Forecasting for Swap Regret for All Downstream Agents

Aaron Roth

Mirah Shi

(Penn)

Forecasting for Decision Making

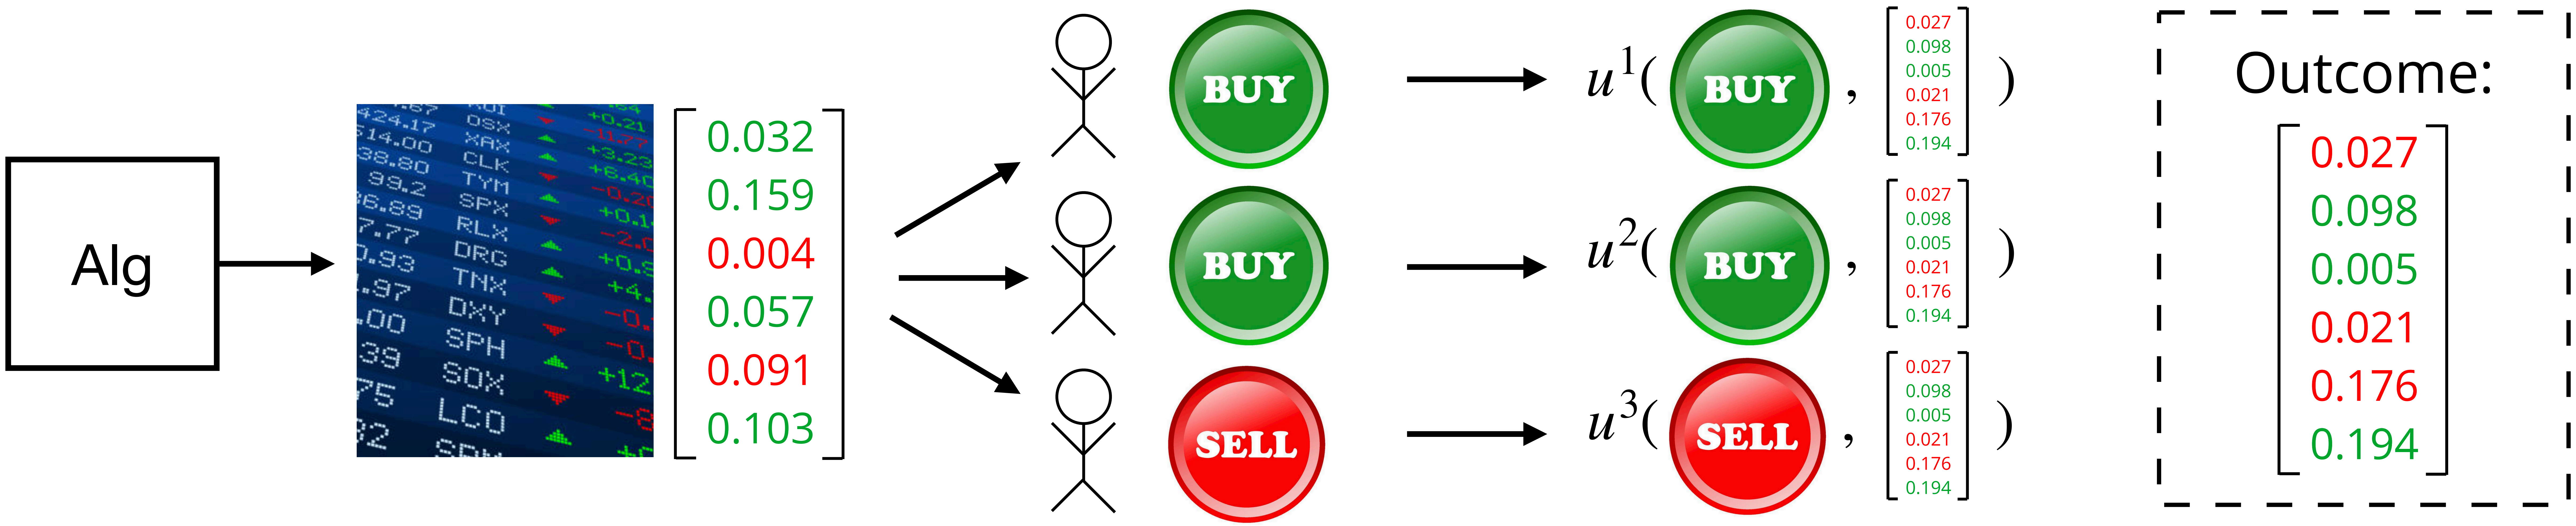

Learning algorithm provides forecast

Downstream agents take some action based on forecast

Agents receive utility as a function of their action and the outcome

How should we make forecasts for downstream decision-makers?

An Online Forecasting Setting

In each round $t = 1,\ldots,T$:

- Learning algorithm outputs forecast $p_t \in [0,1]^d$
- Agent observes p_t and takes action $a_t \in \mathcal{A}$
  - We will focus on agents who ***best respond***:

$$a_t = BR_u(p_t) = \operatorname{argmax}_{a \in \mathcal{A}} u(a, p_t)$$

"Treat predictions as if they are correct"

- Outcome $y_t \in [0,1]^d$ is revealed (possibly adversarially chosen)
- Agent receives utility $u(a_t, y_t)$

Measuring Performance via Regret

Agent wants to maximize their total utility $\sum_{t=1}^{T} u(a_t, y_t)$

- But impossible to maximize against an adversary

Instead, we measure performance through ***regret***: counterfactual guarantee against a class of benchmarks

E.g. all fixed actions (**external regret**):

$$\text{Reg} = \max_{a^* \in \mathcal{A}} \sum_{t=1}^{T} u(a^*, y_t) - u(a_t, y_t)$$

An agent has **no regret** if $\text{Reg} = o(T)$

Swap Regret

Swap regret measures performance against all *mappings* ϕ of actions to actions:

$$\text{SwapReg} = \max_{\phi:\mathcal{A}\to\mathcal{A}} \sum_{t=1}^{T} u(\phi(a_t), y_t) - u(a_t, y_t)$$

"Every time I played action a, I wouldn't have done much better by playing action b"

An agent has **no swap regret** if SwapReg $= o(T)$

Why minimize swap regret? Useful in strategic settings:

- Convergence to correlated equilibrium [Foster-Vohra '97]
- Strategy-robustness in repeated games [DSS '19] [MMSS '22] [ACS '24]

How to Forecast for No Swap Regret?

For **one** agent with a known utility function:

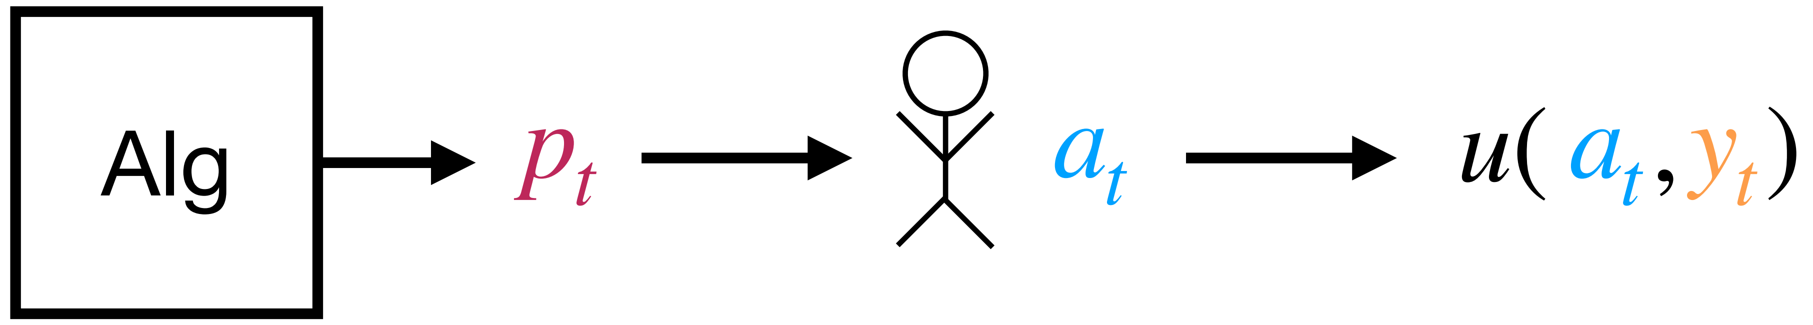

Run standard no swap regret algorithm ⇒

SwapReg = $\tilde{O}(\sqrt{|\mathscr{A}|T})$ (*optimal swap regret bound*)

What if there are many agents?

What if we don't know their utilities?

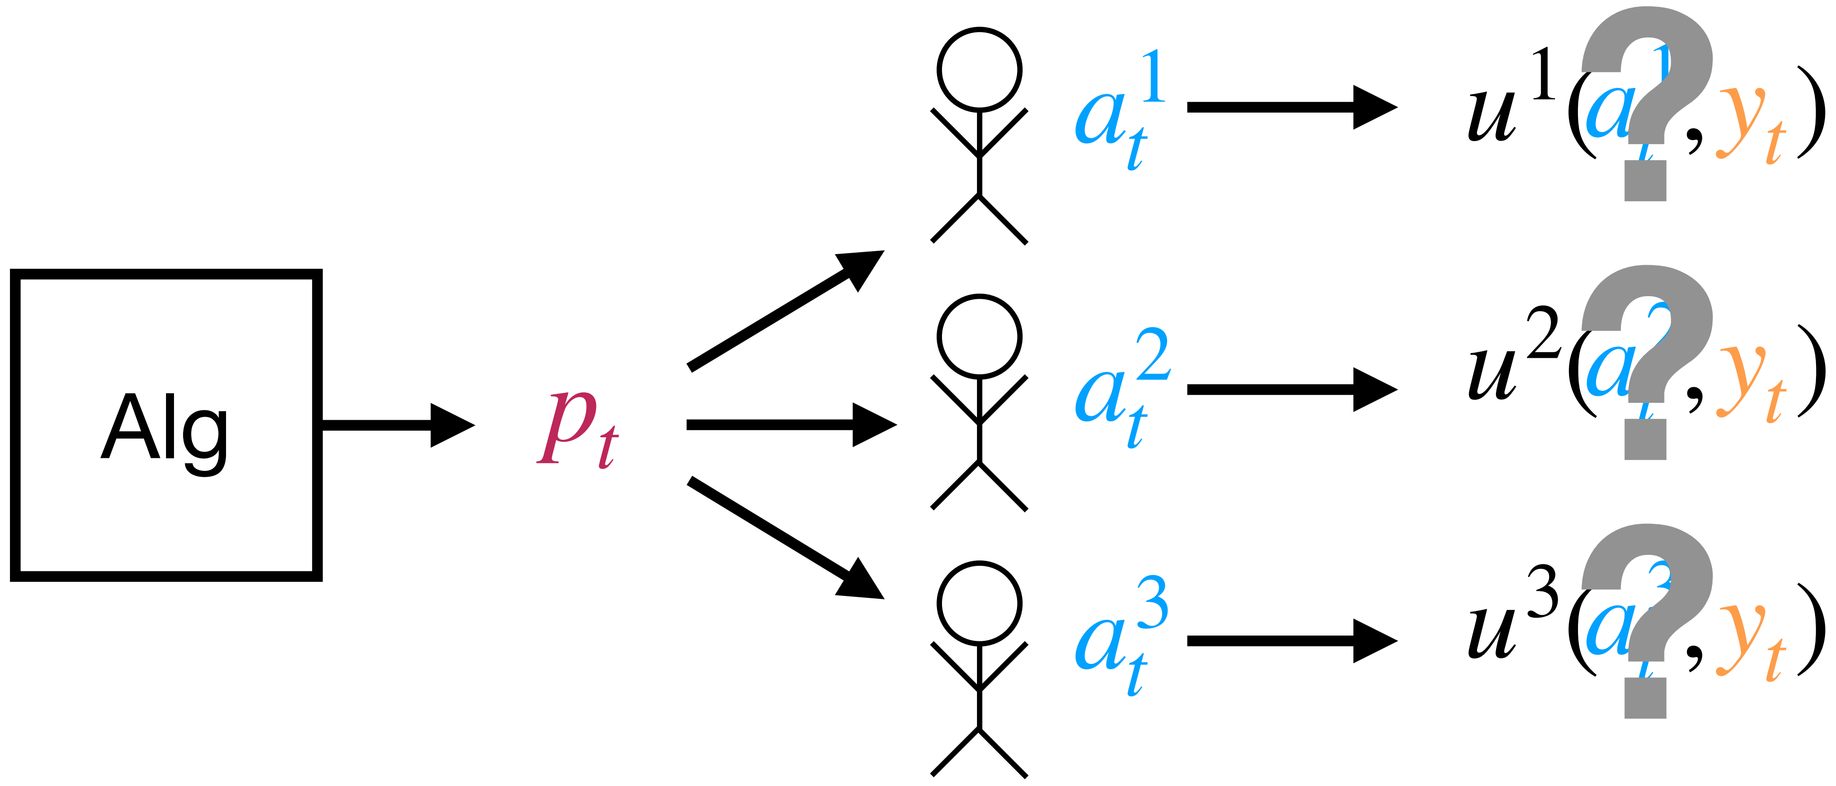

Question: How can we make forecasts that guarantee no swap regret simultaneously to any agent, *regardless of their utility function*?

The Story (Prior to This Work)

Question: How can we make forecasts that guarantee no swap regret simultaneously to any agent, *regardless of their utility function*?

We can do this via ***calibrated forecasts*** [Foster-Vohra ’97]: unbiased conditional on value of forecast itself

Intuition: can view calibration as a “no swap condition” applied to forecasts

But:

- Bad rates: swap regret bound scales poorly with the dimension
- $\Omega(T^{0.54389})$ lower bound on calibrated forecasts in 1 dimension [Qiao-Valiant ’21] [DDFGKO ’24]

The Story (Prior to This Work)

Question: How can we make forecasts that guarantee no swap regret simultaneously to any agent, *regardless of their utility function*?

- [KPST '23] show how to get ***external regret*** $O(\sqrt{T})$ for any agent for 1-dimensional and multi-class forecasts ("U-calibration")
- [NRRX '23] show how to get swap regret $\tilde{O}(\sqrt{T})$ for a ***fixed*** set of agents

Can we circumvent calibration to achieve no *swap regret* for *any* agent?

This work: yes!

This Work (Informally)

We show how to make predictions so that any agent who best responds has swap regret:

- $\tilde{O}(\sqrt{T})$ for 1-dimensional forecasts ← Optimal rate! **Note**: bypasses calibration lower bound
- $\tilde{O}(T^{5/8})$ for 2-dimensional forecasts

and any agent who smoothly best responds has swap regret:

- $\tilde{O}(T^{2/3})$ for d-dimensional forecasts

Main assumption: agent utilities are *linear* in the predictions

- Note: generalizes multi-class prediction setting of [KPST '23]

Remainder

- Forecasts in 1 dimension **(focus of this talk)**
- Forecasts in higher dimensions

Key Ingredient from Previous Work

Theorem [Noarov Ramalingam Roth Xie ’23]:

- Fix an agent with utility function u. If for all a, my forecasts are (on average) unbiased estimates of the outcomes conditional on the days t where $BR_u(p_t) = a$, then the agent has no swap regret.
- We can make conditionally unbiased forecasts.

Takeaway: enough for forecasts to be unbiased conditional on an agent’s ***best response regions***

Best response region for utility u and action a :

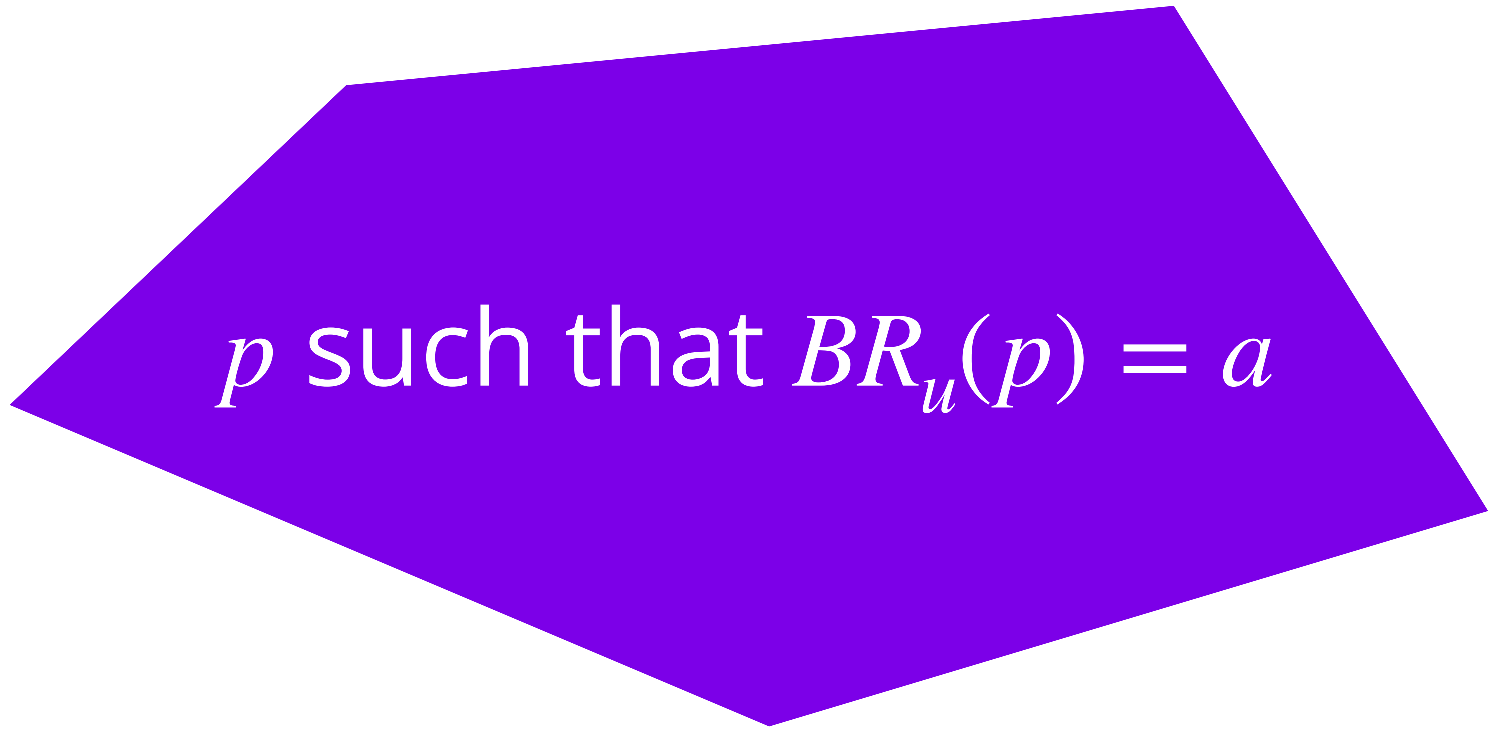

But we don’t know agents’ best response regions…

Forecasts in 1 Dimension (A Geometric Approach)

Structural property: best response regions are convex

Main Idea: require predictions to be unbiased conditional on *every* possible best response region aka convex set

In 1 dimension: sub-intervals of $[0,1]$

- Not too many after discretizing predictions

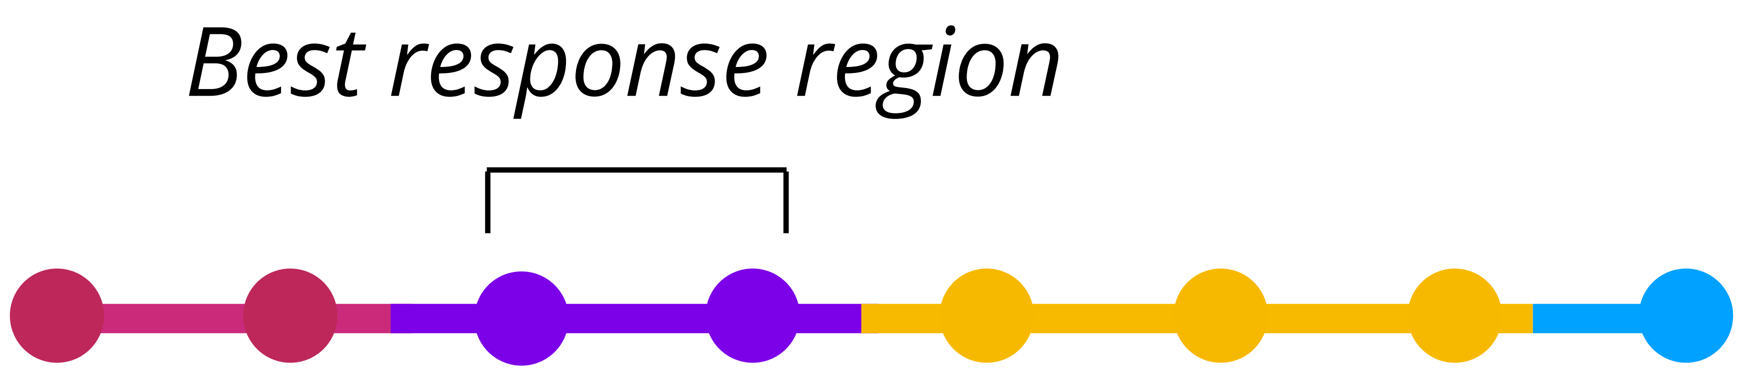

m predictions $\longrightarrow m^2$ possible best response regions

(Need to manage tradeoff with error introduced by discretization)

Forecasts in 1 Dimension

Theorem $(d = 1)$:

Using the conditionally unbiased prediction algorithm of [NRRX '23], we can guarantee any best-responding agent swap regret $\tilde{O}(|\mathcal{A}|\sqrt{T})$.

Same approach gives result for forecasts in 2 dimensions

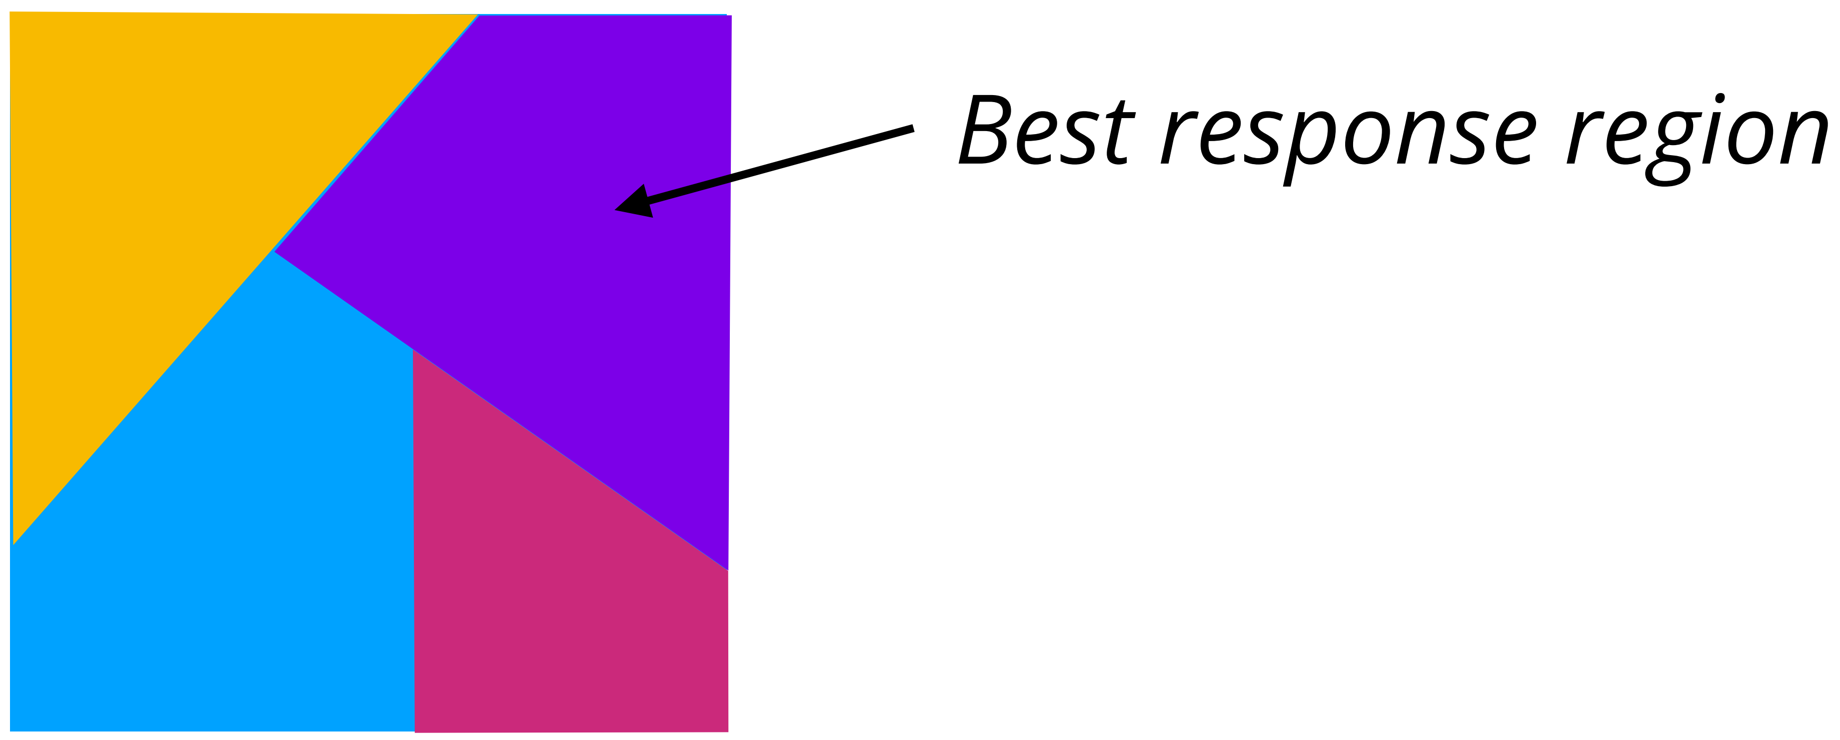

Beyond Low Dimensions

Approach for 1 dimension: enumerate best response regions aka convex sets

In higher dimensions: naively, number of convex sets is doubly exponential in d

New approach: enumerate best response regions of a *discretized* set of utility functions

How many? $\left(\frac{1}{\varepsilon}\right)^{d|\mathscr{A}|}$ to cover any u up to an ε - approximation

Problem: best response function is discontinuous…

…require *smooth* best response (e.g. quantal response)

Theorem (arbitrary d) :

Using the conditionally unbiased prediction algorithm of [NRRX '23], we can guarantee swap regret $\tilde{O}(|\mathscr{A}|d^{1/2}T^{2/3})$ for any agent who *smoothly* best responds.

The Story (Present Day)

Takeaway: can minimize downstream swap regret for any agent *without requiring calibrated forecasts!*

- Optimal bounds in 1 dimension

Follow-up work [Hu-Wu '24]: removes dependence on $|\mathscr{A}|$ (# actions) in 1 dimension

Future directions:

- Action-independent bounds in > 1 dimension?
- Removing linearity assumption on utility functions
- Efficient algorithm (i.e. complexity scaling polynomially with d)

Thanks!

mirahshi@seas.upenn.edu