# Multidimensional and Selective Learning

## Sloan-Nomis Workshop

Srijita Ghosh

New York University

February 22, 2019

# Motivation

- DM needs to choose multiple attributes/inputs
- Any arbitrary correlation allowed between attributes
- Optimal learning strategy: what and how much to learn?
- Examples: agricultural input choices, job assignment etc.

# Why Multidimensional Decision Problem?

- **Selective Learning:** Productivity of only a subset of attribute learned separately (Hanna et al 2014, Bloom et al 2014 etc)
- Belief about correlation affects learning choice

# Research Agenda

- Solve for the optimal learning strategy in the multidimensional choice setting
- Find conditions under which selective learning is optimal
- Policy Implication: optimal information provision
  - extension services
  - management practices training

# Choice Problem: $n = 2$ inputs

- $A = X \times Y$; action/decision space
- $X = \{x_1, x_2, \ldots x_n\}$, $Y = \{y_1, y_2, \ldots y_n\}$: two inputs
- $Y = \{0, 1\}$: output
- $\pi : A \to Y$: payoff function
- $\Omega$: state space (set of all possible payoff function),
- $\omega$: one realization of $\pi$, typical state
- $\mu_0 \in \Delta(\Omega)$ : prior
- $\mu^*$: true state; $supp(\mu^*) \subseteq supp(\mu_0)$

# Learning Technology

| | $y_1$ | $y_2$ | $y_3$ | |
|---|---|---|---|---|
| $x_1$ | 1 | 0 | 0 | 1/3 |
| $x_2$ | 0 | 1 | 1 | 2/3 |
| $x_3$ | 0 | 0 | 0 | 0 |
| | 1/3 | 1/3 | 1/3 | |

Table 1: Example of payoff matrix/state

- DM can uncover any cell for a fixed cost of $c_l$: observe 0 or 1
- DM can uncover any average for a fixed cost of $c_a$: observe the true row or column average
- DM can open any number of cells or average in any order sequentially
- Bayesian updating

# Properties of belief

- Expected payoff given belief $\mu_t$ : $\pi_t$
- Uncertainty given belief $\mu_t$: $H(\mu_t)$

  i. $H(\mu_t)$: Shannon entropy of belief at round $t$

  $$H(\mu_t) = -\sum_{\omega \in \Omega} \mu_t(\omega) \ln \mu_t(\omega)$$

  ii. $H(\mu_t) \in \mathbf{R}_+$

# Decision Problem

## Learning strategy

Learning strategy specifies a conditional sequence of observations ($\gamma(\mathcal{P})$) such that observation of $t^{th}$ round depends on the belief after of the $(t-1)$ observations.

## Decision Problem

DM chooses a learning strategy to maximize his expected payoff subject to the cost of learning,

$$W(\mu_0) = \max_{\gamma(\mathcal{P})} E\left[\pi(\mathbf{a}_{ij}) - c(\mathcal{P}) | \gamma, \mu_0\right] \tag{DP}$$

# Recursive Problem: Main Results

1. When is learning optimal?
   - Uncertainty is in an interval, i.e., not too low or too high
2. Whether to observe a cell of an average?
   - Higher uncertainty $\Rightarrow$ observe a average
   - Lower Uncertainty $\Rightarrow$ observe a cell
3. Which cell or average to observe?
   - Cell: highest one-round ahead expected payoff
   - Average: Reduces most uncertainty

# Recursive Problem: Main Results

1. When is learning optimal?
   - Uncertainty is in an interval, i.e., not too low or too high
2. Whether to observe a cell of an average?
   - Higher uncertainty $\Rightarrow$ observe a average
   - Lower Uncertainty $\Rightarrow$ observe a cell
3. Which cell or average to observe?
   - Cell: highest one-round ahead expected payoff
   - Average: Reduces most uncertainty

# Recursive Problem: Main Results

1. When is learning optimal?
    - Uncertainty is in an interval, i.e., not too low or too high
2. Whether to observe a cell of an average?
    - Higher uncertainty $\Rightarrow$ observe a average
    - Lower Uncertainty $\Rightarrow$ observe a cell
3. Which cell or average to observe?
    - Cell: highest one-round ahead expected payoff
    - Average: Reduces most uncertainty

# Stopping problem: Main results

- **Optimal learning strategy**: start with averages then switch to cell permanently
- **Selective Learning**: optimal if averages are sufficiently informative and learning is costly
- **Policy Implication**: reducing only one cost can decrease learning

# Alternate Mechanism

- Sequential search: informationally inefficient
- Optimal Categorization: no bias-variance trade-off

# Higher Dimensions: General Cost functions?

- For $n > 2$: scalability issues
- **Research Question:** does there exist a cost of learning function that is observationally equivalent to the prescribed learning mechanism?
- Symmetric Prior: Shannon entropy

# Breadth vs Depth: Tree algorithms

- For $n > 2$ tree representation more tractable
- Tree structure assumes sequence of observation