Efficient Competitions and

Online Learning with Strategic Forecasters

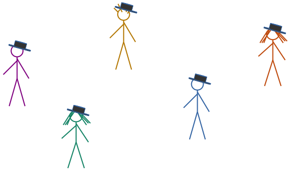

Rafael Frongillo[a], Robert Gomez[b], Anish Thilagar[c], Bo Waggoner[d]

[a]Colleague at CU Boulder
[b]Undergrad thesis student
[c]2nd-3rd year PhD student
[d]me

Classic online learning from expert advice

On rounds $t = 1, \ldots, T$:

- Expert i predicts $p_{it} \in [0, 1]$
- Algorithm chooses an expert
- Outcome $\omega \in \{0, 1\}$; i's loss is $(\omega - p_{it})^2$
- Algorithm's goal: low regret to the **best expert**

Classic online learning from expert advice

On rounds $t = 1, \ldots, T$:

- Expert i predicts $p_{it} \in [0, 1]$
- Algorithm chooses an expert
- Outcome $\omega \in \{0, 1\}$; i's loss is $(\omega - p_{it})^2$
- Algorithm's goal: low regret to the **best expert**

Multiplicative weights (MW):
choose i w.prob. $\propto e^{-\eta(\text{total loss})}$.

Guarantees: Regret $O(\sqrt{T})$.

Strategic experts

Changes to model:

- Experts report some r_{it}, potentially $\neq p_{it}$
- Experts want to be chosen, e.g. max $\mathbb{E}[\#$ times chosen$]$
- *Strategic regret*: to the **best expert's knowledge**

  still according to p_{it}'s

Strategic experts

Changes to model:

- Experts report some r_{it}, potentially $\neq p_{it}$
- Experts want to be chosen, e.g. max $\mathbb{E}[\#\text{ times chosen}]$
- *Strategic regret*: to the **best expert's knowledge**

  still according to p_{it}'s

Question: what is the cost of strategic behavior in online learning?

Prior work

Roughgarden and Schrijvers (Neurips 2017):
experts want to maximize their (unnormalized) weight.

Prior work

Roughgarden and Schrijvers (Neurips 2017):
experts want to maximize their (unnormalized) weight.

Freeman, Pennock, Podimata, Wortman Vaughan (ICML 2020):
Experts want to maximize $\Pr[\text{chosen}]$.

Prior work

Roughgarden and Schrijvers (Neurips 2017):
experts want to maximize their (unnormalized) weight.

Freeman, Pennock, Podimata, Wortman Vaughan (ICML 2020):
Experts want to maximize $\Pr[\text{chosen}]$.

- Myopic experts: $O(\sqrt{T})$ regret truthful algorithm

Prior work

Roughgarden and Schrijvers (Neurips 2017):
experts want to maximize their (unnormalized) weight.

Freeman, Pennock, Podimata, Wortman Vaughan (ICML 2020):
Experts want to maximize $\Pr[\text{chosen}]$.

- Myopic experts: $O(\sqrt{T})$ regret truthful algorithm
- Forward-looking experts: **open problem**
  (truthful algorithm, but no regret guarantee)

Tool: wagering mechanisms

- One event: wagering mechs (scoring rule based)
  FPPV20 $\implies$ truthful $\sqrt{T}$ regret for **myopic agents**

Tool: wagering mechanisms

- One event: wagering mechs (scoring rule based)
  FPPV20 $\implies$ truthful $\sqrt{T}$ regret for **myopic agents**

- Sequence: forecasting competition, ELF (Witkowski et al. 2018)
  FPPV20 $\implies$ truthful for **forward-looking**, regret unknown

Tool: wagering mechanisms

- One event: wagering mechs (scoring rule based)
  FPPV20 $\implies$ truthful $\sqrt{T}$ regret for **myopic agents**

- Sequence: forecasting competition, ELF (Witkowski et al. 2018)
  FPPV20 $\implies$ truthful for **forward-looking**, regret unknown

Suggested approach:

- Find a better truthful forecasting competition
- Use it for online learning

Tool: wagering mechanisms

- One event: wagering mechs (scoring rule based)
  FPPV20 $\implies$ truthful $\sqrt{T}$ regret for **myopic agents**
- Sequence: forecasting competition, ELF (Witkowski et al. 2018)
  FPPV20 $\implies$ truthful for **forward-looking**, regret unknown

Suggested approach:

- Find a better truthful forecasting competition
- Use it for online learning

But: we don't know how.

Our result

Algorithm:

1. Take multiplicative weights

Our result

Algorithm:

1. Take multiplicative weights
2. Run it

Our result

Algorithm:

1. Take multiplicative weights
2. Run it

Theorem

MW achieves $O(\sqrt{T})$ strategic regret.

Our result

Algorithm:

1. Take multiplicative weights
2. Run it

Theorem

MW achieves $\mathrm{O}(\sqrt{\mathrm{T}})$ strategic regret.

Solution concept? In equilibrium?

Our result

Algorithm:

1. Take multiplicative weights
2. Run it

Theorem

MW achieves $O(\sqrt{T})$ strategic regret when experts play undominated strategies.

(more discussion at the end)

Proof Ideas

(1) MW with $\eta = O(1/\sqrt{T})$ has $\sqrt{T}$ regret to **reports**.

Proof Ideas

(1) MW with $\eta = O(1/\sqrt{T})$ has $\sqrt{T}$ regret to **reports**.

Definition

A mechanism is γ-**approximately truthful** if for all undominated strategies r_i, $|r_{it} - p_{it}| \leq \gamma$ for all t.

Proof Ideas

(1) MW with $\eta = O(1/\sqrt{T})$ has $\sqrt{T}$ regret to **reports**.

Definition

A mechanism is γ-**approximately truthful** if for all undominated strategies r_i, $|r_{it} - p_{it}| \leq \gamma$ for all t.

(2) MW is 4η-**approximately truthful**.

Proof Ideas

(1) MW with $\eta = O(1/\sqrt{T})$ has $\sqrt{T}$ regret to **reports**.

Definition

A mechanism is γ-**approximately truthful** if for all undominated strategies r_i, $|r_{it} - p_{it}| \leq \gamma$ for all t.

(2) MW is 4η-**approximately truthful**.

(1) + (2) $\implies$ MW has $\sqrt{T}$ regret to beliefs.

Proof Ideas

(1) MW with $\eta = O(1/\sqrt{T})$ has $\sqrt{T}$ regret to **reports**.

Definition

A mechanism is γ-**approximately truthful** if for all undominated strategies r_i, $|r_{it} - p_{it}| \leq \gamma$ for all t.

(2) MW is 4η-**approximately truthful**.

(1) + (2) $\implies$ MW has $\sqrt{T}$ regret to beliefs.

(Also enables better forecasting competitions.)

Truthfulness of MW

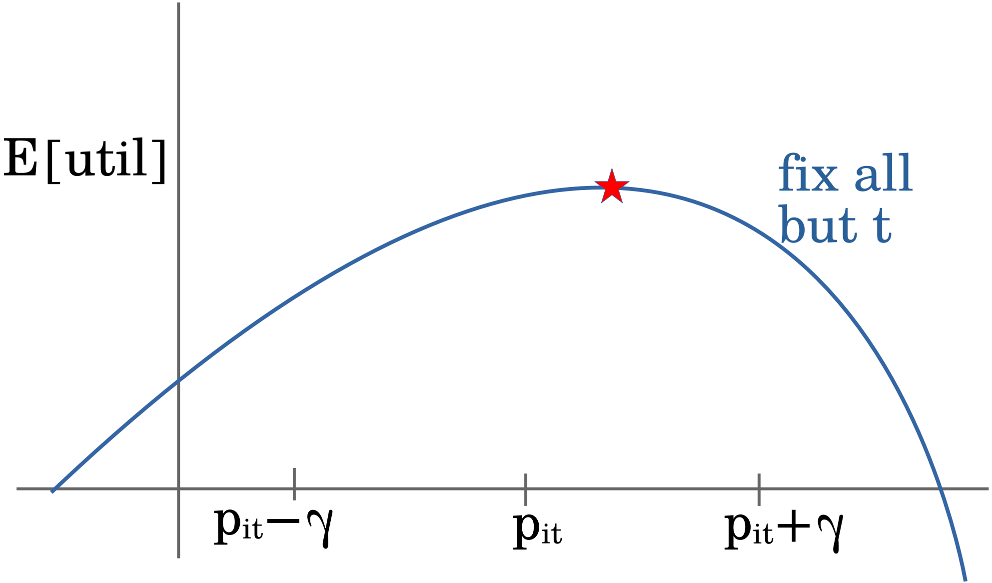

Truthfulness of MW

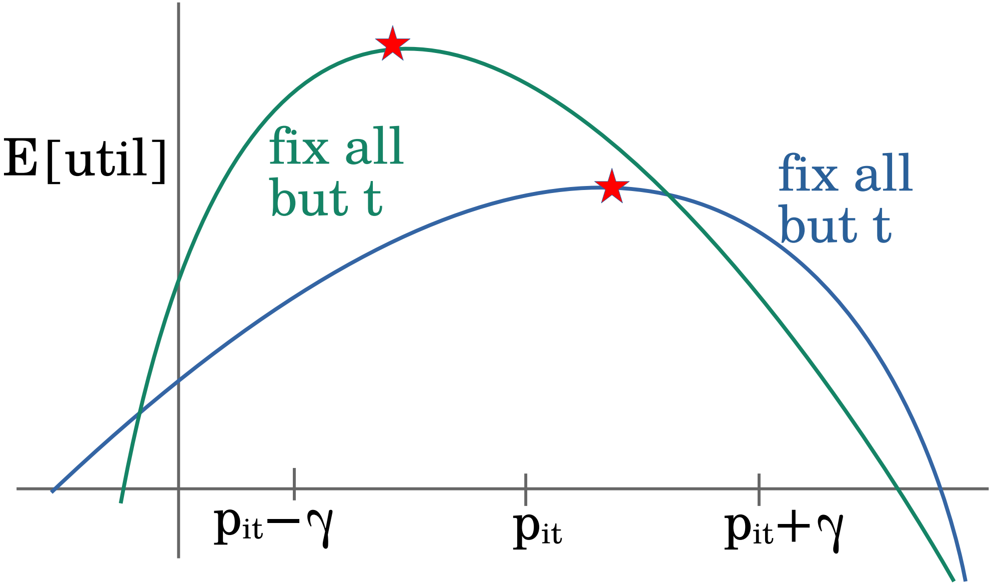

Truthfulness of MW

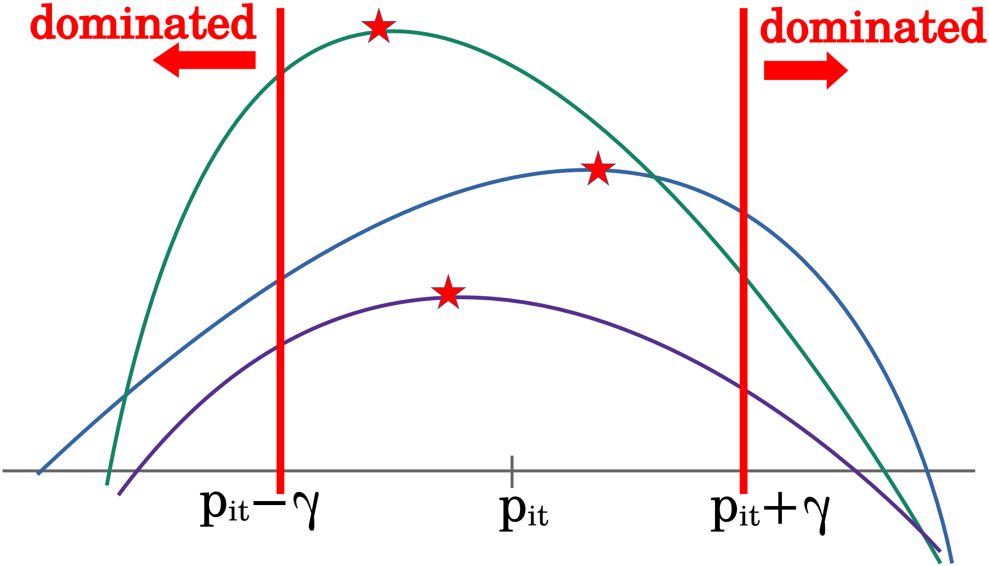

Robustness of result

Result: MW is very **robust**.

Is the *result* robust?

Robustness of result

Result: MW is very **robust**.

Is the *result* robust?

Not really: need specific conditions on FTRL regularizer.
e.g. OGD fails

Robustness of result

Result: MW is very **robust**.

Is the *result* robust?

Not really: need specific conditions on FTRL regularizer.
e.g. OGD fails

But: **Report Noisy Min** is also approximately truthful:

Robustness of result

Result: MW is very **robust**.

Is the *result* robust?

Not really: need specific conditions on FTRL regularizer.
e.g. OGD fails

But: **Report Noisy Min** is also approximately truthful:

1. Let $Y_i = (\text{total loss of } i) + \text{Laplace}(\gamma)$

Robustness of result

Result: MW is very **robust**.

Is the *result* robust?

Not really: need specific conditions on FTRL regularizer.
e.g. OGD fails

But: **Report Noisy Min** is also approximately truthful:

1. Let $Y_i = (\text{total loss of } i) + \text{Laplace}(\gamma)$
2. Choose $\arg\min_i Y_i$.

Robustness of result

Result: MW is very **robust**.

Is the *result* robust?

Not really: need specific conditions on FTRL regularizer.
e.g. OGD fails

But: **Report Noisy Min** is also approximately truthful:

1. Let $Y_i = (\text{total loss of } i) + \text{Laplace}(\gamma)$
2. Choose $\arg\min_i Y_i$.

But but: not true for Gaussian noise!

Truthfulness of Report Noisy Min

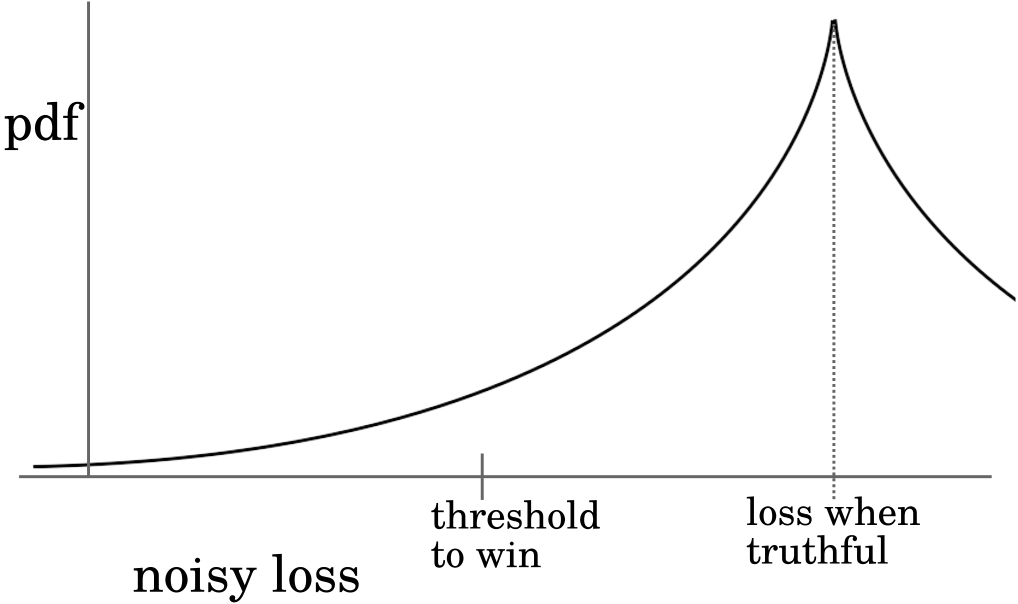

Truthfulness of Report Noisy Min

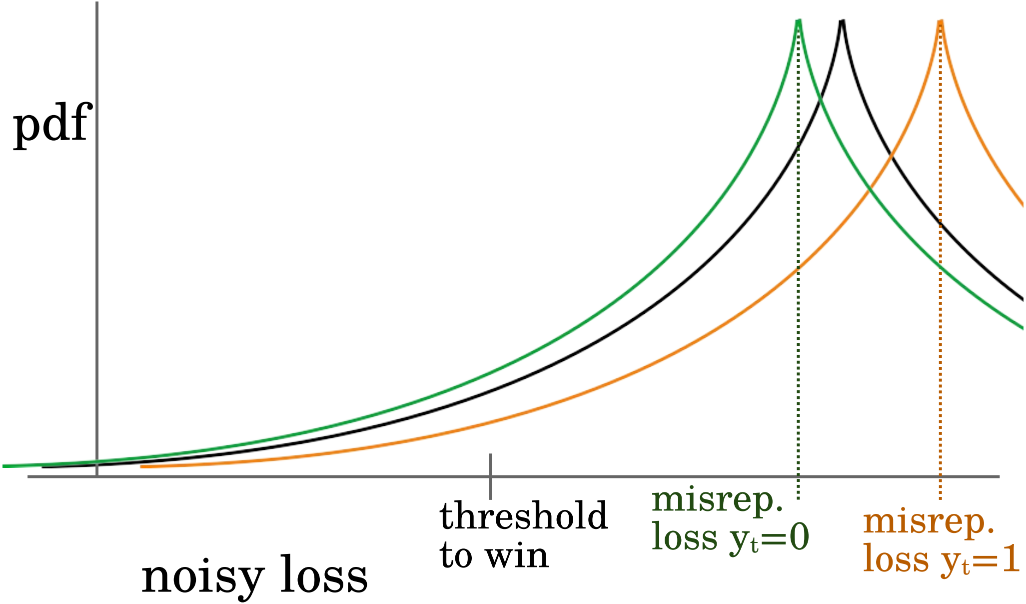

Truthfulness of Report Noisy Min

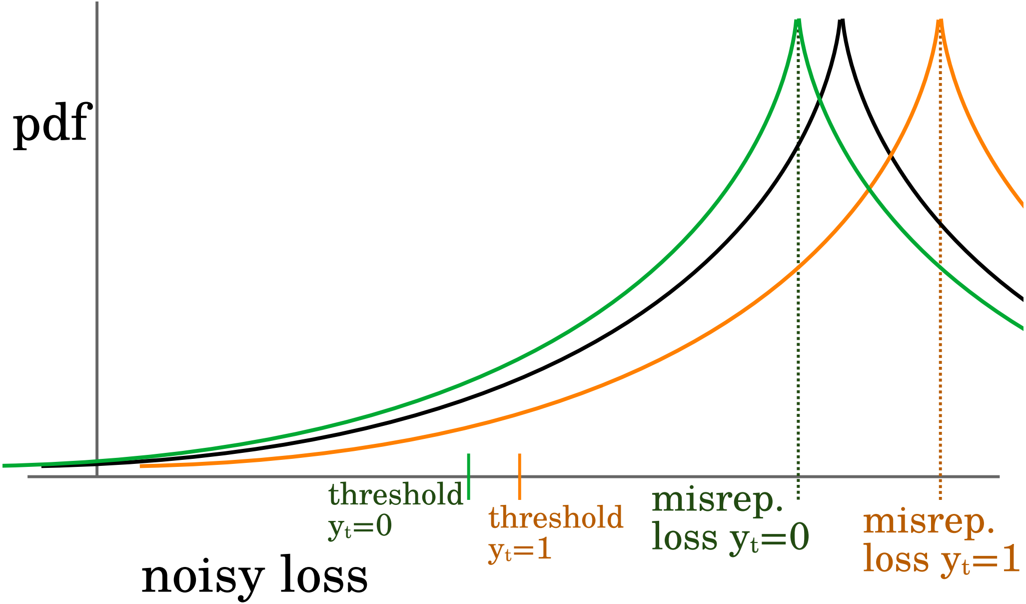

Model of strategic behavior

Our model: **immutable beliefs**

- Participant has beliefs p_{it}, unchanging
- Strategy is a plan of reports $r_{i1}, \dots, r_{iT}$
- Believes rounds are independent, opponents have plans r_{jt}
- Utility is arbitrary positive-weighted sum of $\Pr[\text{chosen}]$

Model of strategic behavior

Our model: **immutable beliefs**

- Participant has beliefs p_{it}, unchanging
- Strategy is a plan of reports $r_{i1}, \dots, r_{iT}$
- Believes rounds are independent, opponents have plans r_{jt}
- Utility is arbitrary positive-weighted sum of $\Pr[\text{chosen}]$

Extensions / questions: Bayesian models, sequential equilibrium, ...

Conjecture

MW achieves strategic regret $O(\sqrt{T})$ in any of the above models.

Conclusion

Setting:

1. Online learning from strategic experts
2. Experts try to maximize expected $\#$ times chosen
3. Immutable belief model

Results:

1. MW has strategic regret $O(\sqrt{T})$ in undominated strategies.
2. exponentially more efficient forecasting competitions (not covered)

Open problems:

1. Truthful no-regret algorithm?
2. Bayesian settings

Conclusion

Setting:

1. Online learning from strategic experts
2. Experts try to maximize expected # times chosen
3. Immutable belief model

Results:

1. MW has strategic regret $O(\sqrt{T})$ in undominated strategies.
2. exponentially more efficient forecasting competitions (not covered)

Open problems:

1. Truthful no-regret algorithm?
2. Bayesian settings

Thanks!