (2017)
  - an even more general model of *attention filters*, following Masatlioglu, Nakajima, and Ozbay (2011)
  - axiomatic characterization, uniqueness result

Introduction

Random Utility/Discrete Choice
Properties of RU
Special Cases
Random Expected Utility (REU)

Learning

Attention
Optimal Attention
Random Attention

Deliberate Randomization

Deliberate Randomization

Idea: The agent directly chooses a probability distribution on actions $\rho \in \Delta(A)$ to maximize some non-linear value function $V(\rho)$

Examples:

- Allais-style lottery preferences
- Implementation Costs
- Hedging against ambiguity
- Regret minimization

Allais-style lottery preferences

- Agent is choosing between lotteries, $X = \Delta(Z)$
- She has a deterministic nonlinear lottery preference $\succsim^\ell$ over $\Delta(Z)$
- If $\succsim^\ell$ is quasiconcave, then the agent likes to toss a "mental coin"

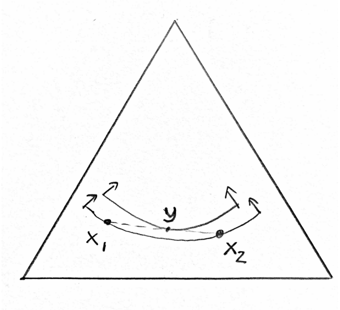

- Example: $x_1 \sim^\ell x_2$
- Strictly prefer y
- To implement this, choice from $A = \{x_1, x_2\}$ is $\rho(x_1, A) = \rho(x_2, A) = \frac{1}{2}$
- what if $B = \{x_1, x_2, y\}$? (Is the "mental coin" better or worse than actual coin?)

Allais-style lottery preferences

- Machina (1985): derives some necessary axioms that follow from maximizing any general $\succsim^\ell$

- Cerreia-Vioglio, Dillenberger, Ortoleva, and Riella (2017):
  - characterize maximization of a general $\succsim^\ell$ $\longrightarrow$ Rational Mixing axiom
  - characterize maximization of a specific $\succsim^\ell$ that belongs to the Cautious Expected Utility class $\longrightarrow$ Rational Mixing + additional axioms

- Other classes of risk preferences $\succsim^\ell$?

Implementation Costs

Idea: The agent implements her choices with an error (trembling hands)

- can reduce error at a cost that depends on the tremble probabilities

- When presented with a menu A choose $\rho \in \Delta(A)$ to maximize

$$V(\rho) = \sum_x v(x)\rho(x) - C(\rho)$$

- $v \in \mathbb{R}^X$ is a deterministic utility function
- C is the cost of implementing ρ
  - zero for the uniform distribution
  - higher as ρ focuses on a particular outcome

- This is called the Perturbed Utility model, used in game theory

Additive Perturbed Utility

Typically used specification: Additive Perturbed Utility

$$C(\rho) = \eta \sum_{x \in A} c(\rho(x))$$

- log cost: $c(t) = -\log(t)$ (Harsanyi, 1973)
- quadratic cost: $c(t) = t^2$ (Rosenthal, 1989)
- entropy cost: $c(t) = t \log t$ (Fudenberg and Levine, 1995),

General C function used in

- Mattsson and Weibull (2002), Hofbauer and Sandholm (2002), van Damme and Weibull (2002)

The Triple Equivalence

Theorem (Anderson, de Palma, and Thisse, 1992): The following are equivalent

(i) ρ has a Luce representation

(ii) ρ has a logit representation

(iii) ρ has an entropy APU representation

Comments:

- Another application to game theory: Quantal Response Equilibrium (McKelvey and Palfrey, 1995, 1998) uses logit

Additive Perturbed Utility

Theorem† (Fudenberg, Iijima, and Strzalecki, 2015): The following are equivalent under Positivity:

(i) ρ has an APU representation

(ii) ρ satisfies Acyclicity†

(iii) ρ satisfies Ordinal IIA†

Comments:

- Weaker forms of Acyclicity if c is allowed to depend on A or on z (Clark, 1990; Fudenberg, Iijima, and Strzalecki, 2014)
- The model explains any ρ if c is allowed to depend on both A and z
- Hedging against ambiguity interpretation (Fudenberg, Iijima, and Strzalecki, 2015)

Evidence

- In experiments (Agranov and Ortoleva, 2017; Dwenger, Kubler, and Weizsacker, 2013) subjects are willing to pay money for an "objective" coin toss

- So "objective" coin better than "mental" coin

- No room in above models for this distinction...

Summary

- Models so far
  - Random Utility
  - Learning
  - Attention
  - Deliberate Randomization
- Lecture 2 uses these as building blocks to study dynamic choices

Lecture 2 on Stochastic Choice

Tomasz Strzalecki

Hotelling Lectures in Economic Theory
Econometric Society European Meeting, Lisbon, August 25, 2017

Plan

Purpose: Overview where the field is and where it seems to be going

Lecture 1: Static Choice

- Random Utility (and Discrete Choice)
- Learning, Attention, Deliberate Randomization

Lecture 2: Dynamic Choice

- Dynamic Random Utility
- Dynamic Discrete Choice
- Drift-Diffusion Models

Basic Model (with two periods)

- In period $t = 0$ have ρ_0 with a RU representation with utility $\tilde{U}_0(x_0)$
- In period $t = 1$ have ρ_1 with a RU representation with utility $\tilde{U}_1(x_1)$
- $\tilde{U}_0$ and $\tilde{U}_1$ are possibly correlated
  - preferences are somewhat stable over time
  - "persistent types" in dynamic games/mechanism design/taxation

Question: What does this assume about behavior?

Answer: Selection on Unobservables/History Dependence

History Dependence

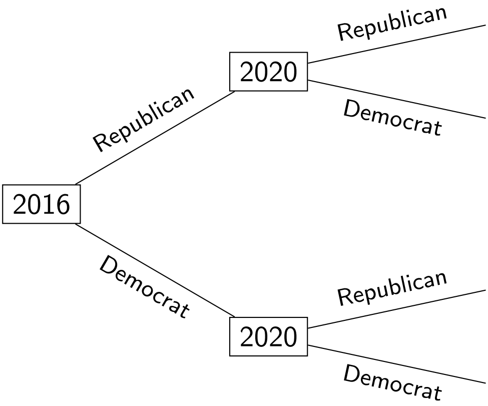

If political preferences persistent over time, expect history dependence:

$$\rho(R_{2020}|R_{2016}) > \rho(R_{2020}|D_{2016})$$

History independence only if preferences completely independent over time

Types of History Dependence (Heckman, 1981)

1. **Choice Dependence**: A consequence of the informational asymmetry between the analyst and the agent
   - Selection on unobservables
   - Utility is serially correlated (past choices partially reveal it)

2. **Consumption Dependence**: Past consumption changes the state of the agent
   - Habit formation or preference for variety (preferences change)
   - Experimentation (beliefs change)

Questions:

- How to put this into the model?
- What happens if we ignore this?
- How to distinguish between the two?

Decision Trees: $x_t = (z_t, A_{t+1})$

– Choice today leads to an immediate payoff and a menu for tomorrow

Examples:

– fertility and schooling choices (Todd and Wolpin, 2006)

– engine replacement (Rust, 1987)

– patent renewal (Pakes, 1986)

– occupational choices (Miller, 1984)

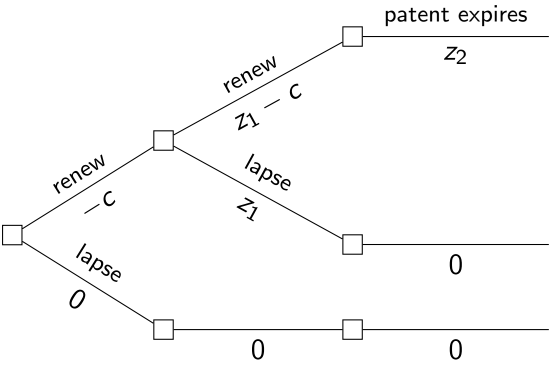

Primitive

- The analyst observes the conditional choice probabilities $\rho_t(\cdot|h_{t-1})$
  - at each node of a decision tree
- Dynamic Discrete Choice literature
  - typically for a fixed tree
- Decision Theory literature
  - typically across decision trees

Full model (with two periods)

In addition, it is often assumed that:

- In period 0 the agent's utility is

$$\tilde{U}_0(z_0, A_1) = \tilde{u}_0(z_0) + \delta\mathbb{E}_0\left[\max_{z_1 \in A_1} \tilde{u}_1(z_1)\right]$$

- $\tilde{u}_0$ is private information in $t = 0$
- $\tilde{u}_1$ is private information in $t = 1$ (so may be unknown in $t = 0$)

Question: What do these additional assumptions mean?

Introduction

Dynamic Random Utility

Dynamic Discrete Choice

Decision Times

Decision Trees

Time: $t = 0, 1$

Per-period outcomes: Z

Decision Nodes: $\mathcal{A}_t$ defined recursively:

- period 1: menu A_1 is a subset of $X_1 := Z$
- period 0: menu A_0 is a subset of $X_0 := Z \times \mathcal{A}_1$

pairs $x_0 = (z_0, A_1)$ of current outcome and continuation menu

Comment: Everything extends to finite horizon by backward induction; infinite horizon—need more technical conditions (a construction similar to universal type spaces)

Conditional Choice Probabilities

ρ is a sequence of **history-dependent** choice distributions:

period 0: for each menu A_0, observe choice distribution

$$\rho_0(\cdot, A_0) \in \Delta(A_0)$$

period 1: for each menu A_1 and history h^0 that leads to menu A_1, observe choice distribution conditional on h^0

$$\rho_1(\cdot, A_1 | h^0) \in \Delta(A_1)$$

$\mathcal{H}_0 \cdots\cdots\cdots\cdots$ period-0 histories

$$\mathcal{H}_0 := \{h^0 = (A_0, x_0) : \rho_0(x_0, A_0) > 0\}$$

$\mathcal{H}_0(A_1) \cdots\cdots\cdots$ is set of histories that lead to menu A_1

$$\mathcal{H}_0(A_1) := \{h^0 = (A_0, x_0) \in \mathcal{H}_0 : x_0 = (z_0, A_1) \text{ for some } z_0 \in Z\}$$

Dynamic Random Utility

Definition: A *DRU* representation of ρ consists of

- a probability space $(\Omega, \mathcal{F}, \mathbb{P})$
- a stochastic process of utilities $\tilde{U}_t : \Omega \to \mathbb{R}^{X_t}$

such that for all $x_0 \in A_0$

$$\rho_0(x_0, A_0) = \mathbb{P}\left[C(x_0, A_0)\right]$$

and for all $x_1 \in A_1$ and histories $h^0 \in \mathcal{H}_0(A_1)$,

$$\rho_1(x_1, A_1 | h^0) = \mathbb{P}\left[C(x_1, A_1) | C(h^0)\right]$$

where $C(x_t, A_t) := \{\omega \in \Omega : \tilde{U}_{t,\omega}(x_t) \geq \tilde{U}_{t,\omega}(y_t) \text{ for all } y_t \in A_t\}$

- for technical reasons allow for ties and use tie-breaking

History Independence

General idea:

- Agent's choice history $h^0 = (A_0, x_0)$ reveals something about his period-0 private information, so expect $\rho_1(\cdot|h^0)$ to depend on h^0
- But dependence cannot be arbitrary: some histories are *equivalent* as far as the private information they reveal
- The axioms of Frick, Iijima, and Strzalecki (2017)
  - Identify two types of equivalence classes of histories
  - Impose history *independence* of ρ_1 within these classes

Contraction History Independence

Definition: History (A_0, x_0) is *contraction equivalent* to (B_0, x_0) if

(i) $A_0 \subseteq B_0$

(ii) $\rho_0(x_0, A_0) = \rho_0(x_0, B_0)$

Axiom (Contraction History Independence): If (A_0, x_0) is contraction equivalent to (B_0, x_0), then

$$\rho_1(\cdot; \cdot | A_0, x_0) = \rho_1(\cdot; \cdot | B_0, x_0)$$

Example

- 2 convenience stores (A & B)
- Stable set of weekly customers with identical preferences
- Week 0 market shares:

milk type	share at A
whole	40%
2%	60%

milk type	share at B
whole	40%
2%	35%
1%	25%

- Alice and Bob buy whole milk in week 0
- **Claim:** If in week 1 all types of milk available at both stores, expect Alice and Bob's choice probabilities to be the same

Example

Why?

- We have same information about Alice and Bob:

- Possible week-0 preferences:
  - **Alice:** $w \succ 2 \succ 1$ or $w \succ 1 \succ 2$ or $1 \succ w \succ 2$
  - **Bob:** $w \succ 2 \succ 1$ or $w \succ 1 \succ 2$

- So in principle learn more about Bob

- But condition *(ii)* of the axiom says

$$\rho_0(w|\{w, 1, 2\}) = \rho_0(w|\{w, 2\}) = 0.4$$

so no customers have ranking $1 \succ w \succ 2$!

Adding Lotteries

Add lotteries: $X_t = \Delta(Z \times \mathcal{A}_{t+1})$, assume each utility function is vNM

- Helps formulate the second kind of history-independence
- Makes it easy to build on the REU axiomatization
- Helps overcome the limited observability problem
  - not all menus observed after a given history; how to impose axioms?
- Helps distinguish choice-dependence from consumption-dependence

$$h^0 = (A_0, x_0) \quad \text{vs} \quad h^0 = (A_0, x_0, z_0)$$

Dynamic Random Expected Utility

First, assume away consumption dependence and allow only for choice dependence

$$\rho_1(\cdot|(A_0, x_0, z_0)) = \rho_1(\cdot|(A_0, x_0, z_0'))$$

Theorem‡ (Frick, Iijima, and Strzalecki, 2017): ρ has a DREU representation if and only it satisfies

- REU axioms in each period
- Contraction History Independence
- Linear History Independence†
- History-Continuity†

How to incorporate Dynamic Optimality?

- In the definition above, no structure on the family $(\tilde{U}_t)$

Definition: ρ has an *Evolving Utility* representation if it has a DREU representation where the process $(\tilde{U}_t)$ satisfies the Bellman equation

$$\tilde{U}_t(z_t, A_{t+1}) = \tilde{u}_t(z_t) + \delta \mathbb{E}\left[\max_{x_{t+1} \in A_{t+1}} \tilde{U}_{t+1}(x_{t+1}) | \mathcal{G}_t\right]$$

for $\delta > 0$ and a $\mathcal{G}_t$-adapted process of vNM utilities $\tilde{u}_t : \Omega \to \mathbb{R}^Z$

Question: What are the additional assumptions?

Answer:

- Option value calculation (Preference for Flexibility)
- Rational Expectations (Sophistication)

Simplifying assumption: No selection

Simplifying Assumption:

1. The payoff in $t = 0$ is fixed
2. There is no private information in $t = 0$

What this means:

- Choices in $t = 0$:
  - are deterministic
  - can be represented by a preference $A_1 \succsim_0 B_1$
- Choices in $t = 1$:
  - are random, represented by ρ_1
  - are history-independent
    - $t = 0$ choices do not reveal any information

Preference for Flexibility

Definition: $\succsim_0$ has an *option-value representation* if there exists a random $u_1 : \Omega \to \mathbb{R}^Z$ such that

$$U_0(A_1) = \mathbb{E}_0 \left[\max_{z_1 \in A_1} \tilde{u}_1(z_1) \right]$$

Axiom (Preference for Flexibility): If $A \supseteq B$, then $A \succsim_0 B$

Theorem† (Kreps, 1979): $\succsim_0$ has an option-value representation iff it satisfies Completeness, Transitivity, Preference for Flexibility, and Modularity†

Comments:

- In econometrics U_0 is called the *consumer surplus*
- To improve the uniqueness properties, Dekel, Lipman, and Rustichini (2001); Dekel, Lipman, Rustichini, and Sarver (2007) specialize to choice between lotteries, $X_1 = \Delta(Z_1)$

Rational Expectations

Specify to $X_1 = \Delta(Z_1)$ and suppose that

- $\succsim_0$ has an option-value representation $(\Omega, \mathcal{F}, \mathbb{P}_0, u_1)$
- ρ_1 has a REU representation with $(\Omega, \mathcal{F}, \mathbb{P}_1, u_1)$

Definition: $(\succsim_0, \rho_1)$ has *Rational Expectations* iff $\mathbb{P}_0 = \mathbb{P}_1$

Axiom (Sophistication)†: For any is a menu without ties† $A \cup \{x\}$

$$A \cup \{x\} \succ_0 A \Longleftrightarrow \rho_1(x, A \cup \{x\}) > 0$$

Theorem‡ (Ahn and Sarver, 2013): $(\succsim_0, \rho_1)$ has Rational Expectations iff it satisfies Sophistication.

Comment: Relaxed Sophistication ($\Rightarrow$ or $\Leftarrow$) pins down either an *unforeseen contingencies* model or a *pure freedom of choice* model

Identification of Beliefs

Theorem‡ (Ahn and Sarver, 2013): If $(\succsim_0, \rho_1)$ has Rational Expectations, then the distribution over cardinal utilities u_1 is uniquely identified.

Comments:

- Just looking at ρ_1 only identifies the distribution over ordinal risk preferences (Gul and Pesendorfer, 2006)
- Just looking at $\succsim_0$ identifies even less (Dekel, Lipman, and Rustichini, 2001)
- But jointly looking at the evaluation of a menu and the choice from the menu helps with the identification

Putting Selection Back In

- In general, want to relax the simplifying assumption
  - in reality there are intermediate payoffs
  - and informational asymmetry in each period
  - choice is stochastic in each period
  - and there is history dependence

- To characterize the evolving utility model need to add Preference for Flexibility and Sophistication
  - but those are expressed in terms of $\succsim_0$
  - when the simplifying assumption is violated we only have access to ρ_0
  - Frick, Iijima, and Strzalecki (2017) find a way to extract $\succsim_0$ from ρ_0

Passive and Active Learning

- Evolving Utility: randomness in choice comes from changing tastes

- Passive Learning: randomness in choice comes from random signals
  - tastes are time-invariant, but unknown $\tilde{u}_t = \mathbb{E}[\tilde{u}|\mathcal{G}_t]$ for some time-invariant vNM utility $\tilde{u} : \Omega \to \mathbb{R}^Z$

- To characterize the passive learning model, need to add a "martingale" axiom

- The paper also relaxes consumption-independence and characterizes habit-formation and active learning (experimentation) models
  - parametric models of active learning used by, e.g., Erdem and Keane (1996), Crawford and Shum (2005)

- Uniqueness of the utility process, discount factor, and information

Related Work

- The Bayesian probit model Natenzon (2016) can be viewed as a model of a sequence of static choice problems where choice probabilities are time dependent

- Cerreia-Vioglio, Maccheroni, Marinacci, and Rustichini (2017) also study a sequence of static choice problems using a Luce-like model

- Gabaix and Laibson (2017) use a model of gradual learning to microfound "as-if" discounting and present bias

- Lu and Saito (2016) study $t = 0$ choices between consumption stream

- Krishna and Sadowski (2012, 2016) characterize a class of models similar to Evolving Utility by looking at menu-preferences

Preference for making choices late

- Positive value of information: desire to delay the choice as late as possible to capitalize on incoming information (unless there is a cost)

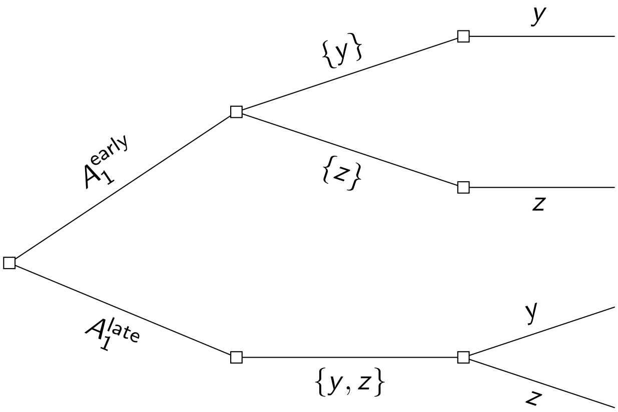

Theorem†: If ρ has an Evolving Utility represetation, then absent ties†

$$\rho_0(A_1^{\text{late}}, \{A_1^{\text{early}}, A_1^{\text{late}}\}) = 1$$

Introduction

Dynamic Random Utility

Dynamic Discrete Choice

Decision Times

DDC model

DDC: There is a process of shocks $\tilde{\epsilon}_t : \Omega \to \mathbb{R}^{X_t}$ s.t.

$$V_t(z_t, A_{t+1}) = \left(v_t(z_t) + \delta \mathbb{E} \left[\max_{x_{t+1} \in A_{t+1}} V_{t+1}(x_{t+1}) | \mathcal{G}_t \right] \right) + \tilde{\epsilon}_t^{(z_t, A_{t+1})}$$

where

- v_t are deterministic
- $\mathcal{G}_t$ is generated by $\tilde{\epsilon}_t$

Special cases of DDC

- $\tilde{\epsilon}_t^{(z_t, A_{t+1})}$ and $\tilde{\epsilon}_\tau^{(y_t, B_{t+1})}$ are i.i.d.
  - **shocks to actions**
  - I will also refer to it as i.i.d. DDC
  - ρ is history independent

- $\tilde{\epsilon}_t^{(z_t, A_{t+1})} = \tilde{\epsilon}_t^{(z_t, B_{t+1})} =: \tilde{\epsilon}_t^{z_t}$
  - **shocks to payoffs**
  - allows for serial correlation of $\tilde{\epsilon}_t$
  - ρ is a special case of evolving utility

Dynamic logit

- A special case of i.i.d. DDC where $\tilde{\epsilon}_t$ are distributed extreme value
- Very tractable due to the "log-sum" expression for "consumer surplus"

$$V_t(A_{t+1}) = \log \left(\sum_{x_{t+1} \in A_{t+1}} e^{v_{t+1}(x_{t+1})} \right)$$

- (This formula is also the reason why nested logit is so tractable)
- Dynamic logit is a workhorse for estimation
  - e.g., Miller (1984), Rust (1989), Hendel and Nevo (2006), Gowrisankaran and Rysman (2012)

Axiomatization (Fudenberg and Strzalecki, 2015)

Axiom (Recursivity):

$$\rho_t((z_t, A_{t+1}), \{(z_t, A_{t+1}), (z_t, B_{t+1})\}) \geq \rho_t((z_t, B_{t+1}), \{(z_t, A_{t+1}), (z_t, B_{t+1})\})$$

$$\Updownarrow$$

$$\rho_{t+1}(A_{t+1}, A_{t+1} \cup B_{t+1}) \geq \rho_{t+1}(B_{t+1}, A_{t+1} \cup B_{t+1})$$

Axiom (Weak Preference for Flexibility): If $A_{t+1} \supseteq B_{t+1}$, then

$$\rho_t((z_t, A_{t+1}), \{(z_t, A_{t+1}), (z_t, B_{t+1})\}) \geq \rho_t((z_t, B_{t+1}), \{(z_t, A_{t+1}), (z_t, B_{t+1})\})$$

Comments:

- Recursivity leverages the "log-sum" expression
- Preference for flexibility is weak because support of $\tilde{\epsilon}_t$ is unbounded
- Also, identification results, including uniqueness of δ

Models that build on Dynamic Logit

- View $\tilde{\epsilon}_t$ as errors, not utility shocks
  - Fudenberg and Strzalecki (2015): errors lead to "choice aversion"
  - Ke (2016): a dynamic model of mistakes
- Dynamic attribute rule
  - Gul, Natenzon, and Pesendorfer (2014)

Questions about DDC

- Characterization of the general i.i.d. DDC model? General DDC?
  - In general, no formula for the "consumer surplus", but the theorem of Williams–Daly–Zachary says that the choice probabilities are the derivative of the "social surplus" (Chiong, Galichon, and Shum, 2016)
  - It is an envelope-theorem result, like the Hotelling lemma
  - It ties together choices in different time periods so conceptually related to Sophistication, Recursivity, and the axiom of Lu (2016)
- There is a vast DDC literature on identification (Manski, 1993; Rust, 1994; Magnac and Thesmar, 2002; Norets and Tang, 2013)
  - δ not identified unless make assumptions about "observable attributes"
  - How does this compare to the "menu variation" approach

Properties of i.i.d. DDC

Key Assumption: Shocks to actions, $\tilde{\epsilon}_t^{(z_t, A_{t+1})}$ and $\tilde{\epsilon}_t^{(z_t, B_{t+1})}$ are i.i.d. regardless of the nature of the menus A_{t+1} and B_{t+1}

Theorem (Fudenberg and Strzalecki, 2015; Frick, Iijima, and Strzalecki, 2017): If ρ has a i.i.d. DDC representation with $\delta < 1$, then $\rho_0(A_1^{\text{late}}, \{A_1^{\text{early}}, A_1^{\text{late}}\}) < \frac{1}{2}$

Intuition:

- The agent gets the ϵ not at the time of consumption but at the time of decision (even if the decision has only delayed consequences)
- So making decisions early allows him to get the $\max \epsilon$ earlier

Question: How much does this result extend beyond i.i.d. ?

- Mixture models: Kasahara and Shimotsu (2009)

Modeling Choices

- DRU: so far few convenient parametrization (Pakes, 1986) but
  - bigger menus w/prob. 1
  - late decisions w/prob. 1

- i.i.d. DDC: statistical tractability, but
  - bigger menus w/prob. $\in (\frac{1}{2}, 1)$
  - late decisions w/prob. $\in (0, \frac{1}{2})$

Comments:

- i.i.d. DDC violates a key feature of Bayesian rationality: positive option value

- Does this mean the model is misspecified?
  - Maybe not as a model of (potentially behavioral) consumers
  - But what about profit maximizing firms?
    - biased parameter estimates?

Modeling Choices

Comments:

- Note that in the static setting i.i.d. DC is a special case of RU
  - though it has its own problems (blue bus/red bus)

- But in the dynamic setting, i.i.d. DDC is outside of DRU!

- Negative option value is not a consequence of history independence
  - no such problem in the Evolving Utility model with i.i.d utility

- It is a consequence of shocks to actions vs shocks to payoffs

Introduction

Dynamic Random Utility

Dynamic Discrete Choice

Decision Times

Decision Times

New Variable: How long does the agent take to decide?

Time: $\mathcal{T} = [0, \infty)$ or $\mathcal{T} = \{0, 1, 2, \ldots\}$

Observe: Joint distribution $\rho \in \Delta(A \times \mathcal{T})$

Question:

- Are fast decisions “better” or “worse” than slow ones?

Decision Times

Intuitions:

- More time $\Rightarrow$ more information $\Rightarrow$ better decisions
- But time is costly, so speed-accuracy tradeoff
  - Want to stop early if get an informative signal $\rightarrow$ selection effect

Comment: These two effects push in opposite directions. Which one wins?

Stylized fact: Decreasing accuracy: fast decisions are "better"

- Well established in perceptual tasks, where "better" is objective
- Also in experiments where subjects choose between consumption items

Experiment of Krajbich, Armel, and Rangel (2010)

- X: 70 different food items
- Step 1: Rate each $x \in X$ on the scale -10, . . . , 10
- Step 2: Choose from $A = \{x, y\}$ (100 different pairs)
  - record choice and decision time
- Step 3: Draw a random pair and get your choice

Decreasing Accuracy

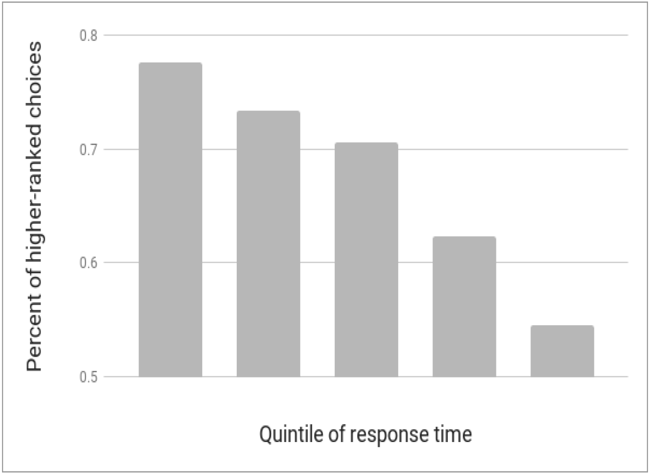

(based on data from Krajbich, Armel, and Rangel, 2010)

Model

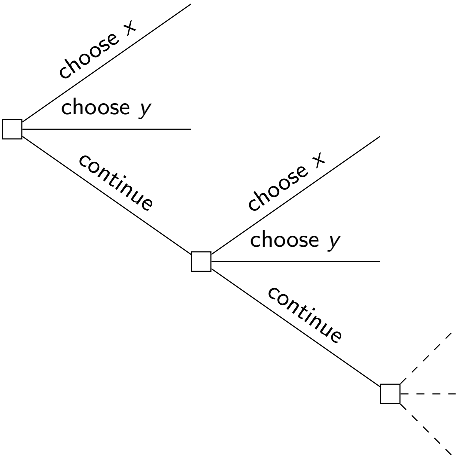

Model

- $S \cdots$ set of unknown states
- $p \in \Delta(S) \cdots$ prior belief
- $v : S \to \mathbb{R}^X \cdots$ state-dependent utility function
- $(\mathcal{G}_t) \cdots$ information of the agent (filtration)
- $\tau \cdots$ stopping time (with respect to $\mathcal{G}_t$)
- Conditional on stopping, the agent maximizes expected utility

$$\text{choice}_\tau = \text{argmax}_{x \in A} \mathbb{E}[v(x)|\mathcal{G}_\tau]$$

- So the only problem is to choose the stopping time

Optimal Stopping Problem

The agent chooses the stopping time optimally

$$\max_{\tau} \mathbb{E}[v(\text{choice}_\tau)] - C(\tau)$$

Comments:

- $(\mathcal{G}_t)$ and τ generate a joint distribution of choices and times
  - conditional on the state $\rho^s \in \Delta(A \times \mathcal{T})$
  - unconditional (averaged out according to p) $\rho \in \Delta(A \times \mathcal{T})$
- Even though $(\mathcal{G}_t)$ is fixed, there is an element of optimal attention
  - Waiting longer gives more information at a cost
  - Choosing τ is like choosing the distribution over posteriors μ
  - How close is this to the static model of optimal attention?

Optimal Stopping: Further Assumptions

- Continuous time, linear cost $C(t) = ct$
- Binary choice $A = \{x, y\}$
- $s = (u(x), u(y)) \in \mathbb{R}^2$
- Signal: $\mathcal{G}_t$ is generated by (G_t^x, G_t^y) where

$$G_t^x = t \cdot u(x) + B_t^x \qquad \text{and} \qquad G_t^y = t \cdot u(y) + B_t^y$$

  and B_t^x, B_t^y are Brownian motions; often look at $G_t := G_t^x - G_t^y$

- Two classes of priors:
  - Binary: the state is $(\lambda, 0)$ or $(0, \lambda)$ with equal probabilities
  - Normal (independent):

$$u(x) \sim N(\mu_0, \sigma_0^2) \qquad \text{and} \qquad u(y) \sim N(\mu_0, \sigma_0^2)$$

Binary Prior

Theorem (Wald, 1945): With binary prior the optimal strategy in the stopping model takes a boundary-hitting form: there exists $b \geq 0$ such that

$$\tau := \inf\{t \geq 0 : |G_t| \geq b\} \qquad \text{choice}_\tau := \begin{cases} x & \text{if} \quad G_\tau = b \\ y & \text{if} \quad G_\tau = -b \end{cases}$$

Binary Prior

Theorem (Wald, 1945): With binary prior the optimal strategy in the stopping model takes a boundary-hitting form: there exists $b \geq 0$ such that

$$\tau := \inf\{t \geq 0 : |G_t| \geq b\} \qquad \text{choice}_\tau := \begin{cases} x & \text{if} \quad G_\tau = b \\ y & \text{if} \quad G_\tau = -b \end{cases}$$

Binary Prior

Theorem (Wald, 1945): With binary prior the optimal strategy in the stopping model takes a boundary-hitting form: there exists $b \geq 0$ such that

$$\tau := \inf\{t \geq 0 : |G_t| \geq b\} \qquad \text{choice}_\tau := \begin{cases} x & \text{if} \quad G_\tau = b \\ y & \text{if} \quad G_\tau = -b \end{cases}$$

Drift-Diffusion Models

- The boundary-hitting model is called a Drift-Diffusion Model
- Most often used as a reduced-form model
  - No optimization problem, just a boundary-hitting exercise
- Brought to the psychology literature by Stone (1960) and Edwards (1965) to study perception; memory retrieval (Ratcliff, 1978)
- Closed-form solutions for choice probabilities (logit) and expected decision time

Accuracy

Definition: *Accuracy* is the probability of making the correct choice

$$\alpha(t) := \mathbb{P}\left[\text{choice}(\tau) = \text{argmax}_{x \in A}\, \tilde{u}(x) | \tau = t\right]$$

Problem: In DDM $\alpha(t)$ is constant in t, so the model does not explain the stylized fact

Intuition:

- Unconditional on stopping:
  - higher $t \Rightarrow$ more information $\Rightarrow$ better accuracy
- But t is not chosen at random: it depends on information
  - stop early after informative signals
- The two effects balance each other out perfectly!

Drift-Diffusion Models

Many ad-hoc extensions, in particular time-varying boundary $b(t)$

$$\tau := \inf\{t \geq 0 : |G_t| \geq b(t)\} \qquad \text{choice}(\tau) := \begin{cases} x & \text{if} \quad G_\tau = b(\tau) \\ y & \text{if} \quad G_\tau = -b(\tau) \end{cases}$$

Theorem (Fudenberg, Strack, and Strzalecki, 2017): Conditional on state

$$\text{accuracy } \alpha \text{ is } \begin{Bmatrix} \text{increasing} \\ \text{decreasing} \\ \text{constant} \end{Bmatrix} \text{ iff boundary } b \text{ is } \begin{Bmatrix} \text{increasing} \\ \text{decreasing} \\ \text{constant} \end{Bmatrix}$$

Comment: Unconditional on state—analogous relation but need to look at a different monotonicity condition on b

Normal Prior

Question: How to microfound such non-constant boundaries? Do they correspond to any particular optimization problem?

Theorem (Fudenberg, Strack, and Strzalecki, 2017): In the Normal optimal stopping problem the optimal behavior leads to decreasing accuracy (unconditional on state)

Intuition: Decreasing Boundary. Suppose $G_t^x \approx G_t^y$ after a long t

- With a binary prior agent thinks: "the signal must have been noisy"
  - so she doesn't learn anything $\Rightarrow$ she continues
- With a Normal prior agent thinks: "I must be indifferent"
  - so she learned a lot $\Rightarrow$ she stops

Other Boundaries

Question: How to microfound other non-constant boundaries? Do they correspond to any particular optimization problem?

Theorem‡ (Fudenberg, Strack, and Strzalecki, 2017): For any b there exists a (nonlinear) cost function C such that b is the optimal solution to the stopping problem

Optimal Attention

- Pure optimal stopping problem (given a fixed $(\mathcal{G}_t)$, choose τ):

$$\max_{\tau} \mathbb{E}\left[\max_{x \in A} \mathbb{E}[\tilde{u}(x)|\mathcal{G}_\tau]\right] - C(\tau)$$

- Pure optimal attention (given a fixed τ, choose $(\mathcal{G}_t)$)

$$\max_{(\mathcal{G}_t)} \mathbb{E}\left[\max_{x \in A} \mathbb{E}[\tilde{u}(x)|\mathcal{G}_\tau]\right] - C(\mathcal{G}_t)$$

- Joint optimization

$$\max_{\tau,(\mathcal{G}_t)} \mathbb{E}\left[\max_{x \in A} \mathbb{E}[\tilde{u}(x)|\mathcal{G}_\tau]\right] - C(\tau, \mathcal{G}_t)$$

Optimal Attention

- In the pure optimal attention problem information choice is more flexible than in the pure stopping problem
  - The agent can focus on one item, depending on what she learned so far
- Woodford (2014) solves a pure optimal attention problem
  - with a constant boundary
  - shows that optimal behavior leads to a decreasing choice accuracy
- Joint optimization puts the two effects together
- In experiments eye movements are often recorded (Krajbich, Armel, and Rangel, 2010; Krajbich and Rangel, 2011; Krajbich, Lu, Camerer, and Rangel, 2012)
  - Do the optimal attention models predict them?

Optimal Attention

- Liang, Mu, and Syrgkanis (2017) study the pure attention as well as joint optimization models
  - Find conditions under which the dynamically optimal strategy is close to the myopic strategy
- Che and Mierendorff (2016) study the joint optimization problem in a Poisson environment with two states; find that coexistence of two strategies is optimal:
  - Contradictory strategy that seeks to challenge the prior
  - Confirmatory strategy that seeks to confirm the prior

Other Models

- Ke and Villas-Boas (2016) joint optimization with two states per alternative in the diffusion environment
- Steiner, Stewart, and Matějka (2017) optimal attention with the mutual information cost and evolving (finite) state
- Branco, Sun, and Villas-Boas (2012); Ke, Shen, and Villas-Boas (2016) application to consumers searching for products
- Epstein and Ji (2017): ambiguity averse agents may never learn
- Gabaix and Laibson (2005): a model of bounded rationality

Optimal Stopping vs Optimal Attention

- In the pure optimal stopping problem $(\mathcal{G}_t)$ is fixed like in the passive learning model

- But there is an element of optimal attention
  - Waiting longer gives more information at a cost
  - Choosing τ is like choosing the distribution over posteriors μ
  - Morris and Strack (2017) show all μ can be obtained this way if $|S| = 2$

- So in a sense this boils down to a static optimal attention problem
  - With a specific cost function: Morris and Strack (2017) show that the class of such cost functions is equal to separable cost functions as long as the flow cost depends only on the current posterior

- Hébert and Woodford (2017) show a similar reduction to a static separable problem in the joint optimization problem
  - Converse to their theorem?

Other Questions

Question:

- Are “close” decisions faster or slower?

Intuitions:

- People “overthink” decision problems which don’t matter, “underthink” those with big consequences
- It is optimal to think more when options are closer (higher option value)

Experiment: Oud, Krajbich, Miller, Cheong, Botvinick, and Fehr (2014)

Other questions

Question: How does the decision time depend on the menu size?

- "Hick–Hyman Law:" the decision time increases logarithmically in the menu size
  - At least for perceptual tasks (Luce, 1986)
- Frick and Iijima (2015) introduce a model that explains the monotonic relationship (among other things)
  - The decision maker is "conflicted" about the choice
  - Different "selves" are playing a Poisson competition game

Other questions

Question: Are fast decisions impulsive/instinctive and slow deliberate/cognitive?

- Rubinstein (2007); Rand, Greene, and Nowak (2012); Krajbich, Bartling, Hare, and Fehr (2015); Caplin and Martin (2016)

Question: Use reaction times to understand how people play games?

- Costa-Gomes, Crawford, and Broseta (2001); Johnson, Camerer, Sen, and Rymon (2002); Brocas, Carrillo, Wang, and Camerer (2014)

Final Slide: Some Open Questions

- A general model of learning with menu-dependent information

- Comparison of DRU and DDC
  - Are the parameter estimates indeed biased?
  - Comparison of identification results ("menu" vs "attribute" variation)

- How to extend other static models to decision trees?
  - E.g, Random Attention, Perturbed Utility

- General analysis of DDM and related models
  - Without relying on distributional assumptions
  - Axioms?

Thank you!

Appendix: additional material

Tiebreakers

Random Utility Axioms

Stochastic Transitivity

Fechnerian models

References

Random Utility (with a tiebreaker)

- To break ties, Gul and Pesendorfer (2006) introduce a *tie-breaker* $w : \Omega \to \mathbb{R}^X$, (which is always a strict preference)
- The agent first maximizes u and if there is a tie, it gets resolved using w
- For any $v \in \mathbb{R}^X$ let $M(v, A) = \text{argmax}_{x \in A} v(x)$
- $C^{u,w}(x, A) := \{\omega \in \Omega : x \in M(w_\omega, M(u_\omega, A))\}$

Definition: ρ has a *random utility representation with a tie-breaker* if there exists $(\Omega, \mathcal{F}, \mathbb{P})$, $u, w : \Omega \to \mathbb{R}^X$ such that $\mathbb{P}(Tie^w) = 0$, and

$$\rho(x, A) = \mathbb{P}\left(C^{u,w}(x, A)\right).$$

Equivalence

Theorem: The following are equivalent:

- ρ has a random utility representation
- ρ has a random utility representation with a tiebreaker

Tiebreakers

Random Utility Axioms

Stochastic Transitivity

Fechnerian models

References

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

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

Axiom (McFadden and Richter, 1990) For any n, for any sequence $(x_1, A_1), \ldots, (x_n, A_n)$ such that $x_i \in A_i$

$$\sum_{i=1}^{n} \rho(x_i, A_i) \leq \max_{\omega \in \Omega} \sum_{i=1}^{n} \mathbb{1}_{C^{\succsim}(x_i, A_i)}(\succsim).$$

Axiom (Clark, 1996) For any n, for any sequence $(x_1, A_1), \ldots, (x_n, A_n)$ such that $x_i \in A_i$, and for any sequence of real numbers $\lambda_1, \ldots, \lambda_n$

$$\sum_{i=1}^{n} \lambda_i \mathbb{1}_{C^{\succsim}(x_i, A_i)} \geq 0 \Longrightarrow \sum_{i=1}^{n} \lambda_i \rho(x_i, A_i) \geq 0.$$

Remark: The last two axioms refer to the canonical random preference representation where Ω is the set of all strict preference relations and the mapping $\succsim$ is the identity

Characterization

Theorem: The following are equivalent

(i) ρ has a random utility representation

(ii) ρ satisfies the Block–Marschak axiom

(iii) ρ satisfies the McFadden–Richter axiom

(iv) ρ satisfies the Clark axiom.

Comments:

- The equivalence *(i)–(ii)* was proved by Falmagne (1978) and Barberá and Pattanaik (1986).
- The equivalences *(i)–(iii)* and *(i)–(iv)* were proved by McFadden and Richter (1990, 1971) and Clark (1996) respectively. They hold also when X is infinite (Clark, 1996; McFadden, 2005; Chambers and Echenique, 2016).

Tiebreakers

Random Utility Axioms

Stochastic Transitivity

Fechnerian models

References

Stochastic Preference

Definition: $x \succsim^s y$ iff $\rho(x, A) \geq \rho(y, A)$ for $A = \{x, y\}$

Comments:

- In Fechnerian models, where v is part of the representation
  - the Luce model
  - i.i.d. DC
  - APU

  the following is true $p \succsim^s q$ iff $v(p) \geq v(q)$

- In fact, in Luce we have $x \succsim^s y$ iff $\rho(x, A) \geq \rho(y, A)$ for all A
  - this characterizes the Luce model under a richness condition

(Gul, Natenzon, and Pesendorfer, 2014)

Weak Stochastic Transitivity

Definition: ρ satisfies *Weak Stochastic Transitivity* iff $\succsim^s$ is transitive

Satisfied by:

- Fechnerian models

Can be violated by:

- RU (Marschak, 1959)
- random attention
- deliberate randomization (Machina, 1985)

Stylized Fact: Weak Stochastic Transitivity is typically satisfied in lab experiments (Rieskamp, Busemeyer, and Mellers, 2006)

Forms of Stochastic Transitivity

Let $p = \rho(x, \{x, y\})$, $q = \rho(y, \{y, z\})$, $r = \rho(x, \{x, z\})$.

Definition: Suppose that $p, q \geq 0.5$. Then ρ satisfies

- *Weak Stochastic Transitivity* if $r \geq 0.5$
- *Moderate Stochastic Transitivity* if $r \geq \min\{p, q\}$
- *Strong Stochastic Transitivity* if $r \geq \max\{p, q\}$

Fechnerian Models

Definition: ρ has a *Fechnerian* representation if there exist a utility function $v : X \to \mathbb{R}$ and a strictly increasing transformation function F such that

$$\rho(x, \{x, y\}) = F(v(x) - v(y))$$

Comments:

- This property of ρ depends only on its restriction to binary menus
- The following models are Fechnerian
  - Luce
  - APU
  - i.i.d. DC
- RU in general is not Fechnerian because it violates Weak Stochastic Transitivity (Marschak, 1959)

References: Davidson and Marschak (1959); Block and Marschak (1960); Debreu (1958); Scott (1964)

References I

Agranov, M., and P. Ortoleva (2017): "Stochastic choice and preferences for randomization," *Journal of Political Economy*, 125(1), 40–68.

Ahn, D. S., and T. Sarver (2013): "Preference for flexibility and random choice," *Econometrica*, 81(1), 341–361.

Allais, M. (1953): "Le Comportment de 1'Homme Rational devant 1e Risque, Critique des Postulates et Axiomes de 1'Eco1e Americaine," *Econometrica*, 21, 803–815.

Anderson, S., A. de Palma, and J. Thisse (1992): *Discrete choice theory of product differentiation*. MIT Press.

Apesteguia, J., M. Ballester, and J. Lu (2017): "Single-Crossing Random Utility Models," *Econometrica*.

Apesteguia, J., and M. A. Ballester (2017): "Monotone Stochastic Choice Models: The Case of Risk and Time Preferences," *Journal of Political Economy*.

Ballinger, T. P., and N. T. Wilcox (1997): "Decisions, error and heterogeneity," *The Economic Journal*, 107(443), 1090–1105.

Barberá, S., and P. Pattanaik (1986): "Falmagne and the rationalizability of stochastic choices in terms of random orderings," *Econometrica*, pp. 707–715.

References II

Block, D., and J. Marschak (1960): "Random Orderings And Stochastic Theories of Responses," in *Contributions To Probability And Statistics*, ed. by I. O. et al. Stanford: Stanford University Press.

Brady, R. L., and J. Rehbeck (2016): "Menu-Dependent Stochastic Feasibility," *Econometrica*, 84(3), 1203–1223.

Branco, F., M. Sun, and J. M. Villas-Boas (2012): "Optimal search for product information," *Management Science*, 58(11), 2037–2056.

Brocas, I., J. D. Carrillo, S. W. Wang, and C. F. Camerer (2014): "Imperfect choice or imperfect attention? Understanding strategic thinking in private information games," *The Review of Economic Studies*, p. rdu001.

Caplin, A., and M. Dean (2013): "Behavioral implications of rational inattention with shannon entropy," Discussion paper, National Bureau of Economic Research.

——— (2015): "Revealed preference, rational inattention, and costly information acquisition," *The American Economic Review*, 105(7), 2183–2203.

Caplin, A., and D. Martin (2016): "The Dual-Process Drift Diffusion Model: Evidence from Response Times," *Economic Inquiry*, 54(2), 1274–1282.

Cattaneo, M., and Y. Masatlioglu (2017): "A Random Attention Model: Identification, Estimation and Testing," *mimeo*.

References III

Cerreia-Vioglio, S., D. Dillenberger, P. Ortoleva, and G. Riella (2017): "Deliberately Stochastic," *mimeo*.

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and A. Rustichini (2017): "Multinomial logit processes and preference discovery," .

Chambers, C. P., and F. Echenique (2016): *Revealed preference theory*, vol. 56. Cambridge University Press.

Che, Y.-K., and K. Mierendorff (2016): "Optimal Sequential Decision with Limited Attention," *in prepraration*.

Chiong, K. X., A. Galichon, and M. Shum (2016): "Duality in dynamic discrete-choice models," *Quantitative Economics*, 7(1), 83–115.

Clark, S. (1996): "The random utility model with an infinite choice space," *Economic Theory*, 7(1), 179–189.

Clark, S. A. (1990): "A concept of stochastic transitivity for the random utility model," *Journal of Mathematical Psychology*, 34(1), 95–108.

Costa-Gomes, M., V. P. Crawford, and B. Broseta (2001): "Cognition and behavior in normal-form games: An experimental study," *Econometrica*, 69(5), 1193–1235.

Crawford, G. S., and M. Shum (2005): "Uncertainty and learning in pharmaceutical demand," *Econometrica*, 73(4), 1137–1173.

References IV

van Damme, E., and J. Weibull (2002): "Evolution in games with endogenous mistake probabilities," *Journal of Economic Theory*, 106(2), 296–315.

Davidson, D., and J. Marschak (1959): "Experimental Tests of Stochastic Decision Theory," in *Measurement Definitions and Theories*, ed. by C. W. Churchman. John Wiley and Sons.

De Oliveira, H., T. Denti, M. Mihm, and M. K. Ozbek (2016): "Rationally inattentive preferences and hidden information costs," *Theoretical Economics*, pp. 2–14.

Debreu, G. (1958): "Stochastic Choice and Cardinal Utility," *Econometrica*, 26(3), 440444.

Debreu, G. (1960): "Review of RD Luce, Individual choice behavior: A theoretical analysis," *American Economic Review*, 50(1), 186–88.

Dekel, E., B. Lipman, and A. Rustichini (2001): "Representing preferences with a unique subjective state space," *Econometrica*, 69(4), 891–934.

Dekel, E., B. L. Lipman, A. Rustichini, and T. Sarver (2007): "Representing Preferences with a Unique Subjective State Space: A Corrigendum1," *Econometrica*, 75(2), 591–600.

Dwenger, N., D. Kubler, and G. Weizsacker (2013): "Flipping a Coin: Theory and Evidence," Discussion paper.

References V

Echenique, F., K. Saito, and G. Tserenjigmid (2013): "The Perception Adjusted Luce Model," Discussion paper.

Edwards, W. (1965): "Optimal strategies for seeking information: Models for statistics, choice reaction times, and human information processing," *Journal of Mathematical Psychology*, 2(2), 312–329.

Epstein, L. G., and S. Ji (2017): "Optimal Learning and Ellsberg's Urns," .

Erdem, T., and M. P. Keane (1996): "Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets," *Marketing science*, 15(1), 1–20.

Ergin, H. (2003): "Costly contemplation," *Unpublished paper, Department of Economics, Duke University.[22] Ergin, Haluk and Todd Sarver (2010),A unique costly contemplation representation. Econometrica*, 78, 1285–1339.

Ergin, H., and T. Sarver (2010): "A unique costly contemplation representation," *Econometrica*, 78(4), 1285–1339.

Falmagne, J. (1978): "A representation theorem for finite random scale systems," *Journal of Mathematical Psychology*, 18(1), 52–72.

Fishburn, P. C. (1998): "Stochastic Utility," *Handbook of Utility Theory: Volume 1: Principles*, p. 273.

References VI

Frick, M., and R. Iijima (2015): "A Warring Selves Model of Random Choice and Delay," .

Frick, M., R. Iijima, and T. Strzalecki (2017): "Dynamic Random Utility," *mimeo*.

Fudenberg, D., R. Iijima, and T. Strzalecki (2014): "Stochastic choice and revealed perturbed utility," *working paper version*.

——— (2015): "Stochastic choice and revealed perturbed utility," *Econometrica*, 83(6), 2371–2409.

Fudenberg, D., and D. K. Levine (1995): "Consistency and Cautious Fictitious Play," *Journal of Economic Dynamics and Control*, 19, 1065–1089.

Fudenberg, D., P. Strack, and T. Strzalecki (2017): "Speed, Accuracy, and the Optimal Timing of Choices," .

Fudenberg, D., and T. Strzalecki (2015): "Dynamic logit with choice aversion," *Econometrica*, 83(2), 651–691.

Gabaix, X., and D. Laibson (2005): "Bounded rationality and directed cognition," *Harvard University*.

——— (2017): "Myopia and Discounting," .

Gowrisankaran, G., and M. Rysman (2012): "Dynamics of Consumer Demand for New Durable Goods," *mimeo*.

References VII

Gul, F., P. Natenzon, and W. Pesendorfer (2014): "Random Choice as Behavioral Optimization," *Econometrica*, 82(5), 1873–1912.

Gul, F., and W. Pesendorfer (2006): "Random expected utility," *Econometrica*, 74(1), 121–146.

Gul, F., and W. Pesendorfer (2013): "Random Utility Maximization with Indifference," *mimeo*.

Harsanyi, J. (1973): "Oddness of the number of equilibrium points: A new proof," *International Journal of Game Theory*, 2(1), 235–250.

Hausman, J., and D. McFadden (1984): "Specification tests for the multinomial logit model," *Econometrica: Journal of the Econometric Society*, pp. 1219–1240.

Hébert, B., and M. Woodford (2017): "Rational Inattention with Sequential Information Sampling," *mimeo*.

Heckman, J. J. (1981): "Heterogeneity and state dependence," in *Studies in labor markets*, pp. 91–140. University of Chicago Press.

Hendel, I., and A. Nevo (2006): "Measuring the implications of sales and consumer inventory behavior," *Econometrica*, 74(6), 1637–1673.

Hey, J. D. (1995): "Experimental investigations of errors in decision making under risk," *European Economic Review*, 39(3), 633–640.

References VIII

——— (2001): "Does repetition improve consistency?," *Experimental economics*, 4(1), 5–54.

Hofbauer, J., and W. Sandholm (2002): "On the global convergence of stochastic fictitious play," *Econometrica*, 70(6), 2265–2294.

Huber, J., J. W. Payne, and C. Puto (1982): "Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis," *Journal of Consumer Research*, 9.

Iyengar, S. S., and M. R. Lepper (2000): "When choice is demotivating: Can one desire too much of a good thing?," *Journal of Personality and Social Psychology*, 79(6), 995–1006.

Johnson, E. J., C. Camerer, S. Sen, and T. Rymon (2002): "Detecting failures of backward induction: Monitoring information search in sequential bargaining," *Journal of Economic Theory*, 104(1), 16–47.

Kahneman, D., and A. Tversky (1979): "Prospect theory: An analysis of decision under risk," *Econometrica*, pp. 263–291.

Kamenica, E. (2008): "Contextual Inference in Markets: On the Informational Content of Product Lines," *American Economic Review*, 98, 2127–2149.

Kasahara, H., and K. Shimotsu (2009): "Nonparametric identification of finite mixture models of dynamic discrete choices," *Econometrica*, 77(1), 135–175.

References IX

Ke, S. (2016): "A Dynamic Model of Mistakes," working paper.

Ke, T., and M. Villas-Boas (2016): "Optimal Learning before Choice," *mimeo*.

Ke, T. T., Z.-J. M. Shen, and J. M. Villas-Boas (2016): "Search for information on multiple products," *Management Science*, 62(12), 3576–3603.

Kitamura, Y., and J. Stoye (2016): "Nonparametric analysis of random utility models: testing," .

Krajbich, I., C. Armel, and A. Rangel (2010): "Visual fixations and the computation and comparison of value in simple choice," *Nature neuroscience*, 13(10), 1292–1298.

Krajbich, I., B. Bartling, T. Hare, and E. Fehr (2015): "Rethinking fast and slow based on a critique of reaction-time reverse inference.," *Nature Communications*, 6(7455), 700.

Krajbich, I., D. Lu, C. Camerer, and A. Rangel (2012): "The attentional drift-diffusion model extends to simple purchasing decisions," *Frontiers in psychology*, 3, 193.

Krajbich, I., and A. Rangel (2011): "Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions," *Proceedings of the National Academy of Sciences*, 108(33), 13852–13857.

References X

Kreps, D. (1979): "A representation theorem for" preference for flexibility"," *Econometrica*, pp. 565–577.

Krishna, V., and P. Sadowski (2012): "Dynamic Preference for Flexibility," *mimeo*.

——— (2016): "Randomly Evolving Tastes and Delayed Commitment," *mimeo*.

Liang, A., X. Mu, and V. Syrgkanis (2017): "Optimal Learning from Multiple Information Sources," .

Lin, Y.-H. (2017): "Stochastic Choice and Rational Inattention," *mimeo*.

Lu, J. (2016): "Random choice and private information," *Econometrica*, 84(6), 1983–2027.

Lu, J., and K. Saito (2016): "Random intertemporal choice," Discussion paper, mimeo.

Luce, D. (1959): *Individual choice behavior*. John Wiley.

Luce, R. D. (1986): *Response times*. Oxford University Press.

Luce, R. D., and H. Raiffa (1957): *Games and decisions: Introduction and critical survey*. New York: Wiley.

Machina, M. (1985): "Stochastic choice functions generated from deterministic preferences over lotteries," *The Economic Journal*, 95(379), 575–594.

References XI

Magnac, T., and D. Thesmar (2002): "Identifying dynamic discrete decision processes," *Econometrica*, 70(2), 801–816.

Manski, C. F. (1993): "Dynamic choice in social settings: Learning from the experiences of others," *Journal of Econometrics*, 58(1-2), 121–136.

Manzini, P., and M. Mariotti (2014): "Stochastic choice and consideration sets," *Econometrica*, 82(3), 1153–1176.

Marschak, J. (1959): "Binary Choice Constraints on Random Utility Indicators," Cowles Foundation Discussion Papers 74, Cowles Foundation for Research in Economics, Yale University.

Masatlioglu, Y., D. Nakajima, and E. Y. Ozbay (2011): "Revealed attention," 102, 2183–2205.

Matejka, F., and A. McKay (2014): "Rational inattention to discrete choices: A new foundation for the multinomial logit model," *The American Economic Review*, 105(1), 272–298.

Mattsson, L.-G., and J. W. Weibull (2002): "Probabilistic choice and procedurally bounded rationality," *Games and Economic Behavior*, 41, 61–78.

Matzkin, R. L. (1992): "Nonparametric and distribution-free estimation of the binary threshold crossing and the binary choice models," *Econometrica*, pp. 239–270.

References XII

McFadden, D. (1973): "Conditional logit analysis of qualitative choice behavior," in *Frontiers in Econometrics*, ed. by P. Zarembka. Institute of Urban and Regional Development, University of California.

McFadden, D., and M. Richter (1971): "On the Extension of a Set Function on a Set of Events to a Probability on the Generated Boolean σ-algebra," *University of California, Berkeley, working paper*.

McFadden, D., and M. Richter (1990): "Stochastic rationality and revealed stochastic preference," *Preferences, Uncertainty, and Optimality, Essays in Honor of Leo Hurwicz*, pp. 161–186.

McFadden, D. L. (2005): "Revealed stochastic preference: a synthesis," *Economic Theory*, 26(2), 245–264.

McKelvey, R., and T. Palfrey (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10, 638.

McKelvey, R. D., and T. R. Palfrey (1998): "Quantal response equilibria for extensive form games," *Experimental economics*, 1(1), 9–41.

Miller, R. (1984): "Job matching and occupational choice," *The Journal of Political Economy*, 92, 1086–1120.

Morris, S., and P. Strack (2017): "The Wald Problem and the Equivalence of Sequential Sampling and Static Information Costs," *mimeo*.

References XIII

Morris, S., and M. Yang (2016): "Coordination and Continuous Choice," .

Natenzon, P. (2016): "Random choice and learning," .

Norets, A., and X. Tang (2013): "Semiparametric Inference in dynamic binary choice models," *The Review of Economic Studies*, p. rdt050.

Oud, B., I. Krajbich, K. Miller, J. H. Cheong, M. Botvinick, and E. Fehr (2014): "Irrational Deliberation in Decision Making," *mimeo*.

Pakes, A. (1986): "Patents as options: Some estimates of the value of holding European patent stocks," *Econometrica*, 54, 755–784.

Raiffa, H., and R. Schlaifer (1961): *Applied statistical decision theory*. Boston: Division of Research, Harvard Business School.

Rand, D. G., J. D. Greene, and M. A. Nowak (2012): "Spontaneous giving and calculated greed," *Nature*, 489(7416), 427–430.

Ratcliff, R. (1978): "A theory of memory retrieval.," *Psychological review*, 85(2), 59.

Rieskamp, J., J. R. Busemeyer, and B. A. Mellers (2006): "Extending the Bounds of Rationality: Evidence and Theories of Preferential Choice," *Journal of Economic Literature*, 44(3), 631–661.

Rosenthal, A. (1989): "A bounded-rationality approach to the study of noncooperative games," *International Journal of Game Theory*, 18.

References XIV

RUBINSTEIN, A. (2007): "Instinctive and cognitive reasoning: A study of response times," *The Economic Journal*, 117(523), 1243–1259.

RUST, J. (1987): "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica*, pp. 999–1033.

——— (1989): "A Dynamic Programming Model of Retirement Behavior," in *The Economics of Aging*, ed. by D. Wise, pp. 359–398. University of Chicago Press: Chicago.

——— (1994): "Structural estimation of Markov decision processes," *Handbook of econometrics*, 4, 3081–3143.

SCOTT, D. (1964): "Measurement structures and linear inequalities," *Journal of mathematical psychology*, 1(2), 233–247.

SIMS, C. A. (2003): "Implications of rational inattention," *Journal of monetary Economics*, 50(3), 665–690.

STEINER, J., C. STEWART, AND F. MATĚJKA (2017): "Rational Inattention Dynamics: Inertia and Delay in Decision-Making," *Econometrica*, 85(2), 521–553.

STONE, M. (1960): "Models for choice-reaction time," *Psychometrika*, 25(3), 251–260.

TODD, P. E., AND K. I. WOLPIN (2006): "Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility," *American economic review*, 96(5), 1384–1417.

References XV

TVERSKY, A. (1969): "Intransitivity of Preferences," *Psychological Review*, 76, 31–48.

——— (1972): "Choice by Elimination," *Journal of Mathematical Psychology*, 9, 341–367.

WALD, A. (1945): "Sequential tests of statistical hypotheses," *The Annals of Mathematical Statistics*, 16(2), 117–186.

WILCOX, N. T. (2008): "Stochastic models for binary discrete choice under risk: A critical primer and econometric comparison," in *Risk aversion in experiments*, pp. 197–292. Emerald Group Publishing Limited.

——— (2011): "Stochastically more risk averse:A contextual theory of stochastic discrete choice under risk," *Journal of Econometrics*, 162(1), 89–104.

WOODFORD, M. (2014): "An Optimizing Neuroeconomic Model of Discrete Choice," *Columbia University working paper*.

Table of Contents

Lecture 1

Introduction
Random Utility/Discrete Choice
- Properties of RU
- Special Cases
- Random Expected Utility (REU)

Learning
Attention
- Optimal Attention
- Random Attention

Deliberate Randomization

Lecture 2

Introduction
Dynamic Random Utility
Dynamic Discrete Choice
Decision Times