Revealed Bayesian Learning[1]

Andrew Caplin, Daniel Martin, and Philip Marx

June 2022

[1]We thank the NOMIS Foundation and the Alfred. P. Sloan Foundation for their generous support for research on the Cognitive Foundations of Economic Behavior.

Big Picture

- Bayesian learning models ubiquitous across economics, psychology, and neuroscience
- They provide a shared cognitive foundation for decision-making in informationally rich environments
- Key assumptions: agents acquire informative signals, update using Bayes' rule, and choose the action that maximized expected utility
- Common questions: Are Bayesian learning models appropriate for capturing choices? If so, what can they tell us about what agents like and what they learn?

Three Model Classes

- Most learning models fall into one of three large and nested classes:
    1. Fixed learning: information structures not impacted by payoffs
        - Exogenous learning in markets, observational/social learning, etc.
    2. Capacity-constrained learning: feasible information structure chosen to match payoffs
        - Fixed information-processing capacity (Sims 2003), optimal encoding (Woodford 2014), etc.
    3. Costly learning: costly information structure chosen given net payoffs
        - Search, bandits, rational inattention, etc.
- Note they are nested: fixed implies capacity constrained implies costly learning not vice versa

Revealing Bayesian Learning

- What is the behavioral imprint of each of these model classes?
- Fundamental challenge: Utility is unknown to econometrician and learning is private
- Recent advances: characterizations of fixed learning (Lu 2016) and costly learning (Caplin and Dean 2015 CD15) provide partial answers in different data (Lu, perfect information, CD15 known utility)
- As yet no characterization of capacity-constrained learning

Results

- We characterize all three Bayesian learning model classes within a unified framework
- Our characterizations make testing these models possible
- Our characterizations identify what can be known about utility
- And what can be known about learning
- Definitive results in our framework

Apologia

- Quite a few definitions just to set up the questions
- The theorems are conceptually straight forward but nonetheless intricate
- I will cover what I can in talk
- I expect that it will not all sink in first time
- Talk is being recorded so possible to review finer points
- There is no paper so I am including technicalities in an Appendix for those who want to pursue

What Next?

- Opens the door to applications such as:
  1. Machine learning (presentation in Summer School which closes out characterizations)
  2. Psychometrics (ongoing with Bustamente, Daw, Grahek, Ham, Musslick)
  3. Labor economics (ongoing with Deming, Leth-Petersen, Weidmann)
  4. Judicial, Medical, Educational...
- Demonstrates ubiquity of Bayesian learning models and strong potential for interdisciplinary collaboration

Talk Structure

1. Definitions
2. Motivating example
3. FIR and mean and optimality preserving spreads
4. CCR and the NIAAS Cone
5. CIR and the NIAC Cone
6. Recovery: Example revisited
7. Concluding remarks

1. Definitions

Objects

- Finite states of the world $\omega \in \Omega$
- Fixed prior $\mu \in \Delta(\Omega)$
- Global set of actions $\mathcal{A}$
- Decision problem a finite set of such actions $A \subset \mathcal{A}$
- Finite prize set $Z = \{z_k\}_{k=1}^{K}$ and mapping $z : \mathcal{A} \times \Omega \to Z$ makes it important for the decision maker (DM) to learn before choosing

Strategies

- State dependent stochastic choice (SDSC) $P(a,\omega)$ that reflects prior

$$\sum_{a \in A} P(a,\omega) = \mu(\omega).$$

- As in Caplin and Martin (2015 CM15), characterize by **revealed information structure** which is defined by **unconditional action probabilities**

$$P(a) = \sum_{\omega \in \Omega} P(a,\omega);$$

and **revealed posteriors** for all chosen actions

$$\gamma_P^a \equiv \frac{P(a,\omega)}{P(a)}.$$

Strategies

- As in CD15, DM observed making choices from a finite set of decision problems $\{A^m\}_{m=1}^M$.
- Paired action sets and corresponding SDSC $\{(A^m, P^m)\}_{m=1}^M$ are the objects of analysis.
- As in Kamenica and Gentzkow (2011) learning a Bayes consistent distributions of posteriors $\gamma \in \Delta(\Omega)$:

$$\mathcal{Q} \equiv \{Q \in \Delta(\Delta(\Omega)) \text{ with } |\text{ supp } Q| < \infty | \sum_{\gamma \in \text{ supp } Q} \gamma Q(\gamma) = \mu\}.$$

Strategies

- Post-learning mixed action strategy based on the resulting posterior $q(a|\gamma)$.
- The overall strategy space is $\Lambda(A)$:

$$\Lambda(A) \equiv \{(Q,q)|Q \in \mathcal{Q}, q : \text{supp } Q \longrightarrow \Delta(A)\}.$$

- Questions concern how to interpret strategies that might have produced observed data. Define $P_{(Q,q)}$ as the SDSC that any strategy $(Q,q) \in \Lambda(A)$ would generate,

$$P_{(Q,q)}(a,\omega) = \sum_{\gamma \in \text{supp } Q} q(a|\gamma)Q(\gamma)\gamma(\omega).$$

- RHS: add across possible posteriors how likely is action given posterior, how likely is posterior, how likely is state given posterior.

CIR

- CIR defined as in CD15 but now with **unknown utility**.
- Costs of learning: $K : \mathcal{Q} \to \bar{\mathbb{R}}$ extended real line.
- Associate with each strategy $(Q, q) \in \Lambda(A)$ the corresponding expected prize utility less costs of learning:

$$V(Q, q|u, K) \equiv \sum_{\gamma \in \text{supp Q}} \sum_{a \in A} Q(\gamma) q(a|\gamma) \sum_{\omega \in \Omega} \gamma(\omega) u(z(a, \omega)) - K(Q).$$

- The value function and the corresponding optimal strategies standard:

$$\begin{aligned} \hat{V}(A|u, K) &\equiv \sup_{\{(Q,q) \in \Lambda(A)\}} V(Q, q|u, K) \\ \hat{\Lambda}(A|u, K) &\equiv \left\{ (Q, q) \in \Lambda(A) \,|\, V(Q, q|u, K) = \hat{V}(A|u, K) \right\}. \end{aligned}$$

CIR

- Goal is to identify necessary and sufficient conditions on $\{(A^m, P^m)\}_{m=1}^M$ for there to exist a fixed utility function over prizes and a learning cost function such that the data can be rationalized by a model in which total expected prize utility less learning costs is maximized.
- Logic is sequential. The key is to identify utility functions for which costs of learning and strategies can be identified that together rationalize the data.

Definition

$\{(A^m, P^m)\}_{m=1}^M$ have a **costly information representation** (CIR) if there exist u, K and $(Q^m, q^m) \in \hat{\Lambda}(A^m|u, K)$ for $1 \leq m \leq M$ such that $P^m = P_{(Q^m, q^m)}$. u **admits** a CIR of $\{(A^m, P^m)\}_{m=1}^M$ if there exists K and $(Q^m, q^m) \in \hat{\Lambda}(A^m|u, K)$ that in combination with u provide a CIR.

CCR and FIR

- Technically, data admit a **capacity constrained representation** (CCR) if the learning in each decision problem is optimal for some fixed feasible set of experiments $\mathcal{Q}^* \subset \mathcal{Q}$.
- Nested: CCR is a CIR in which costs of all revealed information structures are equal and those of all others are infinity.
- Again logic sequential. We want to pin down utility functions that admit a CCR.
- A **fixed information representation** (FIR) is a special case of a CCR in which the feasible set of learning strategies is a singleton.

2. Motivating example

Example

- Two choice sets and two equiprobable states, $\Omega = \{\omega_1, \omega_2\}$ with $\mu = (\frac{1}{2}, \frac{1}{2})$.
- Three prizes, z_B, z_M, z_G; data will reveal good, medium, and bad prizes, with $u_G > u_M > u_B$.
- Three actions a_k for $1 \leq k \leq 3$ with corresponding prizes as follows:

Action	**State** ω_1	**State** ω_2
a_1	z_G	z_B
a_2	z_B	z_G
a_3	z_M	z_M

Revealed Posteriors

- Two decision problems are faced: $A^1 = \{a_1, a_2\}$ and $A^2 = \{a_1, a_2, a_3\}$. SDSC recorded in matrix row actions column states

$$P^1 = \begin{array}{c} \\ \\ \\ \end{array}\begin{array}{cc} \omega_1 & \omega_2 \end{array} \qquad$$

$$P^1 = \begin{array}{cc} & \begin{array}{cc} \omega_1 & \omega_2 \end{array} \\ & \begin{pmatrix} .4 & .1 \\ .1 & .4 \end{pmatrix} \end{array} \begin{array}{c} \\ a_1 \\ a_2 \end{array} \quad \& \quad P^2 = \begin{array}{c} \begin{array}{cc} \omega_1 & \omega_2 \end{array} \\ \begin{pmatrix} .25 & .0 \\ .05 & .2 \\ .2 & .3 \end{pmatrix} \end{array} \begin{array}{c} \\ a_1 \\ a_2 \\ a_3 \end{array}$$

- Column sums reflect prior, row sums unconditional action probabilities.
- Data set P^1 reveals symmetric partial understanding.
- Data set P^2 is somewhat asymmetric. It reveals partial understanding when actions a_2, a_3 are chosen, full understanding only when a_1 is.
- Best summary of data is **revealed information structure**.

Revealed Posteriors

$$P^1 = \begin{array}{cc} \omega_1 & \omega_2 \\ \end{array} \quad \begin{pmatrix} .4 & .1 \\ .1 & .4 \end{pmatrix} \begin{array}{c} a_1 \\ a_2 \end{array} \quad \& \quad P^2 = \begin{pmatrix} .25 & .0 \\ .05 & .2 \\ .2 & .3 \end{pmatrix} \begin{array}{c} a_1 \\ a_2 \\ a_3 \end{array}$$

- **Unconditional choice probabilities** (row sums):
  $P^1(a_1) = P^1(a_2) = P^2(a_3) = 0.5$ and $P^2(a_1) = P^2(a_2) = 0.25$
- **Revealed posteriors** (normalized rows):

$$\gamma_{P^1}^{a_1} = \begin{pmatrix} .8 \\ .2 \end{pmatrix} \begin{array}{c} \omega_1 \\ \omega_2 \end{array} \quad \& \quad \gamma_{P^1}^{a_2} = \begin{pmatrix} .2 \\ .8 \end{pmatrix} \begin{array}{c} \omega_1 \\ \omega_2 \end{array}$$

$$\gamma_{P^2}^{a_1} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \begin{array}{c} \omega_1 \\ \omega_2 \end{array} \quad \& \quad \gamma_{P^2}^{a_2} = \begin{pmatrix} .2 \\ .8 \end{pmatrix} \begin{array}{c} \omega_1 \\ \omega_2 \end{array} \quad \& \quad \gamma_{P^2}^{a_3} = \begin{pmatrix} .4 \\ .6 \end{pmatrix} \begin{array}{c} \omega_1 \\ \omega_2 \end{array}$$

BEU and NIAS

- All representations require post-learning choice rationalizable by Bayesian expected utility maximization. Differ in what is learned before deciding
- $u : Z \longrightarrow \mathbb{R}$ can rationalize P in A if there exists a strategy $(\hat{Q}, \hat{q}) \in \Lambda(A)$ that produces the data such that all actions chosen are optimal at each posterior.
- CM 15 establish that this is equivalent to ruling out all holistic action switches as improving according to the data itself using the **no improving action switch (NIAS)** inequalities,

  $$\sum_{\omega \in \Omega} u(a, \omega) P(a, \omega) \geq \sum_{\omega \in \Omega} u(b, \omega) P(a, \omega),$$

  all $a, b \in A$.
- In all that follows such a utility function exists that rationalizes full data $\{(A^m, P^m)\}_{m=1}^{M}$.

NIAS

- NIAS constrains prize utilities. The fact that when a_1 is chosen in A^2 the state is ω_1 for sure is consistent with expected utility maximization only if utilities satisfy

$$u_G \geq \max\{u_M, u_B\}.$$

- a_3 chosen at revealed posterior $(0.4, 0.6)$ in A^2 reveals u_M to be at least as good as action a_2 at that posterior, which would yield z_G with probability 0.6, else z_B

$$u_M \geq 0.6u_G + 0.4u_B.$$

- a_2 chosen at revealed posterior $(0.2, 0.8)$ in both A^1, A^2 shows the resulting lottery which yields z_G with probability 0.8, else z_B, to be at least as good as z_M

$$0.8u_G + 0.2u_B \geq u_M.$$

NIAS

- Putting these together, the only non-vacuous representation requires strict preference $u_G > u_M > u_B$. Further, if we normalize to $u_G = 1$ and $u_B = 0$, there are upper and lower bounds on u_M,

$$0.6 \leq u_M \leq 0.8.$$

CIR

- CIR equivalent to costs K^1, K^2 such that maximized expected utility net of learning costs in action set A^1 is higher in data set P^1 than in alternative data set P^2, and conversely that maximized expected utility net of learning costs in action set A^2 is higher in data set P^2 than in alternative data set P^1.

CIR

- Given that $u_G = 1$ and $u_B = 0$ and $u_M \in [0.6, 0.8]$:

$$
\begin{aligned}
0.8 - K^1 &\geq 0.75 - K^2; \\
0.45 + 0.5u_M - K^2 &\geq 0.8 - K^1;
\end{aligned}
$$

- Top inequality: LHS odds of good prize 80% in P^1. Shifting A^1 to P^2 would be 60% with probability 0.5 (revealed posterior (0.4,0.6)) and 80% and 100% each with probability 0.25 (revealed posteriors (0.2,0.8) and (1,0)). Overall odds 0.75 as on RHS.
- Second inequality LHS: A^2 in P^2 yields z_M probability 0.5, z_G with probability 0.25 with a_1 and 0.2 with a_2 is.
- Second inequality RHS: with $u_M \leq 0.8$, A^2 in P^1 yields z_G with probability 0.8.

CIR

- In combination we have the following:

$$K^1 - K^2 \in [0.35 - 0.5u_M, 0.05].$$

- Note that this set is non-empty for all $u_M \in [0.6, 0.8]$.
- When $u_M = 0.6$, the only rationalizing cost is $K^1 - K^2 = 0.05$.
- When $u_M = 0.7$,

$$K^1 - K^2 \in [0, 0.05].$$

- When $u_M = 0.8$,

$$K^1 - K^2 \in [-0.05, 0.05].$$

CCR

- Not all utility functions that admit a CIR also admit a CCR.
- The easiest way to see this is to note that conditions for a CCR are equivalent to existence of a CIR with equal costs.
- This holds only for $u_M \in [0.7, 0.8]$.

FIR

- Logic for FIR very different and very restrictive
- Here is the simplest possible information structure that allows the data to be rationalized. Defining this structure by the posterior probability of state ω_1 we have:

$$Q(1) = Q(0.6) = 0.25; Q(0.2) = 0.5.$$

- In this case the unique strategies that rationalize the data are:

$$\begin{aligned} q^1(a_1|1) &= q^1(a_1|0.6) = q^1(a_2|0.2) = 1; \\ q^2(a_1|1) &= q^2(a_3|0.6) = 1; \\ q^2(a_2|0.2) &= q^2(a_3|0.2) = 0.5. \end{aligned}$$

FIR

- That the above strategies are optimizing is clear. All deterministic strategies involve actions that are uniquely optimal at the corresponding posterior, while the mixed strategy in data set A^2 in which both a_2 and a_3 are chosen when the probability of state ω_1 is 0.2 reflects the fact that both yield equal utility of $u_M = 0.8$ at this posterior.
- To confirm that this produces the data consider by way of illustration choice of action a_1 in data set P^1:

$$\begin{aligned} P_{(Q,q^1)}(a_1,\omega_1) &= Q(1)q^1(a_1|1)*1 + Q(0.6)q^1(a_1|0.6)*0.6 = 0.4; \\ P_{(Q,q^1)}(a_1,\omega_2) &= Q(0.6)q^1(a_1|0.6)*0.4 = 0.1; \end{aligned}$$

with precisely analogous logic for a_2 in data set P^1. One can confirm same for data set P^2. Figure illustrates

FIR

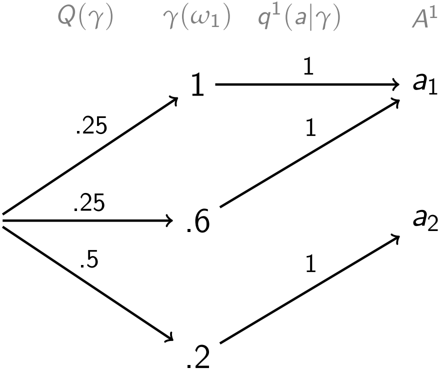

$P^1(a, \omega_1)$

$0.4 = .25 \times 1 \times 1 + .25 \times .6 \times 1$

$0.1 = .5 \times .2 \times 1$

FIR

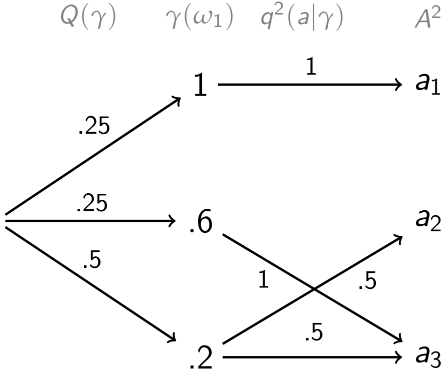

$P^2(a, \omega_1)$

$0.25 = .25 \times 1 \times 1$

$0.05 = .5 \times .2 \times .5$

$0.20 = .25 \times .6 \times 1 + .5 \times .2 \times .5$

FIR

- Why only $u_M = 0.8$ for an FIR?
- The key is the combination of unconditional action probabilities and revealed posteriors of chosen actions.
- Recall data set P^1 has two equiprobable revealed posteriors of state ω_1 which are recorded with their unconditional probabilities as:

$$\gamma_{P^1}^{a_1} \equiv 0.8; \text{ and } P^1(a_1) = 0.5$$
$$\gamma_{P^1}^{a_2} \equiv 0.2; \text{ and } P^1(a_2) = 0.5$$

- In contrast data set P^2 has three revealed posteriors of state ω_1 again recorded with their unconditional probabilities as:

$$\begin{aligned} \gamma_{P^2}^{a_1} &\equiv 1; &\quad \text{and} \quad & P^2(a_1) = 0.25 \\ \gamma_{P^2}^{a_2} &\equiv 0.2; &\quad \text{and} \quad & P^2(a_2) = 0.25; \\ \gamma_{P^2}^{a_3} &\equiv 0.4; &\quad \text{and} \quad & P^2(a_2) = 0.5; \end{aligned}$$

FIR

- Key observation: choice of a_2 in both choice sets A^1 and A^2 has a common revealed posterior probability of state ω_1 of 0.2 but a probability of 0.5 rather than 0.25 in A^1 rather than in A^2.
- Hence there must be some posterior at which a_2 is chosen in A^1 at which a_3 was chosen in A^2. By optimality, these have to be the higher posteriors in the range, if any.
- If indeed there were any possible posteriors in the range $(0.2, 0.5]$ at which a_3 was chosen in A^2 while a_2 was in A^1, their removal would strictly lower the revealed posterior of that action in decision problem A^2, which they do not.

FIR

- Hence 0.2 must be the only possible posterior in the range $[0, 0.5)$ and it must have unconditional probability $Q(0.2) = 0.5$.
- To explain the data we then need the mixed strategy at this posterior to assign equal probability to actions a_2, a_3.
- It is precisely the need for this to be consistent with an optimal strategy that pins down $u_M = 0.8$ as necessary for an FIR.

FIR

- Continuing on with the common information structure, note that $Q(1) = 0.25$.
- It must be at least that high to explain $\gamma_{P^2}^{a_1} \equiv 1$; and $P^1(a_1) = 0.25$. It cannot be higher since it is strictly optimal to choose a_1 at this posterior. In fact there is no possible posterior in the range $(0.8, 1]$ since if there were, optimality would imply a strictly higher probability of choosing action a_1 in data set P^2.
- We are left only to find the final 0.25 uncommitted posterior probabilities.
- In data set P^1 we know that the average posterior when a_1 is chosen is 0.8, and that it is chosen with probability 0.5. Since we also know that the posterior is precisely 1 with probability 0.25, it must average precisely 0.6 otherwise.

FIR

- To summarize the remaining 0.25 probability must all be assigned to posteriors at which a_1 is optimal in data set A^1 and a_3 is optimal in action set A^2. This sets the support as $[0.5, 0.8]$. Further we know that in this range the average posterior must be 0.6.
- Finally the average posterior other than 0.2 at which a_3 is chosen in data set P^2 must be 0.6 to rationalize the revealed posterior of 0.4. But this is already implied above, so adds no new conditions. We conclude that there are indeed no more restrictions.
- What this means is that we can generalize the particular example provided in which $Q(0.6) = Q(1) = 0.25$ and $Q(0.2) = 0.5$ in only one respect. We can replace $Q(0.6)$ with any set of posteriors on the support $[0.5, 0.8]$ that average to 0.6 and then set the corresponding strategy of deterministically choosing a_1 at all such posteriors in A^1 and a_3 at all such posteriors in A^2.

4. Characterization of Fixed Information Representation

Mean and Optimality Preserving Spread

- There is a simple organizing logic for FIR that is a variation of the mean preserving spread characterization of more informative information structures due to Blackwell.
- If we disregard the requirement of optimality, an information structure can rationalize data if and only if it is at least as Blackwell informative as the revealed information structure by the mean preserving spread characterization of the Blackwell order based on a Markov matrix.
- To preserve optimality at all possible posteriors: one needs to start with optimality and then perform an optimality check in each row of the corresponding Markov matrix
- There are technical subtleties outlined in the appendix

Mean and Optimality Preserving Spread

- The Markov matrix B that defines the MPS has row cardinality corresponding to the number of distinct actions chosen in the data with column cardinality corresponding to the support of posterior distribution Q that defined an MPS. Optimality is defined at the level of action choice.
- Further we need to specify a utility function and a particular decision problem in order to define optimality.
- We also need to insist that the original data satisfy NIAS for that utility function so that there is optimality prior to application of the B matrix. If this did not hold optimality would fail at the outset, leaving nothing to preserve.
- With this, maintenance of optimality states that all posteriors that are possible according to the matrix in the row corresponding to chosen action a^i retain its optimality.

Mean and Optimality Preserving Spread

- With all technicalities in place, we know how to define an information structure Q as a **mean and optimality preserving spread** (MOPS) of choice data P
- The characterization theorem is based on identifying a common mean and optimality preserving spread of all observed data.
- The starting point is a utility function for which NIAS is satisfied in all decision problems.
- In addition to identifying precisely when a FIR exists, the theorem identifies mixed action strategies from the matrix.
- The procedure is to set these according to the relative probabilities in column j when unconditional action probabilities in the data are run through the transition matrix.

Characterization

Theorem

Given $\{(A^m, P^m)\}_{m=1}^{M}\}$ any u satisfying NIAS admits a FIR if and only if there exist an information structure Q and for all $1 \leq m \leq M$ Markov matrices B^m that define Q as a MOPS of P^m. Labeling chosen actions in A^m by a^{mi} for $1 \leq i \leq I_m$ the corresponding $q^m : \text{ supp } Q \longrightarrow \Delta(A^m)$ are pinned down as:

$$q^m(a^{mi}|\gamma^j) = \frac{P^m(a^{mi})B^m_{(mi)j}}{Q(\gamma^j)}. \tag{1}$$

5. Characterization of Capacity Constrained Representation

The NIAS Cone

- CM21 introduce an NIAS Cone characterization of utility functions that permit a Bayesian expected utility representation
- This involves ensuring that no counterfactual holistic action switches produce higher expected utility
- Intuitively applied in example: will see again at end.
- All the remaining characterizations extend that idea using richer counterfactual switches

Action and Attention Switches

- The characterization of utility functions that admit a CCR of $\{(A^m, P^m)\}_{m=1}^{M}$ is based on ruling out a set of feasible changes in action choices and information structures as improving.
  1. Start with a given decision problem A^m and select a target data set P^n for $1 \leq m, n \leq M$. It is allowed that $m = n$.
  2. Map chosen actions in $\mathcal{A}_{P^n}$ to any actions in set A^m according to $f : \mathcal{A}_{P^n} \longrightarrow A^m$
  3. Given any triple (m, n, f) define the corresponding SDSC on $a \in A^m$ as:

$$P^{(m,n,f)}(a, \omega) \equiv \sum_{\{b \in A^n | f(b) = a\}} P^n(b, \omega)$$

  4. Define the prize lottery associated with the actual observed data in A^m, $L(P^m)$, and the prize lottery after all action switches, $L(P^{(m,n,f)})$.

Action and Attention Switches

- The defining feature of a CCR is that this alternative data could have been generated and used with action set A^m, so that none of these switches can be improving.
- Index all distinct triples (m, n, f) by $1 \leq h \leq H$ to simplify notation in representing the implied restrictions on utility and define the corresponding change in lottery as D^h,

$$D^h \equiv L(P^{m(h)})) - L(P^{(m(h),n(h),f(h))})$$

Action and Attention Switches

- We call the convex cone defined by all such inequalities the no improving action or attention switch (NIIAS) cone since this is the set of operations that are ruled out as improving.

Definition

We call the convex cone D^{NIAAS} formed by all D^h the *NIAAS cone*

$$\mathcal{D}^{NIAAS} = \{D = \sum_{h=1}^{H} \alpha^h D^h \in \mathbb{R}^K | \alpha^h \in \mathbb{R}_+\}.$$

NIAAS Cone Characterization of CCR

- A clear necessary condition for $u : Z \longrightarrow \mathbb{R}$ to admit a CCR representation is that no such change raise utility,

$$\sum_{k=1}^{K} u_k D_k^h \geq 0$$

Definition

We call the set of all utility functions $u : Z \longrightarrow \mathbb{R}$ such that

$$\sum_{k=1}^{K} u_k D_k \geq 0$$

all $D \in \mathcal{D}^{NIAAS}$ the *NIAAS utility cone*

Theorem

Utility function $u : Z \longrightarrow \mathbb{R}$ admits a CCR if and only if it lies in the NIAAS utility cone.

- This is a corollary of the next result characterizing CIR

6. Characterization of Costly Information Representation

Attention Cycles

- As for the CCR, the characterization of utility functions that admit a CIR of $\{(A^m, P^m)\}_{m=1}^{M}$, is based on ruling out a set of feasible changes in action choices and learning strategies as improving.
- CD15 identify that NIAS and NIAC are necessary and sufficient conditions for a given utility function to provide a CIR.
- For the current context we need to adapt the reasoning to allow for the utility function to be unknown. What we will show is that the CM21 approach can be adapted based on rich lottery comparisons that are implicit in the NIAC constraints.

Attention Cycles

- In formal terms:

Definition

A **cycle of attention and action switches** $(J, \vec{m}, \{f^j\}_{j=1}^{J-1})$ comprises:

1. Cycle length $2 \leq J \leq M+1$
2. A corresponding vector of indices $\vec{m} = (m(1), \cdots, m(J))$ with $1 \leq m(1) = m(J) \leq M$ but $1 \leq m(j) \neq m(j') \leq M$ otherwise.
3. A set of mappings $f^j : \mathcal{A}_{P^{m(j+1)}} \longrightarrow A^{m(j)}$ for $1 \leq j \leq J-1$.

- Note that direct switches in data are included by setting $J = 2$.

Attention Cycles

- At each stage $1 \leq j \leq J-1$ in any such cycle, we follow the definitions for a CIR in the last section to define the lottery in the data, the lottery received under the specified action switches in the ensuing data, and the difference between them.
- We then average up the corresponding lottery changes across the cycle to arrive at the overall lottery change which we denote $D^{(J,\vec{m},\{f^j\}_{j=1}^{J-1})}$
- The defining feature of a CIR is that any cycle of attention and action switches is in principle feasible at no additional cost and so cannot be improving.

NIAC Cone Characterization of CIR

- To simplify the notation for expressing the implied constraint on utility we index the finite number of distinct such triples by h' for $1 \leq h' \leq H'$.
- We call the convex cone formed by all $D^{h'}$ the *NIAC cone*: direct cycles are included hence a subset of NIAS cone.

$$\mathcal{D}^{NIAC} = \{D' = \sum_{h'=1}^{H} \alpha^{h'} D^{h'} \in \mathbb{R}^K | \alpha^{h'} \in \mathbb{R}_+\}.$$

- A clear necessary condition for $u : Z \longrightarrow \mathbb{R}$ to admit a CIR is that no such change raise utility,

$$\sum_{k=1}^{K} u_k D_k^{h'} \geq 0$$

NIAC Cone Characterization of CIR

Definition

We call the set of all utility functions $u : Z \longrightarrow \mathbb{R}$ such that

$$\sum_{k=1}^{K} u_k D_k \geq 0$$

all $D \in \mathcal{D}^{NIAC}$ the *NIAC utility cone*

Theorem

Utility function $u : Z \longrightarrow \mathbb{R}$ admits a CIR if and only if it lies in the NIAC utility cone.

- This follows from CD15 logic.

NIAC Cone Characterization of CIR

- CD15 show how to identify all cost functions that provide a CIR for a given cost function
- CMM22 introduce a simple method of operationalizing
- Will present in ML session of Summer School

7. Recovery In Example

NIAS Cone in Example

- Revisit examples to think in terms of counterfactual prize lotteries (p_G, p_M, p_B).
- NIAS utility cone only counterfactual lotteries from any action switches within each decision problem
- With A^1 the possible action switches are choosing a^2 in place of a^1 and vice versa.
- In both cases the actual choice yields lottery $(0.8, 0, 0.2)$ and the counterfactual yields lottery $(0.2, 0, 0.8)$.
- The difference from sticking with the action rather than switching is $(0.6, 0, -0.6)$, and this cannot be improving, so that the vector of EUs (u_G, u_M, u_B) must have a positive dot product with this vector,

$$0.6u_G \geq 0.6u_B,$$

NIAS Cone in Example

- In choice set A^2, consider first action switches of choosing either a^2 or a^3 in place of a^1.
- When chosen a^1 produces the good prize for sure, $(1, 0, 0)$.
- Switching to a^3 cannot be improving since it yields $u_B \leq u_G$.
- Switching to a^2 would yield z_3 for sure, so that the corresponding NIAS inequality implies that $u_G \geq u_M$.
- Switching from a^2 to a^1 cannot be improving since all it does is lower the probability of receiving z_G rather than z_B.
- For it to be non-improving not to switch from a^2 to a^3 requires that the resulting lottery $(0.8, 0, 0.2)$ be at least as good as the reward of z_M for sure achievable by switching to a^3,

$$0.8u_G + 0.2u_B \geq u_M.$$

NIAS Cone in Example

- Finally, consider possible switching from a^3 to a^2 or a^1. No improvement requires that the actual reward of z_M for sure be at least as good as the best alternative, which would be to switch to a^1 and get the lottery $(0.6, 0, 0.4)$. Hence,

$$u_M \geq 0.6u_G + 0.4u_B.$$

- Overall the only non-vacuous solutions involve $u_G > u_M > u_B$.
- With this we can normalize to $u_G = 1$ and $u_B = 0$ and conclude that the only other restriction is $u_M \in [0.6, 0.8]$.

NIAC Cone in Example

- The inequalities defining the NIAC cone are a superset of those defining the NIAS cone, so all that remains is to rule out improving attention cycles.
- To work out the resulting inequalities we separately analyze the implications of switching A^1 to P^2 and A^2 to P^1.
- As a first step in this direction we compute the full prize lotteries from the actual choices. In choice set A^1 the realized lottery in P^1 is $(0.8, 0, 0.2)$. In choice set A^2 the realized lottery in P^2 is $(0.45, 0.5, 0.05)$.
- The implications of switching A^1 to P^2 are straight forward. At each posterior it is best to choose the action that is more likely to yield prize u_G rather than u_B. The average such probability in P^2 is 0.75 rather than 0.8. Hence sticking with the chosen attention yields a lottery change of $(0.05, 0, -0.05)$ relative to the best alternative. axes corresponding to the second and third prizes, both measuring probabilities and utilities.

NIAC Cone in Example

- The implications of switching A^2 to P^1 are equally simple: in either case the maximum utility derives from picking the action that yields lottery $(0.8, 0, 0.2)$ given that we have already established $u_M \leq 0.8$. Hence the difference between sticking with the actual attention strategy rather than switching is defined by lottery difference $(-0.35, 0.5, -0.15)$.
- To derive the NIAC inequality from the cycle in which A^1 is switched to P^2 and A^2 is switched to P^1 we average the two changes. This averaged lottery change is

$$\frac{(-0.35, 0.5, -0.15) + (0.05, 0, -0.05)}{2} = (-0.15, 0.25, -0.1).$$

- The corresponding inequality on utilities that make this have positive utility given what is already known is,

$$0.25u_M - 0.25 \geq 0 \iff u_M \geq 0.6.$$

- Confirms NIAC cone the same as the NIAS cone

NIAAS Cone in Example

- What of the NIAAS Cone?
- The key difference between the NIAC Cone and the NIAAS Cone lies in the replacement of attention cycles with attention switches.
- So we need to rule out separately any advantage of switching A^1 to P^2 or switching A^2 to P^1.
- As just computed, it is clear that there is a pure loss in switching A^1 to P^2 since it merely lowers the probability of winning prize z_G rather than prize z_B.
- However ensuring that there is no utility gain from switching A^2 to P^1 imposes an additional constraint. Specifically, it must be that the lottery difference $(-0.35, 0.5, -0.15)$ not be improving. It terms of utilities, this imposes the additional constraint

$$0.5u_M - 0.35 \geq 0 \iff u_M \geq 0.7.$$

NIAAS Cone in Example

- Again this confirms the direct logic of section 3 that a CCR exists if and only if $u_M \in [0.7, 0.8]$.
- Normalizing to $u_G = 1$ and $u_B = 0$ we can illustrate all utility cones in a figure with axes corresponding to the second and third prizes, both measuring probabilities and utilities.

Utility Recovery: Example

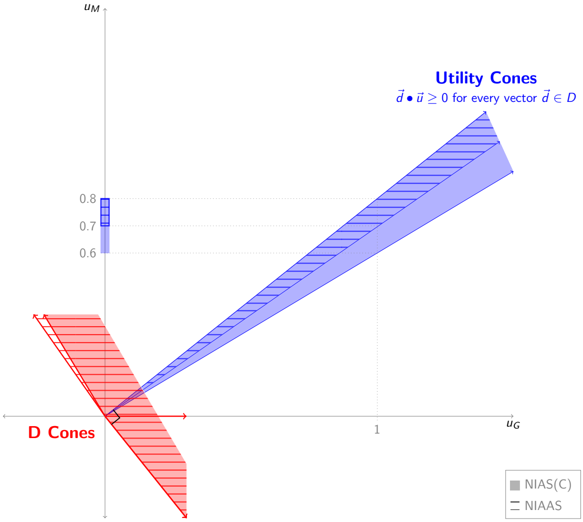

MOPS in Example

- Note that MOPS exists only for $u_M = 0.8$ by argument in the given example.
- To specify the transition matrices, recall rows are chosen actions and columns possible posteriors: we label these $\gamma^1 = (1, 0)$, $\gamma^2 = (0.6, 0.4)$ and $\gamma^3 = (0.2, 0.8)$
- Recall the simplest information structure Q that defines an MOPS of P^1 and P^2 is $Q(\gamma^1) = Q(\gamma^2) = 0.25$ and $Q(\gamma^3) = 0.5$.

MOPS in Example

- The 2×3 transition matrix B^1 is:

$$B^1 = \begin{array}{c} \\ \\ \end{array}\begin{array}{c} \gamma^1 \quad\;\; \gamma^2 \quad\;\; \gamma^3 \\ \begin{pmatrix} 0.5 & 0.5 & 0 \\ 0 & 0 & 1.0 \end{pmatrix} \end{array} \begin{array}{c} \\ a_1 \\ a_2 \end{array}$$

- Indeed a Markov matrix.
- The action strategies that rationalize the data are:

$$q^1(a_1|\gamma^1) = q^1(a_1|\gamma^2) = q^1(a_2|\gamma^3) = 1.$$

MOPS in Example

- The MOPS theorem applied to this case says that for $i = 1, 2, j = 1, 2, 3$

$$q^1(a_i|\gamma^j) = \frac{P^1(a_i)B^1_{ij}}{Q(\gamma^j)}. \tag{2}$$

- In confirmation

$$\begin{aligned}
\frac{P^1(a_1)B^1_{11}}{Q(\gamma^1)} &= \frac{0.25}{0.25} = 1 = q^1(a_1|\gamma^1);\\
\frac{P^1(a_1)B^1_{12}}{Q(\gamma^2)} &= \frac{0.25}{0.25} = 1 = q^1(a_1|\gamma^2);\\
\frac{P^1(a_2)B^1_{23}}{Q(\gamma^3)} &= \frac{0.5}{0.5} = 1 = q^1(a_2|\gamma^3).
\end{aligned}$$

MOPS in Example

- The 3×3 transition matrix B^2 is:

$$
B^2 = \begin{array}{c c}
\begin{array}{ccc} \gamma^1 & \gamma^2 & \gamma^3 \end{array} & \\
\begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0.5 & 0.5 \end{pmatrix} & \begin{array}{c} a_1 \\ a_2 \\ a_2 \end{array}
\end{array}
$$

- Confirmation of formula valuable exercise: in Appendix

7. Conclusion

What We Did

- We characterizes three leading Bayesian learning model classes within a unified framework
- Our characterizations make testing these models possible
- We characterize what can be known about utility
- And what can be known about learning
- Definitive results in our framework

Where We Are Going

- Several applications ongoing others on wish list
  1. Psychometrics
  2. Labor economics
  3. Machine learning
  4. Judicial, Medical
  5. Teaching and Testing
  6. ......
- Demonstrates ubiquity of Bayesian learning models and strong potential for interdisciplinary collaboration

Appendix

Defining mean and optimality preserving spread

- Defining a MOPS requires indexing posteriors revealed in the data and those in any rationalizing information structure.
- With regard to the former, we define I_P as the cardinality of $\mathcal{A}_P$, the set of actions chosen in SDSC P, and index these posteriors as γ_P^i with corresponding action a^i for $1 \leq i \leq I$.
- Correspondingly, we index posteriors in the support of information structure Q as γ^j for $1 \leq j \leq J$.
- With this we define what it means for an information structure Q to be a mean preserving spread of SDSC P.

Defining mean and optimality preserving spread

Definition

Q is a **mean preserving spread** of P if there there exists an $I \times J$ non-negative Markov matrix matrix B with $\sum_{j=1}^{J} B_{ij} = 1$ such that:

1. For each $1 \leq j \leq J$ the unconditional probabilities in Q are generated as P, B imply:
$$\sum_{i=1}^{I} P(a^i) B_{ij} = Q(\gamma^j).$$
2. The posteriors γ_P^i are generated as Q, B imply:
$$\gamma_P^i = \sum_{j=1}^{J} B_{ij} \gamma^j$$

Defining mean and optimality preserving spread

- In comparison with the standard definition in terms of comparing information structures, note that the Markov matrix B that defines the MPS has row cardinality corresponding to the number of distinct actions chosen in the data, I, with column cardinality corresponding to J, the cardinality of the support of posterior distribution Q.
- For a MOPS We want to ensure that optimality is maintained, in the sense that all posteriors that are possible according to the matrix in row i have the feature that action a^i remains optimal. We need to specify a utility function and a particular decision problem in order to define optimality and impose NIAS so that there is optimality prior to application of the B matrix.

Defining mean and optimality preserving spread

Definition

Given $u : Z \longrightarrow \mathbb{R}$ and finite action set A, the corresponding optimal posterior set is $\hat{\Gamma}(a|u, A) \subset \Delta(\Omega)$

$$\hat{\Gamma}(a|u, A) = \{\gamma| \sum_{\omega\in\Omega} u(z(a,\omega))\gamma(\omega) \geq \sum_{\omega\in\Omega} u(z(b,\omega))\gamma(\omega) \text{ all } b \in A\}.$$

Definition

Given (A, P) and u such that NIAS is satisfied, Q is a **mean and optimality preserving spread** of P if there exists a matrix B that defines it as a mean preserving spread in which optimality is preserved;

$$B_{ij} > 0 \implies \gamma^j \in \hat{\Gamma}(a^i|u, A).$$

MOPS in Example

- The 3×3 transition matrix B^2 is:

$$B^2 = \begin{array}{cccc} \gamma^1 & \gamma^2 & \gamma^3 & \\ \left(\begin{array}{c} 1 \\ 0 \\ 0 \end{array}\right. & \begin{array}{c} 0 \\ 0 \\ 0.5 \end{array} & \left.\begin{array}{c} 0 \\ 1 \\ 0.5 \end{array}\right) & \begin{array}{c} a_1 \\ a_2 \\ a_2 \end{array} \end{array}$$

- Indeed a Markov matrix.
- In this case the action strategies that rationalize the data are:

$$\begin{aligned} q^2(a_1|\gamma^1) &= q^2(a_3|\gamma^2) = 1; \\ q^2(a_2|\gamma^3) &= q^2(a_3|\gamma^3) = 0.5. \end{aligned}$$

MOPS in Example

- The MOPS theorem applied to this case says that for $i, j = 1, 2, 3$:

$$q^2(a_i|\gamma^j) = \frac{P^2(a_i)B^2_{ij}}{Q(\gamma^j)}. \tag{3}$$

- Recall that $P^2(a_1) = P^2(a_2) = 0.25$ and $P^2(a_3) = 0.5$.
- Recall also that $Q(\gamma^1) = Q(\gamma^2) = 0.25$ and $Q(\gamma^3) = 0.5$

MOPS in Example

- In confirmation of the general formula:

$$
\begin{aligned}
\frac{P^2(a_1)B_{11}^2}{Q(\gamma^1)} &= \frac{0.25}{0.25} = 1 = q^2(a_1|\gamma^1); \\
\frac{P^2(a_2)B_{23}^2}{Q(\gamma^3)} &= \frac{0.25}{0.5} = 0.5 = q^1(a_2|\gamma^3); \\
\frac{P^2(a_3)B_{32}^2}{Q(\gamma^2)} &= \frac{0.5 * 0.5}{0.25} = 1 = q^2(a_2|\gamma^3); \\
\frac{P^2(a_3)B_{33}^2}{Q(\gamma^3)} &= \frac{0.5 * 0.5}{0.5} = 0.5 = q^2(a_3|\gamma^3);
\end{aligned}
$$