Learning Time Dependent Choice

Zachary Chase   Siddharth Prasad

Caltech

ITCS 2019

Motivation

Intertemporal choice: how people make choices over time.

- Discounted utility: I prefer \$10 today over \$10 a year from today. What about \$10 today versus \$100 a year from today?

Discounted utility models of intertemporal choice widely used in economics and by researchers in other fields:

- Model savings and borrowing decisions
- Evaluating climate change policies
- Self control in humans and animals

Predictions made by these economic models match neurobiological data obtained via MRI scans.

Setup - preferences

- T time periods.
- $x \in \mathbf{R}^T$ are *plans*
- $\succsim \subseteq \mathbf{R}^T \times \mathbf{R}^T$ is a preference
- A preference model $\mathcal{P}$ is a collection of preferences

Setup - learning

- Agent makes choices from pairs $(x, y) \in \mathbf{R}^T \times \mathbf{R}^T$ according to $\succsim \in \mathcal{P}$
- After observing finitely many choices output hypothesis that w.h.p. is very close to $\succsim$.
- PAC Learning:
  - Questions (x, y) drawn from an unknown distribution on $\mathbf{R}^T \times \mathbf{R}^T$, receive the agent's choice for every question drawn.
  - Sample complexity of PAC learning is $O\left(\frac{1}{\varepsilon}\left(VC(\mathcal{P}) + \log\frac{1}{\delta}\right)\right)$.
- Active Learning:
  - Stream model: pairs drawn from unknown distribution, but analyst can choose whether or not to request agent response.
  - Membership queries: analyst directly chooses questions to ask the analyst.

Previous work on learning economic choice parameters

- Kalai (2001). Learnability of choice functions: observe sets of alternatives along with most preferred alternative from each set.
- Beigman and Vohra (2006), Zadimoghaddam and Roth (2012), Balcan et al. (2014). Learning utility functions from revealed preference: observe chosen bundles of goods when faced with prices and budget constraint.
- Basu and Echenique (2018). Learning choice under uncertainty. Preference relations given by $x \succsim y$ iff x yields higher expectation than y, for various notions of subjective expectation.

Discounted utility

Discounted utility model of preferences $\mathcal{P}_{\mathcal{D}}$: preferences $\succsim$ such that there is a decreasing map $D : \{1, \ldots, T\} \to (0, 1)$ where

$$x \succsim y \iff \sum_{t=1}^{T} D(t)x_t \geq \sum_{t=1}^{T} D(t)y_t.$$

Exponential discounting $D(t) = \delta^t$, $\delta \in (0, 1)$.

Hyperbolic discounting $D(t) = \frac{1}{1+t\alpha}$, $\alpha > 0$.

Quasi-hyperbolic discounting $D(t) = \begin{cases} 1 & t = 1 \\ \beta \cdot \delta^{t-1} & t > 1 \end{cases}$

Discounted utility

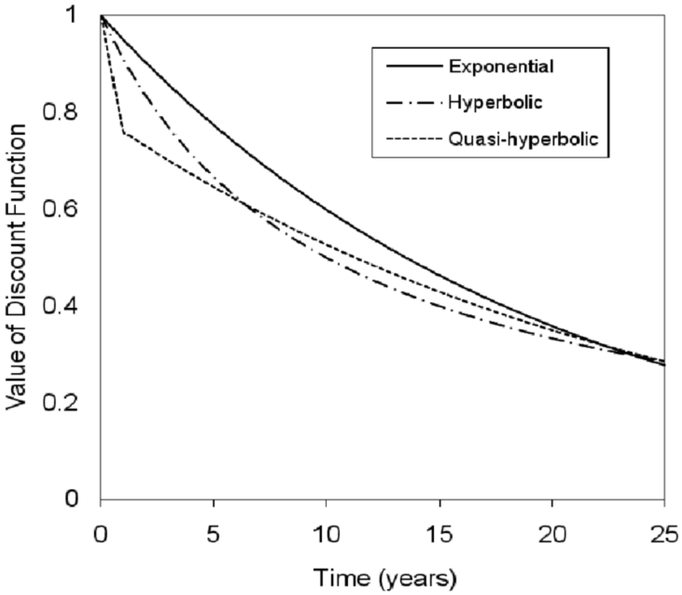

Need structure on discounting for fast learning!

Most general preference model with weights: $\mathcal{P}_{\mathcal{W}}$ where $x \succsim y$ iff $w.x \geq w.y$ for some $w \in \mathbf{R}^T$.

- $T - 1 \leq VC(\mathcal{P}_{\mathcal{W}}) \leq T + 1$.

With no structure to discounting, cannot improve this:

Proposition

$T - 1 \leq VC(\mathcal{P}_{\mathcal{D}}) \leq T + 1.$

Want structural conditions on the preference model that yield better learning results and capture the commonly used discounting models (exponential, hyperbolic, quasi-hyperbolic).

A structural result

$Q_1, \ldots, Q_T$ polynomials of degree $\leq d$.

- Preference model $\mathcal{P}_{\mathcal{PW}}$: $x \succsim y$ iff $\sum_{t=1}^{T} Q_t(\delta)x_t \geq \sum_{t=1}^{T} Q_t(\delta)y_t$.
- Preference model $\mathcal{P}_{\mathcal{BPW}}$: $x \succsim y$ iff $\left(\frac{1}{\beta} - 1\right) \sum_{t=1}^{T} Q_t(0)x_t + \sum_{t=1}^{T} Q_t(\delta)x_t \geq$ $\left(\frac{1}{\beta} - 1\right) \sum_{t=1}^{T} Q_t(0)y_t + \sum_{t=1}^{T} Q_t(\delta)y_t$.

Exponential improvement in learning:

Theorem

- *For all* $\varepsilon > 0$, $VC(\mathcal{P}_{\mathcal{PW}}), VC(\mathcal{P}_{\mathcal{BPW}}) \leq (1+\varepsilon)\log d$ *for large enough* d.
- *If* $d \leq T-1$, *and the* Q_t *span space of polynomials of degree* $\leq T-1$, $VC(\mathcal{P}_{\mathcal{PW}}), VC(\mathcal{P}_{\mathcal{BPW}}) \geq \log(T-1)$.

Exponential, Hyperbolic, and Quasi-hyperbolic discounting models learnable with logarithmic growth in sample size.

A more powerful analyst

Stream-based setting:

- Disagreement methods: request choice for a pair only if there is something to be learned from it.
- In general appears difficult to analyze label complexities due to dependence on underlying distribution.
- Redeeming result for exponential discounting ($x \succsim y \iff \sum_t \delta^t x_t \geq \sum_t \delta^t y_t$):

Theorem

There exists a distribution on $\mathbf{R}^T \times \mathbf{R}^T$ for which the exponential discounting model is learnable with $\widetilde{O}\left(\log T \log \frac{1}{\varepsilon}\right)$ labels (exponential improvement over PAC)

Membership queries: can learn exponential discounting factor to ε-accuracy with $O(\log \frac{1}{\varepsilon})$ queries (by performing a binary search).