Forecasting for Swap Regret for All Downstream Agents

Aaron Roth, Mirah Shi
University of Pennsylvania

Predictions for Decision-Making

Agents take **actions** based on **predictions**

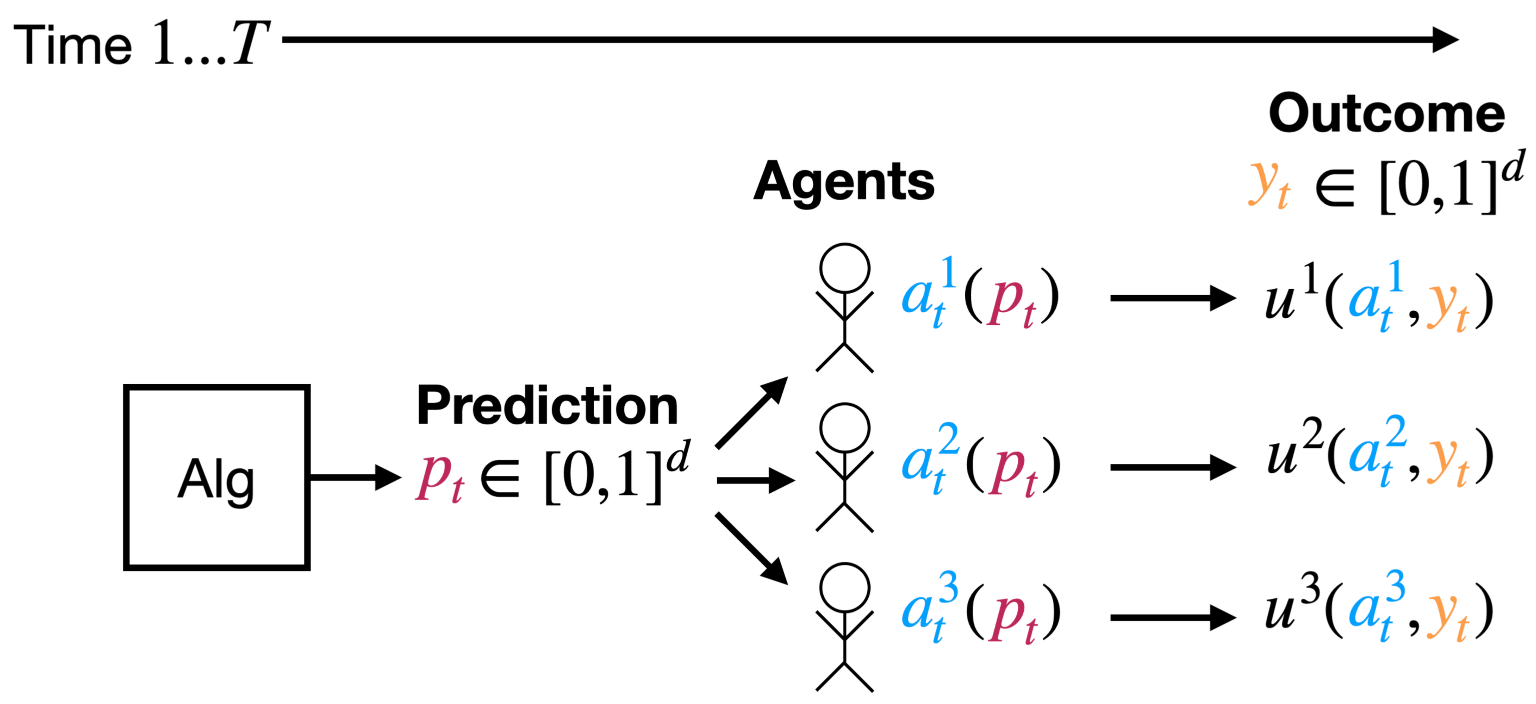

Question: How do we make online predictions that are simultaneously valuable to any downstream decision-making agent?
- Potentially many agents with different utilities
- Learner may not know agents' utility functions

Goal: Guarantee ***low swap regret**** for any downstream agent, *regardless of utility function*

* "Every time I played action a, I wouldn't have wanted to play action b in hindsight."

Why swap regret? Useful in strategic settings, e.g. :
- Convergence to correlated equilibria
- Strategy-robustness in repeated games

The Story (Prior to This Work)

Calibration* is one solution [Foster and Vohra '98]

* Unbiasedness conditional on value of prediction itself

Major Drawbacks...
- Swap regret rates degrade poorly with d
- $\Omega(T^{0.528})$ lower bound on 1-d calibration error [Qiao & Valiant '21] but can achieve $O(\sqrt{T})$ swap regret for a single agent [Blum & Mansour '07]

Other previously known solutions for:
- a weaker benchmark: no *external* regret [KPST '23]
- a *fixed* collection of agents [NRRX '23]

Q: Can we circumvent calibration to achieve no swap regret for any agent, regardless of their utility?
Our work: Yes!

Check out more here:
arxiv.org/abs/2402.08753

Our Results (Informal)

We show how to make predictions so that any downstream agent who best responds has swap regret:
- $\tilde{O}(\sqrt{T})$ for $d = 1$ **(Focus of this poster)**
- $\tilde{O}(T^{5/8})$ for $d = 2$

and any agent who smoothly best responds has swap regret:
- $\tilde{O}(T^{2/3})$ for $d > 2$

no matter what their utility function is.
Takeaway: can minimize downstream swap regret *without requiring calibrated forecasts!*

Key Idea

For a fixed agent, enough for predictions to be unbiased estimates conditional on their ***best response regions**** [NRRX '23]

* Predictions p inducing a best response of a

But we don't know agents' best response regions...

Structural property: Best response regions are convex!
In 1 dimension: convex sets correspond to sub-intervals of $[0,1]$
- Not too many after discretizing predictions

$d = 1$:

m predictions
↓
m^2 possible best response regions

Algorithm: Make predictions that are unbiased conditional on lying in all sub-intervals
- Use unbiased prediction algorithm of NRRX '23

Beyond Low Dimensions

Complexity scales in higher dimensions...
New approach: Consider best response regions of a discretized set of utility functions
- Challenge: best response function is discontinuous ⟶ require smooth approx.