



Full Text

Introduction & Motivation

Automated Decision Systems are becoming mainstream in societal resource allocation:

- ML models predict individuals' behavior, which informs resource allocation
- Recent literature grapples with how prediction and allocation interact [3]

What is the best strategy for prediction [6, 7]? :

- Measure underlying states? / Estimate counterfactual probabilities? / Utilitarian metric optimization?

How does prediction inform intervention allocation? [8]

Domain dependent. For example:

- Education: risk score of dropping out, informs allocation of extra tutoring
- Healthcare: risk score of patients not surviving, informs allocation of ICU space
- Prediction accuracy improvement $\neq$ improvement in social outcome [4,5]
- Prediction systems become part of bureaucratic decision making and affect who gets what.
- Prediction can improve allocation, depending on underlying inequality [3].

Core Question: How do prediction and allocation interact when prediction uncertainty is heterogeneous and outcomes are thresholded?

In scarce societal resource allocation literature, two common types of prioritization [9] :

- **Maximum Marginal Benefit (MMB)** : who benefits the most from allocation?
 - E.g., Battlefield triage, Covid ventilator allocation (Maximize Δ Pr (Survival))
- **Vulnerability First (VF)**. Commonly used in homelessness service allocation

We show that **even when underlying welfare states are identical**:

- Under a **scarce resource** budget, MMB favors individuals with low uncertainty
- **Abundance of resources flips prioritization**: Higher uncertainty individuals become favored

New ethical dilemma:

- Suppose two individuals have the same expected change in underlying state from an intervention. Is it OK to take predictive uncertainty into account in allocation?


Uncertainty Aware MMB (UA-MMB)

Under resource budget α , the policy $a(x, \sigma)$ can be described as:

$$\begin{aligned} \max_{a(\cdot, \cdot)} \quad & \mathbb{E}_{x, \sigma} [a(x, \sigma) \Pr(x + L + I \geq \theta) + (1 - a(x, \sigma)) \Pr(x + L \geq \theta)] \\ \text{s.t.} \quad & \mathbb{E}_{x, \sigma} [a(x, \sigma)] \leq \alpha. \end{aligned}$$

$$\text{Equivalently: } \max_{a(\cdot, \cdot)} \mathbb{E}_{x, \sigma} [a(x, \sigma) \Delta(x, \sigma)]; \quad \text{s.t.: } \mathbb{E}_{x, \sigma} [a(x, \sigma)] \leq \alpha$$

We call $\Delta(x, \sigma)$ the '**Utility gain function**'



$\Delta(x, \sigma)$ is defined as:


- Increase in probability of crossing θ if intervention is allocated:
 $\Pr[x + L + I \geq \theta] - \Pr[x + L \geq \theta]$

Nature of Δ as life uncertainty σ_s varies: (Theorem 1 in full text)

- Individuals with high x may face **harm** from intervention (negative Δ)
- When $\sigma_s = 0$, the 'benefit' to 'harm' crossover happens at exactly θ
 - Individuals **already above threshold** see increased chance of **falling into a negative welfare state** upon receiving intervention
- As σ_s increases, this crossover moves more to the right
 - Those **same individuals** above the threshold **see increased marginal benefit** when they have high predictive uncertainty

UU-MMB in action

- High uncertainty individuals receive allocation iff budget $\alpha > \text{Cutoff}$
- Below this cutoff, only low uncertainty individuals receive allocation
- UA-MMB performs triage (Theorem 2 in main text)
 - **Scarce** resource $\rightarrow$ **only low uncertainty** population receives intervention
 - **Plentiful** resource $\rightarrow$ **both** groups receive allocation
- Past this threshold, as budget increases, treated population mass grows faster for the high uncertainty group. (Theorem 3 in main text)



Acknowledgements

We are grateful for support from NSF Award 2533162.

References

- [1] OECD. *PISA 2022 Technical Report*. PISA, OECD Publishing, Paris, 2024.
- [2] OECD. *PISA 2022 Database*. 2023. <https://www.oecd.org/en/data/datasets/pisa-2022-database.html>.
- [3] Ali Shirali, Rediet Abebe, and Moritz Hardt. Allocation requires prediction only if inequality is low. ICML 2024.
- [4] Juan Carlos Perdomo, Tolani Britton, Moritz Hardt, and Rediet Abebe. Difficult lessons on social prediction from Wisconsin public schools. FAcT 2025.
- [5] Bryan Wilder, Bistra Dilikina, and Milind Tambe. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. AAAI 2019
- [6] Susan Athey. Beyond prediction: Using big data for policy problems. Science, 355(6324):483–485, 2017
- [7] Chelsea Barabas, Madars Virza, Karthik Dinakar, Joichi Ito, and Jonathan Zittrain. Interventions over predictions: Reframing the ethical debate for actuarial risk assessment. *1st Conference on Fairness, Accountability and Transparency 2018*.
- [8] Lydia T. Liu, Solon Barocas, Jon Kleinberg, and Karen Levy. On the actionability of outcome prediction. AAAI 2024
- [9] Jon Elster. Local Justice: How institutions allocate scarce goods and necessary burdens. Russell Sage 1992

Mathematical Model

For each individual, the allocation mechanism has a belief including:

- Scalar estimate of a baseline welfare state x ;
- Uncertainty level σ (think s.d.)

A threshold θ exists on underlying welfare states of individuals:

- If $x \geq \theta$, the individual is in a positive welfare state (e.g.: passed a test)
- If $x < \theta$, the individual is in a negative welfare state (e.g.: failed a test)

Prediction uncertainty is modelled $L = \sigma Z$

- We assume Z is a centered, symmetric, log-concave, location-scale family distribution
- Parameter σ defined convex order. Higher $\sigma \Leftrightarrow$ more uncertainty

Interventions (I) are additive and can introduce additional uncertainty:

- $I = \xi + \tau W$;
- Allocating intervention adds ξ to the underlying welfare state
- τW corresponds to introduced uncertainty:
 - W is a centered, symmetric, log-concave, location-scale family distribution
 - Higher $\tau \Leftrightarrow$ more uncertainty introduced

$$x' \mid \text{no intervention} = x + L = x + \sigma Z$$

$$x' \mid \text{intervention} = x + L + I = x + \xi + \sigma Z + \tau W;$$

x' is the underlying welfare state measured after interventions are distributed

Uncertainty Unaware MMB (UU-MMB), Vulnerability-First policies and Efficiency Loss


UU-MMB:

- The allocation algorithm follows maximum marginal gain
 - But only has access to population level estimate of prediction uncertainty
 - No individual estimates of prediction uncertainty.

VF: Intervention goes to individuals with lowest x values

Relative Efficiency Loss: By being oblivious to uncertainty, how much efficiency is lost?

- UU-MMB loses efficiency at low resources
- VF always suffers heavy efficiency loss, until population close to θ starts receiving intervention



Real World Study: Educational Achievement

We use the PISA 2022 dataset released by OECD [1,2]. The dataset contains:

- 15 year old students' cognitive assessments in math, reading, and science
- 613,744 records in total
- Questionnaire responses from students, school staff and parents
- 'Plausible Values' (PV) reported instead of raw score for standardization


- Students divided into low and high socioeconomic (SES) groups using variable **ESCS** (1st and 3rd tertile)

- **XGBoost** trained to predict math PVs using sociodemographic and prior academic performance features.

- **Training set:** 10,000 students (5,000 low-SES; 5,000 high-SES), randomly sampled from the full population.


- **Test set:** 1,000 students (500 low-SES; 500 high-SES), randomly sampled.

- We use prediction error RMS as prediction uncertainty
 - **High SES group** : 74.24; **Low SES group**: 63.65; **Global** : 69.16
 - The high SES group is also the higher uncertainty group!



Experiment:

- Intervention: Increases student PV by 10 points
- Goal: Increase PV past a proficiency level cutoff. We separately test 420 and 482 (basic and more proficient levels)
- Intervention budget varies from 0 to 300
- UA-MMB and UU-MMB both try to maximize total number of students crossing proficiency threshold



Key Takeaways:

- **At low resource**, UA-MMB's triage strategy is evident:
 - **Low uncertainty/low socioeconomic group prioritized**
- When resources increase, the high-uncertainty group receives allocation at a higher rate.