Lectures on Stochastic Choice

Tomasz Strzalecki

this version: March 8, 2019

Table of Contents

Disclaimer

- I won't get too deeply into any one area

- The 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

$x, y, z \in X$ ····················· 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 (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?

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

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

Why is individual choice random?

- Randomly fluctuating tastes
- Noisy signals
- Attention is random
- Experimentation (experience goods)

} Agent does not see his choice as random. Stochastic choice is a result of informational asymmetry between agent and analyst.

- People just like to randomize
- Trembling hands

} Analyst and agent on the same footing

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 (2018) 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 mostly abstract from this here

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

Random Utility (RU)

Idea: Choice is random because:

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

Story (Data Generating Process):

- The menu A is drawn at random
  - maybe by the analyst or the experimenter
- A utility function is drawn at random (with probability distribution $\mathbb{P}$)
  - independently of the menu
- Agent chooses $x \in A$ whenever x maximizes the utility on A

Questions

- The most we can obtain from the data (with infinite sample) is
  - The distribution over menus
  - The conditional choice probability of choosing x from A: $\rho(x, A)$
- So you can think of the likelihood function $\mathbb{P} \mapsto \rho(x, A)$
  - Is this mapping one-to-one? (Identification/ Partial Identification)
  - What is the image of this mapping? (Axiomatization)

Random Utility (RU)

How to model a random utility function on X?

Depending on the context it will be either:

- a probability distribution μ over utility functions living in $\mathbb{R}^X$
- a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a $\mathcal{F}$-measurable random utility function $\tilde{U} : \Omega \to \mathbb{R}^X$

Notes:

- Given μ we can always take the canonical state space where $\Omega = \mathbb{R}^X$, $\tilde{U}$ the identity mapping, and $\mathbb{P} = \mu$.
- Or even $\Omega =$ all strict preferences
- Ω is a subjective state space, related to Kreps (1979) and Dekel, Lipman, and Rustichini (2001) $\rightsquigarrow$ Lecture 3

Random Utility (RU)

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

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

This is the event in which the utility is maximized at $x \in A$

Definition: ρ has a *random utility representation* if there exists $\tilde{U}$ 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 $\tilde{U}(x) = v(x) + \tilde{\epsilon}(x)$

Remark: This is sometimes called the *additive random utility* model

Ties

- $T^{\tilde{U}}$ is the event in which there is a tie

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

- Notice that RU implies that $\mathbb{P}(T^{\tilde{U}}) = 0$
  - this is because $\sum_{x \in A} \rho(x, A) = 1$
  - in DC guaranteed by assuming that $\tilde{\epsilon}$ has a density

- So not every $\tilde{U}$ leads to a legitimate ρ

Ways to deal with ties

- Sometimes convenient to allow ties (esp. when X is infinite)
- For example, randomize uniformly over $\text{argmax}_{x \in A} \tilde{U}_\omega(x)$
- A more general idea of tie-breaking was introduced by Gul and Pesendorfer (2006)
- A different approach is to change the primitive (stochastic choice correspondence: Barberá and Pattanaik, 1986; Lu, 2016; Gul and Pesendorfer, 2013)

Random Utility (with a tiebreaker)

- A *tie-breaker* is a random utility function $\tilde{W} : \Omega \to \mathbb{R}^X$, (which is always a strict preference)

- The agent first maximizes $\tilde{U}$ and if there is a tie, it gets resolved using $\tilde{W}$

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

$$\rho(x, A) = \mathbb{P}\left(\{\omega \in \Omega : \tilde{W}_\omega(x) \geq \tilde{W}_\omega(y) \text{ for all } y \in \operatorname{argmax}_{x \in A} \tilde{U}_\omega(x)\}\right).$$

Equivalence

Theorem: The following are equivalent when X is finite:

- ρ has a RU representation
- ρ has a RU representation with uniform tie breaking
- ρ has a RU representation with a tiebreaker

Thus, even though the representation is more general, the primitive is not.

When X is infinite (for example lotteries) these things are different.

Positivity

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

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?

i.i.d. DC

- It is often assumed that $\tilde{\epsilon}_x$ are i.i.d. across $x \in X$
  - logit, where ϵ has a mean zero extreme value distribution
  - probit, where ϵ has a mean zero Normal distribution

- In i.i.d. DC the binary choice probabilities are given by

$$\begin{aligned}\rho(x, \{x, y\}) &= \mathbb{P}(v(x) + \tilde{\epsilon}_x \geq v(y) + \tilde{\epsilon}_y) \\ &= \mathbb{P}(\tilde{\epsilon}_y - \tilde{\epsilon}_x \leq v(x) - v(y)) = F(v(x) - v(y)),\end{aligned}$$

where F is the cdf of $\tilde{\epsilon}_y - \tilde{\epsilon}_x$ (such models are called Fechnerian)

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

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

Comments:

- This property of ρ depends only on its restriction to binary menus
- RU in general is not Fechnerian because it violates Weak Stochastic Transitivity (Marschak, 1959)
- Some models outside of RU are Fechnerian, e.g., APU $\rightsquigarrow$ Lecture 2

References: Davidson and Marschak (1959); Block and Marschak (1960); Debreu (1958); Scott (1964); Fishburn (1998)

The Luce Model

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

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

Intuition 1: 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.

Intuition 2: $w(x)$ is the "response strength" associated with x. Choice probability is proportional to the response strength.

Equivalence

Theorem (McFadden, 1973): The following are equivalent

(i) ρ has a logit representation with v

(ii) ρ has a Luce representation with $w = e^v$

Proof: This is a calculation you all did in 1^{st} year metrics $\square$

DC with characteristics

- Typically menu is fixed, $A = X$
- $\xi \in \mathbb{R}^n$ vector of observable characteristics
- $\tilde{U}(x;\xi) = v(x;\xi) + \tilde{\epsilon}(x)$
- Observed choices $\rho(x, X; \xi)$

Note: Allen and Rehbeck (2019) also study $\rho(x, X; \xi)$, but with a perturbed utility representation instead of discrete choice representation $\rightsquigarrow$ Lecture 2

Generalizations

- Removing Positivity (Echenique and Saito, 2015; Cerreia-Vioglio, Maccheroni, Marinacci, and Rustichini, 2018; Ahumada and Ülkü, 2018)

- Nested logit (Train, 2009, for axioms see Kovach and Tserenjigmid, 2019)

- GEV (generalized extreme value; Train, 2009)

- Multivariate probit (Train, 2009)

- Mixed logit (McFadden and Train, 2000; Gul, Natenzon, and Pesendorfer, 2014; Saito, 2018)

Generalizations

- 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, 2018)
- Imbalanced Luce (Kovach and Tserenjigmid, 2018)
- Threshold Luce (Horan, 2018)

Regularity

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.

Theorem (Block and Marschak, 1960). If ρ has a random utility representation, then it satisfies Regularity.

Proof:

Step 1: If x maximizes u on B, then x maximizes u on A (because A is smaller). Thus, For any $x \in A \subseteq B$ we have $C(x, A) \supseteq C(x, B)$.

Step 2: $\mathbb{P}$ is set-monotone, so $\mathbb{P}(C(x, A)) \geq \mathbb{P}(C(x, B))$ □

Violations of Regularity

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

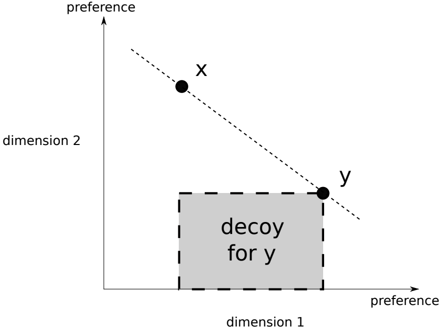

Axiomatic Characterizations

Theorem (Block and Marschak, 1960). Suppose that $|X| \leq 3$. If ρ satisfies Regularity, then ρ has a random utility representation.

Proof Idea:

- For each menu A sets $C(x, A)$ form a partition of Ω
- ρ defines a probability distribution over each partition
- Need to ensure that they are consistent with a single $\mathbb{P}$

Proof

Wlog $\Omega = \{xyz, xzy, yxz, yzx, zxy, zyx\}$. To define $\mathbb{P}(xyz)$ note that $C(y, \{y, z\}) = \{xyz, yxz, yzx\}$ and $C(y, X) = \{yxz, yzx\}$, so define

$$\mathbb{P}(xyz) := \rho(y, \{y, z\}) - \rho(y, X).$$

Likewise,

$$\begin{aligned}
\mathbb{P}(xzy) &:= \rho(z, \{y, z\}) - \rho(z, X) \\
\mathbb{P}(yxz) &:= \rho(x, \{x, z\}) - \rho(x, X) \\
\mathbb{P}(yzx) &:= \rho(z, \{x, z\}) - \rho(z, X) \\
\mathbb{P}(zxy) &:= \rho(x, \{x, y\}) - \rho(x, X) \\
\mathbb{P}(zyx) &:= \rho(y, \{x, y\}) - \rho(y, X)
\end{aligned}$$

By Regularity, they are nonnegative. They sum up to $3 - 2 = 1$. Finally, $\rho(x, A) = \mathbb{P}(C(x, A))$ follows from the above definitions as well. $\square$

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

Block and Marschak

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

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

Comments:

- Generalizes the idea that we get information from looking at difference between $\rho(x, A)$ and $\rho(x, B)$
- Inclusion-Exclusion formula (Möbus transform)

Other Axioms

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: These 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

Axiomatic Characterizations

Theorem: The following are equivalent for a finite X

(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).

Axioms for Luce/Logit

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

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

Axiom (Luce's Choice Axiom). For all $x \in A \subseteq B$

$$\rho(x, B) = \rho(x, A)\rho(A, B).$$

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

(i) ρ satisfies Positivity and Luce's IIA

(ii) ρ satisfies Positivity and Luce's Choice Axiom

(iii) ρ has a Luce representation

(iv) ρ has a logit representation

Proof

Luce $\Rightarrow$ **IIA** is straightforward:

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

Luce $\Rightarrow$ **Positivity** is also straightforward since $w(x) > 0$ for all $x \in X$

Proof

To show **Luce's IIA+Positivity** $\Rightarrow$ **Luce** for X finite, define $w(x) := \rho(x, X)$.

Fix A and $x^* \in A$. By IIA,

$$\rho(x, A) = \rho(x, X)\frac{\rho(x^*, A)}{\rho(x^*, X)} = w(x)\frac{\rho(x^*, A)}{w(x^*)}.$$

Summing up over $y \in A$ and rearranging we get $\frac{\rho(x^*,A)}{w(x^*)} = \frac{1}{\sum_{y \in A} w(y)}$.
When X is infinite, need to modify the proof slightly.

Proof

To show **IIA+Positivity** $\Rightarrow$ **Luce** for X finite, define $w(x) := \rho(x, X)$.

Fix A. By Luce's Choice Axiom,

$$\begin{aligned}\rho(x, X) &= \rho(x, A)\rho(A, X) \\ &= \rho(x, A) \sum_{y \in A} \rho(y, X)\end{aligned}$$

so $w(x) = \rho(x, A) \sum_{y \in A} w(y)$. When X is infinite, need to modify the proof slightly. $\square$

Remark: Luce's IIA is equivalent to Luce's Choice Axiom even without Positivity, see Cerreia-Vioglio, Maccheroni, Marinacci, and Rustichini (2018).

Other forms of IIA

Remark: IIA has a cardinal feel to it (we require ratios of probabilities to be equal to each other). Consider the following ordinal axiom.

Axiom (GNP IIA). If $A \cup B$ and $C \cup D$ are disjoint, then

$$\rho(A, A \cup C) \geq \rho(B, B \cup C) \Longrightarrow \rho(A, A \cup D) \geq \rho(B, B \cup D)$$

Theorem† (Gul, Natenzon, and Pesendorfer, 2014). In the presence of Richness†, ρ satisfies GNP IIA iff it has a Luce representation.

Blue bus/red bus paradox for i.i.d. DC

Example: Transportation choices are: train, or bus. There are two kinds of buses: blue bus and red bus. So $X = \{t, bb, rb\}$. Suppose that we observed that

$$\rho(t, \{t, bb\}) = \rho(t, \{t, rb\}) = \rho(bb, \{bb, rb\}) = \frac{1}{2}.$$

If ρ is i.i.d. DC, then $\rho(t, X) = \frac{1}{3}$. But this doesn't make much sense if you think that the main choice is between the modes of communication (train or bus) and the bus color is just a tie breaker. In that case we would like to have $\rho(t, X) = \frac{1}{2}$.

If you are still not convinced, imagine that there n colors of buses. Would you insist on $\rho(t, X) \to 0$ as $n \to \infty$?

Solution to the blue bus/red bus paradox

- Don't use i.i.d. DC, but some RU model
  - for example put equal probability on orders
    $bb \succ rb \succ t,\ rb \succ bb \succ t,\ t \succ bb \succ rb,\ t \succ rb \succ bb$
  - or use the attribute rule of Gul, Natenzon, and Pesendorfer (2014)
  - or use parametric DC families (nested logit, GEV)
- But no need to go outside of the RU class. This is not a paradox for RU but for i.i.d. DC

Weak Stochastic Transitivity

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

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

Satisfied by: Fechnerian models because $x \succsim^s y$ iff $v(x) \geq v(y)$

Can be violated by:

- RU (Marschak, 1959)
- random attention (Manzini and Mariotti, 2014)
- 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 $r \geq \max\{p, q\}$

Forms of Stochastic Transitivity

Tversky and Russo (1969) characterize the class of binary choice models that satisfy (a slightly stronger version of) Strong Stochastic Transitivity. Under Positivity, they are the models that have a *simple scalability* representation.

Definition: ρ has a *simple scalability* representation if $\rho(x, \{x, y\}) = F(v(x), v(y))$ for some $v : X \to \mathbb{R}$ and $F : \mathbb{R}^2 \to [0, 1]$, defined on an appropriate domain, is strictly increasing in the first argument and strictly decreasing in the second argument.

Note: Fechnerian is a special case where $F(v(x), v(y)) = F(v(x) - v(y))$.

Forms of Stochastic Transitivity

He and Natenzon (2018) characterize the class of models that satisfy (a slightly stronger version of) Moderate Stochastic Transitivity. These are the models that are represented by *moderate utility*.

Definition: ρ has a *moderate utility* representation if $\rho(x, \{x, y\}) = F\left(\frac{u(x)-u(y)}{d(x,y)}\right)$ for some $u : X \to \mathbb{R}$, distance metric $d : X \times X \to \mathbb{R}_+$, and $F : \mathbb{R} \to [0, 1]$ strictly increasing transformation, defined on an appropriate domain, that satisfies $F(t) = 1 - F(1 - t)$

Identification of Utilities

- 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.

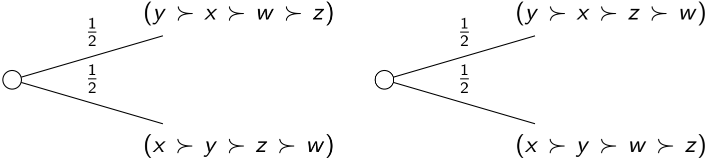

Note that these two distributions have disjoint supports!

Identification of "Marginal" Preferences

Theorem (Falmagne, 1978). If $\mathbb{P}_1$ and $\mathbb{P}_2$ are RU 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 $k = 1, \ldots, |X|$.

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"

i.i.d. DC with known distribution of $\tilde{\epsilon}$

Theorem: If $(v_1, \tilde{\epsilon})$ and $(v_2, \tilde{\epsilon})$ are i.i.d DC representations of ρ that share the distribution of $\tilde{\epsilon}$, then there exists $k \in \mathbb{R}$ such that $v_2(x) = v_1(x) + k$ for all $x \in X$.

Proof: Fix $x^* \in X$ and normalize $v(x^*) = 0$. Let F be the cdf of the ϵ difference. By Fechnerianity

$$\rho(x, \{x, x^*\}) = F(v(x)),$$

so $v(x) = F^{-1}(\rho(\{x, \{x, x^*\}))$. $\square$

Remark: If we know F, this gives us a recipe for identifying v from data.

Unknown distribution of $\tilde{\epsilon}$; observable attributes

- Under appropriate assumptions can identify $v(x; \xi)$ and the distribution of $\tilde{\epsilon}$
- Matzkin (1992) and the literature that follows

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
- Typical items are now $p, q, r \in X$

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

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

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

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 (modulo ties)

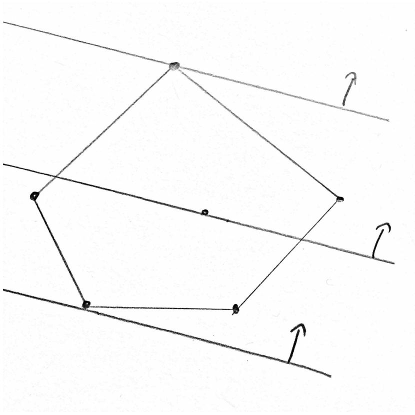

REU—Axioms

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: The vNM Independence axiom applied utility by utility

$$\tilde{u}_w \in C(p, A) \Longleftrightarrow u_w \in C\left(\alpha p + (1-\alpha)q, \alpha A + (1-\alpha)q\right)$$

Linearity

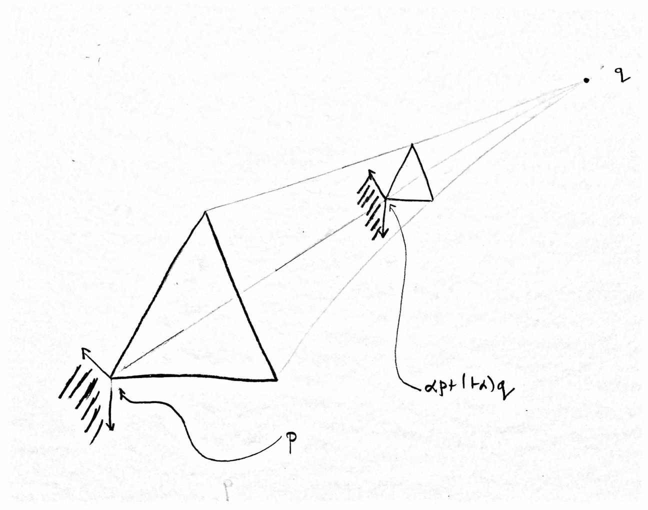

$C(p, A)$ is the normal cone of A at p
same for the mixture with q
they are equal because they are "corresponding angles"

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 identified.

REU—Comments

- Simpler axioms
- Better identification results
- Stronger assumptions: vNM relaxed $\rightsquigarrow$ Lecture 2
  - Allais (1953) paradox is a rejection of Linearity
  - Agranov and Ortoleva (2017) is a rejection of Extremeness
- Model used as a building block for a lot to come
- This is only one possible specification of risk preferences ...

REU—Comments

- Gul and Pesendorfer (2006) introduce tiebreakers
  - weakening of Continuity, tiebreaker are finitely additive
  - Extremeness hinges on tiebreakers being EU themselves (uniform tiebreaking violates Extremenes)
- finite support: Ahn and Sarver (2013)
  - Add a Finiteness axiom to get finitely many $\tilde{U}$
  - Useful in dynamic model to avoid conditioning on zero-probability events

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

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

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

Measuring Risk Preferences

Notation:

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

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

Theorem‡: (Wilcox, 2008, 2011; Apesteguia and Ballester, 2017) If p SOSD q, then $\rho_\theta^{DC}(p, \{p, q\})$ is strictly decreasing for large enough θ.

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

Lecture 2 on Stochastic Choice

Tomasz Strzalecki

Recap of Random Utility

- $(\Omega, \mathcal{F}, \mathbb{P})$ $\cdots\cdots\cdots\cdots\cdots\cdots\cdots$ probability space

- $\tilde{U} : \Omega \to \mathbb{R}^X$ $\cdots\cdots\cdots\cdots\cdots\cdots$ random utility

- $C(x, A) := \{\omega \in \Omega : \tilde{U}_\omega(x) \geq \tilde{U}_\omega(y) \text{ for all } y \in A\}$
  - agent learns the state (his utility) and chooses optimally

- $\rho(x, A) = \mathbb{P}(C(x, A))$
  - analyst does not see the state; the observed choice frequency of x from A is the probability that x is the argmax of the agent's utility on A

Learning

- In RU choice is stochastic because preferences are fluctuating
- Another possible reason: choices are driven by agent's noisy signals
  - Agent does not learn the state perfectly but gets a signal of it
- What kinds of ρ does this lead to?
  - If information is independent of the menu, this is a special case of RU
  - Strict subset of RU if model is rich enough
  - What if information can depend on the menu? $\rightsquigarrow$ later today

Learning—probabilistic model

- Let $\mathcal{G}$ represent *menu-independent* information the agent is learning.
- Conditional on the signal the agent maximizes $\mathbb{E}[\tilde{U}(x)|\mathcal{G}]$

Example: Hiring an applicant based on an interview

- Interview is a noisy signal
- Interview goes well $\Longrightarrow \mathbb{E}[\tilde{U}(\mathit{hire})|\mathcal{G}] > \mathbb{E}[\tilde{U}(\mathit{not})|\mathcal{G}]$
- Interview goes badly $\Longrightarrow \mathbb{E}[\tilde{U}(\mathit{hire})|\mathcal{G}] < \mathbb{E}[\tilde{U}(\mathit{not})|\mathcal{G}]$

Learning—probabilistic model

Comment: Choices are random because they depend on the signal realization

- No information ($\mathcal{G}$ trivial) $\Rightarrow$ choices are deterministic (agent maximizes ex ante expected utility)
- 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

Learning—probabilistic model

Proposition: ρ has a probabilistic learning representation iff it has a RU representation

Proof:

- For any $\mathcal{G}$ the induced ρ has a RU representation $(\Omega, \mathcal{G}, \mathbb{P}, \tilde{V})$ with $\tilde{V} := \mathbb{E}[\tilde{U}|\mathcal{G}]$
- Any RU has a learning representation where $\mathcal{G} = \mathcal{F}$ (signal is fully revealing) $\square$

Comment: Need to enrich the model to get a strictly special case

- Separation of tastes and beliefs $\rightsquigarrow$ next slide
- Strictly special case of RU in a dynamic setting (Frick, Iijima, and Strzalecki, 2019) $\rightsquigarrow$ Lecture 3

Learning—statistical model

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

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

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

$\mathbb{E}_p u(x)$ ··················· (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 $\Rightarrow$ random posterior $\tilde{q}$
  - For each posterior $\tilde{q}$ the agent maximizes $\max_{x \in A} \mathbb{E}_{\tilde{q}} u(x)$

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
- This is a special case of observable attributes ξ from Lecture 1

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

- That's when the analyst does not observe s

Learning—statistical model

Example (Classical experimental design in perception literature):

- s% of dots on the screen are moving left, $100 - s$% are moving right
- subject has to guess where most dots are moving
- imperfect perception, so noisy guesses
- experimenter controls s, observes ρ^s

Learning—statistical model

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 $s \rightsquigarrow$ Caplin and Martin (2015)
- 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
  - Typical acts are denoted $f, g, h \in X$
- Random Utility $\tilde{U}(f) = \sum_{s \in S} u(f(s))\tilde{q}(s)$, where
  - u is a (deterministic) linear utility 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

Let $A(s) := \{f(s) : f \in A\}$.

Axiom (S-monotonicity): If $\rho(f(s), A(s)) = 1$ for all $s \in S$ then $\rho(f, A) = 1$.

Axiom (C-determinism): If A is a menu of constant acts, then $\rho(f, A) = 1$ for some $f \in A$.

Axiom (Non-degeneracy): $\rho(f, A) \in (0, 1)$ for some $f \in A$.

Learning—the Lu (2016) model

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

‡ (Ties dealt with by stochastic choice correspondence)

Theorem‡ (Uniqueness). Under Non-degeneracy the the information structure μ is unique and the utility function u is cardinally-unique.

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

Theorem‡ (Comparative Statics). Fix u 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'.

Menu-dependent 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 depend on the menu (so we would have $\mathcal{G}^A$ or μ^A):
  - if new items provide more information
  - or if there is limited attention $\rightsquigarrow$ later today

Example

	$\tilde{U}_\omega(\textit{steak tartare})$	$\tilde{U}_\omega(\textit{chicken})$	$\tilde{U}_\omega(\textit{fish})$
$\omega = \textit{good chef}$	10	7	3
$\omega = \textit{bad chef}$	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 (2018) 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
- **Question**: What is the family of ρ that has a general menu-dependent learning representation? What is the additional bite of Blackwell monotonicity? What if Blackwell is violated (information overload)?

Learning so far

- Information independent of the menu (RU or special case of RU)
- Information dependent on the menu (more general than RU)

In both cases, the true state was uncorrelated with the menu. What if there is such a correlation? $\rightsquigarrow$ in general can explain any ρ

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, 2018)
    - correlation between menu and state (Kamenica, 2008)
  - General analysis? Axioms?

Learning

Attention

Optimal Attention

Random Attention

Controlled 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

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

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:

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

- 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

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	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)
- $\rho(x, \{x, y\}) = 0$, $\rho(x, \{x, y, z\}) = \frac{1}{2}$

Optimal Attention

Special cases of the cost function:

- Separable cost functions $C(\mu) = \int \phi(\tilde{q})\mu(d\tilde{q})$
  - for some function $\phi : \Delta(S) \to \mathbb{R}$
- Mutual information: separable where $\phi(q)$ is the relative entropy (Kullback-Leibler divergence) of q with respect to the prior $\int \tilde{q}\mu(d\tilde{q})$
- General cost functions: C is just Blackwell-monotone and convex

Optimal Attention

Question: Is it harder to distinguish "nearby" states than "far away" states?

- In the dots example, is it harder to distinguish $s = 49\%$ from $s = 51\%$ or $s = 1\%$ from $s = 99\%$?
- 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)
  - show the optimal solution leads to weighted-Luce choice probabilities ρ^s
  - can be characterized by two Luce IIA-like axioms on ρ^s

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

- Denti (2018) and Caplin, Dean, and Leahy (2018) characterize separable cost functions (and mutual information)
  - additional axioms beyond the two acyclicity-like axioms on ρ^s

- Chambers, Liu, and Rehbeck (2018) characterize a more general model without the $V^A(\mu) - C(\mu)$ separability
  - like Caplin and Dean (2015), they assume that the analyst knows the agent's utility function and the prior

Optimal Attention

- Lin (2017) 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

- Duraj and Lin (2019a) the agent can buy a fixed signal
  - either at a cost, or experimentation takes time
  - axiomatic characterization and uniqueness results

Learning

Attention

Optimal Attention

Random Attention

Controlled 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

- So this could be called Random Consideration

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

- Imagine that $\tilde{\Gamma}(A) = \tilde{\Gamma} \cap A$ for some random set $\tilde{\Gamma}$. Items outside of $\tilde{\Gamma}$ have their utility set to $-\infty$; inside of $\tilde{\Gamma}$ utility is unchanged. If $\tilde{\Gamma}$ is independent of the menu, then this is a special case of RU.

Random Attention

- Brady and Rehbeck (2016): allow for correlation
  - axiomatic characterization, uniqueness result
  - now can violate Regularity
- Cattaneo, Ma, Masatlioglu, and Suleymanov (2018): even more general
  - *attention filters*, following Masatlioglu, Nakajima, and Ozbay (2011)
  - axiomatic characterization, uniqueness result

Random Attention

- Suleymanov (2018) provides a clean axiomatic classification of these models
- Aguiar, Boccardi, Kashaev, and Kim (2018)
  - theoretical and statistical framework to test limited and random consideration at the population level
  - experiment designed to tease them apart

Random Attention

- Abaluck and Adams (2017): model with characteristics ξ
  - a version of Manzini and Mariotti (2014) where $\gamma(x)$ depend only on ξ^x
  - a version where the probability of being asleep (only looking at status quo) depends only on $\xi^{\text{status quo}}$
  - identification results and experimental proof of concept

Satisficing

- Aguiar, Boccardi, and Dean (2016): agent is Satisficing
  - draws a random order
  - goes through items till an item is "good enough"
  - randomness in orders generates randomness in choice

Learning

Attention

Optimal Attention

Random Attention

Controlled Randomization

Controlled 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:

- Trembling hands with implementation costs
- Allais-style lottery preferences
- Hedging against ambiguity
- Regret minimization

Trembling Hands

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

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

- 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)

a recent decision theoretic study is Allen and Rehbeck (2019)

The Quadruple Equivalence

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

(i) ρ satisfies Positivity and IIA

(ii) ρ has a Luce representation

(iii) ρ has a logit representation

(iv) 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

Axiom (Acyclicity): For any n and bijections $f, g : \{1, ..., n\} \to \{1, ..., n\}$,

$$\rho(x_1, A_1) > \rho(x_{f(1)}, A_{g(1)})$$
$$\rho(x_k, A_k) \geq \rho(x_{f(k)}, A_{g(k)}) \quad \text{for } 1 < k < n$$

implies

$$\rho(x_n, A_n) < \rho(x_{f(n)}, A_{g(n)}).$$

Condition (Ordinal IIA): For some continuous and monotone $\phi : (0, 1) \to \mathbb{R}_+$

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

for each menu $A, B \in A$ and $x, y \in A \cap B$.

Additive Perturbed Utility

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

(i) ρ has an APU representation with steep cost†

(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)

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"

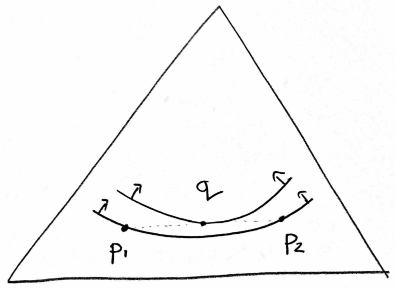

- Example: $p_1 \sim^\ell p_2$
- Strictly prefer q
- To implement this, choice from $A = \{p_1, p_2\}$ is $\rho(p_1, A) = \rho(p_2, A) = \frac{1}{2}$
- what if $B = \{p_1, p_2, q\}$? (Is the "mental coin" better or worse than actual coin?)

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
  - show that violations of Regularity obtain iff $\succsim^\ell$ has a point of strict convexity
  - characterize maximization of a specific $\succsim^\ell$ that belongs to the Cautious Expected Utility class $\longrightarrow$ Rational Mixing + additional axioms

- Lin (2019) shows lack of uniqueness for other classes of risk preferences
  - betweenness
  - also can rationalize REU as betweenness

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...

Lecture 3 on Stochastic Choice

Tomasz Strzalecki

Recap of Static Random Utility

- $(\Omega, \mathcal{F}, \mathbb{P})$ ··················· probability space
- $\tilde{U} : \Omega \to \mathbb{R}^X$ ················ random utility
- $C(x, A) := \{\omega \in \Omega : \tilde{U}_\omega(x) = \max_{y \in A} \tilde{U}_\omega(y)\}$
  - agent learns the state (his utility) and chooses optimally
- $\rho(x, A) = \mathbb{P}(C(x, A))$
  - analyst does not see the state; the observed choice frequency of x from A is the probability that x is the argmax of the agent's utility on A

Dynamic Random Utility (DRU)

In every period ρ_t has a RU representation with utility $\tilde{U}_t(x_0)$

Conditional choice probability (given a history of choices h^t):

$$\rho_t(x_t, A_t|h^t) = \mathbb{P}\Big[C(x_t, A_t)\Big|h^t\Big]$$

Two main dynamic effects that connect ρ_t and ρ_{t+1}

- **Backward Looking**: (if $\tilde{U}_t$ and $\tilde{U}_{t+1}$ are correlated)
  - History-Dependence, Choice-Persistence
- **Forward Looking**: (if $\tilde{U}_t$ satisfies the Bellman Equation)
  - Agent is Bayesian, has rational expectations, and correctly calculates option value

History Dependence and Selection on Unobservables

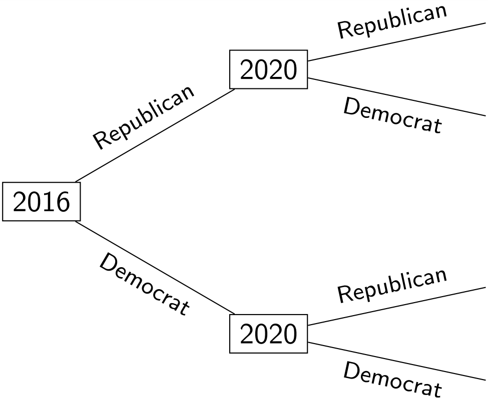

If political preferences persistent over time, expect history dependence:

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

History independence only if preferences completely independent over time.

History Dependence is a result of informational asymmetry between agent

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 distinguish between the two?
- How much history-dependence can there be?
- What are the axioms that link ρ_t and ρ_{t+1}?

Dynamic Decisions

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)

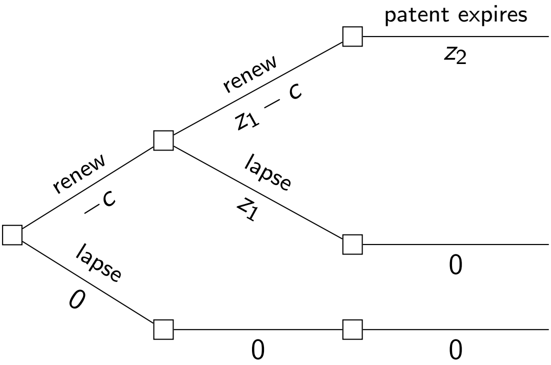

Primitive

- 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, but have covariates ξ
- Decision Theory literature
  - typically across decision trees

Bellman Equation

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?

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 $(A_0, x_0) \in \mathcal{H}_0(A_1)$,

$$\rho_1(x_1, A_1 | h^0) = \mathbb{P}\left[C(x_1, A_1) | C(x_0, A_0)\right]$$

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

- 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 (2019)
  - Identify two types of equivalence classes of histories
  - Impose history *independence* of ρ_1 within these classes

Contraction History Independence

Axiom (Contraction History Independence): If

(i) $A_0 \subseteq B_0$

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

then

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

Example

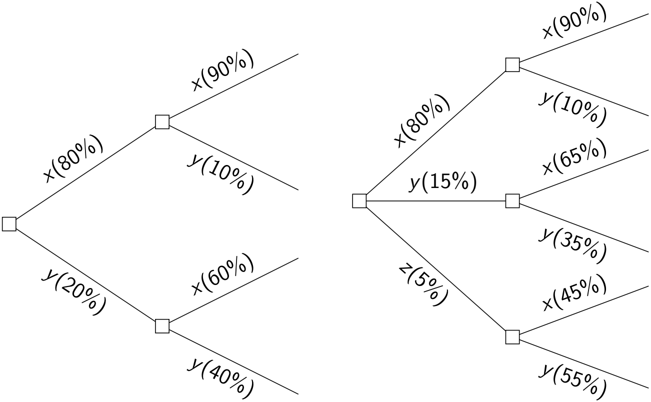

- z does not steal any customers from x in period $t = 0$
- so what people do in $t = 1$ after choosing x should be the same
- (note that z steals from y, so we have a mixture)

Adding Lotteries

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

- Denote lotteries by $p_t \in X_t$
- 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

$$h^0 = (A_0, x_0) \ \text{ vs } \ h^0 = (A_0, p_0, z_0)$$

Consumption History Independence

For now, assume away consumption dependence and allow only for choice dependence

Axiom (Consumption Independence): For any $p_0 \in A_0$ with $p_0(z_0), p_0(z_0') > 0$

$$\rho_1(\cdot|(A_0, p_0, z_0)) = \rho_1(\cdot|(A_0, p_0, z_0'))$$

Weak Linear History Independence

Idea: Under EU-maximization, choosing p_0 from A_0 reveals the same information as choosing option $\lambda p_0 + (1-\lambda) q_0$ from menu $\lambda A_0 + (1-\lambda)\{q_0\}$.

Axiom (Weak Linear History Independence)

$$\rho_1(\cdot,\cdot|A_0,p_0) = \rho_1\big(\cdot,\cdot|\lambda A_0 + (1-\lambda) q_0, \lambda p_0 + (1-\lambda) q_0\big).$$

Example

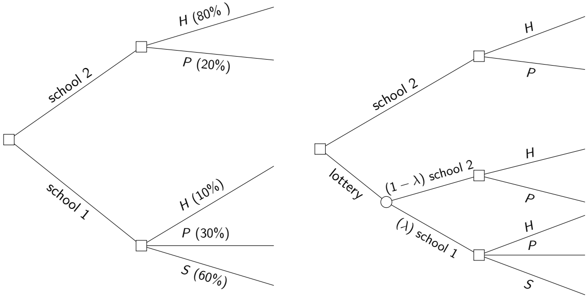

- school 2 offers two after-school programs, school 1 offers three
- different parents self-select to different schools
- how would school-1 parents choose between $\{H, P\}$?
- lottery to get in to the school
- Axiom says choice between $\{H, P\}$ independent of λ

Linear History Independence

Axiom (Weak Linear History Independence)

$$\rho_1(\cdot,\cdot|A_0,p_0)=\rho_1\big(\cdot,\cdot|\lambda A_0+(1-\lambda)q_0,\lambda p_0+(1-\lambda)q_0\big).$$

Idea was to mix-in a lottery q_0. Next we mix-in a set of lotteries B_0

Axiom (Linear History Independence)

$$\rho_1(\cdot,\cdot|A_0,p_0)\rho_0(p_0,A_0)$$

$$=\sum_{q_0\in B_0}\rho_1\Big(\cdot,\cdot|\lambda A_0+(1-\lambda)B_0,\lambda p_0+(1-\lambda)q_0\Big)\cdot\rho_0\Big(\lambda p_0+(1-\lambda)q_0,\lambda A_0+(1-\lambda)B_0\Big)$$

Dynamic Random Expected Utility

Theorem 1: ρ has a DREU representation if and only it satisfies

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

Remark: For REU axioms we use the approach of Gul and Pesendorfer (2006); Ahn and Sarver (2013). We need to extend their result to infinite spaces because X_1 is infinite (our Theorem 0).

Consumption Persistence

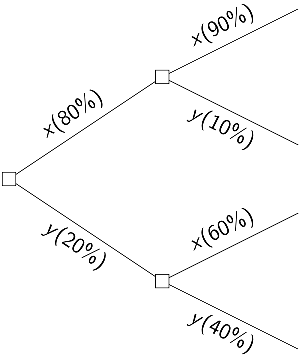

- $\rho_1(x|x) > \rho_1(x|y)$
- again, there is no habit here
- but serially correlated utility
- widely studied in marketing literature
- Frick, Iijima, and Strzalecki (2019) formulate behavioral notions of persistence and relate them to the serial correlation of utility

How to incorporate Dynamic Optimality?

- In the definition above, no structure on the family $(\tilde{U}_t)$
- But typically $\tilde{U}_t$ satisfies the Bellman equation

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

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

for $\delta > 0$ and a $\mathcal{F}_t$-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

Suppose that there are no lotteries, so $X_1 = Z_1$ is finite and $X_0 = M(X_1)$.

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): Preference $\succsim_0$ has an option-value representation iff it satisfies Completeness, Transitivity, Preference for Flexibility, and Modularity[†]

Comments:

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

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 menu without ties† $A \cup \{p\}$

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

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

Analogues in econometrics

- Analogue of Sophistication is the Williams-Daly-Zachary theorem
  - ρ_1 is the gradient of U_0 (in the space of utilities)
  - see, e.g., Koning and Ridder (2003); Chiong, Galichon, and Shum (2016)
- Hotz and Miller (1993) and the literature that follows exploits this relationship
- Sophistication is in a sense a "test" of this property

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 BEU 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 (2019) find a way to extract $\succsim_0$ from ρ_0
  - Impose stochastic versions of time-separability, DLR, and Sophistication

Passive and Active Learning

- BEU: 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

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

- Frick, Iijima, and Strzalecki (2019) relax consumption-independence and characterize habit-formation and active learning (experimentation) models

  - parametric models of active learning used by, e.g., Erdem and Keane (1996), Crawford and Shum (2005)

Related Work

- The Bayesian probit model of Natenzon (2018) 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 (2018) 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

Related Work

- Lu and Saito (2019)
  - study a random utility model where separability is violated, as in Epstein and Zin (1989)
  - show that even on simple domains where the continuation menu is fixed the analysts estimates of the function u are biased because they are contaminated by the nonlinear continuation utility.

- Duraj (2018) adds an objective state space to DREU and studies dynamic stochastic choice between Anscombe–Aumann acts
  - new primitive: augmented stochastic choice function
  - direct test of whether the agent's beliefs reflect the true data generating process

Preference for making choices late

- Suppose you got admitted to PhD programs at Harvard and MIT
- Do you make your decision before the visit days or after?

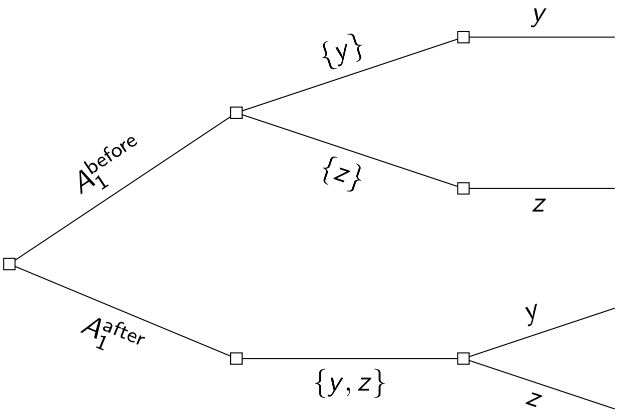

Preference for making choices late

Theorem†: If ρ has a BEU representation, then absent ties†

$$\rho_0(A_1^{\text{after}}, \{A_1^{\text{before}}, A_1^{\text{after}}\}) = 1$$

Comments:

- BEU has positive value of information: desire to delay the choice as late as possible to capitalize on incoming information (unless there is a cost)
- Here delaying decision does not delay consumption. The situation is different in optimal stopping models $\rightsquigarrow$ later today

DDC model

Definition: The *DDC model* is a restriction of DREU to deterministic decision trees that additionally satisfies the Bellman equation

$$\tilde{U}_t(z_t, A_{t+1}) = v_t(z_t) + \delta\mathbb{E}\left[\max_{y_{t+1}\in A_{t+1}} \tilde{U}_{t+1}(y_{t+1})|\mathcal{F}_t\right] + \tilde{\epsilon}_t^{(z_t, A_{t+1})},$$

with deterministic utility functions $v_t : \Omega \to \mathbb{R}^Z$; discount factor $\delta \in (0,1)$; and $\mathcal{F}_t$-adapted zero-mean *payoff shocks* $\tilde{\epsilon}_t : \Omega \to \mathbb{R}^{Y_t}$.

Special cases of DDC

- **BEU** is a special case, which can be written by setting $\tilde{\epsilon}_t^{(z_t, A_{t+1})} = \tilde{\epsilon}_t^{(z_t, B_{t+1})}$
  - **shocks to consumption**

- **i.i.d. DDC** where $\tilde{\epsilon}_t^{(z_t, A_{t+1})}$ and $\tilde{\epsilon}_\tau^{(y_t, B_{t+1})}$ are i.i.d.
  - **shocks to actions**

Remarks:

- i.i.d. DDC displays history-independence because $\tilde{U}_t$ are independent
- BEU can also be history-independent
- but these two models are different

Other special cases of DDC

- **permanent unobserved heterogeneity**: $\tilde{\epsilon}_t^{(z_t, A_{t+1})} = \tilde{\pi}_t^{z_t} + \tilde{\theta}_t^{(z_t, A_{t+1})}$, where
  - $\tilde{\pi}_t^{z_t}$ is a "permanent" shock that is measurable with respect to $\mathcal{F}_0$
  - $\tilde{\theta}_t^{(z_t, A_{t+1})}$ is a "transitory" shock, i.i.d. conditional on $\mathcal{F}_0$
  - Kasahara and Shimotsu (2009)
- **transitory but correlated shocks to actions**: $\tilde{\epsilon}_t^{(z_t, A_{t+1})}$ and $\tilde{\epsilon}_\tau^{(x_\tau, B_{\tau+1})}$ are i.i.d. whenever $t \neq \tau$, but might be correlated within any fixed period $t = \tau$
- **unobservable serially correlated state variables**: almost no structure on ϵ
  - Norets (2009); Hu and Shum (2012)

Dynamic logit

- A special case of i.i.d. DDC where $\tilde{\epsilon}_t$ are distributed extreme value
- Dynamic logit is a workhorse for estimation
  - e.g., Miller (1984), Rust (1989), Hendel and Nevo (2006), Gowrisankaran and Rysman (2012)
- 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)

Axiomatization (Fudenberg and Strzalecki, 2015)

Notation: $x \succsim_t^s y$ iff $\rho_t(x, A) \geq \rho_t(y, A)$ for $A = \{x, y\}$

Axiom (Recursivity):

$$(z_t, A_{t+1}) \succsim_t^s (z_t, B_{t+1})$$
$$\Updownarrow$$
$$\rho_{t+1}(A_{t+1}, A_{t+1} \cup B_{t+1}) \geq \rho_{t+1}(B_{t+1}, A_{t+1} \cup B_{t+1})$$

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

$$(z_t, A_{t+1}) \succsim_t^s (z_t, B_{t+1})$$

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" (each menu is penalized by a function of its size)
  - Ke (2016): a dynamic model of mistakes (agent evaluates each menu by the expectation of the utility under her own stochastic choice function)
- 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 Williams-Daly-Zachary theorem may be useful here?
- 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

Understanding the role of i.i.d. ϵ

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

Let $A_0 := \{(z_0, A_1^{\text{small}}), (z_0, A_1^{\text{big}})\}$ where $A_1^{\text{small}} = \{z_1\}$ and $A_1^{\text{big}} = \{z_1, z_1'\}$.

Proposition (Frick, Iijima, and Strzalecki, 2019): If ρ has a i.i.d. DDC representation, then

$$0 < \rho_0\left((z_0, A_1^{\text{small}}), A_0\right) < 0.5.$$

Moreover, if the ε shocks are scaled by $\lambda > 0$, then this probability is strictly increasing in λ whenever $v_1(z_1') > v_1(z_1)$.

Understanding the role of i.i.d. ϵ

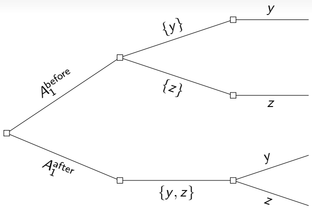

Proposition (Fudenberg and Strzalecki, 2015; Frick, Iijima, and Strzalecki, 2019): If ρ has a i.i.d. DDC representation with $\delta < 1$, then

$$0.5 < \rho_0\left((x, A_1^{\text{early}}), A_0\right) < 1.$$

Moreover, if ε is scaled by $\lambda > 0$, then $\rho_0((x, A_1^{\text{early}}), A_0)$ is strictly increasing in λ (modulo ties).

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

Beyond i.i.d. DDC

- This result extends in a straightforward way to DDC with permanent unobserved heterogeneity
  - this is just a mixture of i.i.d DDC models, so inherits this property
- Also to DDC with transitory but correlated shocks to actions
- serially correlated unobservable heterogeneity:
  - the result may not hold in general
  - example: in the mixture model of i.i.d. DDC with BEU there is a horse race between the two effects

Other Decision Problems

- So far, looked at pure manifestations of option value
  - direct choice between nested menus
  - costless option to defer choice
- DDC models typically not applied to those
- But these forces exist in "nearby" choice problems
- So specification of shocks matters more generally

Modeling Choices

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

- i.i.d. DDC: widely used because of statistical tractability, but
  - smaller menus w/prob. $\in (0, \frac{1}{2})$
  - early decisions w/prob. $\in (\frac{1}{2}, 1)$

Comments:

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

- Biased Parameter Estimates

- Model Misspecification
  - Maybe a fine model of (behavioral) consumers
  - But what about profit maximizing firms?

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 BEU!
- Negative option value is not a consequence of history independence
  - e.g., no such problem in the i.i.d. BEU
- It is a consequence of shocks to actions vs shocks to payoffs

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?

Are quick decisions better than slow ones?

Informational Effect:

- More time ⇒ more information ⇒ better decisions
  - if forced to stop at time t, make better choices for higher t
  - seeing more signals leads to more informed choices

Selection Effect:

- Time is costly, so you decide to stop depending on how much you expect to learn (option value of waiting)
  - Want to stop early if get an informative signal
  - Want to continue if get a noisy signal

- This creates dynamic selection
  - stop early after informative signals
  - informative signals more likely when the problem is easy

Decreasing accuracy

The two effects push in opposite directions. Which one wins?

Stylized fact: Decreasing accuracy: fast decisions are “better”

- Well established in perceptual tasks (dots moving on the screen), where “better” is objective
- Also in experiments where subjects choose between consumption items

When are decisions "more accurate?"

In cognitive tasks, **accurate = correct**

In choice tasks, **accurate = preferred**

$\boldsymbol{p(t)}$:= probability of making the correct/preferred choicce conditional conditional on deciding at t

Definition:

$$P \text{ displays } \begin{Bmatrix} \text{increasing} \\ \text{decreasing} \\ \text{constant} \end{Bmatrix} \text{ accuracy iff } p(t) \text{ is } \begin{Bmatrix} \text{increasing} \\ \text{decreasing} \\ \text{constant} \end{Bmatrix}$$

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 $A = \{\ell, r\}$ (100 different pairs)
  - record choice and decision time
- Step 3: Draw a random pair and get your choice

Remark: Here we are using the ratings as a proxy for true preferences. Of course, this is imperfect, as the ratings themselves are probably stochastic as well, so this approach should be treated only as the first step.

Decreasing Accuracy

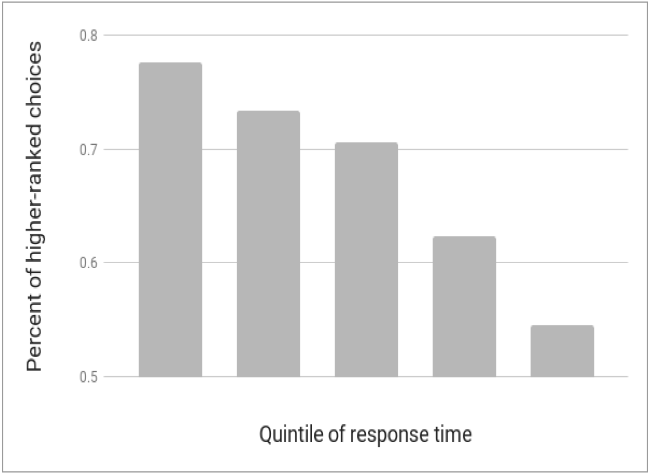

(based on data from Krajbich, Armel, and Rangel, 2010)

Domain

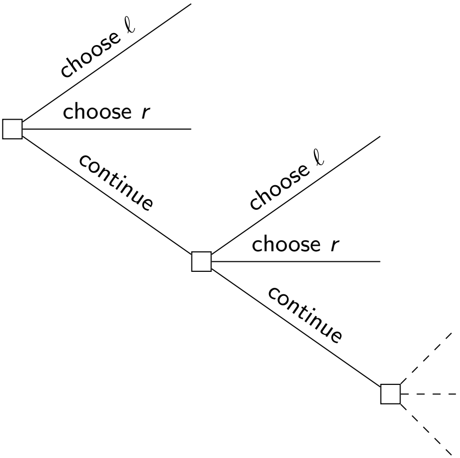

Models

There are two kinds of models:

1. **Optimal stopping models**: The agent is optimally choosing when to stop (and what to choose). The benefit of waiting is that it gives more information. But there is a cost of waiting too, so the optimal decision balances the two.

   – Wald's model (with a binary prior)

   – Chernoff's model (with a Gaussian prior)

2. **Hitting-time models**: The agent stops when some accumulation process hits a certain boundary. This is a heuristic model, there is no optimization here.

   – Drift-Diffusion models

Under certain conditions, 2 is a solution to 1.

Optimal Stopping 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)
- τ $\cdots$ stopping time, $\{\tau \leq t\} \in \mathcal{G}_t$
- 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

Interpretation of the Signal Process

- recognition of the objects on the screen
- retrieving pleasant or unpleasant memories
- coming up with reasons pro and con
- introspection
- signal strength depends on the utility difference or on the ease of the perceptual task

In animal experiments, some neuroscientists record neural firing and relate it to these signals

We don't do this, treat signals as unobserved by the analyst

Exogenous vs Endogenous Stopping

Example: If stopping is exogenous (τ is independent of signal G_t), and prior is symmetric, there is **increasing** accuracy: waiting longer gives better information so generates better decisions

- Key assumption above: stopping independent of signal
- If stopping is conditional on the signal, this could get reversed
- Intuition: with endogenous stopping you
  - *#1* stop early after informative signals (and make the right choice); wait longer after noisy signals (and possibly make a mistake)
  - *#2* probably faced an easier problem if you decided quickly

Optimal Stopping Problem

The agent chooses the stopping time optimally

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

Comments:

- Assume first that cost is linear $C(t) = ct$
- $(\mathcal{G}_t)$ 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? $\rightsquigarrow$ later

Further Assumptions

- Binary choice $A = \{x, y\}$
- $s = \left(u^\ell, u^r\right) \in \mathbb{R}^2$; utility function is identity
- Continuous time
- Signal: $\mathcal{G}_t$ is generated by (G_t^ℓ, G_t^r) where

$$G_t^i = t \cdot u^i + B_t^i$$

  and B_t^ℓ, B_t^r are Brownian motions; often look at $G_t := G_t^\ell - G_t^r$

Examples of Prior/Posterior Families

- "certain difference" (Wald's model)
  - binomial prior: either $s = (1, 0)$ or $s = (0, 1)$
  - binomial posterior: either $s = (1, 0)$ or $s = (0, 1)$
- "uncertain difference" (Chernoff's model)
  - Gaussian prior: $u^i \sim N(X_0^i, \sigma_0^2)$, independent
  - Gaussian posterior: $u^i \sim N(X_t^i, \sigma_t^2)$, independent

The "certain difference" model

* Assumptions:
  - binomial prior: either $s = (1, 0)$ or $s = (0, 1)$
  - binomial posterior: either $s = (1, 0)$ or $s = (0, 1)$

* Key intuition: **stationarity**
  - suppose that you observe $G_t^\ell \approx G_t^r$ after a long t
  - you think to yourself: "the signal must have been noisy"
  - so you don't learn anything $\Rightarrow$ you continue

* Formally, the option value is constant in time

The "certain difference" model

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

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

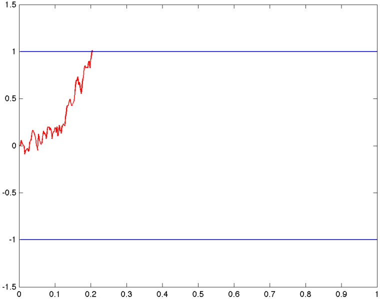

The "certain difference" model

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

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

The "certain difference" model

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

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

Comments

- The solution to the optimal stopping problem is a hitting-time model
- Can use this as a reduced-form model to generate $\rho \in \Delta(A \times \mathcal{T})$
  - 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)
- Used extensively for perception tasks since the 70's; pretty well established in psych and neuroscience
- Closed-form solutions for choice probabilities (logit) and expected decision time

Comments

- More recently used to study choice tasks by a number of teams of authors including Colin Camerer and Antonio Rangel

- Many versions of the model
  - ad-hoc tweaks (not worrying about optimality)
    - assumptions about the process G_t
    - functional forms for the time-dependent boundary
  - much less often, optimization used:
    - time-varying costs (Drugowitsch, Moreno-Bote, Churchland, Shadlen, and Pouget, 2012)
    - endogenous attention (Woodford, 2014)

Hitting Time Models

Definition:

- Stochastic "stimulus process" G_t starts at 0
- Time-dependent boundary $b : \mathbb{R}_+ \to \mathbb{R}_+$
- Hitting time $\tau = \inf\{t \geq 0 : |G_t| \geq b(t)\}$
- Choice $= \begin{cases} \ell & \text{if} \quad G_\tau = +b(\tau) \\ r & \text{if} \quad G_\tau = -b(\tau) \end{cases}$

Anything Goes

Proposition (Fudenberg, Strack, and Strzalecki, 2018): Any Borel $\rho \in \Delta(A \times T)$ has a hitting time representation where the stochastic process G_t is a time-inhomogeneous Markov process and the barrier is constant

Remarks:

- This means that the general model is without loss of generality
  - for a fixed binary menu (but maybe not across menus?)
- In particular, it is without loss of content to assume that b is independent of time
- However, in the general model the process G_t may have jumps
- From now on we focus on the DDM special cases

Drift Diffusion Model (DDM)

Special case where the stimulus process G_t is a diffusion with constant drift and volatility

$$G_t = \delta t + B_t$$

Definition: ρ has a **DDM representation** if it can be represented by a stimulus process $G_t = \delta t + B_t$ and a time-dependent boundary b. We write this as $\rho = DDM(\delta, b)$

Remarks:

- The optimal solution to the certain difference model is a DDM with $\delta = u^\ell - u^r$ and constant b.
- The Brownian assumption has bite. A partial axiomatization was obtained by Baldassi, Cerreia-Vioglio, Maccheroni, and Marinacci (2018) but they only look at expected decision times, so ignore the issue of correlation of times and decisions

Average DDM

Definition: ρ has an **average DDM representation** $DDM(\mu, b)$ with $\mu \in \Delta(\mathbb{R})$ if $\rho = \int \rho(\delta, b) d\mu(\delta)$.

- In an average DDM model the analyst does not know δ, but has a correct prior
- Intuitively, it is unknown to the analyst how hard the problem is for the agent
- This is the unconditional choice function ρ (the average of ρ^s)

Accuracy

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

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

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!

Accuracy in DDM

Theorem (Fudenberg, Strack, and Strzalecki, 2018): Suppose that $\rho = DDM(\delta, b)$.

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

Intuition for decreasing accuracy: this is our selection effect #1

- higher bar to clear for small t, so if the agent stopped early, G must have been very high, so higher likelihood of making the correct choice

Accuracy in DDM models

Theorem (Fudenberg, Strack, and Strzalecki, 2018): Suppose that $\rho = DDM(\mu, b)$, with $\mu = \mathcal{N}(0, \sigma_0)$

$$\text{accuracy } \alpha \text{ is } \begin{Bmatrix} \text{increasing} \\ \text{decreasing} \\ \text{constant} \end{Bmatrix} \text{ iff } b(t) \cdot \sigma_t \text{ is } \begin{Bmatrix} \text{increasing} \\ \text{decreasing} \\ \text{constant} \end{Bmatrix}$$

where $\sigma_t^2 := \frac{1}{\sigma_0^{-2} + t}$

Intuition for decreasing accuracy: this is our selection effect #2

- σ_t is a decreasing function; this makes it an easier bar to pass

selection effect #2

Proposition (Fudenberg, Strack, and Strzalecki, 2018): Suppose that $\mu = \mathcal{N}(0, \sigma_0)$, and $b(t) \cdot \sigma_t$ non-increasing. Then $|\delta|$ **decreases in** τ in the sense of FOSD, i.e. for all $d > 0$ and $0 < t < t'$

$$\mathbb{P}\left[\,|\delta| \geq d \mid \tau = t\right] > \mathbb{P}\left[\,|\delta| \geq d \mid \tau = t'\right] .$$

- larger values of $|\delta|$ more likely when the agent decides quicker
- problem more likely to be "easy" when a quick decision is observed
- this is a selection coming from the analyst not knowing how hard the problem is

Extended DDM

- Constant DDM cannot explain decreasing accuracy
- **Extended DDM** adds more parameters to match the data better:
  - random starting point of Z_t
  - random drift δ
  - random initial latency T
- Sometimes this is also called **full DDM**
- See, e.g., Ratcliff and McKoon (2008); Bogacz, Brown, Moehlis, Holmes, and Cohen (2006); Ratcliff and Smith (2004)

Example: random starting point

- random starting point may seem ad-hoc, but it sometimes makes sense:
  - example: there is a window of time in which the agent gathers information but cannot act yet
- Let T_0 be the length of this window and the drift be $\lambda(u^\ell - u^r)$
- Then the starting point of Z will be distributed $N(T_0\lambda(u^\ell - u^r), T_0)$
- Chiong, Shum, Webb, and Chen (2018) study non-skippable video ads in apps (estimate the model and simulate skippable counterfactual)

Microfounding the Boundary

* So far, only the constant boundary b was microfounded

* Do any other boundaries come from optimization?

* Which boundaries should we use?

* We now derive the optimal boundary

The "uncertain difference" model

* Assumptions:
  - Gaussian prior: $u^i \sim N(X_0^i, \sigma_0^2)$
  - Gaussian posterior: $u^i \sim N(X_t^i, \sigma_t^2)$

* Key intuition: **nonstationarity**
  - suppose that you observe $G_t^l \approx G_t^r$ after a long t
  - you think to yourself: "I must be indifferent"
  - so you have learned a lot $\Rightarrow$ you stop

* Formally $\sigma_t^2 = \frac{1}{\sigma_0^{-2}+t}$ so option value is decreasing in time

* Intuition for the difference between the two models:
  - interpretation of signal depends on the prior

The "uncertain difference" model

Theorem (Fudenberg, Strack, and Strzalecki, 2018): In the uncertain difference model the optimal behavior has a DDM representation. Moreover, unconditional on the state, accuracy is decreasing.

Other Boundaries

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

Theorem‡ (Fudenberg, Strack, and Strzalecki, 2018): 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

- Fudenberg, Strack, and Strzalecki (2018) show that in their model it is always optimal to pay equal attention to both alternatives
- 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
- Zhong (2018) shows that in a broad class of models Poisson signals are optimal.

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
- Duraj and Lin (2019b): decision-theoretic foundations for optimal sampling

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
- But it is optimal to think more when options are closer
  - The option value of thinking is higher
  - Would you like to spend more time thinking about the choice "Harvard vs MIT" or "Harvard vs Alabama State"?

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

Other questions

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

- Kahneman (2011); Rubinstein (2007); Rand, Greene, and Nowak (2012); Krajbich, Bartling, Hare, and Fehr (2015); Caplin and Martin (2016)

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)

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)

Thank you!

References I

Abaluck, J., and A. Adams (2017): "What Do Consumers Consider Before They Choose? Identification from Asymmetric Demand Responses," .

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

Aguiar, V. H., M. J. Boccardi, and M. Dean (2016): "Satisficing and stochastic choice," *Journal of Economic Theory*, 166, 445–482.

Aguiar, V. H., M. J. Boccardi, N. Kashaev, and J. Kim (2018): "Does Random Consideration Explain Behavior when Choice is Hard? Evidence from a Large-scale Experiment," *arXiv preprint arXiv:1812.09619*.

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

Ahumada, A., and L. Ülkü (2018): "Luce rule with limited consideration," *Mathematical Social Sciences*, 93, 52–56.

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.

Allen, R., and J. Rehbeck (2019): "Revealed Stochastic Choice with Attributes," Discussion paper.

References II

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*.

Baldassi, C., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci (2018): "Simulated Decision Processes An axiomatization of the Drift Diffusion Model and its MCMC extension to multi-alternative choice," .

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.

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.

Bogacz, R., E. Brown, J. Moehlis, P. Holmes, and J. D. Cohen (2006): "The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks.," *Psychological review*, 113(4), 700.

References III

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., M. Dean, and J. Leahy (2018): "Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy," Discussion paper.

Caplin, A., and D. Martin (2015): "A testable theory of imperfect perception," *The Economic Journal*, 125(582), 184–202.

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

Cattaneo, M., X. Ma, Y. Masatlioglu, and E. Suleymanov (2018): "A Random Attention Model: Identification, Estimation and Testing," *mimeo*.

References IV

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," .

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and A. Rustichini (2018): "Multinomial logit processes and preference discovery: inside and outside the black box," Discussion paper.

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

Chambers, C. P., C. Liu, and J. Rehbeck (2018): "Costly Information Acquisition," Discussion paper.

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

Chiong, K., M. Shum, R. Webb, and R. Chen (2018): "Split-second Decision-Making in the Field: Response Times in Mobile Advertising," *Available at SSRN*.

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

References V

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.

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), 440–444.

References VI

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.

Denti, T. (2018): "Posterior-Separable Cost of Information," Discussion paper.

Drugowitsch, J., R. Moreno-Bote, A. K. Churchland, M. N. Shadlen, and A. Pouget (2012): "The cost of accumulating evidence in perceptual decision making," *The Journal of Neuroscience*, 32(11), 3612–3628.

Duraj, J. (2018): "Dynamic Random Subjective Expected Utility," Discussion paper.

Duraj, J., and Y.-H. Lin (2019a): "Costly Information and Random Choice," Discussion paper.

——— (2019b): "Identification and Welfare Analysis in Sequential Sampling Models," Discussion paper.

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

Echenique, F., and K. Saito (2015): "General Luce Model," *Economic Theory*, pp. 1–16.

References VII

Echenique, F., K. Saito, and G. Tserenjigmid (2018): “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,” .

Epstein, L. G., and S. Zin (1989): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57(4), 937–969.

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.

References VIII

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

FRICK, M., R. IIJIMA, AND T. STRZALECKI (2019): "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 (2018): "Speed, Accuracy, and the Optimal Timing of Choices," *American Economic Review*, 108, 3651–3684.

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 IX

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.

He, J., and P. Natenzon (2018): "Moderate Expected Utility," Discussion paper, Working Paper, Washington University in Saint Louis.

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.

References X

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

——— (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.

Horan, S. (2018): "Threshold Luce Rules," Discussion paper.

Hotz, V. J., and R. A. Miller (1993): "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 60(3), 497–529.

Hu, Y., and M. Shum (2012): "Nonparametric identification of dynamic models with unobserved state variables," *Journal of Econometrics*, 171(1), 32–44.

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.

References XI

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. (2011): *Thinking, fast and slow*. Macmillan.

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.

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 (2018): "Nonparametric analysis of random utility models," *Econometrica*, 86(6), 1883–1909.

Koning, R. H., and G. Ridder (2003): "Discrete choice and stochastic utility maximization," *The Econometrics Journal*, 6(1), 1–27.

References XII

Kovach, M., and G. Tserenjigmid (2018): "The Imbalanced Luce Model," Discussion paper.

——— (2019): "Behavioral Foundations of Nested Stochastic Choice and Nested Logit," Discussion paper.

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.

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*.

——— (2019): "Random Non-Expected Utility: Non-Uniqueness," *mimeo*.

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

Lu, J., and K. Saito (2018): "Random intertemporal choice," *Journal of Economic Theory*, 177, 780–815.

Lu, J., and K. Saito (2019): "Repeated Choice," *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.

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

References XIV

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.

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.

References XV

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., and K. Train (2000): "Mixed MNL models for discrete response," *Journal of applied Econometrics*, pp. 447–470.

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, 6–38.

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*.

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

References XVI

Natenzon, P. (2018): "Random choice and learning," *Journal of Political Economy*.

Norets, A. (2009): "Inference in dynamic discrete choice models with serially orrelated unobserved state variables," *Econometrica*, 77(5), 1665–1682.

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.

Ratcliff, R., and G. McKoon (2008): "The diffusion decision model: Theory and data for two-choice decision tasks," *Neural computation*, 20(4), 873–922.

Ratcliff, R., and P. L. Smith (2004): "A comparison of sequential sampling models for two-choice reaction time.," *Psychological review*, 111(2), 333.

References XVII

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.

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.

Saito, K. (2018): "Axiomatizations of the Mixed Logit Model," .

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.

References XVIII

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.

Suleymanov, E. (2018): "Stochastic Attention and Search," *mimeo*.

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.

Train, K. (2009): *Discrete choice methods with simulation*. Cambridge university press, 2nd edn.

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

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

Tversky, A., and J. E. Russo (1969): "Substitutability and similarity in binary choices," *Journal of Mathematical psychology*, 6(1), 1–12.

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*.

ZHONG, W. (2018): "Optimal Dynamic Information Acquisition," Discussion paper.