The Hidden Cost of Waiting for Accurate Predictions

Ali Shirali
UC Berkeley

Ariel D. Procaccia
Harvard University

Rediet Abebe
Harvard University

Abstract

Algorithmic predictions are increasingly informing societal resource allocations by identifying individuals for targeting. Policymakers often build these systems with the assumption that by gathering more observations on individuals, they can improve predictive accuracy and, consequently, allocation efficiency. An overlooked yet consequential aspect of prediction-driven allocations is that of timing. The planner has to trade off relying on earlier and potentially noisier predictions to intervene before individuals experience undesirable outcomes, or they may wait to gather more data to make more precise allocations. We examine this tension using a simple mathematical model, where the planner collects observations on individuals to improve predictions over time. We analyze both the *ranking* induced by these predictions and optimal *resource allocation*. We show that though prediction accuracy may improve over time, counter-intuitively, the average ranking loss can worsen. As a result, the planner's ability to improve social welfare can decline. We identify inequality as a driving factor behind this phenomenon. Our findings provide a nuanced perspective and challenge the conversational wisdom that it is preferable to wait for more accurate predictions to ensure the most efficient allocations.

1 Introduction

Algorithmic predictions are playing a central role in societal resource allocation. Policymakers and organizations are increasingly turning to algorithmically-driven systems in contexts where resources are scarce in order to target resources with greater precision [1–6]. Underpinning this growing reliance on predictions is the assumption that by gathering more observations about individuals over time, we can improve prediction accuracy and, consequently, targeting and allocation efficiency.

In practice, decisions around the *timing* of predictions and how they inform allocations reveal consequential trade-offs that the planner must navigate when designing prediction-driven allocation systems. On the one hand, the planner may wait to collect extensive data to refine their predictions before intervening. On the other hand, they can intervene early by relying on coarser data and noisier predictions. The potential advantage of the latter is that, in a fixed-horizon setting where the planner wants to prevent individuals from experiencing undesirable outcomes, the "window of opportunity" for this undesirable outcome to be realized closes. Furthermore, the underlying population changes with time, as those at greatest risk of experiencing such outcomes are more likely to "fail out" of the population early if they do not receive resources [7–9]. These factors pull in different directions, and it is not immediately apparent which factor dominates.

We examine this tension using a simple, versatile model where the planner predicts and intervenes on a population over time. Modeling a generic resource allocation problem, we assume the planner has a fixed budget of resources to prevent individuals from experiencing undesirable outcomes, such as eviction, job loss, poor health, or dropping out of school [4, 5, 10–14].[1] At each time step, the planner collects observations about individuals to improve their estimate of their underlying failure

[1] We present an extensive discussion on motivating and related work in Appendix B.

probability. The planner then uses the rankings induced by these estimates to allocate resources. Specifically, we ask:

1. *Ranking:* How does the ranking loss change as the planner collects more data, but some individuals fail out of the population?
2. *Allocation:* For a given instance of this problem, what is the optimal time to allocate resources? When is early intervention justified?

We present our results for two allocation problems: First, in a stylized setting, the planner is tasked with allocating all resources at once but can choose when to do so. We then use this as a building block to study the case where the planner can allocate resources over time. For both the ranking and allocation problems, we examine the role of inequality—as measured by the variance in the underlying failure probabilities—and surface it as a driving factor behind the optimal solutions.

We show that although prediction accuracy improves with more observations, counter-intuitively, this does not translate into improvements in the average ranking loss. To observe this, we decompose ranking loss into two counteracting effects: one due to improvements in prediction from additional observations and the second due to the change in population as individuals fail out of the active pool. We identify fundamental statistics that drive these two effects. We show that the change-in-population effect negatively impacts ranking performance and that this effect grows at least proportionally to the variance in the failure probabilities.

We then address both instantiations of the resource allocation problem. For the setting where the planner must allocate all resources simultaneously, we derive an upper bound on the optimal allocation time when targeting is broad, and the pool of active individuals is small. We show that, in this setting, allocating resources earlier yields greater social welfare. For the setting where the designer can allocate the budget over time, we design a provably optimal algorithm whose running time is independent of the number of individuals. Using this algorithm, we then demonstrate that the optimal solution can concentrate the allocation around any time-point t, depending on the prior distribution of failure probability among individuals.

Our results provide a nuanced perspective on the role of timing in prediction-driven allocations. In settings where the planner observes and intervenes on a population over time, they must balance the desire for more accurate predictions with the necessity for timely interventions. In the presence of significant inequality within the population, more accurate predictions do not necessarily lead to better ranking or improved allocations, providing a potential justification for early resource allocation.

2 Model and preliminaries

In this section, we first introduce the notations necessary to present our basic models. We provide further notation, as needed, throughout the paper and summarize the key notations along with their interpretation in Table 1.

We model the population over which the planner acts. We assume there is an initial population of N individuals and consider a finite horizon setting where $t \in [1, T]$.[2] Each individual i has some failure probability $p_i \in [0, 1]$, which captures their likelihood of dropping out of the population between time steps. This failure probability remains the same across time in the absence of interventions. Once an individual fails, they are no longer in the active pool of the population. We denote this active pool at time t by $\mathcal{A}^t$ and define $N^t := |\mathcal{A}^t|$ and $n^t := N^t/N$.

Prediction and ranking. At each time step t, the planner observes a signal o_i^t from each active individual $i \in \mathcal{A}^t$. These signals are analogous to observing loan or rental payments in housing and credit scoring, exam scores in education, and medical check-ups and tests in clinical settings. In our working model, these signals are drawn independently from a noisy Bernoulli process

$$o_i^t \sim \mathrm{Ber}(f(p_i)) \oplus \mathrm{Ber}(\epsilon)\,, \tag{1}$$

where we flip the Bernoulli random variable's value with probability $\epsilon < 1/2$. This noise parameter ϵ controls the level of uncertainty in the observations. The function $f : [0, 1] \to [0, 1]$ is an increasing

[2]Though we primarily consider the finite horizon setting, we also show that the key insights hold in the infinite horizon setting.

function: The more likely an individual is to fail, the more likely we are to observe signals indicating this possibility. When $\epsilon = 0$, individual i will leave $f(p_i)/p_i$ observations in expectation. Therefore, a larger f corresponds to more observations from individuals before they fail.

We note that $\mathrm{Ber}(f(p)) \oplus \mathrm{Ber}(\epsilon)$ is itself a Bernoulli random variable with effective parameter $\tilde{p}(p) \coloneqq (1-\epsilon)f(p) + (1-f(p))\epsilon$. For $\epsilon < 1/2$, this parameter $\tilde{p}$ is increasing in p since f is increasing. We omit the dependence of $\tilde{p}(\cdot)$ on p when it is clear from the context.

The planner is interested in the predictions as a means to rank individuals. Given observations drawn from Eq. (1) and a prior over the failure probability, we examine how the pairwise ranking risk of the Bayes optimal ranking, measured on the active population, changes over time.

Targeting and allocation. The planner has a budget B of resources to allocate to individuals in the active pool. Examples include housing vouchers, food assistance programs, unemployment insurance, or preventative health screenings. We assume a fixed unit cost in the budget to assign a resource to an individual at any time t. Once an individual i receives such a resource, they will not fail out of the population at subsequent time steps.

When the planner allocates a resource to an individual, the utility of this intervention is equivalent to the probability that this individual would have failed by time T without this resource. That is,

$$u^t(p) \coloneqq 1 - (1-p)^{T-t}\,. \tag{2}$$

The optimal allocation maximizes the overall utility over the whole population.

We consider two allocation problems: We first study a stylized version, which we call *one-time allocation*, where the planner is tasked with finding the optimal time t to allocate their entire budget to maximize the total utility. Using this as a building block, we then consider an *over-time allocation* problem, where the planner can spend their budget over time. The planner is tasked with finding the optimal distribution of the budget across time. For both versions, the planner allocates resources to the individuals with the highest predicted rank at that time.

3 Ranking over time

The planner uses predictions as risk scores to rank individuals, with the goal of prioritizing the most vulnerable individuals. We are therefore interested in the ranking loss. We decompose the ranking problem into finding a pairwise ranker that, given observations from two individuals, predicts which individual has a higher failure probability.

Define $y_i^t \coloneqq \sum_{t' \in [t]} o_i^{t'}$ as the number of positive observations from an active individual i up to time t. Let $\mathcal{P}^t$ be the posterior over p for an individual active at t and assume that $\mathcal{P}^t$ has no point mass. Proposition E.4 shows that, under the observation model in Eq. (1), ranking individuals based on their y^t is Bayes optimal in terms of the zero-one pairwise ranking loss. Formally, the zero-one risk of optimal (pairwise) ranking at time t is

$$R^t = \mathrm{Pr}_{1,2}^t\big(y_2^t < y_1^t \mid p_2 \geq p_1\big)\,.$$

Here, $\mathrm{Pr}_{1,2}^t(\cdot)$ is the probability involving two independently chosen active individuals at t. For a given $\tilde{p}$, the term y^t follows $\mathrm{Binomial}(t, \tilde{p})$. For analytic tractability, we approximate this distribution with $\mathcal{N}(t \cdot \tilde{p}, t \cdot \tilde{\sigma}^2)$, where $\tilde{\sigma}^2 \coloneqq \tilde{p} \cdot (1-\tilde{p})$. The independence of the draws implies that $\frac{y_2^t - y_1^t}{t} \mid (\tilde{p}_1, \tilde{p}_2) \sim \mathcal{N}\big(\tilde{p}_2 - \tilde{p}_1, \frac{\tilde{\sigma}_{12}^2}{t}\big)$, where $\tilde{\sigma}_{12}^2 \coloneqq \tilde{\sigma}_1^2 + \tilde{\sigma}_2^2$. Denoting the cumulative distribution function of the standard normal distribution with $G(\cdot)$, it is then straightforward to simplify R^t as

$$R^t \approx \mathbb{E}_{1,2}^t\Big[G\big(-\frac{|\tilde{p}_2 - \tilde{p}_1|}{\tilde{\sigma}_{12}}\sqrt{t}\big)\Big]\,.$$

The dependence of R^t on t appears in two places: inside $G(\cdot)$ and in $\mathbb{E}_{1,2}^t$. The former captures the effect of gathering more observations over time, and the latter models a change-in-population effect.

We are interested in understanding the change in R^t over one time step: $\Delta R^t \coloneqq R^{t+1} - R^t$. The sign of ΔR^t determines whether ranking improves over time. Using the above observations, we can decompose ΔR^t into two parts: one capturing the change of $G(\cdot)$'s argument and another capturing the change of $\mathbb{E}^t_{1,2}$ due to the change of $\mathcal{P}^t(p_1)\,\mathcal{P}^t(p_2)$. We can approximate the former by taking the derivative with respect to t. To find the latter, since an active individual at time t with a failure probability of p survives until $t+1$ with a probability of $1-p$, we can write

$$\mathcal{P}^{t+1}(p) = \big(\frac{1-p}{1-\mu^t}\big)\,\mathcal{P}^t(p)\,, \tag{3}$$

where $\mu^t := \mathbb{E}^t[p]$. This allows us to write $\mathbb{E}^{t+1}_{1,2}[\cdot] = \mathbb{E}^t_{1,2}[(1-p_1)\,(1-p_2)\cdot(\cdot)]/(1-\mu^t)^2$. By approximating $G(x)$ with $\frac{1}{\sqrt{2\pi x}}\exp(-x^2/2)$, we obtain

$$\Delta R^t \approx \mathbb{E}^t_{1,2}\left[\frac{1}{\sqrt{2\pi t}}\exp\Big(\frac{-(\tilde{p}_2-\tilde{p}_1)^2}{2\tilde{\sigma}_{12}^2}t\Big)\cdot\Big\{\frac{\tilde{\sigma}_{12}}{|\tilde{p}_2-\tilde{p}_1|}\big[\frac{(1-p_1)(1-p_2)}{(1-\mu^t)^2}-1\big]-\frac{|\tilde{p}_2-\tilde{p}_1|}{2\tilde{\sigma}_{12}}\Big\}\right]. \tag{4}$$

Intuitively, we can think of the exponential multiplier as a kernel enforcing similarity of $\tilde{p}_1$ and $\tilde{p}_2$. Such a filter pushes ΔR^t towards positive values, which means that as the filter becomes stricter for larger t, the effect of losing vulnerable individuals dominates the gain from collecting more observations. We formalize this intuition in the following proposition.

Proposition 3.1. *Suppose there exists a constant $\alpha \in (0,1)$ bounded away from zero such that plugging the approximation $\exp(-x^2) \approx \mathbb{1}\{\exp(-x^2) \geq \alpha\}$ into Eq. (4) does not change the value of ΔR^t. Denote the mean and variance of p for an active individual at t by $\mu^t := \mathbb{E}^t[p]$ and $\mathrm{Var}^t[p]$, respectively. If f^{-1} in Eq. (1) is $O(1)$-Lipschitz, we can bound ΔR^t as*

$$\Delta R^t \geq \sqrt{\frac{t}{2\ln(1/\alpha)}}\cdot\Big\{\underbrace{\frac{\mathrm{Var}^t[p]}{(1-\mu^t)^2} - O\big(\frac{1}{\sqrt{t}}\big)}_{\text{loss of vulnerable individuals}} \underbrace{-\frac{\ln(1/\alpha)}{\sqrt{2\pi}\,t}}_{\text{gain in observations}}\Big\}. \tag{5}$$

See proof on page 19. The proof presents the exact form of $O\left(1/\sqrt{t}\right)$, which we skip here for clarity of exposition. The lower bound of Eq. (5) decomposes into two terms: the first due to vulnerable individuals dropping out of the population and the second due to information gain. Factoring the common terms, the positive term grows proportionally to $\mathrm{Var}^t[p]/(1-\mu^t)^2$ and the negative term scales with $1/\sqrt{t}$. So, a higher variance of p and a larger μ^t (corresponding to the average failure probability) contribute positive cost to waiting. On the other hand, the gain from collecting more observations declines rapidly. Therefore, in the presence of high inequality or a high average failure probability, the cost of waiting for accurate predictions may not be justified.

4 One-time allocation

In this instantiation of the problem, the planner must select a time t where they allocate their entire budget. This stylized problem will provide preliminary insights into the allocation problem and provide a building block for the more general problem studied in Section 5.

Our main observation in this section is that under weak conditions, the optimal allocation time to maximize an *eligibility-based utility* function occurs before targeting becomes *broad*.

Theorem 4.1 (Informal)**.** *Assuming a Lipschitz condition on f in the observation model of Eq. (1), if targeting is broad, deferring the allocation is not justified.*

An eligibility-based utility is a lower bound for the utility u^t in Eq. (2). It assumes that a fixed utility will be realized if an individual with $u^t(p)$ larger than a threshold u^* is treated, i.e., $\underline{u}^t(p) := u^*\cdot\mathbb{1}\{u^t(p)\geq u^*\}$. Focusing on this notion of utility allows us to establish a clearer connection between the ranking and allocation problems and to theoretically analyze the optimal allocation.

Roughly speaking, targeting becomes broad when there are more eligible individuals than the available budget. Targeting may become broad for two particular reasons: As we approach the horizon, u^t decreases, and therefore, even for a fixed population, the pool of eligible individuals shrinks due to a stricter eligibility requirement. On the other hand, even if the eligibility criterion is fixed in terms of the minimum value of p, the relative number of eligible individuals above a threshold decreases as the individuals with a higher p are more likely to fail. We will connect this latter phenomenon to inequality, as measured by the Gini index of the p values.

To prove Theorem 4.1, we first derive the expected utility of allocating the budget at time t to maximize an eligibility-based utility in the limit of many individuals. We then formally define broad targeting and its relation with inequality. By examining how overall utility changes over time, we prove that broad targeting justifies early allocation.

4.1 Derivation of the total utility

For a chosen time t, we assume the planner ranks the active individuals at t, denoted by $\mathcal{A}^t$, and targets the first B of them. Denote such a ranking by injection $r^t : [N^t] \to \mathcal{A}^t$, where, for example, $r^t(1)$ is the most eligible individual from the planner's view. We assume the planner uses a Bayes optimal ranking. Assuming the prior over failure probability p has no point mass, Proposition E.4 implies that under the observation model of Eq. (1) the optimal ranking would be to simply rank individuals in descending order in terms of their number of positive observations thus far, denoted by y^t. In case of a tie, we can assume any arbitrary individual will be chosen.

For a ranking r^t, the total utility of allocating budget B at t is $\sum_{i=1}^{B} u^t\big(p_{r^t(i)}\big)$. To make the connection between the allocation problem and the ranking problem clearer and to make the solution analytically easier to follow, we next derive a lower bound for this utility using a threshold function that assesses eligibility. We then redefine the planner's problem based on this lower bound.

Eligibility-based utility function. Suppose an individual i is eligible for receiving intervention if their utility $u^t(p_i)$ exceeds a certain threshold u^*. This eligibility criterion can define a new utility function $\underline{u}^t(p) := u^* \cdot \mathbb{1}\{u^t(p) \geq u^*\}$ that serves as a lower bound for $u^t(p)$. The monotonicity of u^t allows us to simplify $\underline{u}^t$ as an eligibility criterion over p instead of u^t. More precisely, excluding $t = T$ since spending budget at T is not rational ($u^T(p) = 0$), $u^t(p)$ of Eq. (2) is increasing in p, and we can write $\underline{u}^t(p) = u^* \cdot \mathbb{1}\{p \geq c^t\}$, where

$$c^t := (u^t)^{-1}(u^*) = 1 - (1 - u^*)^{1/(T-t)} \,. \tag{6}$$

Using $\underline{u}^t(p) = u^* \cdot \mathbb{1}\{p \geq c^t\}$ in place of u^t, the planner's problem becomes

$$\arg\max_t \sum_{i=1}^{N^t} \mathbb{1}\{p_{r^t(i)} \geq c^t\} \cdot \mathbb{1}\{i \leq B\} = \sum_{i \in \mathcal{A}^t} \mathbb{1}\{p_i \geq c^t\} \cdot \mathbb{1}\{(r^t)^{-1}(i) \leq B\} \,. \tag{7}$$

Substituting r^t by the Bayes optimal ranking, we can further upper bound $(r^t)^{-1}(i)$ by $\sum_{j \in \mathcal{A}^t} \mathbb{1}\{y_j^t \geq y_i^t\}$. Plugging this into Eq. (7) further bounds the total utility of the allocation. We assume the planner maximizes this lower bound, referred to as U^t, in order to decide on the optimal time of intervention.

Total utility in the limit of many individuals. As the number of individuals grows ($N \to \infty$) while the ratio B/N remains fixed, the planner must solve

$$\arg\max_t \; U^t := N\, n^t \Pr_1^t(p_1 \geq c^t) \cdot \Pr_1^t\Big(n^t \Pr_2^t(y_2^t \geq y_1^t) \leq B/N \mid p_1 \geq c^t\Big) \,. \tag{8}$$

In this context, $\Pr_i^t$ denotes the probability when $p_i \sim \mathcal{P}^t$ and $y_i^t = \sum_{t' \in [t]} o_i^{t'}$ is drawn according to the observation model of Eq. (1). Our focus in the following will be on this asymptotic case.

Introducing $\tilde{b}^t$ and simplifying U^t. Define $\mathrm{ER}^t(k) := N^t \Pr^t(y^t \geq k)$ as the expected pessimistic rank for an individual with k positive observations at t. Note that $\mathrm{ER}^t(k)$ is a decreasing function of k. Suppose the budget is neither excessively large, so that $B \leq N$, nor very small, ensuring that it can always treat those individuals who have consistently shown positive observations until t: $B \geq \sum_{i \in \mathcal{A}^t} \mathbb{1}\{y_i^t = t\}$. Then

$$\tilde{b}^t := \frac{1}{t} \cdot \min\big\{k \mid \mathrm{ER}^t(k) \leq B\big\} \tag{9}$$

is well-defined. Intuitively, $\tilde{b}^t$ represents the minimum proportion of times an agent must send a positive signal to be expected among the B targeted individuals. As individuals with higher $\tilde{p}$ are more likely to fail at each time step, the threshold to be expected among the top B individuals decreases, so $\tilde{b}^t$ decreases over time. We also define $b^t := \tilde{p}^{-1}(\tilde{b}^t)$ as the failure probability corresponding to $\tilde{b}^t$. Using $\tilde{b}^t$, we can further simplify U^t in Eq. (8) as

$$U^t = N^t \Pr^t(p \geq c^t) \cdot \Pr^t(y^t \geq t \cdot \tilde{b}^t \mid p \geq c^t) \,. \tag{10}$$

4.2 Broad targeting

Intuitively, if there are not many eligible individuals with $p \geq c^t$ compared to the available budget, deferring allocation may not be justified, as it risks losing those already rare eligible ones. We call this scenario a case of *broad targeting*. Intuitively, the targeting may become broad due to either a high value of c^t as we approach the horizon or the lack of individuals with a high p. We next formally define broad targeting and the factors that may constitute it.

First, it is helpful to introduce some shorthand notation that will be used throughout the rest of this section: Let $\tilde{\sigma}(p)^2 := \tilde{p}(p) \cdot (1 - \tilde{p}(p))$ denote the variance of observations from an individual with failure probability p, and let $\tilde{c}^t := \tilde{p}(c^t)$. With these definitions, broad targeting formally means:

Definition 4.2 (Broad targeting)**.** We say targeting is broad at t if $c^t \geq 1/2$ and

$$\tilde{c}^t - \tilde{b}^t \geq \frac{\tilde{\sigma}^t(c^t)}{\sqrt{t}} \,. \tag{11}$$

Targeting becomes more broad as time progresses: c^t, and so $\tilde{c}^t$, increases over time while $\tilde{b}^t$ decreases. Therefore, the left-hand side of Eq. (11) increases while the right-hand side decreases. Note that a smaller value of $\tilde{b}^t$ contributes to a broader targeting. In the extreme case of a negligible $\tilde{b}^t$, the condition for broad targeting reduces to $\tilde{c}^t \geq \tilde{\sigma}^t(c^t)/\sqrt{t}$.

We next show that high inequality necessarily implies a small $\tilde{b}^t$. More precisely, denoting the Gini index of $\tilde{p}$ at time t by $\widetilde{G}^t$, in Proposition 4.3, we show an upper bound on $\tilde{b}^t$ that scales with $(1 - \widetilde{G}^t)^2$. Therefore, high inequality at time t implies a small value of $\tilde{b}^t$ and consequently pushes the allocation towards a broader targeting regime.

Proposition 4.3. *Suppose the probability density function of $\tilde{p}$ at time t is non-increasing. In the limit of many individuals, denote the (population) Gini index of $\tilde{p}$ at time t by $\widetilde{G}^t$. If $\widetilde{G}^t \geq 1/3$, then $\tilde{b}^t$ (Eq. (9)) is bounded from above:*

$$\tilde{b}^t \leq \frac{9}{8} \cdot \frac{N^t}{B} \cdot (1 - \widetilde{G}^t)^2 + \frac{1}{t} \,.$$

See proof on page 19. Complementing this observation, Proposition E.6 shows that $\widetilde{G}^t$ can only increase over time. So, an initially high inequality pushes the allocation towards broader targeting at every subsequent step. We next show allocation should not be postponed if the targeting is broad.

4.3 Optimal intervention time in case of broad targeting

At the heart of Theorem 4.1 is the following lemma, which provides an upper bound on the change in U^t over a single time step. If this change is negative at time t, then the optimal allocation should occur at t or earlier. The lemma decomposes the change of U^t into two counteracting effects: a positive effect due to gathering more information and a negative effect caused by losing easy-to-predict vulnerable individuals. Using this decomposition, we will then show that when targeting becomes broad, the latter negative effect becomes dominant, which will prove Theorem 4.1.

Lemma 4.4 (Decompose and upper bound $\frac{\mathrm{d}U^t}{\mathrm{d}t}$)**.** *Let t^* be the earliest time when targeting is broad according to $Definition$ 4.2 and suppose $\mathcal{P}^{t^*}(p)$ is a non-increasing function of p. By approximating the binomial distribution with a Gaussian distribution according to the central limit theorem, for $t \geq t^*$, we can bound the change in U^t over one step as*

$$\frac{1}{N^t}\frac{\mathrm{d}U^t}{\mathrm{d}t} \leq \Pr{}^t(p \geq c^t) \cdot \Big\{ \underbrace{\frac{\tilde{c}^*}{2\tilde{\sigma}^*} \cdot \frac{1}{\sqrt{t}}\, g\big(\sqrt{t/t^*}\big)}_{\text{gain in observations}} \underbrace{- G\big(\sqrt{t/t^*}\big) \cdot \Big[\frac{c^t - b^t}{2} + \frac{\mathrm{d}c^t}{\mathrm{d}t}\Big]}_{\text{loss of vulnerable individuals}} \Big\} \,. \tag{12}$$

Here $G(\cdot)$ and $g(\cdot)$ denote the cumulative distribution and density function of the standard normal distribution, and $\tilde{c}^*$ and $\tilde{\sigma}^*$ are shorthands for $\tilde{c}^{t^*}$ and $\tilde{\sigma}(c^{t^*})$.

See proof on page 20. This theorem suggests that while the marginal gain from extra information diminishes exponentially, the marginal cost of losing vulnerable individuals from the active pool consistently increases. We next present the formal statement of Theorem 4.1.

Theorem 4.1. *Let t^* be the earliest time when targeting is broad according to $Definition$ 4.2 and suppose $\mathcal{P}^{t^*}(p)$ is a non-increasing function of p. Approximating the binomial distribution with a Gaussian distribution, if f in Eq. (1) is L-Lipschitz with $L \leq \frac{\sqrt{2\pi}\, G(1)}{1-2\epsilon}$, then $\frac{\mathrm{d}U^t}{\mathrm{d}t} \leq 0$ for any $t \geq t^*$.*

See proof on page 23.

5 Over-time allocation

We now study the problem of allocating B over time. In general, the state at every time step t can be described by the number of individuals with $y^t = k$, denoted by N_k^t, for $k \leq t$. An allocation policy then specifies how many individuals from each group k will be treated at each state. The policy and the current state, along with the prior distribution over p, are sufficient to specify a distribution over the next state. Therefore, we can think of the optimal dynamic allocation problem as a Markov decision process (MDP). We first characterize the optimal over-time allocation as the optimal policy for this MDP. Based on this characterization, we then devise an efficient algorithm to find the optimal allocation. Using this algorithm, we simulate different scenarios and intuitively present the factors that can lead to early intervention being optimal.

5.1 Characterizing the optimal over-time allocation

Lemma E.2 implies that individuals with a higher y^t yield a higher expected utility. Therefore, the optimal allocation at every time t should treat those with the highest y^t. Given a budget of B, a rollout of such a policy in the described MDP can be specified by the budget spent at each time step. Therefore, there are $\binom{B+1}{T-1}$ possibilities. In the case of many agents, and so a large B, the MDP dynamics become almost deterministic due to the law of large numbers, and the optimal policy converges to a single fixed rollout. However, naively iterating through all $\binom{B+1}{T-1}$ possibilities to find the optimal rollout is computationally infeasible. The next theorem further characterizes the optimal allocation in the limit of many individuals.

Theorem 5.1 (Characterize optimal over-time allocation)**.** *Consider any utility function $u^t(p)$ that is non-decreasing in p and non-increasing in t. Assuming the prior over failure probability has no point mass, the optimal over-time allocation in the limit of many individuals follows a specific pattern: For a non-decreasing sequence $q : [T] \to \{0, 1, \ldots, T\}$, and an exceptional time step $\hat{t} \in [T]$,*

- *At $t \neq \hat{t}$, everyone with $y^t \geq q(t)$ will be treated, and $q(t+1) \in \{q(t), q(t)+1\}$.*
- *At $t = \hat{t}$, everyone with $y^t > q(t)$ and some with $y^t = q(t)$ will be treated, and $q(t+1) \in \{q(t)+1, q(t)+2\}$.*

See proof on page 23.

5.2 Designing an efficient algorithm

Theorem 5.1 leaves three parameters of the optimal allocation unspecified: $\hat{t}$, $q(\cdot)$, and the portion of individuals with $y^{\hat{t}} = q(\hat{t})$ to be treated at $\hat{t}$. We next explain how we can search this space.

There are T possibilities for $\hat{t}$. A valid sequence $q(\cdot)$ can then be determined by specifying $q(1)$ and a binary sequence of length $T-1$. The t^{th} binary value in this sequence determines whether $q(t+1) - q(t)$ will be: 0 or 1 if $t \neq \hat{t}$, or 1 or 2 if $t = \hat{t}$. Since the maximum y^t at each time t is t, one can verify that for any sequence that starts with $q(1) > 2$, there exists another sequence with $q(1 \leq 2)$ that treats similar individuals. Therefore, there are effectively only three choices for $q(1)$. Thus far, we counted $3T \cdot 2^{T-1}$ possibilities.

Given $\hat{t}$ and a valid sequence $q(\cdot)$, the allocation at $\hat{t}$ can be determined based on the available budget. Let $\mathcal{A}_k^t$ be the set of active individuals at t with $y^t = k$. Denote the proportion of $\mathcal{A}_{q(\hat{t})}^{\hat{t}}$ who are not treated by ρ. One can verify that for a fixed $\hat{t}$ and $q(\cdot)$, both the spending and the expected total utility are linear in ρ. Therefore, to find ρ, we only need to simulate an allocation for two distinct values of ρ. This allows us to identify the linear relationship and determine the optimal ρ under the budget constraint. In Algorithm 1, we do this by simulating the trajectories for $\rho = 0$ and $\rho = 1$.

We use Algorithm 2 to simulate a trajectory. At the heart of this algorithm is the backup formula that updates $N_k^t =:= |\mathcal{A}_k^t|$ based on N_k^{t-1} and N_{k-1}^{t-1}. Algorithm 2 also requires calculating expectation with respect to $p \sim \mathcal{P}^t(\cdot \mid y^t = k)$. In our simulation, this posterior has a closed form. However, in general, since p is a bounded scalar, the posterior calculation can be approximated by a constant number of operations. The expectation of utility under this posterior will also be used in calculating the expected total utility in Algorithm 1.

Since simulating a trajectory in Algorithm 2 requires $O(T^2)$ operations, the overall running time of Algorithm 1 is $O(T^3 \cdot 2^T)$. In practical applications, such as a tenant facing eviction or a student at risk of dropping out, time steps are typically on the scale of a year or a month. Therefore, T is usually a small constant. The significant improvement of Algorithm 1 compared to the naive iteration over $\binom{B+T-1}{T-1}$ possible trajectories is its independence from the number of individuals and B. With the help of this algorithm, we next simulate the optimal over-time allocation in multiple scenarios.

Algorithm 1 Optimal over-time allocation

$U_{\text{opt}} \leftarrow 0$
for $\hat{t} = 1$ to T and $q(\cdot) \in$ valid sequences according to Theorem 5.1 **do**
  Simulate as $\mathcal{A}^{\hat{t}}_{q(\hat{t})}$ *are all treated:*
    $\{N^t_{q(t)}\}_{t=1}^T \leftarrow \textsc{SimulateTraj}(1, \{N_0^1, N_1^1\}, q(\cdot))$
    $E_{\max} \leftarrow \mathbb{1}\{q(1) = 0\} \cdot N_1^1 + \sum_{t=2}^T N^t_{q(t)}$ ▷ maximum expenditure
  Simulate the difference as if no one in $\mathcal{A}^{\hat{t}}_{q(\hat{t})}$ *was treated:*
    $\{\Delta N^t_{q(t)}\}_{t=\hat{t}}^T \leftarrow \textsc{SimulateTraj}(\hat{t}, \{0, \ldots, 0, N^{\hat{t}}_{q(\hat{t})}, 0, \ldots, 0\}, q(\cdot) + \mathbb{1}\{(\cdot) = \hat{t}\})$
    $\Delta E \leftarrow N^{\hat{t}}_{q(\hat{t})} - \sum_{t=\hat{t}}^T \Delta N^t_{q(t)}$ ▷ decrease from the max. expenditure
  $\rho \leftarrow \frac{E_{\max} - B}{\Delta E}$ ▷ proportion of $\mathcal{A}^{\hat{t}}_{q(\hat{t})}$ to be left untreated
  if $\rho \geq 1$ or $\rho \leq 0$ **then** continue to the next possible $q(\cdot)$ **end if**
  Calculate the total expected utility:
    for $t = 1$ to T and $k = q(t)$ **do** $u_k^t \leftarrow \mathbb{E}_{p \sim \mathcal{P}^t(\cdot \mid y^t = k)}[u^t(p)]$ **end for**
    $U \leftarrow (1 - \rho)\, N^{\hat{t}}_{q(\hat{t})}\, u^{\hat{t}}_{q(\hat{t})} + \sum_{t \neq \hat{t}} (N^t_{q(t)} + \rho\, \Delta N^t_{q(t)})\, u^t_{q(t)} + \mathbb{1}\{q(1) = 0\} \cdot N_1^1\, u_1^1$
  if $U > U_{\text{opt}}$ **then** $U_{\text{opt}} \leftarrow U$, $q_{\text{opt}} \leftarrow q$, $\hat{t}_{\text{opt}} \leftarrow \hat{t}$ **end if** ▷ check for optimality
end for
return $U_{\text{opt}}, q_{\text{opt}}, \hat{t}_{\text{opt}}$

5.3 Simulations

Algorithm 1 enables us to tractably find the optimal over-time allocation in various simulated settings. Using this algorithm, we next present the effect of the prior distribution of p and the relative budget size. For simplicity, suppose $f(p) = p$ and $\epsilon = 0$, so $\tilde{p} = p$. Suppose the failure probabilities of individuals before entering the process are drawn from $\mathcal{P}^0 = \text{Beta}(\alpha, \beta)$. Since the Beta distribution is a conjugate prior for the Binomial distribution, the posterior over failure probability has a closed form: $\mathcal{P}(\cdot \mid y^t = k) = \text{Beta}(\alpha + k, \beta + 2t - k)$.

Fig. 1 shows the optimal allocation along with the optimal $q(\cdot)$ over $T = 5$ time steps for three different prior distributions. One can see that as the population concentrates around smaller values of p, distinguishing individuals becomes harder, and the call to predict and act becomes less pressing. Therefore, more of the budget is allocated at later times. Such deferral requires a significant population close to $p = 0$ though, and typically early interventions are preferred.

We study the effect of budget size in Fig. 2 (in Appendix D). For a fixed distribution, we change the relative size of the budget compared to the initial number of individuals. As the figure suggests, prediction will be postponed only when the relative size of the budget becomes very small. This is consistent with our observation in Section 4 that deferring prediction and allocation when the budget is large and targeting is broad cannot be justified.

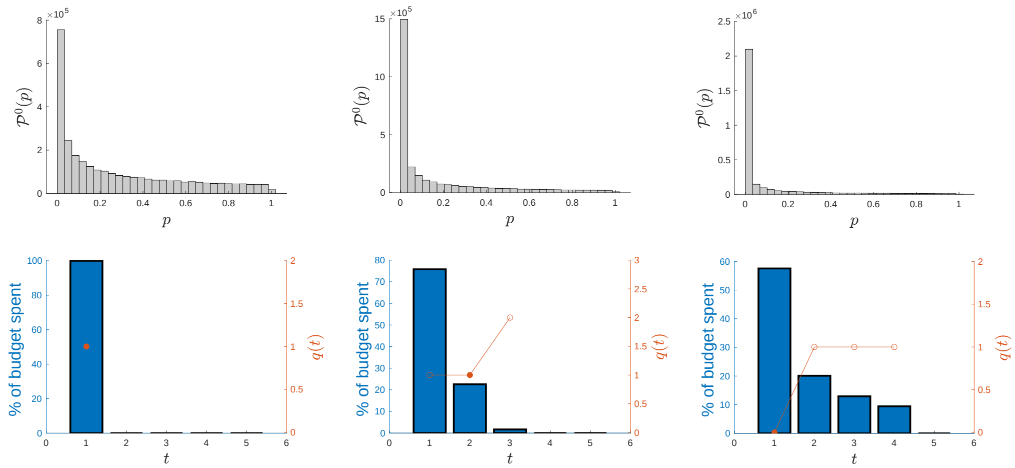

(a) $\mathcal{P}^0 = \text{Beta}(\alpha = 0.4, \beta = 1)$ (b) $\mathcal{P}^0 = \text{Beta}(\alpha = 0.2, \beta = 1)$ (c) $\mathcal{P}^0 = \text{Beta}(\alpha = 0.1, \beta = 1)$

Figure 1: Optimal over-time allocation for three different priors and a fixed budget. The orange curve depicts the optimal $q(\cdot)$ and the filled circle corresponds to $t = \hat{t}$.

6 Discussion

Our work contributes a timing dimension to an emerging body of research on evaluating prediction-driven allocation. Predictive systems are introduced with the promise of minimizing waste and increasing efficiency. This existing predisposition, amplified further by the traditional focus on maximizing predictive accuracy, encourages practices that favor waiting to collect more information over acting early with noisier signals. Our study presents a simple model that challenges this practice.

Our work opens numerous lines of inquiry. For instance, we assume the planner has a fixed budget B, corresponding to a fixed unit-cost intervention, that they can allocate all at once or over time. There are various natural variations worth exploring: For instance, we could consider heterogeneity in cost across time or different p_i values. In the same spirit as Perdomo [15], we can also consider trading off this B with other interventions or parameters in the problem.

We make generic assumptions about the failure probabilities and collection of observations. In settings motivating our study, the failure probabilities change over time, favoring increasing inequality in the absence of interventions. Likewise collecting observations for vulnerable individuals may be more costly, contain less signal, or may otherwise be undesirable [16]. Finally, though the individuals have heterogeneous values of p_i, we do not assume that they have different "starting points." Enriching the model we study to include such insights would only further justify early interventions in the presence of high inequality, though it would be interesting to examine the extent to which it does so.

Our work introduces a potential lens through which to examine tradeoffs incurred by waiting to improve prediction accuracy. Our results, on their own, do not endorse early or late allocations for any specific setting. Each policy problem should be examined empirically, and policymakers must consider various community, policy, and practical considerations. Indeed, targeting as an effective means of improving welfare—which has fueled the use of predictive systems—is, itself, an actively debated policy concept. Nonetheless, we believe that the machine learning community can contribute to discussions around how to best evaluate predictive systems in such policy settings.

References

[1] Virginia Eubanks. *Automating inequality: How high-tech tools profile, police, and punish the poor.* St. Martin's Press, 2018.

[2] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. Prediction policy problems. *American Economic Review*, 105(5):491–495, 2015.

[3] Amanda R Kube, Sanmay Das, and Patrick J Fowler. Fair and efficient allocation of scarce resources based on predicted outcomes: Implications for homeless service delivery. *Journal of Artificial Intelligence Research*, 76:1219–1245, 2023.

[4] Tasfia Mashiat, Alex DiChristofano, Patrick J Fowler, and Sanmay Das. Beyond eviction prediction: Leveraging local spatiotemporal public records to inform action. *arXiv preprint arXiv:2401.16440*, 2024.

[5] Juan Carlos Perdomo, Tolani Britton, Moritz Hardt, and Rediet Abebe. Difficult lessons on social prediction from Wisconsin public schools. *arXiv preprint arXiv:2304.06205*, 2023.

[6] Halil Toros and Daniel Flaming. Prioritizing homeless assistance using predictive algorithms: an evidence-based approach. *Cityscape*, 20(1):117–146, 2018.

[7] Rediet Abebe, Jon Kleinberg, and S Matthew Weinberg. Subsidy allocations in the presence of income shocks. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, pages 7032–7039, 2020.

[8] Karen Bierman. The implementation of the fast track program: An example of a large-scale prevention science efficacy trial. *Journal of Abnormal Child Psychology*, 30:1–17, 2002.

[9] Matthew J Salganik, Ian Lundberg, Alexander T Kindel, Caitlin E Ahearn, Khaled Al-Ghoneim, Abdullah Almaatouq, Drew M Altschul, Jennie E Brand, Nicole Bohme Carnegie, Ryan James Compton, et al. Measuring the predictability of life outcomes with a scientific mass collaboration. *Proceedings of the National Academy of Sciences*, 117(15):8398–8403, 2020.

[10] Sebastian Zezulka and Konstantin Genin. From the fair distribution of predictions to the fair distribution of social goods: Evaluating the impact of fair machine learning on long-term unemployment. *arXiv preprint arXiv:2401.14438*, 2024.

[11] Carri W Chan, Vivek F Farias, Nicholas Bambos, and Gabriel J Escobar. Optimizing intensive care unit discharge decisions with patient readmissions. *Operations research*, 60(6):1323–1341, 2012.

[12] Vinithasree Subbhuraam, editor. *Predictive Analytics in Healthcare*. IOP Publishing, 2021.

[13] 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.

[14] 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*, 12(3):363–390, 2019.

[15] Juan Carlos Perdomo. The relative value of prediction in algorithmic decision making. *arXiv preprint arXiv:2312.08511*, 2023.

[16] Lucas Monteiro Paes, Carol Long, Berk Ustun, and Flavio Calmon. On the epistemic limits of personalized prediction. *Advances in Neural Information Processing Systems*, 35:1979–1991, 2022.

[17] Emily Aiken, Suzanne Bellue, Dean Karlan, Chris Udry, and Joshua E Blumenstock. Machine learning and phone data can improve targeting of humanitarian aid. *Nature*, 603(7903):864–870, 2022.

[18] Ziad Obermeyer and Ezekiel J Emanuel. Predicting the future—big data, machine learning, and clinical medicine. *The New England journal of medicine*, 375(13):1216, 2016.

[19] Debopam Bhattacharya and Pascaline Dupas. Inferring welfare maximizing treatment assignment under budget constraints. *Journal of Econometrics*, 167(1):168–196, 2012.

[20] Toru Kitagawa and Aleksey Tetenov. Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica*, 86(2):591–616, 2018.

[21] Susan Athey and Stefan Wager. Policy learning with observational data. *Econometrica*, 89(1): 133–161, 2021.

[22] Susan Athey. Beyond prediction: Using big data for policy problems. *Science*, 355(6324): 483–485, 2017.

[23] Karen Levy, Kyla E Chasalow, and Sarah Riley. Algorithms and decision-making in the public sector. *Annual Review of Law and Social Science*, 17:309–334, 2021.

[24] Angelina Wang, Sayash Kapoor, Solon Barocas, and Arvind Narayanan. Against predictive optimization: On the legitimacy of decision-making algorithms that optimize predictive accuracy. *ACM Journal on Responsible Computing*, 1(1):1–45, 2024.

[25] Chelsea Barabas, Madars Virza, Karthik Dinakar, Joichi Ito, and Jonathan Zittrain. Interventions over predictions: Reframing the ethical debate for actuarial risk assessment. In *Proceedings of the 1st ACM Conference on Fairness, Accountability, and Transparency*, pages 62–76, 2018.

[26] Xavier de Souza Briggs, Susan J Popkin, and John Goering. *Moving to opportunity: The story of an American experiment to fight ghetto poverty*. Oxford University Press, 2010.

[27] Jens Ludwig, Jeffrey B Liebman, Jeffrey R Kling, Greg J Duncan, Lawrence F Katz, Ronald C Kessler, and Lisa Sanbonmatsu. What can we learn about neighborhood effects from the moving to opportunity experiment? *American Journal of Sociology*, 114(1):144–188, 2008.

[28] Lisa A Gennetian, Lisa Sanbonmatsu, Lawrence F Katz, Jeffrey R Kling, Matthew Sciandra, Jens Ludwig, Greg J Duncan, and Ronald C Kessler. The long-term effects of moving to opportunity on youth outcomes. *Cityscape*, pages 137–167, 2012.

[29] Jens Ludwig, Greg J Duncan, Lisa A Gennetian, Lawrence F Katz, Ronald C Kessler, Jeffrey R Kling, and Lisa Sanbonmatsu. Long-term neighborhood effects on low-income families: Evidence from moving to opportunity. *American Economic Review*, 103(3):226–231, 2013.

[30] Raj Chetty, Nathaniel Hendren, and Lawrence F Katz. The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *American Economic Review*, 106(4):855–902, 2016.

[31] Moritz Hardt and Michael P Kim. Is your model predicting the past? In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, pages 1–8, 2023.

[32] Thomas M Shapiro. *The hidden cost of being African American: How wealth perpetuates inequality*. Oxford University Press, 2004.

[33] Ellora Derenoncourt. Can you move to opportunity? Evidence from the great migration. *American Economic Review*, 112(2):369–408, 2022.

[34] Ayan Mukhopadhyay and Yevgeniy Vorobeychik. Prioritized allocation of emergency responders based on a continuous-time incident prediction model. In *International Conference on Autonomous Agents and MultiAgent Systems*, 2017.

[35] Bryan Wilder, Bistra Dilkina, and Milind Tambe. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1658–1665, 2019.

[36] Adam N Elmachtoub, Jason Cheuk Nam Liang, and Ryan McNellis. Decision trees for decision-making under the predict-then-optimize framework. In *International Conference on Machine Learning*, pages 2858–2867. PMLR, 2020.

[37] Adam N Elmachtoub and Paul Grigas. Smart "predict, then optimize". *Management Science*, 68(1):9–26, 2022.

[38] Shipra Agrawal and Nikhil Devanur. Linear contextual bandits with knapsacks. *Advances in Neural Information Processing Systems*, 29, 2016.

[39] Aleksandrs Slivkins, Karthik Abinav Sankararaman, and Dylan J Foster. Contextual bandits with packing and covering constraints: A modular lagrangian approach via regression. In *Proceedings of the 36th Annual Conference on Learning Theory*, pages 4633–4656, 2023.

A Notational conventions

The variables related to an individual i are indexed with a subscript i. The variables related to time t are indexed with a superscript t. The variables transformed with $\tilde{p}(\cdot)$ are denoted with a tilde.

Symbol	Notion
t	Time step which takes a value from 1 to T
T	Horizon
p_i	Failure probability of individual i
o_i^t	Binary observation from an active individual at time t
ϵ	Noise on observation model of Eq. (1)
$\tilde{p}$	$\tilde{p}(p) \coloneqq (1-\epsilon)f(p) + (1-f(p))\epsilon$
y_i^t	Number of positive observations from an active individual i up to time t: $y_i^t \coloneqq \sum_{t' \in [t]} o_i^{t'}$
$\mathcal{A}^t$	Set of active individuals at time t
$\mathcal{A}_k^t$	Set of active individuals at t with $y^t = k$: $\mathcal{A}_k^t \coloneqq \{i \mid y_i^t = k\}$
$\mathcal{I}^t$	Set of individuals treated at t
$\bar{\mathcal{A}}_k^t$	$\mathcal{A}_k^t$ excluding those who will be treated at t: $\bar{\mathcal{A}}_k^t \coloneqq \mathcal{A}_k^t \setminus \mathcal{I}^t$
N	Number of initial individuals at time $t = 1$
N^t	Number of individuals who made it to time t: $N^t \coloneqq \lvert\mathcal{A}^t\rvert$
n^t	Proportion of initial individuals who made it to time t: $n^t \coloneqq N^t/N$
n_k^t	$n_k^t \coloneqq \lvert\mathcal{A}_k^t\rvert/N$
$\bar{n}_k^t$	$\bar{n}_k^t \coloneqq \lvert\bar{\mathcal{A}}_k^t\rvert/N$
$\mathcal{P}^t(\cdot)$	Posterior over p for an active individual at t
$\mathcal{P}^t(p \mid y^t = k)$	Posterior over p given an individual has made it to t and $y^t = k$
$\Pr^t(\cdot)$	Probability measure over active individuals at time t
$\Pr_{1,2}^t(\cdot)$	Probability measure over two independent active individuals 1 and 2 at time t
$\Pr^t(y^t = k)$	Probability that an active individual at t has $y^t = k$
$\Pr^t(y^t = k \mid p)$	Likelihood that an active individual with failure probability p shows $y^t = k$
$\mathbb{E}^t[\cdot]$	Expectation over active individuals at time t
$\mathbb{E}_{1,2}^t[\cdot]$	Expectation over two independent active individuals 1 and 2 at time t
μ^t	Mean of failure probability at time t: $\mu^t \coloneqq \mathbb{E}^t[p]$
$\mathrm{Var}^t[\cdot]$	Variance under $\Pr^t(\cdot)$
$\tilde{\sigma}_i^2$	$\tilde{\sigma}_i^2 \coloneqq \tilde{p}_i \cdot (1 - \tilde{p}_i)$
$\tilde{\sigma}_{12}^2$	$\tilde{\sigma}_{12}^2 \coloneqq \tilde{\sigma}_1^2 + \tilde{\sigma}_2^2$
$G(\cdot)$	Cumulative distribution function (CDF) of the standard normal distribution
$g(\cdot)$	Probability density function (PDF) of the standard normal distribution
$\tilde{b}^t$	See Eq. (9)

Table 1: Glossary

B Related work

Algorithmic predictions are increasingly employed to identify individuals who are most in need of limited resources. In many of these applications, the timing of prediction and allocation is of the utmost importance. Examples of such applications include directing assistance to tenants at risk of eviction based on their predicted risk, which is informed by their past records [4], or prioritizing homelessness assistance while considering potential population dynamics, such as housing stability [6] and the likelihood of re-entry into homelessness [3], or making ICU discharge decisions based on readmission probability [11], or improve targeting of humanitarian aids [17]. Another example that aligns with our model is the use of early warning systems to identify students at risk of dropping out, allowing for timely interventions to support their academic success and retention within the targeted educational division [13, 14, 5]. For a discussion on the role of machine learning in clinical medicine in particular, refer to Obermeyer and Emanuel [18].

Our work is closely related to subsidy allocation in the presence of income shocks, as studied by Abebe et al. [7]. Their model captures a more general dynamic where individuals fail after experiencing potentially multiple shocks, and their reserve to resist failure increases over time. However, it assumes a full information setting. In contrast, our focus is on understanding the accumulation of information through time and the risks associated with waiting for more information.

Welfare maximizing assignment of treatments under a budget constraint is well-studied in economics. However, it has often focused on estimating a static heterogeneous treatment effect of a fixed population with finite samples neglecting the dynamic aspect of the problem [19, 20]. A large body of this literature is concerned with the strategies to estimate the treatment effect from observational data [21]. In contrast, our simplified model of observations hides the complexity in learning from observational data and allows us to directly study the often overlooked aspect of timing in making predictions and allocations.

Historically, policy planning has relied on aggregate data; however, the promise of improved resource allocation, reduced costs, and more preventative interventions has led to the widespread adoption of algorithmic systems on an individual basis in governments [22, 23]. Our work contributes to this discussion, as our insights have direct implications for policy planning in evolving social contexts. While causal inference can inform policy-making, it is not always necessary, as discussed by Kleinberg et al. [2]. Our framework falls under the category of prediction policy problems where accurate predictions and ranking of individuals are sufficient for effective allocation policies. Related to this topic, Wang et al. [24] raise concerns about the legitimacy of decision-making based on predictive optimization.

The debate surrounding risk assessment tools has largely centered around their inherently predictive nature. However, as emphasized by Barabas et al. [25] in the context of the criminal justice system, the focus of risk assessment should be on guiding interventions rather than merely making predictions. Our research aligns with this perspective by studying prediction not as an isolated task but as an integral part of the resource allocation process.

In our framework, we assume observation noise and the number of observations are fixed, so the planner's estimate can only improve with further observations over time. However, in a broader framework, observation frequency and accuracy might change with further investment. Perdomo [15] asks what the relative value of investing in improving the predictor is compared to expanding access to resources.

The Moving to Opportunity (MTO) experiment, sponsored by the U.S. Department of Housing and Urban Development, exemplifies an early intervention aimed at improving life outcomes by providing low-income families with children living in disadvantaged urban public housing the opportunity to relocate to less distressed private-market housing communities[26–30]. The mixed findings of the MTO experiment across different age groups and the contrast between interim and long-term analyses highlight the crucial role that timing and the considered time horizon play in determining the intervention's effect. The differential impact of the MTO experiment by gender reveals a heterogeneity of effect not captured by our model. Although for some age groups, this intervention can be considered very early, the positive outcomes on certain aspects demonstrate the long-lasting influence that early interventions and environmental factors can have over time. In our model, we consider the extreme case of this when individuals subject to intervention are no longer vulnerable at any future time point. Hardt and Kim [31] discuss how these long-lasting effects inform future predictions.

Another discussion relevant to our setting is the significant role an individual's initial conditions play compared to what they can earn over time. These initial conditions may include wealth, family members, environment, or more abstractly, the probability of failure, as in our model. Shapiro [32] argues that these initial differences, rather than wage disparities, are the primary drivers of persistent inequality in the United States. Consistently, Derenoncourt [33] show that while moving to areas with better economic opportunities theoretically provided improved prospects, local responses counteracted many of the potential benefits. Such observations are compatible with our simplified model of population dynamics that abstracts the complexity of an individual's initial conditions into a fixed probability of failure, allowing us to isolate the effect of timing on prediction and allocation. While this abstraction helps us focus on specific aspects, we acknowledge that it does not capture the full range of dynamics in the real world.

Our discussion is related to decision-focused learning [34–37] in the Operations Research community, where predictions are informed by their downstream applications. In our work, we employ a simplified observation model that allows us to consistently obtain a posterior distribution over hidden variables. This approach circumvents the challenges that could arise from inaccurate or biased predictions and allows us to shift our focus on the prediction dynamic in the problem. There is also a direct connection between our Algorithm 1 and decision-focused learning. If we consider prediction as the ranking of individuals at *all* time points, then this algorithm effectively identifies the optimal prediction tailored for the subsequent allocation step.

Our over-time allocation is also related to multi-armed bandit problems with resource constraints in addition to reward (or utility) generation [38, 39]. Unlike the standard bandit problems, in our problem observations are available from all individuals and not only those treated. The exploration/exploitation tradeoff then lies in waiting for further information or treating those already estimated to be vulnerable.

C Supplementary algorithms

Algorithm 2 Simulate trajectory

inputs:
  t_0 : starting time of simulation
  $\{N_k^{t_0}\}_{k=0}^{t_0}$: number of individuals with $y^{t_0} = k$ for $k = 0, \dots, t_0$
  $q : [T] \to \{0, 1, \dots, T\}$: a valid non-decreasing sequence according to Theorem 5.1
output:
  $\{N_{q(t)}^t\}_{t=t_0}^T$: number of individuals reaching $y^t = q(t)$ for $t > t_0$
function SIMULATETRAJ(t_0, $\{N_k^{t_0}\}_{k=0}^{t_0}$, $q(\cdot)$)
  for $t = t_0 + 1$ to T **do**
    for $k = 0$ to t **do**
      Find the remaining number of individuals from the previous step:
$$\overline{N}_k^{t-1} \leftarrow N_k^{t-1} \cdot \mathbb{1}\{q(t-1) > k\}, \quad \overline{N}_{k-1}^{t-1} \leftarrow N_{k-1}^{t-1} \cdot \mathbb{1}\{q(t-1) > k-1\}$$
      Backup formula:
$$N_k^t \leftarrow \overline{N}_k^{t-1}\, \mathbb{E}_{p\sim\mathcal{P}^{t-1}(\cdot|y^{t-1}=k)}\big[(1-p)(1-\tilde{p})\big] + \overline{N}_{k-1}^{t-1}\, \mathbb{E}_{p\sim\mathcal{P}^{t-1}(\cdot|y^{t-1}=k-1)}\big[(1-p)\,\tilde{p}\big]$$
    end for
  end for
  return $\{N_{q(t)}^t\}_{t=t_0}^T$
end function

D Supplementary figures

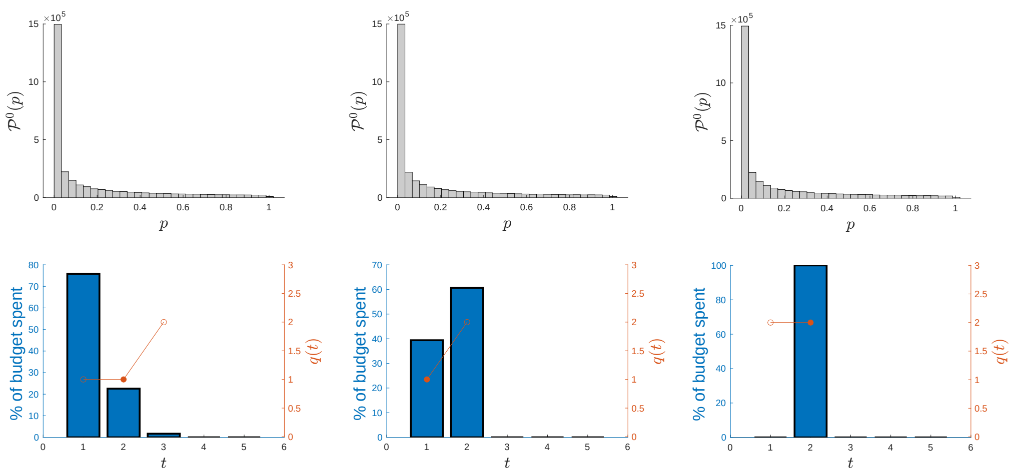

(a) $B/N = 10\%$ (b) $B/N = 2\%$ (c) $B/N = 1\%$

Figure 2: Optimal over-time allocation for three sizes of the budget and a fixed prior. The orange curve depicts the optimal $q(\cdot)$ and the filled circle corresponds to $t = \hat{t}$

E Supplementary statements

E.1 General statements

Lemma E.1 (Bayes optimal ranking)**.** *Consider a statistical model $P = \{p_\theta : \theta \in \Theta = [a, b]\}$ that induces a family of continuous probability distributions over a sample space $\mathcal{X}$. Assume P has a univariant sufficient statistics $T : \mathcal{X} \to \mathbb{R}$. Consider samples drawn independently from two probability distributions with distinct parameters: $X_1 \sim p_{\theta_1}, X_2 \sim p_{\theta_2}$. For a ranking function $\delta : \mathcal{X} \times \mathcal{X} \to \{-1, 1\}$, define the ranking loss as*

$$\operatorname{loss}((\theta_1, \theta_2); \delta(x_1, x_2)) \coloneqq \mathbb{1}\{\delta(x_1, x_2)(\theta_2 - \theta_1) < 0\}\,.$$

Consider Θ_1 and Θ_2 independently drawn from a prior $\mathcal{P}$ over Θ. If for any $\theta_2 \geq \theta_1$,

$$T(x_2) \geq T(x_1) \iff p_{\theta_1}(x_1)\, p_{\theta_2}(x_2) \geq p_{\theta_1}(x_2)\, p_{\theta_2}(x_1)\,, \quad \textit{a.e.}\,, \tag{13}$$

then for any choice of $\mathcal{P}$ that has no point mass, the Bayes optimal ranking rule is $\delta^*(x_1, x_2) = \chi\{T(x_2) \geq T(x_1)\}$.

Proof. For Θ_1 and Θ_2 independently drawn from $\mathcal{P}$, the Bayes risk of ranking is

$$R(\mathcal{P}^{\otimes 2}; \delta) = \mathbb{E}_{\substack{\Theta_1 \sim \mathcal{P} \\ \Theta_2 \sim \mathcal{P}}}\Big[\mathbb{E}_{\substack{X_1 \sim p_{\Theta_1} \\ X_2 \sim p_{\Theta_2}}}\left[\operatorname{loss}((\Theta_1, \Theta_2); \delta(X_1, X_2))\right]\Big]\,.$$

The independence also allows us to decompose the posterior over Θ_1 and Θ_2 given $X_1 = x_1$ and $X_2 = x_2$ as $\mathcal{P}(\Theta_1 \mid x_1)\,\mathcal{P}(\Theta_2 \mid x_2)$. It is well-known that the minimizer of the Bayes risk is

$$\delta^*(x_1, x_2) = \operatorname*{arg\,min}_{\delta(\cdot,\cdot)} R(\mathcal{P}^{\otimes 2}; \delta) \in \operatorname*{arg\,min}_{\delta \in \{-1,1\}} \mathbb{E}_{\substack{\Theta_1 \sim \mathcal{P}(\cdot|x_1) \\ \Theta_2 \sim \mathcal{P}(\cdot|x_2)}}\left[\operatorname{loss}((\Theta_1, \Theta_2); \delta) \mid X_1 = x_1, X_2 = x_2\right].$$

Plugging the ranking loss into this, we can further simplify the conditional expectation and obtain

$$\begin{aligned}
\delta^*(x_1, x_2) &\in \operatorname*{arg\,min}_{\delta \in \{-1,1\}} \mathbb{E}_{\substack{\Theta_1 \sim \mathcal{P}(\cdot|x_1) \\ \Theta_2 \sim \mathcal{P}(\cdot|x_2)}}\left[\operatorname{loss}((\Theta_1, \Theta_2); \delta) \mid x_1, x_2\right] \\
&= \operatorname*{arg\,min}_{\delta \in \{-1,1\}} \frac{1+\delta}{2}\Pr(\Theta_1 > \Theta_2 \mid x_1, x_2) + \frac{1-\delta}{2}\Pr(\Theta_1 \leq \Theta_2 \mid x_1, x_2) \\
&= \operatorname*{arg\,min}_{\delta \in \{-1,1\}} \delta\left[\Pr(\Theta_1 > \Theta_2 \mid x_1, x_2) - \Pr(\Theta_1 \leq \Theta_2 \mid x_1, x_2)\right] \\
&= \operatorname{sign}\big(\Pr(\Theta_1 \leq \Theta_2 \mid x_1, x_2) - \Pr(\Theta_1 > \Theta_2 \mid x_1, x_2)\big)\,.
\end{aligned}$$

Now, using a change of variable trick and the Bayes' rule, we have

$$\begin{aligned}
&\Pr(\Theta_1 \leq \Theta_2 \mid x_1, x_2) - \Pr(\Theta_1 > \Theta_2 \mid x_1, x_2) \\
&\quad = \int_a^b \int_a^{\theta_2} \mathcal{P}(\theta_1 \mid x_1)\,\mathcal{P}(\theta_2 \mid x_2)\,\mathrm{d}\theta_1\,\mathrm{d}\theta_2 - \int_a^b \int_a^{\theta_1} \mathcal{P}(\theta_1 \mid x_1)\,\mathcal{P}(\theta_2 \mid x_2)\,\mathrm{d}\theta_2\,\mathrm{d}\theta_1 \\
&\quad = \int_a^b \int_a^{\theta_2} \left[\mathcal{P}(\theta_1 \mid x_1)\,\mathcal{P}(\theta_2 \mid x_2) - \mathcal{P}(\theta_2 \mid x_1)\,\mathcal{P}(\theta_1 \mid x_2)\right] \mathrm{d}\theta_1\,\mathrm{d}\theta_2 \\
&\quad = \int_a^b \int_a^{\theta_2} \frac{\mathcal{P}(\theta_1)\,\mathcal{P}(\theta_2)}{Z(x_1, x_2)}\left[p_{\theta_1}(x_1)\,p_{\theta_2}(x_2) - p_{\theta_1}(x_2)\,p_{\theta_2}(x_1)\right] \mathrm{d}\theta_1\,\mathrm{d}\theta_2\,,
\end{aligned}$$

where $Z(x_1, x_2)$ is the partition function. The integral bound enforces $\theta_2 \geq \theta_1$. Then if the condition of Eq. (13) holds, we can conclude

$$\operatorname{sign}\big(\Pr(\Theta_1 \leq \Theta_2 \mid x_1, x_2) - \Pr(\Theta_1 > \Theta_2 \mid x_1, x_2)\big) = \operatorname{sign}(T(x_2) - T(x_1))\,.$$

□

Lemma E.2. *If*

1. *for every t and $k \leq t$, $\frac{\Pr^t(y^t = k+1 \mid p)}{\Pr^t(y^t = k \mid p)}$ is a non-decreasing continuous function of p, and*
2. *for every t, the utility $u^t(p)$ is a non-decreasing function of p,*

then $\mathbb{E}_{p\sim\mathcal{P}^t(\cdot|y^t=k+1)}[u^t(p)] \geq \mathbb{E}_{p\sim\mathcal{P}^t(\cdot|y^t=k)}[u^t(p)]$.

Proof. Using the Bayes' rule of $\mathcal{P}^t(p \mid y^t = k) = \mathcal{P}^t(p)\,\mathrm{Pr}^t(y^t = k \mid p)/\mathrm{Pr}^t(y^t = k)$, we can write $\mathcal{P}^t(p \mid y^t = k+1)$ as an update to $\mathcal{P}^t(p \mid y^t = k)$:

$$\mathcal{P}^t(p \mid y^t = k+1) = \mathcal{P}^t(p \mid y^t = k)\frac{\mathrm{Pr}^t(y^t = k+1 \mid p)}{\mathrm{Pr}^t(y^t = k \mid p)}\frac{\mathrm{Pr}^t(y^t = k)}{\mathrm{Pr}^t(y^t = k+1)}\,.$$

Define $\sigma(p) \coloneqq \mathcal{P}^t(p \mid y^t = k+1)/\mathcal{P}^t(p \mid y^t = k)$. The above update and a non-decreasing $\frac{\mathrm{Pr}^t(y^t=k+1|p)}{\mathrm{Pr}^t(y^t=k|p)}$ imply $\sigma(p)$ is a non-decreasing continuous function. Since $\int_0^1 \mathcal{P}^t(p \mid y^t = k)\,\mathrm{d}p = \int_0^1 \mathcal{P}^t(p \mid y^t = k)\,\sigma(p)\,\mathrm{d}p = 1$, and σ is continuous, there should exist a critical value p^* such that $\sigma(p) \geq 1$ for $p \geq p^*$, and $\sigma(p) \leq 1$ for $p < p^*$. Using this critical value to decompose the expectations and the fact that $u^t(\cdot)$ is increasing, we obtain

$$\begin{aligned}
&\mathbb{E}_{p\sim\mathcal{P}^t(\cdot|y^t=k+1)}[u^t(p)] - \mathbb{E}_{p\sim\mathcal{P}^t(\cdot|y^t=k)}[u^t(p)] \\
&= \int_0^1 u^t(p)\,\mathcal{P}^t(p \mid y^t = k)\,(\sigma(p) - 1)\,\mathrm{d}p \\
&= \int_{p^*}^1 u^t(p)\,\mathcal{P}^t(p \mid y^t = k)\,(\sigma(p) - 1)\,\mathrm{d}p - \int_0^{p^*} u^t(p)\,\mathcal{P}^t(p \mid y^t = k)\,(1 - \sigma(p))\,\mathrm{d}p \\
&\geq u^t(p^*)\int_{p^*}^1 \mathcal{P}^t(p \mid y^t = k)\,(\sigma(p) - 1)\,\mathrm{d}p - u^t(p^*)\int_0^{p^*} \mathcal{P}^t(p \mid y^t = k)\,(1 - \sigma(p))\,\mathrm{d}p = 0\,.
\end{aligned}$$

□

Lemma E.3. *Consider two non-increasing functions* $a : \mathbb{R} \to [0, 1]$ *and* $b : \mathbb{R} \to [0, 1]$. *If* $\int_{-\infty}^{\infty} b(x)\,\mathrm{d}x$ *is finite and non-zero, the following inequality always holds:*

$$\Big(\int_{-\infty}^{\infty} b(x)^2\,\mathrm{d}x\Big) \cdot \Big(\int_{-\infty}^{\infty} a(x)\,b(x)\,\mathrm{d}x\Big) \geq \Big(\int_{-\infty}^{\infty} a(x)^2\,b(x)^2\,\mathrm{d}x\Big) \cdot \Big(\int_{-\infty}^{\infty} b(x)\,\mathrm{d}x\Big)\,. \tag{14}$$

Proof. Define the difference between the left-hand side and the right-hand side of the inequality given in Eq. (14) as Δ. For simplicity, consider integrals as a discrete sum with a step size of δ. Increasing the value of $a(x')$ would change the value of Δ by

$$\frac{1}{\delta}\,\frac{\mathrm{d}\Delta}{\mathrm{d}a(x')} = b(x') \cdot \int_{-\infty}^{\infty} b(x)^2\,\mathrm{d}x - 2a(x')\,b(x')^2 \cdot \int_{-\infty}^{\infty} b(x)\,\mathrm{d}x\,.$$

This implies that for any x' such that $b(x') > 0$, increasing $a(x')$ will decrease Δ if and only if

$$a(x')\,b(x') > \frac{1}{2}\,\frac{\int_{-\infty}^{\infty} b(x)^2\,\mathrm{d}x}{\int_{-\infty}^{\infty} b(x)\,\mathrm{d}x}\,.$$

Since both a and b are non-increasing non-negative functions, their multiplication is also a non-increasing non-negative function. Therefore, increasing $a(x')$ will decrease Δ if and only if x' is larger than a critical value x^*. The non-increasing constraint on a then implies that for a fixed b, the minimum value of Δ corresponds to a constant function $a(x) = a_0$. In this case,

$$\Delta = \big(a_0 - a_0^2\big) \cdot \Big(\int_{-\infty}^{\infty} b(x)^2\,\mathrm{d}x\Big) \cdot \Big(\int_{-\infty}^{\infty} b(x)\,\mathrm{d}x\Big)$$

For $a_0 \in [0, 1]$, the above equation is always greater than or equal to zero, which completes the proof. □

E.2 Statements specific to the observation model in Eq. (1)

Proposition E.4. *Under the observation model in Eq.* (1), *for any prior* $\mathcal{P}$ *over* p *that has no point mass, ranking individuals based on their* y^t *is Bayes optimal.*

Proof. At any time t, define the statistical model $P = \{p_\theta = (\mathrm{Ber}(f(\theta)) \oplus \mathrm{Ber}(\epsilon))^{\otimes t} : \theta \in [0,1]\}$ where we can think of the model parameter θ as the failure probability p. All the observations from individual i until t can be interpreted as a sample from a model in P: $X = [o^1, \ldots, o^t] \sim p_\theta$.

Recall that $\mathrm{Ber}(f(\theta)) \oplus \mathrm{Ber}(\epsilon)$ is itself a Bernoulli random variable with effective parameter $\tilde{\theta} = \tilde{p}(\theta)$. So, it is straightforward to see $y^t = \sum_{t' \in [t]} o^t$ is a sufficient statistic for P and $p_\theta(x) = \tilde{\theta}^{y^t}(1-\tilde{\theta})^{t-y^t}$. The strictly increasing property of $\tilde{p}$ also implies that $\theta_2 \geq \theta_1 \iff \tilde{\theta}_2 \geq \tilde{\theta}_1$.

For $\theta_2 \geq \theta_1$, plugging p_θ into the condition of Eq. (13) gives

$$p_{\theta_1}(x_1)\, p_{\theta_2}(x_2) \geq p_{\theta_1}(x_2)\, p_{\theta_2}(x_1) \iff \Big(\frac{\tilde{\theta}_2}{\tilde{\theta}_1} \frac{1-\tilde{\theta}_1}{1-\tilde{\theta}_2}\Big)^{y_2^t - y_1^t} \geq 1\,, \quad \text{a.e.}$$

Since for $1 > \tilde{\theta}_2 \geq \tilde{\theta}_1 > 0$, we have $\frac{\tilde{\theta}_2}{\tilde{\theta}_1}\frac{1-\tilde{\theta}_1}{1-\tilde{\theta}_2} \geq 1$, we can conclude

$$p_{\theta_1}(x_1)\, p_{\theta_2}(x_2) \geq p_{\theta_1}(x_2)\, p_{\theta_2}(x_1) \iff y_2^t \geq y_1^t\,, \quad \text{a.e.}$$

Therefore, P meets the sufficient condition given in Eq. (13) of Lemma E.1, and ranking based on its sufficient statistic y^t is Bayes optimal. □

Proposition E.5. *Under the observation model in Eq. (1), for any utility function that is non-decreasing in p, we have*

$$\mathbb{E}_{p\sim\mathcal{P}^t(\cdot|y^t=k+1)}[u^t(p)] \geq \mathbb{E}_{p\sim\mathcal{P}^t(\cdot|y^t=k)}[u^t(p)]\,.$$

Proof. Recall that $\mathrm{Ber}(f(p)) \oplus \mathrm{Ber}(\epsilon)$ is itself a Bernoulli random variable with effective parameter $\tilde{p}$. Under this parameterization, the likelihood $\Pr^t(y^t = k \mid p)$ has a closed-form of $\binom{t}{k}\tilde{p}^k(1-\tilde{p})^{t-k}$. Plugging this into the likelihood ratio, we obtain

$$\frac{\Pr^t(y^t = k+1 \mid p)}{\Pr^t(y^t = k \mid p)} = \frac{\tilde{p}}{1-\tilde{p}}\Big(\frac{t-k}{k+1}\Big)\,.$$

This is a non-decreasing function of $\tilde{p}$ and, consequently, of p, for every $k \leq t$. Using Lemma E.2, this non-decreasing continuous likelihood ratio along with the non-decreasing utility function is sufficient to complete the proof . □

Proposition E.6. *Under the observation model of Eq. (1), suppose the probability density function of $\tilde{p}$ at time t is non-increasing and $\Pr^t(\tilde{p} > 0) > 0$. Denote the Gini index of $\tilde{p}$ at time t by $\widetilde{G}^t$. In the limit of many individuals, we have*

$$\widetilde{G}^{t+1} \geq \widetilde{G}^t\,.$$

Proof. Denote the cumulative distribution function of p and $\tilde{p}$ at time t by $\mathcal{F}^t$ and $\widetilde{\mathcal{F}}^t$, respectively. Using the update rule of $\mathcal{P}^t$ from Eq. (3), we can write an update rule for $\mathcal{F}^t$:

$$1 - \mathcal{F}^{t+1}(p') = \int_{p'}^1 \Big(\frac{1-p}{1-\mu^t}\Big)\,\mathcal{P}^t(p)\,\mathrm{d}p = (1-\mathcal{F}^t(p')) \cdot \frac{1-\mathbb{E}^t[p \mid p > p']}{1-\mu^t}\,.$$

Define $a(p') := \frac{1-\mathbb{E}^t[p|p>p']}{1-\mu^t}$. Note that $a(p')$ is a non-increasing function, implying that $(1-\gamma F^{(p')})$ decreases at a higher rate for larger values of p'.

Since $\tilde{p}(\cdot)$ is an increasing function, its inverse is well-defined, and $\tilde{a} := a \circ \tilde{p}^{-1}$ is also a non-increasing function. It also allows us to write $\widetilde{\mathcal{F}}^t(\tilde{p}') = \mathcal{F}^t(\tilde{p}^{-1}(\tilde{p}'))$. Then by plugging $p' = \tilde{p}^{-1}(\tilde{p}')$ into the update rule of $\mathcal{F}^t$, we obtain an update for $\widetilde{\mathcal{F}}^t$:

$$1 - \widetilde{\mathcal{F}}^{t+1}(\tilde{p}') = (1 - \widetilde{\mathcal{F}}^t(\tilde{p}')) \cdot \tilde{a}(\tilde{p}')\,.$$

The Gini coefficient of $\tilde{p}$ at t can be calculated by

$$\widetilde{G}^t = 1 - \frac{\int_0^1 (1-\widetilde{\mathcal{F}}^t(\tilde{p}))^2\,\mathrm{d}\tilde{p}}{\int_0^1 (1-\widetilde{\mathcal{F}}^t(\tilde{p}))\,\mathrm{d}\tilde{p}}\,.$$

Note that $\Pr^t(\tilde{p} > 0) > 0$ implies $\widetilde{G}^t$ is well-defined. Moreover, a non-increasing density function for $\tilde{p}$ implies $\mu^t < 1$, so $\tilde{a}$ and consequently $\widetilde{G}^{t+1}$ are well-defined. Using the update rule of $\widetilde{\mathcal{F}}^t$, the change of $\widetilde{G}^t$ over time is

$$\widetilde{G}^{t+1} - \widetilde{G}^t = \frac{\int_0^1 (1-\widetilde{\mathcal{F}}^t(\tilde{p})) \cdot \tilde{a}(\tilde{p}) \,\mathrm{d}\tilde{p} \cdot \int_0^1 (1-\widetilde{\mathcal{F}}^t(\tilde{p}))^2 \,\mathrm{d}\tilde{p} - \int_0^1 (1-\widetilde{\mathcal{F}}^t(\tilde{p})) \,\mathrm{d}\tilde{p} \cdot \int_0^1 (1-\widetilde{\mathcal{F}}^t(\tilde{p}))^2 \cdot \tilde{a}(\tilde{p})^2 \,\mathrm{d}\tilde{p}}{\int_0^1 (1-\widetilde{\mathcal{F}}^t(\tilde{p})) \,\mathrm{d}\tilde{p} \cdot \int_0^1 (1-\widetilde{\mathcal{F}}^t(\tilde{p})) \cdot \tilde{a}(\tilde{p}) \,\mathrm{d}\tilde{p}} .$$

Since $\tilde{a}(\tilde{p})$ and $(1-\widetilde{\mathcal{F}}^t(\tilde{p}))$ are both non-increasing functions of $\tilde{p}$, Lemma E.3 implies $\widetilde{G}^{t+1} \geq \widetilde{G}^t$.

□

F Missing proofs

Proof of Proposition 3.1. Using the step function approximation, the exponential multiplier in Eq. (4) is only non-zero if

$$\frac{|\tilde{p}_2 - \tilde{p}_1|}{\tilde{\sigma}_{12}} \leq \sqrt{\frac{2\ln(1/\alpha)}{t}} .$$

This implies the following lower bound on ΔR^t:

$$\Delta R^t \geq \sqrt{\frac{t}{2\ln(1/\alpha)}} \cdot \mathbb{E}^t_{1,2}\Big[\frac{(1-p_1)(1-p_2)}{(1-\mu^t)^2} - 1 \mid \frac{|\tilde{p}_2 - \tilde{p}_1|}{\tilde{\sigma}_{12}} \leq \sqrt{\frac{2\ln(1/\alpha)}{t}}\Big] - \sqrt{\frac{\ln(1/\alpha)}{4\pi t}} .$$

Suppose f^{-1} is L^{-1}-Lipschitz continuous. Using this and $\tilde{\sigma}_{12}^2 \leq 1/2$, we can further bound ΔR^t by

$$\begin{aligned}
\Delta R^t &\geq \sqrt{\frac{t}{2\ln(1/\alpha)}} \cdot \mathbb{E}^t\Big[\frac{(1-p)(1-p-\frac{L^{-1}}{1-2\epsilon}\sqrt{\frac{2\ln(1/\alpha)}{t}})}{(1-\mu^t)^2} - 1\Big] - \sqrt{\frac{\ln(1/\alpha)}{4\pi t}} \\
&= \sqrt{\frac{t}{2\ln(1/\alpha)}} \cdot \frac{\mathrm{Var}^t[p] - (1-\mu^t)\frac{L^{-1}}{1-2\epsilon}\sqrt{\frac{2\ln(1/\alpha)}{t}}}{(1-\mu^t)^2} - \sqrt{\frac{\ln(1/\alpha)}{4\pi t}} .
\end{aligned}$$

□

Proof of Proposition 4.3. Denote the cumulative distribution function of $\tilde{p}$ at time t by $\widetilde{\mathcal{F}}$ (we drop dependency on t for brevity). A non-increasing density function requires a concave $\widetilde{\mathcal{F}}$. The Gini coefficient of $\tilde{p}$ at t can be calculated by

$$\widetilde{G}^t = 1 - \frac{\int_0^1 (1-\widetilde{\mathcal{F}}(\tilde{p}))^2 \,\mathrm{d}\tilde{p}}{\int_0^1 (1-\widetilde{\mathcal{F}}(\tilde{p})) \,\mathrm{d}\tilde{p}} .$$

In order to upper bound $\tilde{b}^t$, we set an upper bound on ER^t and use this upper bound in the definition of $\tilde{b}^t$ (Eq. (9)). Consider the following upper bound on ER^t:

$$\mathrm{ER}^t(k) = N^t \Pr^t(y^t \geq k) \leq N^t\, \mathbb{E}^t\big[\frac{y^t}{k}\big] = N^t\, \frac{t}{k} \cdot \mathbb{E}^t[\tilde{p}] .$$

To further find a distribution-agnostic bound on $\mathrm{ER}^t(k)$, regardless of k, we solve

$$\begin{aligned}
\max_{\widetilde{\mathcal{F}}} \quad & \mathbb{E}^t[\tilde{p}] = \int_0^1 (1-\widetilde{\mathcal{F}}(\tilde{p})) \,\mathrm{d}\tilde{p} \\
\text{s.t.} \quad & c_G(\widetilde{\mathcal{F}}) \coloneqq 1 - \frac{\int_0^1 (1-\widetilde{\mathcal{F}}(\tilde{p}))^2 \,\mathrm{d}\tilde{p}}{\int_0^1 (1-\widetilde{\mathcal{F}}(\tilde{p})) \,\mathrm{d}\tilde{p}} \geq \widetilde{G}^t , \qquad \text{(Relaxed Gini index constraint)} \\
& \widetilde{\mathcal{F}} \text{ is concave} .
\end{aligned}$$

Next, we show that the optimal solution to this problem follows a specific form. Suppose $\widetilde{\mathcal{F}}^*$ is the maximizer of the above problem. Consider integrals in the objective and the definition of c_G as the limit of a discrete sum with step size $\delta \to 0$. If there exist $1 > \tilde{p}_2 > \tilde{p}_1 > 0$ such that

$$\widetilde{\mathcal{F}}^*(\tilde{p}_2) > (\widetilde{\mathcal{F}}^*(\tilde{p}_2 + \delta) + \widetilde{\mathcal{F}}^*(\tilde{p}_2 - \delta))/2\,,$$
$$\widetilde{\mathcal{F}}^*(\tilde{p}_1) > (\widetilde{\mathcal{F}}^*(\tilde{p}_1 + \delta) + \widetilde{\mathcal{F}}^*(\tilde{p}_1 - \delta))/2\,,$$

then define a new solution

$$\widetilde{\mathcal{F}}(\tilde{p}) := \begin{cases} \widetilde{\mathcal{F}}^*(\tilde{p}) - h\,, & \tilde{p} = \tilde{p}_2\,, \\ \widetilde{\mathcal{F}}^*(\tilde{p}) + h\,, & \tilde{p} = \tilde{p}_1\,, \\ \widetilde{\mathcal{F}}^*(\tilde{p}) & \text{o.w.} \end{cases}$$

where $h \to^+ 0$. This solution preserves concavity while maintaining a similar value for the objective. Moreover, it only improves the relaxed Gini index constraint:

$$\frac{\mathrm{d}c_G}{\mathrm{d}h} = 2\delta\, \frac{\widetilde{\mathcal{F}}^*(\tilde{p}_2) - \widetilde{\mathcal{F}}^*(\tilde{p}_1)}{\int_0^1 (1 - \widetilde{\mathcal{F}}^*(\tilde{p}))\,\mathrm{d}\tilde{p}} \geq 0\,.$$

Therefore, if $\widetilde{\mathcal{F}}^*$ was optimal, $\widetilde{\mathcal{F}}$ should also be optimal. By repetitively applying this operation, we get an optimal solution that follows

$$\widetilde{\mathcal{F}}^*(\tilde{p}) = \widetilde{\mathcal{F}}^*(0) + (1 - \widetilde{\mathcal{F}}^*(0)) \cdot \tilde{p}\,.$$

The objective value for the optimal solution is $\frac{1}{2}(1 - \widetilde{\mathcal{F}}^*(0))^2$. The relaxed Gini index constraint also requires

$$c_G(\widetilde{\mathcal{F}}^*) = 1 - \frac{2}{3}(1 - \widetilde{\mathcal{F}}^*(0)) \geq \widetilde{G}^t\,.$$

Hence, for $\widetilde{G}^t \geq 1/3$, the objective is bounded by $\frac{9}{8}(1 - \widetilde{G}^t)^2$. This gives the following upper bound on ER^t:

$$\mathrm{ER}^t(k) \leq \frac{9}{8}\, N^t \cdot \frac{t}{k} \cdot (1 - \widetilde{G}^t)^2\,.$$

So, for a budget of B we can conclude

$$\tilde{b}^t \leq \frac{1}{t} \cdot \min\big\{k \mid \frac{9}{8}\, N^t \cdot \frac{t}{k} \cdot (1 - \widetilde{G}^t)^2 \leq B\big\} = \frac{1}{t} \cdot \big\lceil \frac{9}{8}\, t \cdot \frac{N^t}{B} \cdot (1 - \widetilde{G}^t)^2 \big\rceil\,.$$

□

Proof of Lemma 4.4. Recall that the observation distribution in Eq. (1) is itself a Bernoulli distribution with an effective parameter $\tilde{p}$. So, for an individual with a failure probability of p, we have $y^t \sim \mathrm{Binomial}(t, \tilde{p})$. We approximate this distribution using the central limit theorem (CLT): $\frac{y^t}{t} \sim \mathcal{N}\big(\tilde{p}, \frac{\tilde{\sigma}^2}{t}\big)$, where $\tilde{\sigma}^2 := \tilde{p} \cdot (1 - \tilde{p})$. Using this approximation, Eq. (10) can be written as

$$\begin{aligned} U^t &= N^t\, \mathbb{E}^t[\mathbb{1}\{p \geq c^t\}] \cdot \mathbb{E}^t[\mathbb{1}\{y^t \geq t \cdot \tilde{b}^t\} \mid p \geq c^t] \\ &= N^t\, \mathbb{E}^t\big[\mathbb{1}\{p \geq c^t\} \cdot \mathbb{1}\{y^t \geq t \cdot \tilde{b}^t\}\big] \\ &\approx N^t\, \mathbb{E}^t\big[\mathbb{1}\{p \geq c^t\} \cdot G\big(\frac{\tilde{p} - \tilde{b}^t}{\tilde{\sigma}} \sqrt{t}\big)\big]\,. \end{aligned}$$

Note that both $\tilde{p}$ and $\tilde{\sigma}$ depend implicitly on p. We also revisit ER^t and $\tilde{b}^t$ using the CLT approximation:

$$\mathrm{ER}^t(k) \approx N^t\, \mathbb{E}^t\big[G\big(\frac{\tilde{p} - k/t}{\tilde{\sigma}} \sqrt{t}\big)\big]\,,$$
$$\tilde{b}^t \approx \frac{1}{t}\, (\mathrm{ER}^t)^{-1}(B)\,.$$

Next, we calculate $\frac{\mathrm{d}U^t}{\mathrm{d}t}$ and bound it from above.

Deriving $\frac{\mathrm{d}U^t}{\mathrm{d}t}$. Calculating $\frac{\mathrm{d}U^t}{\mathrm{d}t}$ requires taking the derivative with respect to t for N^t, $\mathbb{E}^t$, c^t, and $\tilde{b}^t$. In particular, taking the derivative of $\mathbb{E}^t$ requires the derivative of $\mathcal{P}^t$ with respect to t, where we use the approximation of Eq. (3):

$$\frac{\mathrm{d}\mathcal{P}^t}{\mathrm{d}t} \approx \big(\frac{\mu^t - p}{1-\mu^t}\big)\, \mathcal{P}^t(p)\,.$$

This dynamic is valid since the intervention takes place at a single time point, and no prior intervention has altered the distribution $\mathcal{P}^t$. A similar approximation gives $\frac{\mathrm{d}N^t}{\mathrm{d}t} \approx N^{t+1} - N^t = -N^t\,\mu^t$. To calculate $\frac{\mathrm{d}\tilde{b}^t}{\mathrm{d}t}$ we use the property

$$\begin{aligned}\frac{\mathrm{dER}^t(t\cdot\tilde{b}^t)}{\mathrm{d}t} = 0 = -\mu^t\,\mathrm{ER}^t(t\cdot\tilde{b}^t) \\ &+ N^t\,\mathbb{E}^t\Big[\big(\frac{\mu^t-p}{1-\mu^t}\big)\cdot G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\Big] \\ &+ \frac{1}{2\sqrt{t}}\,N^t\,\mathbb{E}^t\big[\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\big)\cdot g\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big] \\ &- \frac{\mathrm{d}\tilde{b}^t}{\mathrm{d}t}\,N^t\,\mathbb{E}^t\big[\frac{\sqrt{t}}{\tilde{\sigma}}\cdot g\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]\,,\end{aligned}$$

and solve for $\frac{\mathrm{d}\tilde{b}^t}{\mathrm{d}t}$, which we omit for brevity. The derivative of c^t is also straightforward to calculate from Eq. (6):

$$\frac{\mathrm{d}c^t}{\mathrm{d}t} = -\ln(1-u^*)\,\frac{1}{(T-t)^2}\,(1-u^*)^{1/(T-t)}\,.$$

With all the necessary elements now available, we can proceed to calculate $\frac{\mathrm{d}U^t}{\mathrm{d}t}$. Defining

$$\delta \coloneqq \frac{\mathbb{E}^t\big[\mathbb{1}\{p\geq c^t\}\cdot\frac{1}{\tilde{\sigma}}\,g\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]}{\mathbb{E}^t\big[\frac{1}{\tilde{\sigma}}\,g\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]}\,, \tag{15}$$

a straightforward calculation shows that we can decompose $\frac{\mathrm{d}U^t}{\mathrm{d}t}$ as

$$\frac{\mathrm{d}U^t}{\mathrm{d}t} = \big(\frac{\mathrm{d}U^t}{\mathrm{d}t}\big)^+ - \big(\frac{\mathrm{d}U^t}{\mathrm{d}t}\big)^-\,,$$

where

$$\begin{aligned}\frac{1}{N^t}\big(\frac{\mathrm{d}U^t}{\mathrm{d}t}\big)^- &\coloneqq \mathbb{E}^t\Big[\mathbb{1}\{p\geq c^t\}\cdot\big(\frac{p-(\mu^t)^2}{1-\mu^t}\big)\cdot G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\Big] - \delta\cdot\mathbb{E}^t\big[\big(\frac{p-(\mu^t)^2}{1-\mu^t}\big)\cdot G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big] \\ &\quad + \mathcal{P}^t(c^t)\cdot G\big(\frac{\tilde{c}^t-\tilde{b}^t}{\tilde{\sigma}(c^t)}\sqrt{t}\big)\cdot\frac{\mathrm{d}c^t}{\mathrm{d}t}\,, &(16)\\ \frac{1}{N^t}\big(\frac{\mathrm{d}U^t}{\mathrm{d}t}\big)^+ &\coloneqq \frac{1}{2\sqrt{t}}\,\mathbb{E}^t\Big[\mathbb{1}\{p\geq c^t\}\cdot\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\big)\cdot g\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\Big] - \frac{\delta}{2\sqrt{t}}\cdot\mathbb{E}^t\big[\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\big)\cdot g\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]\,. &(17)\end{aligned}$$

Next, we derive an upper bound for $\big(\frac{\mathrm{d}U^t}{\mathrm{d}t}\big)^+$ and a lower bound for $\big(\frac{\mathrm{d}U^t}{\mathrm{d}t}\big)^-$ in order to obtain an upper bound for $\frac{\mathrm{d}U^t}{\mathrm{d}t}$.

Lower bound $\big(\frac{\mathrm{d}U^t}{\mathrm{d}t}\big)^-$. Let us define the relative efficiency of budget allocation as U^t/B. Given that the targeting intervenes on individuals with higher expected utility, the relative efficiency is a non-increasing function of B. This results in a lower bound on δ:

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}B}\big(\frac{U^t}{B}\big) &= \frac{\mathrm{d}}{\mathrm{d}B}\Big(\frac{N^t}{B}\,\mathbb{E}^t\big[\mathbb{1}\{p \geq c^t\}\cdot G\big(\frac{\tilde{p}-(\mathrm{ER}^t)^{-1}(B)/t}{\tilde{\sigma}}\sqrt{t}\big)\big]\Big) \\
&= \frac{N^t}{B}\,\mathbb{E}^t\big[\mathbb{1}\{p \geq c^t\}\cdot\frac{1}{\tilde{\sigma}\sqrt{t}}\cdot g\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]\cdot\Big(N^t\,\mathbb{E}^t\big[\frac{1}{\tilde{\sigma}\sqrt{t}}\cdot g\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]\Big)^{-1} \\
&- \frac{N^t}{B^2}\,\mathbb{E}^t\big[\mathbb{1}\{p \geq c^t\}\cdot G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big] \\
&= \frac{N^t}{B}\Big(\delta - \frac{\mathbb{E}^t\big[\mathbb{1}\{p \geq c^t\}\cdot G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]}{\mathbb{E}^t\big[G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]}\Big) \leq 0 \\
&\iff \delta \leq \frac{\mathbb{E}^t\big[\mathbb{1}\{p \geq c^t\}\cdot G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]}{\mathbb{E}^t\big[G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]}\,.
\end{aligned}$$

Plugging this bound into Eq. (16), we obtain

$$\frac{1}{N^t}\big(\frac{\mathrm{d}U^t}{\mathrm{d}t}\big)^{-} \geq \mathbb{E}^t\Big[\mathbb{1}\{p \geq c^t\}\cdot\big(\frac{p-(\mu^t)^2}{1-\mu^t}\big)\cdot G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\Big] \tag{18}$$

$$- \mathbb{E}^t\big[\mathbb{1}\{p \geq c^t\}\cdot G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]\cdot\frac{\mathbb{E}^t\big[\big(\frac{p-(\mu^t)^2}{1-\mu^t}\big)\cdot G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]}{\mathbb{E}^t\big[G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]} \tag{19}$$

$$+ \mathcal{P}^t(c^t)\cdot G\big(\frac{\tilde{c}^t-\tilde{b}^t}{\tilde{\sigma}(c^t)}\sqrt{t}\big)\cdot\frac{\mathrm{d}c^t}{\mathrm{d}t}\,. \tag{20}$$

Next, we will further lower bound each of the three terms in the above bound.

- First of all, the $(\mu^t)^2$ terms cancel out. When $\tilde{p} \geq \tilde{c}^*$, we have $\tilde{p} \geq 1/2$ and $\tilde{p} \geq \tilde{b}^t$. These are sufficient to argue $G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)$ is non-decreasing in $\tilde{p}$ when $\tilde{p} \geq \tilde{c}^*$:

  $$\frac{\mathrm{d}}{\mathrm{d}\tilde{p}}G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big) = \sqrt{t}\,g\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\cdot\frac{\tilde{\sigma}-(\tilde{p}-\tilde{b}^t)\,\frac{1-2\tilde{p}}{2\tilde{\sigma}}}{\tilde{\sigma}^2} \geq 0\,.$$

  Now since $\tilde{p}$ is also an increasing function of p, one can verify that $G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)$ is also a non-decreasing function of p when $p \geq c^*$. This is particularly true for $p \geq c^t$ and $t \geq t^*$ as c^t is increasing in time. Chebyshev's sum inequality then allows us to bound the inner product of two non-decreasing functions and derive a lower bound on the first term:

  $$\Pr^t(p \geq c^t)\cdot\big(\frac{\mathbb{E}^t[p \mid p \geq c^t]-(\mu^t)^2}{1-\mu^t}\big)\cdot\mathbb{E}^t\big[G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big) \mid p \geq c^t\big]\,.$$

- In order to lower bound the second term (Eq. (19)), one can verify

  $$\frac{\mathbb{E}^t\big[p\cdot G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{\tau}\big)\big]}{\mathbb{E}^t\big[G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{\tau}\big)\big]}$$

  is non-decreasing in τ. This ratio intuitively represents the expected probability of failure of those found eligible when the eligibility cutoff is fixed and τ controls the amount of available information. As τ increases, allowing for more information, we expect the treatment allocation to become more efficient. Using this, we can write

  $$\frac{\mathbb{E}^t\big[p\cdot G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]}{\mathbb{E}^t\big[G\big(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\big)\big]} \leq \mathbb{E}^t[p \mid \tilde{p} \geq \tilde{b}^t]\,.$$

- Using the update rule of Eq. (3) and the fact that c^t is increasing, it is straightforward to conclude that if $\mathcal{P}^t(p)$ is non-increasing at $t = t^*$ for $p \geq c^*$, then $\mathcal{P}^t(p)$ is also non-increasing in p for $p \geq c^t$ and $t \geq t^*$. This allows us to lower bound $\mathcal{P}^t(c^t)$ in the third term (Eq. (20)) by $\Pr^t(p \geq c^t)$.

Putting these all together, we can further lower bound $\left(\frac{\mathrm{d}U^t}{\mathrm{d}t}\right)^-$ by

$$\frac{1}{N^t}\left(\frac{\mathrm{d}U^t}{\mathrm{d}t}\right)^- \geq \Pr^t(p \geq c^t) \cdot G\left(\frac{\tilde{c}^t - \tilde{b}^t}{\tilde{\sigma}(c^t)}\sqrt{t}\right) \cdot \left[\mathbb{E}^t[p \mid p \geq c^t] - \mathbb{E}^t[p \mid p \geq b^t] + \frac{\mathrm{d}c^t}{\mathrm{d}t}\right].$$

For a distribution $\mathcal{P}^t(p)$ non-increasing in p, the difference $\mathbb{E}^t[p \mid p \geq c^t] - \mathbb{E}^t[p \mid p \geq b^t]$ is bounded by $(c^t - b^t)/2$ from above. Since 1) we showed $G\left(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\right)$ is non-decreasing in $\tilde{p}$ when $\tilde{p} \geq \tilde{c}^*$, 2) c^t is increasing over time, and 3) $\tilde{b}^t$ is decreasing, we can further lower bound $G\left(\frac{\tilde{c}^t-\tilde{b}^t}{\tilde{\sigma}(c^t)}\sqrt{t}\right)$ by its value at t^*. This gives

$$\frac{1}{N^t}\left(\frac{\mathrm{d}U^t}{\mathrm{d}t}\right)^- \geq \Pr^t(p \geq c^t) \cdot G\left(\sqrt{t/t^*}\right) \cdot \left[\frac{c^t - b^t}{2} + \frac{\mathrm{d}c^t}{\mathrm{d}t}\right].$$

Upper bound $\left(\frac{\mathrm{d}U^t}{\mathrm{d}t}\right)^+$. Plugging δ from Eq. (15) into Eq. (17) and performing a straightforward calculation shows that the $\frac{\tilde{b}^t}{\tilde{\sigma}} \cdot g\left(\frac{\tilde{p}-\tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\right)$ terms will cancel out, so one can verify

$$\frac{1}{N^t}\left(\frac{\mathrm{d}U^t}{\mathrm{d}t}\right)^+ \leq \frac{1}{2\sqrt{t}}\Pr^t(p \geq c^t) \cdot \mathbb{E}^t\left[\frac{\tilde{p}}{\tilde{\sigma}} \cdot g\left(\frac{\tilde{p} - \tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\right) \mid \tilde{p} \geq \tilde{c}^t\right].$$

Observe that for $\tilde{p} \geq \tilde{b}^t$,

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}\tilde{p}}\left(\frac{\tilde{p}}{\tilde{\sigma}} \cdot g\left(\frac{\tilde{p} - \tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\right)\right) &= \left(\frac{\tilde{p}}{2\tilde{\sigma}^3} - \frac{\tilde{p}\,(\tilde{p} - \tilde{b}^t)\,(\tilde{p} + \tilde{b}^t - 2\tilde{b}^t\tilde{p})}{2\tilde{\sigma}^5}t\right) \cdot g\left(\frac{\tilde{p} - \tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\right) \\ &\leq t\frac{\tilde{p}}{2\tilde{\sigma}^5} \cdot \left(\frac{\tilde{\sigma}^2}{t} - (\tilde{p} - \tilde{b}^t)^2\right) \cdot g\left(\frac{\tilde{p} - \tilde{b}^t}{\tilde{\sigma}}\sqrt{t}\right). \end{aligned} \tag{21}$$

Therefore, for any $\tilde{p} \geq \tilde{b}^t + \tilde{\sigma}/\sqrt{t}$, this derivative is non-positive. It is straightforward to verify that if this condition holds for any $\tilde{p}_1 \geq 0.5$, it will also hold for $\tilde{p}_2 \geq \tilde{p}_1$. Now, since 1) this condition holds at $t = t^*$ for $\tilde{p} = \tilde{c}^*$, and 2) $\tilde{c}^* \geq 0.5$, the derivative in Eq. (21) is non-positive for any $\tilde{p} \geq \tilde{c}^*$. In particular, since c^t is increasing and $\tilde{b}^t$ is decreasing over time, we have

$$\begin{aligned}\frac{1}{N^t}\left(\frac{\mathrm{d}U^t}{\mathrm{d}t}\right)^+ &\leq \frac{1}{2\sqrt{t}}\Pr^t(p \geq c^t) \cdot \frac{\tilde{c}^t}{\tilde{\sigma}(c^t)} \cdot g\left(\frac{\tilde{c}^t - \tilde{b}^t}{\tilde{\sigma}(c^t)}\sqrt{t}\right) \\ &\leq \Pr^t(p \geq c^t) \cdot \frac{\tilde{c}^*}{2\tilde{\sigma}^*} \cdot \frac{1}{\sqrt{t}}\, g\left(\sqrt{t/t^*}\right).\end{aligned}$$

□

Proof of Theorem 4.1. The condition of broad targeting at t^* implies that at any time $t \geq t^*$, we have $\tilde{c}^t - \tilde{b}^t \geq \tilde{\sigma}^t/\sqrt{t}$. Then, for an L-Lipschitz f we obtain

$$\frac{\tilde{\sigma}^t}{\sqrt{t}} \leq \tilde{c}^t - \tilde{b}^t = \tilde{p}(c^t) - \tilde{p}(b^t) \leq L\,(1 - 2\epsilon) \cdot (c^t - b^t).$$

Since c^t is increasing and b^t is decreasing in time, for every $t \geq t^*$, we can further lower bound $c^t - b^t$ by $\frac{1}{L\,(1-2\epsilon)} \cdot \frac{\tilde{\sigma}^*}{\sqrt{t^*}}$. Plugging this into Eq. (12), for every $t \geq t^*$, we have

$$\operatorname{sign}\left(\frac{\mathrm{d}U^t}{\mathrm{d}t}\right) \leq \operatorname{sign}\left(\frac{\tilde{c}^*}{2\tilde{\sigma}^*\sqrt{t^*}}\left[\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{t}{2t^*}\right) - \frac{1}{L}\,G(1)\,\frac{1 - \tilde{c}^*}{1 - 2\epsilon}\right]\right).$$

So, for a Lipschitz constant $L \leq \frac{\sqrt{2\pi}\,G(1)}{1-2\epsilon}$, one can verify $\frac{\mathrm{d}U^t}{\mathrm{d}t} \leq 0$ for $t \geq t^*$. □

Proof of Theorem 5.1. Consider N individuals initially at time $t = 1$. Denote the subset of $\mathcal{A}^t$ with $y^t = k$ by $\mathcal{A}^t_k$. Define $n^t_k := |\mathcal{A}^t_k|/N$. Denote the set of individuals treated at t by $\mathcal{I}^t$. Excluding $\mathcal{I}^t$ from $\mathcal{A}^t_k$, denote the remaining by $\bar{\mathcal{A}}^t_k := \mathcal{A}^t \setminus \mathcal{I}^t$, and define $\bar{n}^t_k := |\bar{\mathcal{A}}^t_k|/N$. In the limit of $N \to \infty$, we can treat n^t_k and $\bar{n}^t_k$ as continuous variables taking any value in $[0, 1]$.

Step 1. The following property of the problem dynamics allows us to infer whether $\mathcal{A}^t_k$ is empty, depending on the previous step.

Lemma F.1. *Defining $\bar{n}_{-1}^t = 0$, at any time t and for any k, we have*

$$n_k^{t+1} > 0 \iff \bar{n}_k^t > 0 \;\; or \;\; \bar{n}_{k-1}^t > 0\,.$$

See proof on page 24.

Step 2. Consider two active individuals i and j at time t. If $y_i^t > y_j^t$, for our observation model in Eq. (1) and any utility function non-decreasing in p, Proposition E.5 implies treating i yields more utility than j in expectation. Therefore, the optimal allocation at any point should not target individuals with lower y^t while there are active individuals with a higher y^t.

Step 3. We show that, except for the first time step, at each time t on the optimal path, the treated individuals at t should have similar values of y^t.

Lemma F.2. *For $\mathcal{A}^t \neq \emptyset$ and $t \geq 2$ on the optimal path, for any $i, j \in \mathcal{I}^t$, we should have $y_i^t = y_j^t = \max\{i' \mid i' \in \mathcal{A}^t\}$.*

See proof on page 24.

Step 4. When an individual i with $y_i^t = k \geq 1$ is treated at t, not only is $\mathcal{A}_k^t$ non-empty, but $\mathcal{A}_{k-1}^t$ is also non-empty.

Lemma F.3. *For $k \geq 1$, if there exists $i \in \mathcal{I}^t$ on the optimal path such that $y_i^t = k$, then $n_{k-1}^t > 0$.*

See proof on page 25.

Step 5. Using the structure imposed on the optimal solution in the previous steps, we next restrict $\mathcal{I}^{t+1}$ based on $\mathcal{I}^t$.

Lemma F.4. *Suppose $\mathcal{A}^t \neq \emptyset$ and let $k = \max\{y_i^t \mid i \in \mathcal{A}^t\}$. On the optimal path,*

- *If $\mathcal{I}^t = \mathcal{A}_k^t$, either $\mathcal{I}^{t+1} \subseteq \mathcal{A}_k^{t+1}$ or $\mathcal{I}^{t+1} = \mathcal{A}_{k+1}^{t+1} = \emptyset$.*
- *If $\mathcal{I}^t \subset \mathcal{A}_k^t$, either $\mathcal{I}^{t+1} \subseteq \mathcal{A}_{k+1}^{t+1}$ or $\mathcal{I}^{t+1} = \mathcal{A}_{k+2}^{t+1} = \emptyset$.*

See proof on page 25.

Step 6. Except for one time step, at every t on the optimal path, either $\mathcal{I}^t =$ or $\mathcal{I}^t$ treats everyone with $y^t \geq k$ for some k. If there were two time steps t and t' violating this, because of the linearity of the expected utility in $|\mathcal{I}^t|$ and $|\mathcal{I}^{t'}|$, optimally, one would become zero or treat everyone above a cutoff. This and Lemma F.4 complete the proof.

□

Proof of Lemma F.1. For $k \geq 1$, $\mathcal{A}_k^{t+1}$ will consist of those in $\bar{\mathcal{A}}_k^t$ who survive and have $o^{t+1} = 0$, or those in $\bar{\mathcal{A}}_{k-1}^t$ who survive and have $o^{t+1} = 1$:

$$n_k^{t+1} = \bar{n}_k^t \,\mathbb{E}_{p \sim \mathcal{P}^t(\cdot|y^t=k)}\big[(1-p)(1-\tilde{p})\big] + \bar{n}_{k-1}^t \,\mathbb{E}_{p \sim \mathcal{P}^t(\cdot|y^t=k-1)}\big[(1-p)\,\tilde{p}\big]\,.$$

Here, we implicitly used the fact that the targeting cannot distinguish people with the same y^t. The same update rule works for $k = 0$ if we set $\bar{n}_{-1}^t = 0$. One can also verify that since the prior over p has no point mass, the expectations above are non-zero. This completes the proof. □

Proof of Lemma F.2. The proof is by contradiction with optimality and has multiple steps:

- Let $k = \max\{y_i^t \mid i \in \mathcal{A}^t\}$. Since $n_{k+1}^t = 0$, Lemma F.1 requires $\bar{n}_{k+1}^{t-1} = \bar{n}_k^{t-1} = 0$. On the other hand, when $n_k^t > 0$, Lemma F.1 requires either $\bar{n}_k^{t-1}$ or $\bar{n}_{k-1}^{t-1}$ to be non-zero. Since we just argued $\bar{n}_k^{t-1} = 0$, it is required to have $\bar{n}_{k-1}^{t-1} > 0$.
- Since $\bar{n}_{k-1}^{t-1} > 0$, Lemma F.1 implies $n_{k-1}^t > 0$. Then Step 2 requires that if $\mathcal{I}^t$ contains individuals with different y^t, there should be two individuals $i, j \in \mathcal{I}^t$ such that $y_i^t = k$ and $y_j^t = k - 1$.

- Consider the following tie-breaking when treating individuals with a similar y^t: Assign a random priority value $z_i^0 \in [0, 1)$ to each individual i in the initial pool. At time t, update the priority value by $z_i^t = z_i^{t-1} + o_i^t 2^{t-1}$. If any two individuals are at a tie to receive the treatment, choose the one with the lowest priority value. This tie-breaking does not change the optimality of an allocation rule.
- So far we showed $\bar{n}_{k-1}^{t-1} > 0$, i.e., there exist some individuals in $\mathcal{A}_{k-1}^{t-1}$ left untreated at $t-1$, but there exists $j \in \mathcal{A}_{k-1}^t$ who is treated at t. We argue this is suboptimal as the budget to treat j could have been spent earlier to treat those in $\mathcal{A}_{k-1}^{t-1}$, yielding higher utility. To see this, consider a counterfactual allocation rule that treats one more individual from $\mathcal{A}_{k-1}^{t-1}$ at $t-1$, to be referred to as individual j'. This individual has either failed or made it to t. If failed, the counterfactual treatment could be maximally effective in preventing a failure. If j' is active at t, she will either have $y_{j'}^t = k$ or $y_{j'}^t = k-1$. If $y_{j'}^t = k$, then she is treated with others in $\mathcal{A}_k^t$. If $y_{j'}^t = k-1$, still j' is treated. This is because j' maintains the lowest priority value among $\mathcal{A}_{k-1}^t$ by the construction of the priority values. Since j' is treated in any case if she makes it to t, she could be treated earlier at $t-1$. This could yield a higher or similar utility as u^t is non-increasing in t. This contradiction shows $y_i^t = y_j^t$. Then Step 2 implies $y_i^t = y_j^t = \max\{i' \mid i' \in \mathcal{A}^t\}$.

□

Proof of Lemma F.3. The proof is obvious for $t = 1$ since k can only be 1 and unless the budget is excessively large to treat everyone, we have $n_0^1 > 0$. For $t \geq 2$, Lemma F.2 requires $n_k^t > 0$ and $n_{k+1}^t = 0$. Then Lemma F.1 implies $\bar{n}_k^{t-1} = 0$ and $\bar{n}_{k-1}^{t-1} > 0$, so $n_{k-1}^t > 0$. □

Proof of Lemma F.4. We first prove the first part the lemma. If $\mathcal{I}^t = \mathcal{A}_k^t$, there are two possibilities for k. If $k = 0$, then $\bar{\mathcal{A}}_k^t = \emptyset$ and Step 2 implies no one is left untreated at t. So, $\mathcal{I}^{t+1} = \mathcal{A}_{k+1}^{t+1} = \mathcal{A}^{t+1} = \emptyset$. If $k \geq 1$, then Lemma F.3 implies $n_{k-1}^t > 0$. Then $\mathcal{I}_k^t = \mathcal{A}_k^t$ implies $\bar{n}_{k-1}^t > 0$ and $\bar{n}_k^t = 0$. Applying Lemma F.1 gives $n_k^{t+1} > 0$ and $n_{k+1}^{t+1} = 0$. Therefore, using Lemma F.2, treated individuals at $t+1$ should be among $\mathcal{A}_k^{t+1}$ or no one will be treated.

We next prove the second part of the lemma. If $\mathcal{I}^t \subset \mathcal{A}_k^t$, we have $\bar{n}_k^t > 0$ and $n_{k+1}^t = \bar{n}_{k+1}^t > 0$. Then Lemma F.1 implies $n_{k+1}^{t+1} > 0$ and $n_{k+2}^{t+1} = 0$. Therefore, using Lemma F.2, treated individuals at $t+1$ should be among $\mathcal{A}_{k+1}^{t+1}$ or no one will be treated. □