ICLR'25

The Hidden Cost of Waiting for Accurate Predictions

Presented by: Jiduan Wu — ETH, MPI for Intelligent Systems, & Tübingen AI Center

Ali Shirali, — UC Berkeley

Ariel Procaccia, — Harvard

Rediet Abebe, — ELLIS Institute, MPI for Intelligent Systems, & Tübingen AI Center

April 2025

Predictive systems are widely used to efficiently target individuals and allocate resources in society.

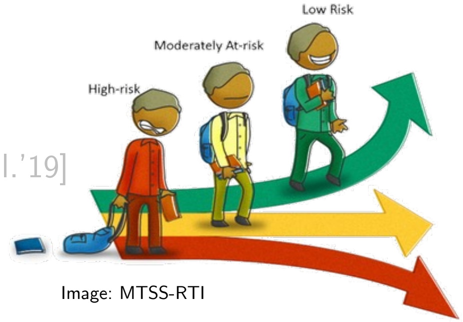

Image: MTSS-RTI

Allocate educational resources based on predicted risk of dropout [Faria et al.'17; Mac Iver et al.'19]

Image: Eric Marks

Allocate homelessness funds based on predicted risk of re-entry into homelessness [Kube et al.'23]

... and many more!

Predictive systems are widely used to efficiently target individuals and allocate resources in society.

Image: MTSS-RTI

Allocate educational resources based on predicted risk of dropout [Faria et al.'17; Mac Iver et al.'19]

Image: Eric Marks

Allocate homelessness funds based on predicted risk of re-entry into homelessness [Kube et al.'23]

… and many more!

The underlying assumption is that more accurate predictions lead to better targeting/allocation.

We examine the extent to which this assumption holds using a well-motivated, stylized model.

Next: Let's work through a concrete example …

Suppose the planner aims to prevent poor outcomes (e.g., eviction, school dropout).

There are limited resources to prevent such outcomes (e.g., housing vouchers, tutors).

Suppose the planner aims to prevent poor outcomes (e.g., eviction, school dropout).

There are limited resources to prevent such outcomes (e.g., housing vouchers, tutors).

Individuals differ in how much they benefit from the same resource.

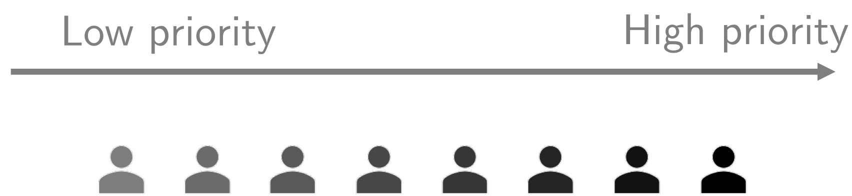

Suppose the planner aims to prevent poor outcomes (e.g., eviction, school dropout).

There are limited resources to prevent such outcomes (e.g., housing vouchers, tutors).

Individuals differ in how much they benefit from the same resource.

Each individual can receive only one unit of the resource and only once.

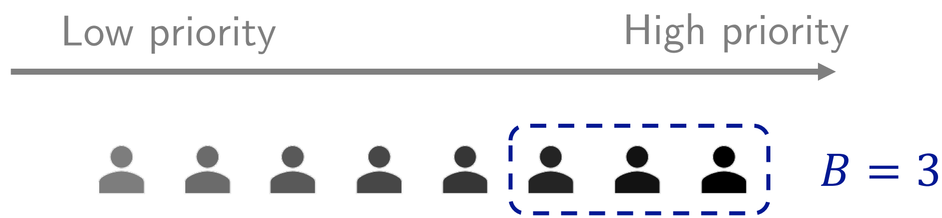

Suppose the planner aims to prevent poor outcomes (e.g., eviction, school dropout).

There are limited resources to prevent such outcomes (e.g., housing vouchers, tutors).

Individuals differ in how much they benefit from the same resource.

Each individual can receive only one unit of the resource and only once.

Solution. Obtain an individual-level risk predictor, ...

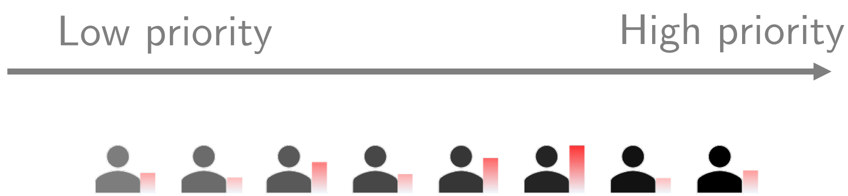

Suppose the planner aims to prevent poor outcomes (e.g., eviction, school dropout).

There are limited resources to prevent such outcomes (e.g., housing vouchers, tutors).

Individuals differ in how much they benefit from the same resource.

Each individual can receive only one unit of the resource and only once.

Solution. Obtain an individual-level risk predictor, then prioritize.

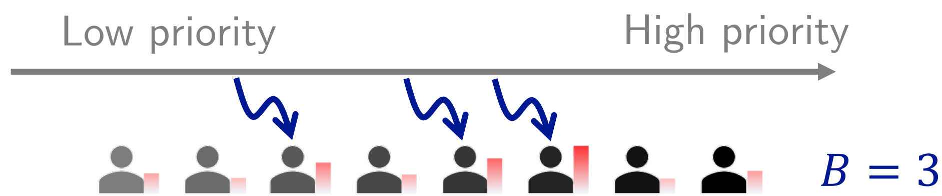

Suppose the planner aims to prevent poor outcomes (e.g., eviction, school dropout).

There are limited resources to prevent such outcomes (e.g., housing vouchers, tutors).

Individuals differ in how much they benefit from the same resource.

Each individual can receive only one unit of the resource and only once.

Solution. Obtain an individual-level risk predictor, then prioritize.

Assumption. Higher accuracy always leads to more efficient targeting.

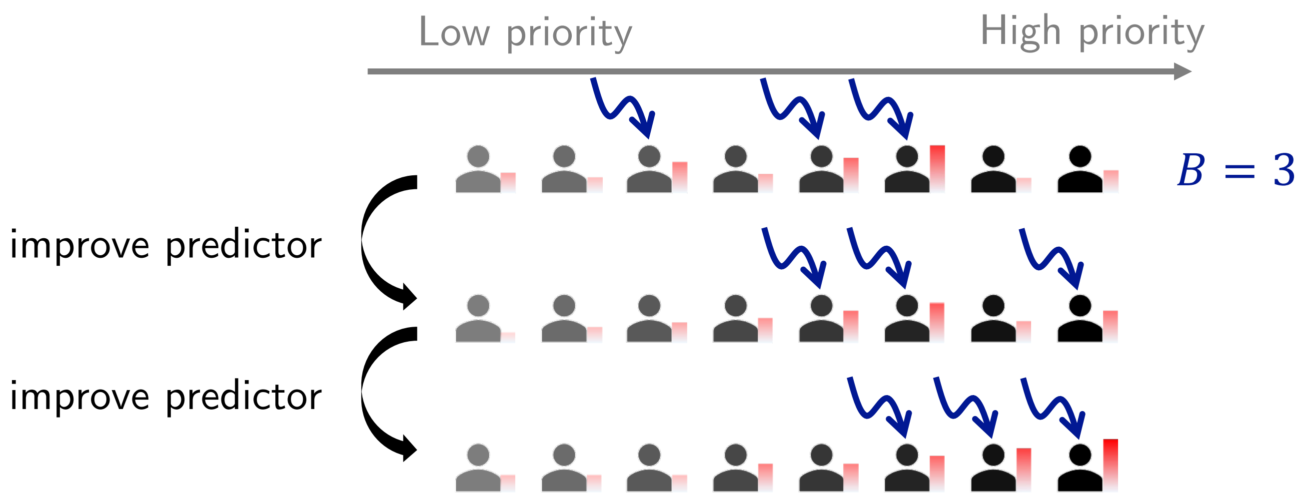

Suppose the planner aims to prevent poor outcomes (e.g., eviction, school dropout).

There are limited resources to prevent such outcomes (e.g., housing vouchers, tutors).

Individuals differ in how much they benefit from the same resource.

Each individual can receive only one unit of the resource and only once.

Solution. Obtain an individual-level risk predictor, then prioritize.

Assumption. Higher accuracy always leads to more efficient targeting.

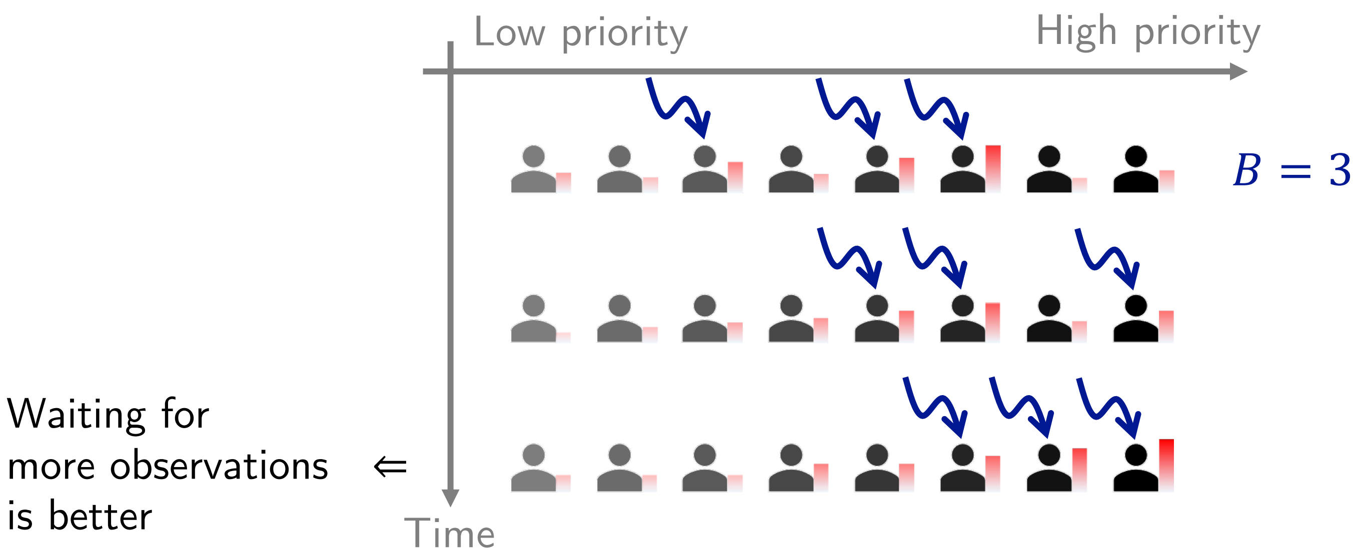

Suppose the planner aims to prevent poor outcomes (e.g., eviction, school dropout).

There are limited resources to prevent such outcomes (e.g., housing vouchers, tutors).

Individuals differ in how much they benefit from the same resource.

Each individual can receive only one unit of the resource and only once.

Solution. Obtain an individual-level risk predictor, then prioritize.

Assumption. Higher accuracy always leads to more efficient targeting.

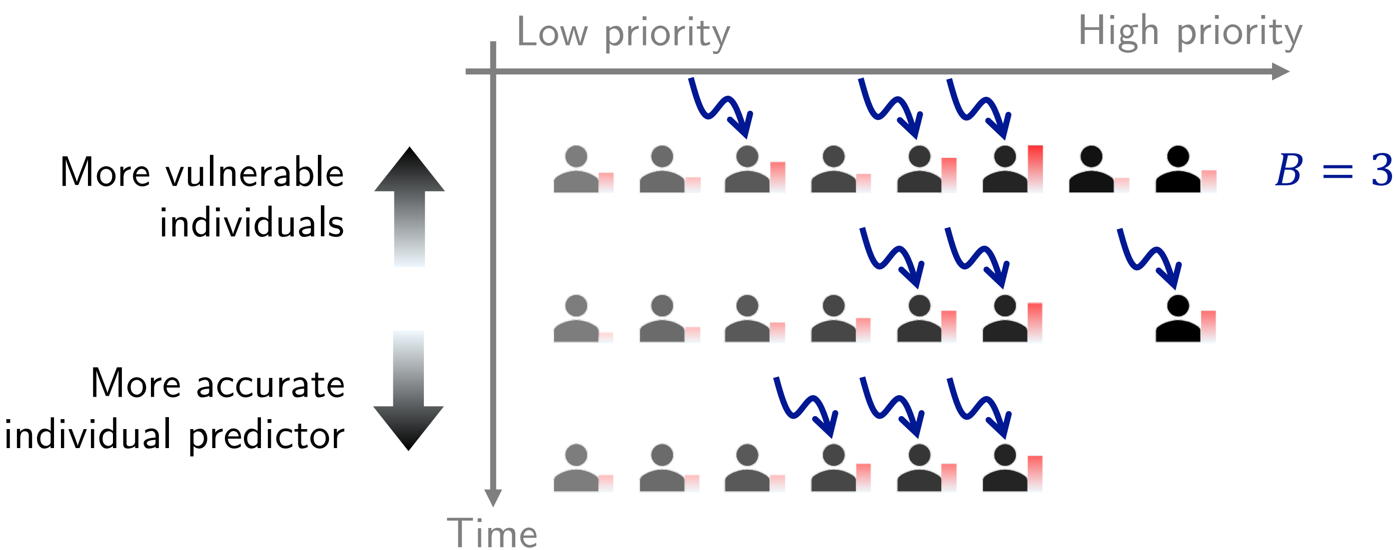

Suppose the planner aims to prevent poor outcomes (e.g., eviction, school dropout).

There are limited resources to prevent such outcomes (e.g., housing vouchers, tutors).

Individuals differ in how much they benefit from the same resource.

Each individual can receive only one unit of the resource and only once.

Solution. Obtain an individual-level risk predictor, then prioritize.

Assumption. Higher $\underbrace{\text{accuracy}}_{\text{at what cost?}}$ always leads to more efficient targeting.

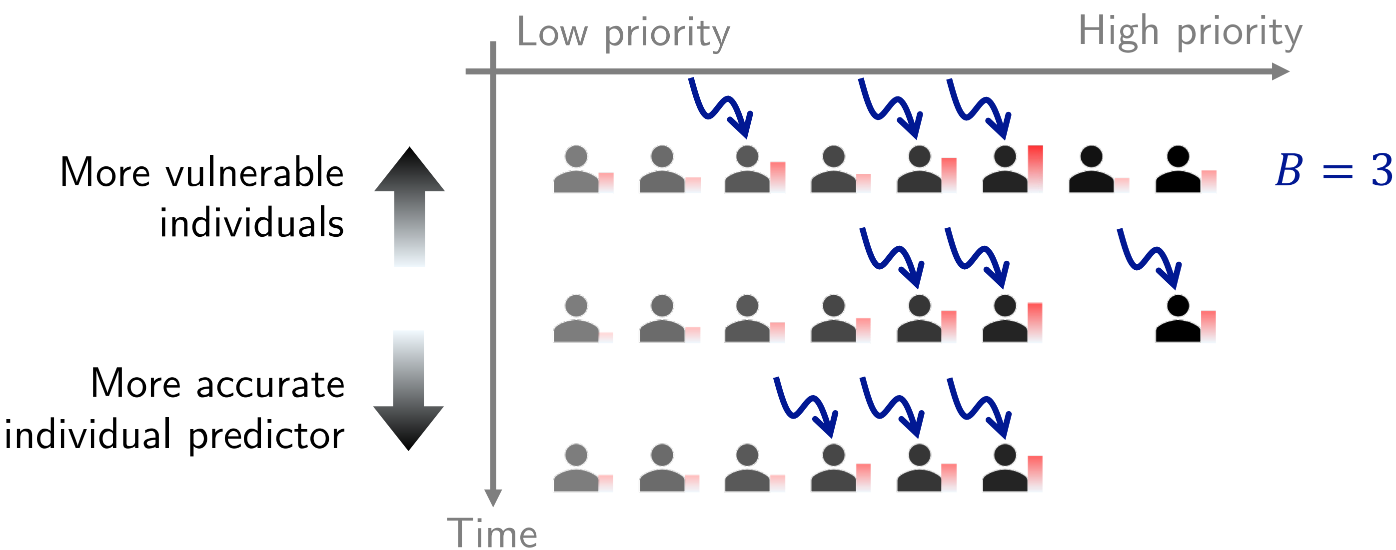

Model

- Individual i may fail between time steps with probability p_i .

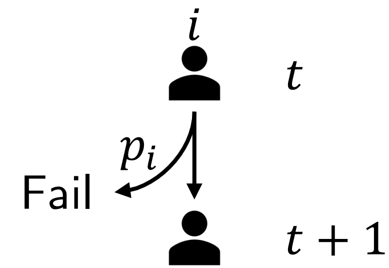

Model

- Individual i may fail between time steps with probability p_i .
- Treating i yields an effect of $\tau_i^t = \tau(p_i, t)$, where, τ is increasing in p and decreasing in t.
- Planner can treat B individuals before time T.
- Planner wants to maximize the sum of treatment effects for those treated.

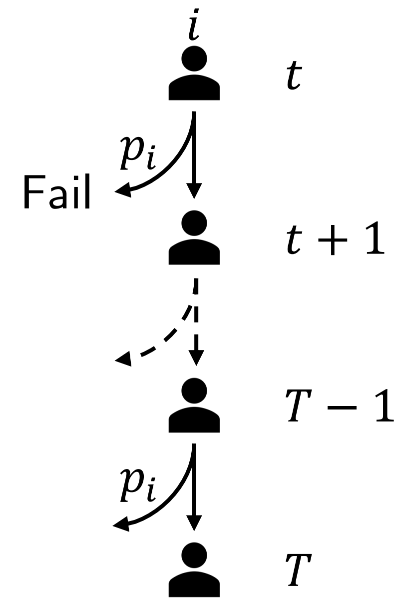

If p_is were known:
Sort individuals in decreasing order of p at $t = 1$ and treat the top B.

Model

- Individual i may fail between time steps with probability p_i .
- Treating i yields an effect of $\tau_i^t = \tau(p_i, t)$, where, τ is increasing in p and decreasing in t.
- Planner can treat B individuals before time T.
- Planner wants to maximize the sum of treatment effects for those treated.
- Planner observes a signal from each active i at t:

$$o_i^t \sim \mathrm{Ber}(\tilde{p}_i) .$$

  Here, $\tilde{p}_i$ is an increasing function of p_i.

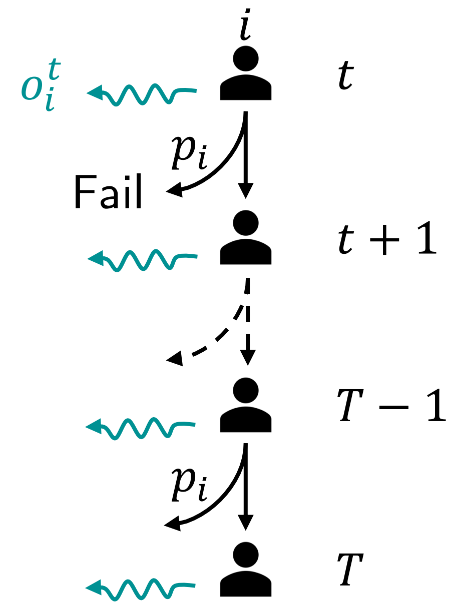

Model

- Individual i may fail between time steps with probability p_i .
- Treating i yields an effect of $\tau_i^t = \tau(p_i, t)$, where, τ is increasing in p and decreasing in t.
- Planner can treat B individuals before time T.
- Planner wants to maximize the sum of treatment effects for those treated.
- Planner observes a signal from each active i at t:
$$o_i^t \sim \mathrm{Ber}(\tilde{p}_i)\ .$$
Here, $\tilde{p}_i$ is an increasing function of p_i.

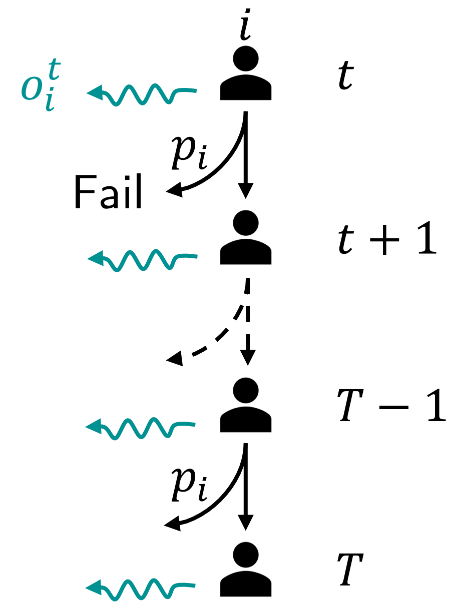

Research Questions

Ranking problem:
Would the planner's predicted ranking of active individuals improve over time?

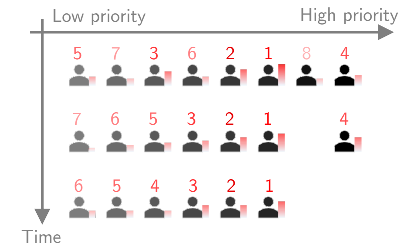

Research Questions

Ranking problem:
Would the planner's predicted ranking of active individuals improve over time?

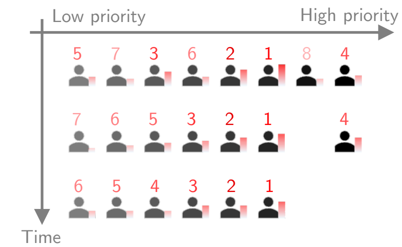

One-time allocation:

When is the best time to act?

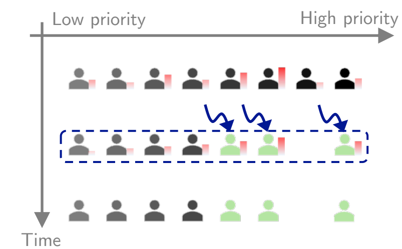

Over-time allocation:

What is the optimal distribution of the budget across time?

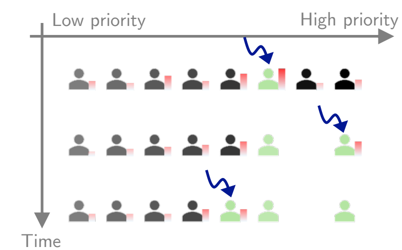

Ranking Over Time

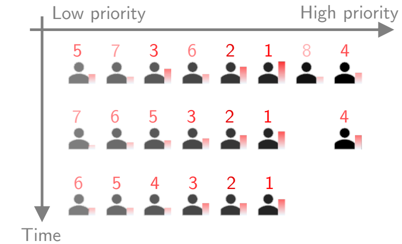

Would the (optimal) predicted ranking of active individuals improve over time?

Thm. 3.1 (informal). Assuming that the inverse of $\tilde{p}(\cdot)$ is $O(1)$-Lipschitz, the optimal pairwise ranking risk at time t, denoted by R^t, can only improve if

$$0 > R^{t+1} - R^t \propto \underbrace{\frac{\overset{\text{inequality}}{\boxed{\mathrm{Var}^t[p]}}}{(1-\mathbb{E}^t[p])^2}}_{\text{change-in-population effect}} - \underbrace{\frac{C_1}{\sqrt{t}} - \frac{C_2}{t}}_{\text{gain in observations}}$$

One-Time Allocation

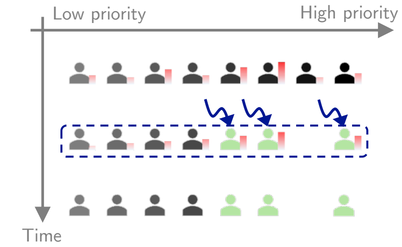

Suppose the planner can only intervene once.

When is the best time to act?

One-Time Allocation

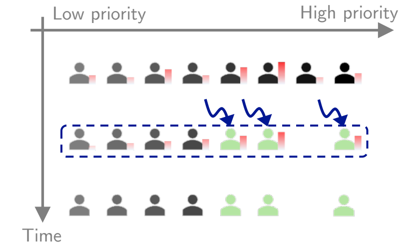

Suppose the planner can only intervene once.

When is the best time to act?

In this talk: Fully effective treatment

$$\tau_i^t = 1 - (1 - p_i)^{T-t}$$

$$\Pr(\text{survive} \mid \text{treated}) - \Pr(\text{survive} \mid \text{untreated})$$

One-Time Allocation

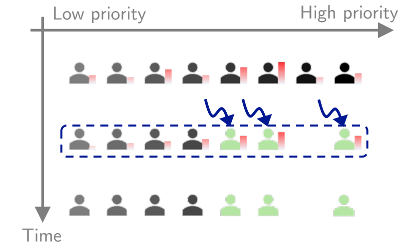

When is the best time to act?

Thm. 4.3 (informal). If the initial distribution over failure probabilities is G-decaying, using a Bayes optimal ranking of individuals to prioritize resources, it is never optimal to wait beyond time $\frac{T}{2} + C\,G\ln\frac{N}{B}$.

One-Time Allocation

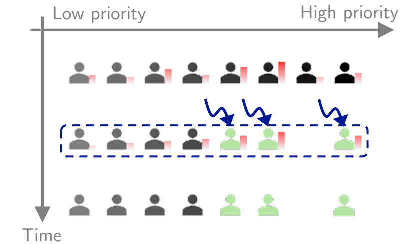

When is the best time to act?

Thm. 4.3 (informal). If the initial distribution over failure probabilities is G-decaying, using a Bayes optimal ranking of individuals to prioritize resources,

Bounds how fast the distribution declines. Lower G, higher inequality.

it is never optimal to wait beyond time $\frac{T}{2} + C\, \boxed{G \ln \frac{N}{B}}$.

Captures everything about the observation model.

Act earlier if

- Budget B is larger (if B is small, we can wait to confidently spend it.)
- G is lower (inequality is higher)

Over-Time Allocation

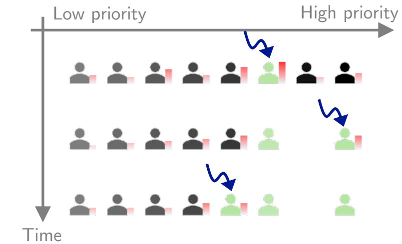

What is the optimal distribution of B across time?

Optimal has some structure: We propose an efficient algorithm that finds optimal over-time allocation with the complexity independent of N and B.

Visualize the effect of B using National Education Longitudinal Study'88:

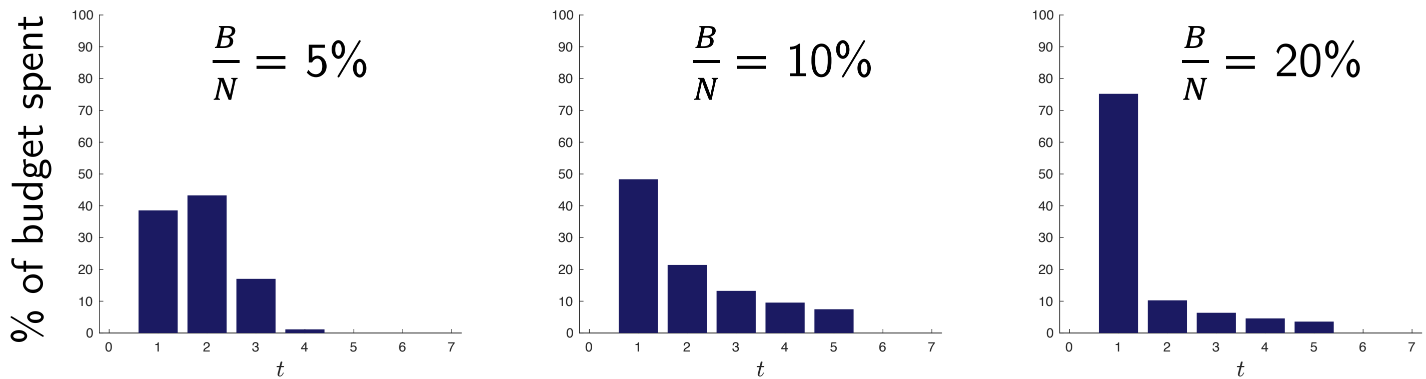

Over-Time Allocation

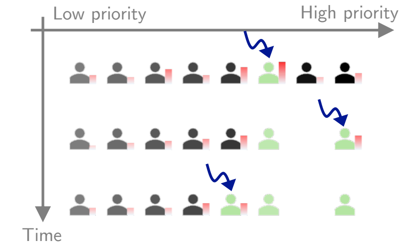

What is the optimal distribution of B across time?

Optimal has some structure: We propose an efficient algorithm that finds optimal over-time allocation with the complexity independent of N and B.

Visualize the effect of G using National Education Longitudinal Study'88:

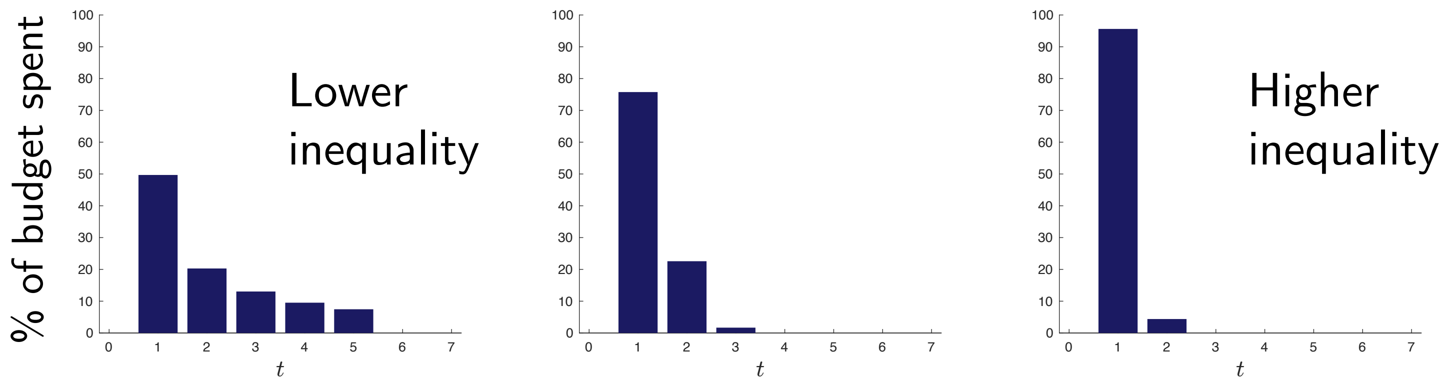

Future Works

- Heterogeneous costs: Collecting observations for vulnerable individuals may be more costly, contain less signal, or may otherwise be undesirable [Paes et al.'22].

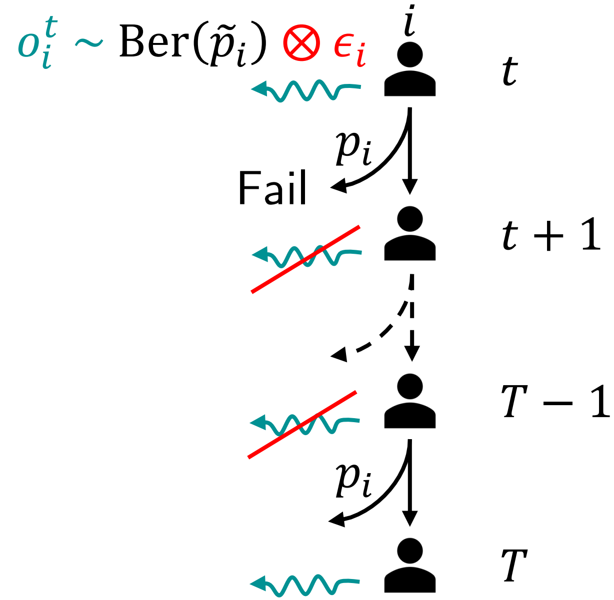

Future Works

- Heterogeneous costs: Collecting observations for vulnerable individuals may be more costly, contain less signal, or may otherwise be undesirable [Paes et al.'22].
- Individual "contexts": While individuals have heterogeneous values of p, we do not account for variations in their initial conditions.

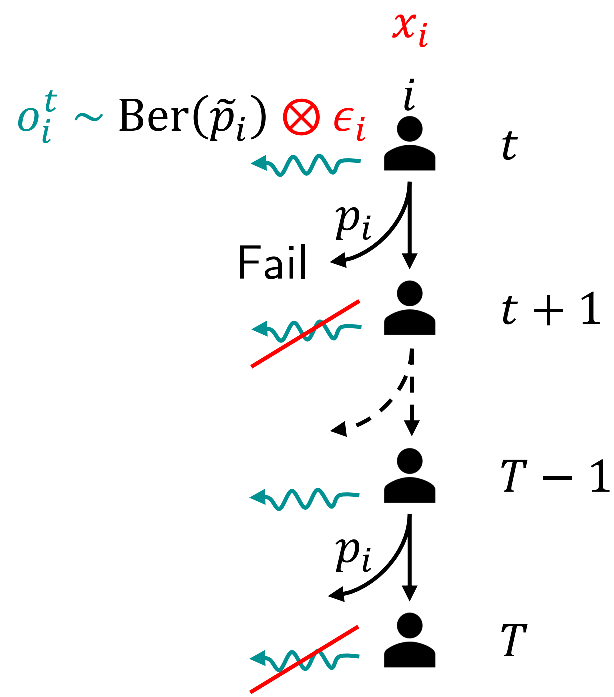

Conclusion

- Waiting to gather more information to improve prediction accuracy is not always preferred.

Ex. Reduce dropout rate by directing educational resources to students [Mac Iver et al'17].

Ex. Reduce eviction rate by directing assistance to renters [Mashiat'24].

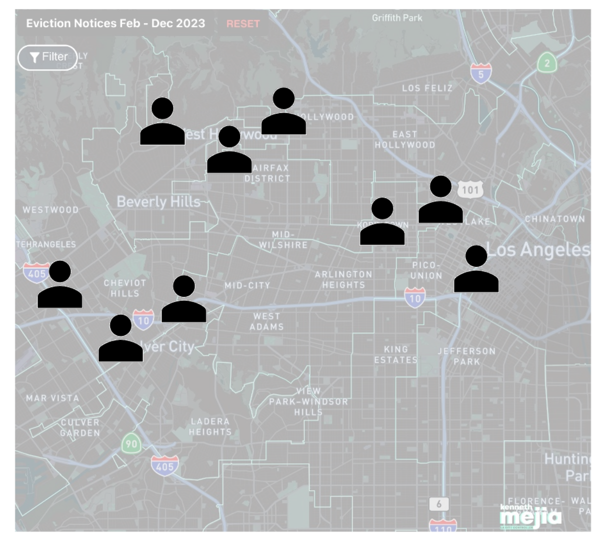

Map: LA City Controller

Conclusion

- Waiting to gather more information to improve prediction accuracy is not always preferred.
- Always ask: What is the cost of improving (individual) predictions?
  - In our work: The opportunity cost treating vulnerable individuals
  - Related prior work:
    - Cheaper, coarser predictions are enough when individuals are under some units [Shirali et el'24]

Ex. Prioritize schools [Mac Iver et al'17, Perdomo'24]

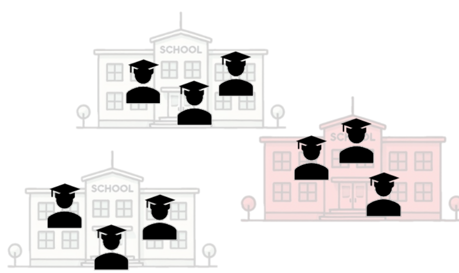

Ex. Prioritize neighborhoods [Mashiat'24]

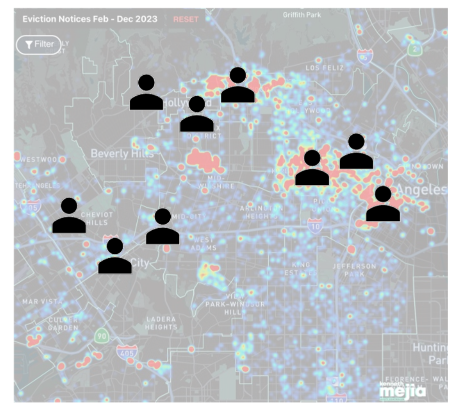

Conclusion

- Waiting to gather more information to improve prediction accuracy is not always preferred.
- Always ask: What is the cost of improving (individual) predictions?
  - In our work: The opportunity cost treating vulnerable individuals
  - Related prior work:
    - Cheaper, coarser predictions are enough when individuals are under some units [Shirali et el'24]
    - Expanding access vs. improving predictors [Perdomo'24]

Discussion

- Evaluating predictive systems for efficient decision-making in societal contexts presents technical challenges, as the **assumptions underlying standard ML evaluations often fail to hold**.
- Our work helps identify key deviations from these assumptions and lays the groundwork for more context-aware evaluation of predictive systems in societal settings.

Thank you!

References

Ann-Marie Faria, Nicholas Sorensen, Jessica Heppen, Jill Bowdon, Suzanne Taylor, Ryan Eisner, and Shandu Foster. Getting students on track for graduation: Impacts of the early warning intervention and monitoring system after one year. Regional Educational Laboratory Midwest, 2017.

Martha Abele Mac Iver, Marc L Stein, Marcia H Davis, Robert W Balfanz, and Joanna Hornig Fox. An efficacy study of a ninth-grade early warning indicator intervention. Journal of Research on Educational Effectiveness, 2019.

Kube, A. R., Das, S., and Fowler, P. J. Fair and efficient allocation of scarce resources based on predicted outcomes: implications for homeless service delivery. Journal of Artificial Intelligence Research, 2023.

Mashiat T, DiChristofano A, Fowler PJ, Das S. Beyond Eviction Prediction: Leveraging Local Spatiotemporal Public Records to Inform Action. FAccT, 2024.

Perdomo, J. C. The Relative Value of Prediction in Algorithmic Decision Making. ICML, 2024.

Shirali, A., Abebe, R., & Hardt, M. Allocation Requires Prediction Only if Inequality Is Low. ICML, 2024.