The Limits of AI-Driven Allocation: Optimal Screening under Aleatoric Uncertainty

Santiago Cortes-Gomez[*]
Machine Learning Department
Carnegie Mellon University
scortesg@cs.cmu.edu

Mateo Dulce Rubio[*]
Center for Data Science
New York University
mateo.d@nyu.edu

Carlos Patiño
University of Amsterdam
carlos.patino.paz@student.uva.nl

Bryan Wilder
Machine Learning Department
Carnegie Mellon University
bwilder@cs.cmu.edu

Abstract

The rise of machine learning has shifted targeted resource allocation in policy and humanitarian settings toward algorithmic targeting based on predicted risk scores. This approach is typically cheaper and faster than traditional screening procedures that directly observe the latent vulnerability status through physical verification. Yet, even access to the true conditional vulnerability probability cannot eliminate misallocation: aleatoric uncertainty over individual vulnerability status is irreducible, and probabilistic targeting inevitably misallocates some resources. In this work, we study how screening and algorithmic targeting should be optimally combined in a two-stage allocation framework, in which a screening stage observes true outcomes for a subset of units before a final allocation stage assigns the resource under a fixed coverage budget. We show that the optimal strategy screens units *at the margin of algorithmic allocation*, while directly targeting the highest-risk units. We empirically characterize when screening and algorithmic targeting act as complements or substitutes: efficiency gains from screening grow as the aleatoric uncertainty in the population increases. Finally, we characterize the relative value of screening and direct allocation, which yields a simple rule for splitting a program budget between the two: the optimal cutoff depends only on their cost ratio. We illustrate our framework with applications in income-based social protection programs and humanitarian demining in Colombia, where the tension between screening costs and allocation efficiency is operationally consequential.

1 Introduction

Resource allocation in high-stakes humanitarian and social settings requires distributing a scarce resource—cash transfers, medical treatment, humanitarian aid—across a population with heterogeneous and unobserved needs. In principle, every individual with a latent vulnerability condition should receive the resource, but universal screening for direct verification of eligibility is often prohibitively expensive. For instance, convening medical panels, conducting household visits, or performing field surveys [11, 22, 12] all involve large logistical demands that make universal screening infeasible under realistic budgets.

The rise of machine learning has offered a practical alternative to physical verification. Risk scores estimated from observable covariates can approximate individual vulnerability at scale: a model

[*]Equal contribution.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

trained on demographic, geographic, or contextual features assigns each unit a predicted probability of being in need, and decision-makers can allocate resources to the highest-scoring units without any direct outcome measurement. This *pre-hoc* algorithmic targeting circumvents the need to measure vulnerability directly, making it substantially cheaper and faster to deploy than traditional screening strategies. As predictive models have grown more accurate and data more abundant, AI-driven allocation has become the default strategy in many humanitarian and social policy contexts, from poverty targeting via proxy means tests [9, 1] to risk prioritization in humanitarian demining [6, 3].

Yet, probabilistic targeting cannot fully substitute for screening, even when the true conditional vulnerability probability is available. Aleatoric uncertainty over individual vulnerability remains irreducible as there typically is inherent variation in the actual vulnerability status of units with the same observable features. Some misallocation is therefore unavoidable under any *ex-ante* targeting strategy, even under a perfectly specified model. Screening, by contrast, resolves this uncertainty directly by observing the true status of screened individuals, allowing certain allocation for those units. Whether the efficiency gain from screening justifies its cost depends on how much aleatoric uncertainty the population exhibits and where in the risk distribution that uncertainty is concentrated.

This tension motivates a natural two-stage design that we formalize and study in this work. In the first stage, a decision-maker screens a subset of units directly observing their true vulnerability status. In the second stage, allocation targets vulnerable screened units and remaining resources are allocated using the vulnerability risk model for the unscreened units. The central question is then: *which units should be screened to maximize allocation efficiency?* Notably, the optimal screening policy does not target the highest- or lowest-risk units, but those *at the margin of algorithmic allocation* (Figure 1). This result characterizes when screening and algorithmic targeting act as complements or substitutes, with efficiency gains from screening growing as population-level aleatoric uncertainty increases.

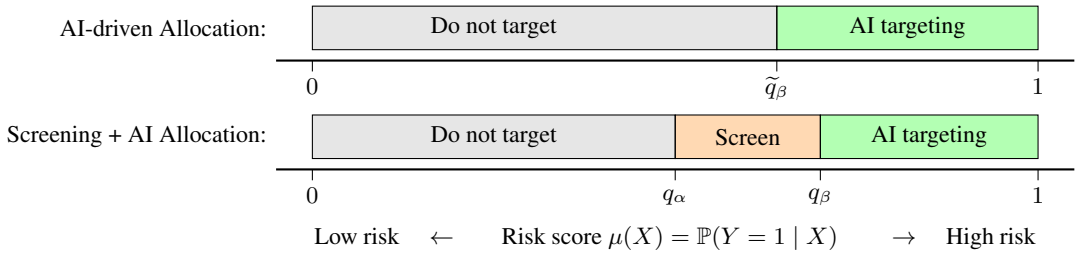

Figure 1: The population is ordered left to right by increasing vulnerability risk score $\mu(X) = \mathbb{E}[Y|X]$. *Top:* AI-driven allocation directly targets units with $\mu(X) > \widetilde{q}_\beta$. *Bottom:* Two-stage allocation with optimal screening band $[q_\alpha, q_\beta]$, where units are allocated only if their true vulnerability status $Y = 1$ is confirmed, and remaining budget is allocated to the highest-risk unscreened units, $\mu(X) > q_\beta$. The optimal screening set $[q_\alpha, q_\beta]$ sits at the margin of algorithmic allocation around $\widetilde{q}_\beta$.

Our contributions are as follows:

- We characterize the optimal screening strategy in a two-stage allocation framework with Bayes-optimal risk. We show that, regardless of the underlying risk distribution, the optimal screening set is an interval targeting units *at the margin of algorithmic allocation*, with endpoints that depend jointly on the risk distribution and the resource constraints, and develop a fixed-point algorithm that recovers the optimal band for arbitrary distributions.
- We establish a distribution-free characterization of the value of screening in allocation efficiency, formalizing a law of diminishing returns, and obtain sharp bounds on the worst-case loss of any fixed heuristic screening interval with respect to the optimal screening set.
- We derive the relative value of screening and direct allocation, leading to a simple, distribution-free, and scale-invariant cutoff for splitting a program budget between the two.
- We provide an empirical characterization of the relationship between aleatoric uncertainty and screening: the higher the irreducible uncertainty in the population, the greater the marginal gains from screening over purely algorithmic targeting.
- We demonstrate the practical effectiveness of our framework through real-world applications in humanitarian demining operations and income-based social programs, showing substantial gains over purely algorithmic targeting across budget regimes and risk estimation methods.

Our work characterizes the limits of AI-driven allocation under aleatoric uncertainty and provides decision-makers with concrete tools for optimally combining algorithmic targeting and screening. We release our code at `https://github.com/cmpatino/optimal-screening` so practitioners can readily apply the framework to their own risk scores, recovering the optimal screening and allocation policy under any budget constraints α and β.

2 Related work

Predictive machine learning is increasingly used to target scarce resources, including proxy means tests for poverty alleviation [9, 1], risk-based prioritization in humanitarian demining [6, 3], AI-assisted triage and health-care resource allocation [23, 8], and decision support in education [16, 19]. These applications typically frame prediction as a lower-cost substitute for direct verification [9, 16, 6]. We instead study screening and algorithmic targeting as complementary instruments, with their complementarity driven by aleatoric uncertainty. Prior work distinguishes aleatoric from epistemic uncertainty, with the former reflecting stochasticity inherent in the data-generating process [14, 13], while recent work argues that this dichotomy can be insufficiently expressive for some decision-relevant uncertainty quantities [20, 10]. We contribute to this discussion by linking aleatoric uncertainty to resource allocation directly, formalizing how its population-level magnitude determines the marginal value of screening and thereby when screening is most needed.

Other levers for improving allocation have also been studied. The idea of screening before allocating has been explored as a form of information acquisition: Cai et al. [2] study a budget-constrained decision-maker who screens selected individuals to acquire additional features that refine their risk estimates before allocating a resource. In contrast, we take the risk to be Bayes-optimal given the available covariates, so that screening observes true outcomes and fully resolves the residual aleatoric uncertainty for screened units, which allows us to characterize the optimal policy analytically. Relatedly, Perdomo [18] and Fischer-Abaigar et al. [7] compare the relative value of improving the predictive model versus expanding access to the resource. We instead hold the model fixed at the Bayes-optimal limit and study targeted screening as a complementary lever that policymakers have long used in practice [9, 22]. Our analysis of the budget split similarly endogenizes resources, but between screening and direct allocation rather than between prediction and access.

3 Setup: Resource Allocation under Aleatoric Uncertainty

Consider a centralized decision-maker that must allocate a scarce resource across a population to maximize allocation value. Each unit is characterized by covariates $X \in \mathcal{X}$ and a binary *vulnerability status* $Y \in \{0, 1\}$, where $Y = 1$ indicates that the unit would benefit from the resource. Crucially, Y is a fixed but latent characteristic of each unit, encoding a pre-existing condition rather than a response to the allocation decision. For instance, $Y = 1$ may indicate the presence of a disease, landmine contamination, or household poverty.

The decision-maker observes X for all units but may observe Y only for those selected for physical verification, which we refer to as *screening*. Units that are not screened must be allocated based on X alone. A resource allocation rule is a measurable function $T : \mathcal{X} \to \{0, 1\}$ that maps covariates to allocation decisions, where $T(X) = 1$ indicates that a unit with features X receives the resource. The *vulnerability risk score* summarizes all available information about a unit's latent status:

$$\mu(X) = \mathbb{E}[Y \mid X] = \mathbb{P}(Y = 1 \mid X). \tag{1}$$

We assume μ is Bayes-optimal given the available covariates unless explicitly stated otherwise.

Let $\beta \in (0, 1)$ denote the allocation budget as a fixed fraction of the population. The objective is to maximize allocation efficiency, which we define here as the probability of a correct allocation to vulnerable units:

$$\mathbb{P}(T(X) = 1, Y = 1) = \mathbb{E}[T(X)Y] = \mathbb{E}[T(X)\mu(X)], \tag{2}$$

subject to the budget constraint $\mathbb{E}[T(X)] \leqslant \beta$. Appendix A summarizes the notation used throughout.

Remark 1. *We assume, for now, that the allocation budget is fixed and the risk score is Bayes-optimal, reflecting short-term constraints where resources cannot be rapidly expanded and the model cannot be improved, so that all gains in allocation efficiency come from screening alone. This allows us to characterize, in full generality, the policy that maximizes allocation value for any given screening budget. Section 4.4 relaxes the first assumption, splitting a total budget between screening and allocation.*

Aleatoric Uncertainty. In our setting, a unit's vulnerability status follows $Y \mid X \sim \text{Bernoulli}(\mu(X))$. That is, outcomes are inherently stochastic for any unit with $\mu(X) \in (0,1)$. This irreducible randomness is the *aleatoric uncertainty* of the problem, reflecting stochasticity in the outcome itself rather than estimation error [14, 13, 10], and cannot be reduced by collecting more data or improving the risk model.

Crucially, aleatoric uncertainty imposes a fundamental limit on the allocation efficiency of purely algorithmic targeting. Any *ex ante* rule $T(X)$ based solely on $\mu(X)$ targets units that have $Y = 1$ only in expectation, so some misallocation is inevitable. The realized allocation value $\mathbb{E}[T(X)Y]$ will therefore always be smaller than the *ex post* optimal allocation where Y is known before the decision is made, even when μ is Bayes-optimal.

The magnitude of this gap depends on the distribution of $\mu(X)$ and in particular on its aleatoric uncertainty. When $\mu(X) \sim \delta_{0.5}$, risk scores are uninformative with high aleatoric uncertainty, and optimal allocation precision is at most 50%. When $\mu(X)$ follows a perfectly separating bimodal distribution, risk is well-separated with low aleatoric uncertainty, and algorithmic targeting can achieve full allocation efficiency. This motivates screening as a mechanism to close this gap.

Screening. We define *screening* as a procedure to directly observe the realized vulnerability status of any unit of interest, resolving aleatoric uncertainty at a cost. Formally, a screening rule is a function $S : \mathcal{X} \to \{0,1\}$, where $S(X) = 1$ indicates that a unit with covariates X is selected for physical verification. A decision-maker with a screening budget $\alpha \in (0, \beta)$ can observe the true outcome Y for a fraction α of the population before the allocation decision is made, converting a uncertain allocation into a certain one for those units. In the next section, we formalize this two-stage framework for optimally combining screening and algorithmic targeting to maximize allocation efficiency.

4 Resource Allocation with Screening

While universal screening would fully resolve aleatoric uncertainty, it is prohibitively costly in practice. Instead, we consider a two-stage framework in which a decision-maker operates under a screening budget $\alpha \in (0, \beta)$ and allocation budget $\beta \in (0,1)$, jointly optimizing over the screening S and allocation T rules to maximize allocation efficiency. The framework proceeds as follows:

1. **Screening stage**: Select a subset of units $S(X) = 1$ to observe their true vulnerability status Y, subject to a screening budget $\mathbb{E}[S(X)] \leqslant \alpha$.
2. **Allocation stage**: Assign the resource based on both covariates X and observed outcomes for screened units, subject to the allocation budget β. Screened units with $Y = 1$ are allocated with certainty, which is always ex-post optimal, and the remaining budget is deployed to unscreened units based on $\mu(X)$ alone.

This defines the two-stage allocation rule

$$\pi(X, Y) = S(X)Y + (1 - S(X))T(X), \tag{3}$$

where Y is only observed for screened units $S(X) = 1$, and $T(X)$ is the allocation rule for unscreened units based on $\mu(X)$ alone. A decision-maker chooses S and T to maximize allocation efficiency as the probability of allocation to vulnerable units, subject to screening and allocation budgets:

$$\begin{aligned} \max_{S,T:\ \mathcal{X}\to\{0,1\}} \quad & \mathbb{E}[\pi(X,Y)Y] \\ \text{subject to} \quad & \mathbb{E}[S(X)] \leqslant \alpha, \qquad \mathbb{E}[S(X)Y] + \mathbb{E}[(1 - S(X))T(X)] \leqslant \beta. \end{aligned} \tag{4}$$

We study this framework by first characterizing the optimal policy in the absence of screening, then deriving the optimal screening rule for a given α. We finalize with a distribution-free characterization of the value of screening and a formal law of diminishing returns to screening in allocation efficiency.

4.1 Optimal Algorithmic Allocation

We first consider the case of purely algorithmic allocation, where $\alpha = 0$ and screening is unavailable. Thus, all allocation decisions are made ex ante based solely on $\mu(X)$, and the problem reduces to

$$\max_{T:\ \mathcal{X}\to\{0,1\}} \quad \mathbb{E}[T(X)\mu(X)] \quad \text{subject to} \quad \mathbb{E}[T(X)] \leqslant \beta. \tag{5}$$

The solution has a simple and intuitive form: allocate to units in decreasing order of their risk score $\mu(X)$ until the budget is exhausted. Note that analogous threshold rules have been derived in related constrained allocation settings [15, 17, 7, 2]. All proofs are deferred to Appendix B.

Proposition 1 (Optimal Allocation Rule without Screening)**.** *The optimal allocation rule for* (5) *is* $T^*(X) = \mathbf{1}\{\mu(X) > \widetilde{q}_\beta\}$*, where* $\widetilde{q}_\beta$ *is the* $(1-\beta)$*-quantile of the distribution of* $\mu(X)$ *for* $\beta \in (0,1)$*.*

4.2 Optimal Screening Rule

We now characterize the optimal solution to the two-stage allocation problem. A critical distinction from the purely algorithmic allocation case is that screening resolves aleatoric uncertainty before allocation so that the realized vulnerability status Y is observed for screened units *before final allocation decisions are made*. From (3), this yields the allocation rule $\pi(X, Y) = Y$ for screened units ($S(X) = 1$), where Y is a deterministic observed value rather than a random variable. Intuitively, screening improves allocation efficiency in two ways: screened units with $Y = 1$ are allocated with certainty, improving targeting precision, while screened units with $Y = 0$ are excluded from allocation, avoiding wasteful use of the fixed budget β.

We revisit the two-stage optimization problem (4) by substituting the two-stage allocation rule (3) into the objective, decomposing $\mathbb{E}[\pi(X, Y)Y]$ by whether units are screened or not:

$$\begin{aligned}\mathbb{E}[\pi(X,Y)Y] &= \mathbb{E}[S(X)Y] + \mathbb{E}[(1 - S(X))T(X)Y] \\ &= \mathbb{E}[S(X)Y] + \mathbb{E}[(1 - S(X))T(X)\mu(X)].\end{aligned} \tag{6}$$

The first term captures allocation efficiency among screened and targeted units, and the second from algorithmic targeting of unscreened units. Equality follows from iterated expectations.

The allocation budget constraint also decomposes accordingly: screened units are allocated if and only if $Y = 1$, using $\mathbb{E}[S(X)Y]$ of the budget, while unscreened units are allocated via $T(X)$, using $\mathbb{E}[(1 - S(X))T(X)]$. The full two-stage optimization problem is therefore

$$\begin{aligned}\max_{S,T:\ \mathcal{X}\to\{0,1\}} \quad & \mathbb{E}[S(X)Y] + \mathbb{E}[(1 - S(X))T(X)\mu(X)] \\ \text{subject to} \quad & \mathbb{E}[S(X)] \leqslant \alpha, \qquad \mathbb{E}[S(X)Y] + \mathbb{E}[(1 - S(X))T(X)] \leqslant \beta,\end{aligned} \tag{7}$$

where α and β are the screening and allocation budgets respectively.

The central question is: *which units should be screened to optimize allocation efficiency?* Unlike the no-screening case, where the optimal rule is a simple quantile threshold (Proposition 1), the interaction between screening and allocation budgets creates a more complex optimization problem. Each screened unit uses allocation budget with probability $\mu(X)$, since it is allocated only when $Y = 1$, so screening the highest-risk units is inefficient and these are better served by direct allocation. Screening the lowest-risk units is equally inefficient since finding $Y = 1$ among them is rare, yielding few confirmed allocations. The optimal screening set therefore targets units in between: those with sufficient uncertainty that observing Y meaningfully improves allocation decisions, *and* sufficient risk to be competitive candidates given the budget β.

The characterization of the optimal screening rule builds on the intuition of Proposition 1: optimal policies are naturally expressed as sets defined by quantile thresholds in the risk score space. Since quantile notation will be central to all subsequent results, we establish it here alongside the relevant distributional objects. Let F denote the CDF of $\mu(X)$ and $F^{-1}(q)$ its q-th quantile, so that an interval $[q_a, q_b]$ refers to all units whose risk score falls between the q_a-th and q_b-th quantiles of F, with higher quantiles corresponding to higher risk. We now state our main theoretical result.

Theorem 1 (Optimal Screening in Two-Stage Allocation)**.** *Let* $\alpha \in (0, \beta)$ *and* $\beta \in (0, 1)$ *with* $\alpha + \beta < 1$*, and assume the risk distribution has full support on* $[0, 1]$*. The optimal screening rule is* $S^*(X) = \mathbb{1}\{\mu(X) \in [q_\alpha, q_\beta]\}$ *and the optimal direct allocation rule is* $T^*(X) = \mathbb{1}\{\mu(X) > q_\beta\}$ *where the thresholds* $0 < q_\alpha \leqslant q_\beta < 1$ *satisfy the screening and allocation budget constraints:*

$$F(q_\beta) - F(q_\alpha) = \alpha \qquad \textit{and} \qquad \int_{q_\alpha}^{q_\beta} \mu \, dF(\mu) + 1 - F(q_\beta) = \beta. \tag{8}$$

Proof sketch. The proof proceeds in three steps. Full details are presented in Appendix B.

Step 1: No directly allocated unit lies below a screened one. Suppose unit z with $\mu(z) < \mu(z')$ is directly allocated while unit z' is screened. Swapping them preserves the objective value and screening budget but reduces the allocation budget used by screening, since screening z uses allocation budget with probability $\mu(z) < \mu(z')$. This extra saved budget allows extending the direct allocation band to additional units, strictly improving the objective and contradicting optimality.

Step 2: Given S, the optimal allocation rule is a threshold in $\mu(X)$. The "second-stage" unscreened allocation problem reduces to the no-screening problem of Proposition 1 with residual budget $\beta - \mathbb{E}[S(X)Y]$. Since Step 1 implies the screening set lies entirely below the direct allocation set, the optimal unscreened allocation rule is a threshold $T^*(X) = \mathbf{1}\{\mu(X) > q_\beta\}$ in the original risk score space, with the highest-risk units always directly allocated.

Step 3: The optimal screening interval is $[q_\alpha, q_\beta]$. Among all screening intervals of mass α, shifting the screening interval upward has two effects: a gain in allocation efficiency from screening higher-risk units and a loss from the induced tightening of the direct allocation budget. The gain exceeds the loss because the reduction in the direct allocation term is proportional to $q_\beta < 1$, while the gain is proportional to the risk of the newly screened units. The net effect is always positive, so the objective is strictly increasing in the upper boundary. The optimal screening interval is therefore pushed up until q_β solves the budget constraints (8), with lower boundary $q_\alpha = F^{-1}(F(q_\beta) - \alpha)$ determined by the screening budget constraint. ■

Theorem 1 formally shows that the optimal screening set targets the highest-uncertainty units among relevant allocation candidates. The screening band $[q_\alpha, q_\beta]$ sits immediately below the direct allocation band $[q_\beta, 1]$, where units above q_β have sufficiently high risk to warrant direct allocation, while units below q_α are too unlikely to have $Y = 1$ to justify the screening cost. Corollary 1 further shows that the screening set targets units *at the margin of algorithmic allocation*: the purely algorithmic targeting threshold $\widetilde{q}_\beta$ from Proposition 1 is contained in the optimal screening interval $[q_\alpha, q_\beta]$.

Corollary 1. *The purely algorithmic allocation threshold $\widetilde{q}_\beta$ falls within the optimal screening band, $q_\alpha < \widetilde{q}_\beta < q_\beta$, for any $\alpha \in (0, \beta)$. Moreover, q_β is increasing in α and q_α is decreasing in α, expanding the screening band toward higher- and lower-risk units as the screening budget grows.*

Notably, the structure of the optimal policy does not depend on the underlying risk distribution F directly. The screening and allocation bands are always contiguous intervals in the risk score space, with boundaries determined solely by the budget constraints (8). The distribution F enters only implicitly through the quantiles q_α and q_β.[2] In the next section, we leverage this result to establish a distribution-free characterization of the value of screening.

Remark 2. *Since the optimal thresholds q_α and q_β are defined in risk score space, the result holds regardless of the nature of X, whether structured tabular data, images, text, or any other modality.*

4.3 Value of Screening

To study how allocation efficiency improves as the screening budget α grows, first note that, under the optimal two-stage allocation rule, the optimal allocation value takes the form

$$V^*(\alpha, \beta) = \int_{q_\alpha}^{1} \mu \, dF(\mu), \tag{9}$$

where $q_\alpha = F^{-1}(F(q_\beta) - \alpha)$ is determined by the screening budget constraint, and q_β solves the allocation budget constraint (8).

Theorem 2 (Marginal Value of Screening)**.** *The marginal value of the screening budget is $\frac{dV^*(\alpha,\beta)}{d\alpha} = \frac{q_\alpha(1-q_\beta)}{1-q_\beta+q_\alpha} > 0$, and $V^*(\alpha, \beta)$ is concave in α.*

Theorem 2 formally establishes that screening always improves optimal allocation efficiency, with $dV^*/d\alpha > 0$ for any feasible α. Moreover, the marginal value of screening is decreasing in α, formalizing a law of diminishing returns to physical verification. This result is distribution-free: the marginal value depends only on the optimal thresholds q_α and q_β, with the risk distribution F entering only implicitly through the budget constraints that determine the screening set.

[2] For the special case of uniform risk, the optimal thresholds can be solved analytically; see Corollary 3 in Appendix B.

The value of screening depends not only on *how much* is screened but also on *where*. Appendix E derives a sharp worst-case bound showing that fixed heuristic screening intervals, such as one centered at the allocation margin, can lose a substantial share of this value relative to the optimal band.

4.4 Relative Value of Screening and Direct Allocation

Having characterized the value of the optimal screening policy, we now study when screening is beneficial relative to direct allocation and derive the optimal split of a shared budget between screening and allocation. That is, we endogenize the budgets (α, β), which so far have been treated as exogenous parameters of the model. We begin with a cost-free characterization of the trade-off between the two.

Theorem 3 (Marginal Rate of Substitution between Screening and Allocation)**.** *For any $0 < \alpha < \beta < 1$, the marginal value of the allocation budget is $\frac{dV^*(\alpha,\beta)}{d\beta} = \frac{q_\alpha}{1-q_\beta+q_\alpha}$, and the marginal rate of substitution between screening and allocation is $\frac{dV^*/d\alpha}{dV^*/d\beta} = 1 - q_\beta$.*

Theorem 3 states that one unit of screening is worth $1 - q_\beta$ units of direct allocation. Intuitively, the marginal unit at the direct allocation threshold has risk q_β, so allocating to it directly wastes the resource with probability $1 - q_\beta$, whereas screening it first allocates only if it is positive. When q_β is low, direct allocation at the margin reaches low-risk units and screening is relatively valuable; as $q_\beta \to 1$, the marginal directly allocated unit is high-risk and screening it reveals little. Notably, this rate depends neither on q_α nor on the risk distribution F, except through q_β.

Next, let c_s and c_a denote the per-unit costs of screening and allocation, respectively, and let C denote the total budget per member of the population, consistent with (α, β) being fractions of the same population. The decision-maker then chooses the budgets that maximize the optimal allocation value:

$$\max_{\alpha,\beta} \; V^*(\alpha, \beta) \quad \text{subject to} \quad c_s\alpha + c_a\beta \leqslant C. \tag{10}$$

Since the marginal values of screening and allocation are positive (Theorems 2 and 3), the solution to (10) exhausts the full budget. Moreover, at an interior solution, where both instruments are used ($0 < \alpha^* < \beta^* < 1$), the marginal value per dollar must equalize across the two instruments, $\frac{dV^*/d\alpha}{c_s} = \frac{dV^*/d\beta}{c_a}$. Equivalently, the marginal rate of substitution must equal the cost ratio.

Corollary 2 (Optimal Budget Split between Screening and Allocation)**.** *Consider problem* (10) *with $0 < c_s < c_a$. If $C \geqslant c_a\,(1 - F(1 - c_s/c_a))$, the optimal direct allocation threshold is $q^*_\beta = 1 - \frac{c_s}{c_a}$, which determines the optimal screening and allocation budgets (α^*, β^*) via the binding total budget constraint. Otherwise, the optimal policy allocates the entire budget directly, $\alpha^* = 0$ and $\beta^* = C/c_a$.*

Corollary 2 gives a simple, distribution-free, and scale-invariant condition: the optimal boundary between the screening band and direct allocation depends only on the cost ratio c_s/c_a. Intuitively, the cost ratio determines *where* to draw the line in risk space, while C and the risk distribution determine *how much* this translates into in absolute terms. The budget condition simply requires that there is enough budget to allocate directly to every unit above the cutoff, whose mass is $1 - F(q^*_\beta)$. Otherwise, the budget is exhausted before the threshold q^*_β is reached.

The cost ratio also clarifies when screening is worthwhile. If $c_s \geqslant c_a$, q^*_β is not a valid quantile and the optimum is not to screen at all, consistent with a unit of screening never being worth more than a unit of direct allocation (Theorem 3). As $c_s/c_a \to 0$, screening extends to high-risk levels, with $\alpha \to \beta$. This is the regime considered in our experiments below, where screening is cheap relative to allocation. To decide whether additional screening is the best use of available funds, a decision-maker need only compute c_s/c_a for their setting, and the framework determines how aggressively to screen.

5 Computing the Optimal Screening Quantiles

We develop a fixed-point algorithm to recover the optimal screening quantiles via Algorithm 1 in Appendix C. While Theorem 1 characterizes the structure of the optimal screening rule, it does not directly yield a constructive method for computing the band thresholds, as q_β requires solving a nonlinear integral equation with no closed-form solution in general. The algorithm partitions the population into three contiguous bands ordered by risk. The top band, containing the highest-risk fraction $\beta - \alpha$ of the population, is fixed and directly allocated without screening. Below it, the

algorithm iterates over two bands: a screening band of mass α and a residual allocation band of directly allocated units using the budget saved by excluding confirmed $Y = 0$ units within the screening band. The key insight is that the saved budget depends on the average risk $\rho = \frac{1}{\alpha}\int_{q_\alpha}^{q_\beta} \mu \, dF(\mu)$ within the screening band: the lower the average risk of screened units, the more $Y = 0$ units are excluded and the larger the saved budget $\alpha(1-\rho)$. Given a candidate $\rho^{(k)}$, the algorithm fixes the residual allocation band mass at $\alpha(1-\rho^{(k)})$, places the screening band immediately below, and recomputes the average risk over the new screening band to obtain $\rho^{(k+1)} = g(\rho^{(k)})$. Theorem 4 in Appendix C proves that g is a contraction, guaranteeing convergence to a fixed point ρ^*; q_β is then recovered as the boundary between the residual allocation and screening bands as $q_\beta = F^{-1}(1-\beta+\alpha\rho^*)$.

6 Experiments

We validate our theoretical findings and illustrate the practical relevance of optimal screening in resource allocation using both synthetic and real-world data. We first characterize the gains from optimal screening across risk distributions with varying levels of aleatoric uncertainty, then present two real-world allocation settings: humanitarian demining operations in Colombia and identifying low-income individuals from the U.S. Census demographic data. All the code used for our experimental analysis is available at `https://github.com/cmpatino/optimal-screening`.

6.1 Empirical Value of Screening under Aleatoric Uncertainty

Our theoretical results are distribution-free: regardless of the underlying risk distribution, the structure of the optimal screening set and the concavity of the value function with respect to the screening budget α hold. To illustrate this property and characterize how the value of screening varies with aleatoric uncertainty, we sweep over the one-parameter family $\mu \sim \text{Beta}(t,t)$, $t \geqslant 0$. All distributions in this family are symmetric around $\mathbb{E}[\mu] = 0.5$, so differences in targeting outcomes are driven purely by distributional shape. Moreover, $\text{Beta}(t,t)$ continuously interpolates between canonical risk distributions: bimodal ($0 < t < 1$), uniform ($t = 1$) and unimodal bell-shaped ($1 < t < \infty$), representing increasing aleatoric uncertainty regimes. In the limit $t \to 0$ the distribution concentrates around binary risk $\mu \in \{0,1\}$, and as $t \to \infty$ it converges to the point mass $\delta_{1/2}$, where all units share the same risk $\mu = 1/2$. These represent the two extremes: one where underlying risk perfectly determines vulnerable units (no aleatoric uncertainty), and one where risk is entirely uninformative about vulnerability status (maximum aleatoric uncertainty).

For a given t, we generate synthetic populations with underlying risk $\mu_i \sim \text{Beta}(t,t)$ and realized outcomes $Y_i \mid \mu_i \sim \text{Bernoulli}(\mu_i)$, and compute the optimal screening policy via Algorithm 1.

Figure 2 fixes $\beta = 35\%$ and shows screening gains for $\alpha \in [0,\beta]$ across four reference distributions: bimodal ($t = 0.1$), uniform ($t = 1$), unimodal ($t = 10$), and $\delta_{0.5}$. Critically, without screening, optimal allocation precision ranges from 50% to 100%, demonstrating that even with access to the true risk distribution, targeting efficiency is fundamentally limited by aleatoric uncertainty. In particular, the bimodal distribution achieves perfect precision *without any screening*, while optimal screening with $\alpha = \beta = 35\%$ raises precision from 82% to 98% for the uniform, from 61% to 83% for the unimodal, and from 50% to 75% for $\delta_{0.5}$. In the latter case, screening yields large efficiency gains even though the risk scores carry no targeting information whatsoever. Crucially, since the total number of targeted units is fixed at β, all gains come from better targeting through observed outcomes via screening. Appendix D and E benchmark against random screening and fixed heuristic, showing that gains depend not only on how much screening is performed but on where it is concentrated.

Figure 3 fixes budgets $\alpha = \beta = 35\%$ and reports the *utility gap* $\Delta V^*(\alpha) = V^*(\alpha) - V^*(0)$, the percentage-point improvement in precision from optimal screening relative to no-screening (Proposition 1), as a function of t. The experimental results demonstrate a systematic relationship between aleatoric uncertainty and the value of screening: $\Delta V^*(\alpha)$ increases monotonically with t, tracing an S-shaped curve with three distinct phases. (i) As $t \to 0$, the marginal value of screening is flat and negligible: risk scores are highly informative and algorithmic targeting acts as a near-perfect substitute for screening. (ii) For $t \in [0.1, 10]$, spanning bimodal, uniform, and unimodal bell-shaped distributions, gains in allocation efficiency steeply increase as screening acts as a complement to algorithmic allocation. (iii) As $t \to \infty$, the curve flattens again as risk becomes highly non-informative and gains from screening saturate at their maximum. Together, Figures 2 and

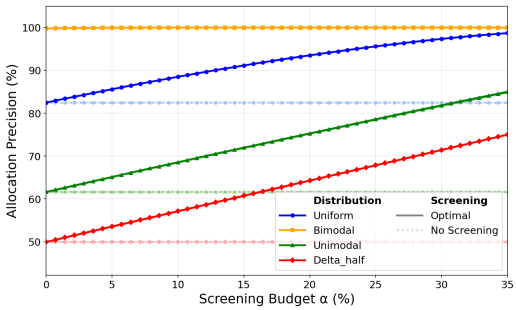

Figure 2: Allocation precision by screening budget α under optimal screening vs. no-screening across risk distributions ($\beta = 35\%$). Results are averaged over 10 simulations; standard deviations are negligible for every α (see Table 2).

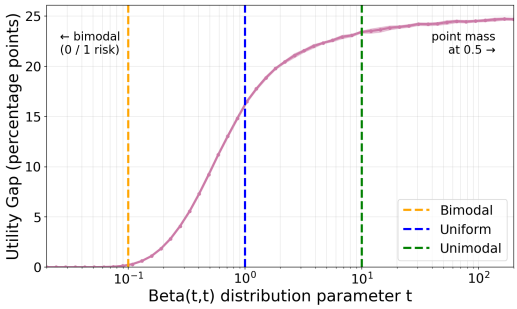

Figure 3: Utility gap $\Delta V^*(\alpha) = V^*(\alpha) - V^*(0)$ between optimal screening and no-screening as a function of the Beta(t, t) concentration parameter t ($\alpha = \beta = 35\%$, averaged over 10 simulations $\pm$ one standard deviation).

3 characterize the conditions under which screening and algorithmic targeting act as substitutes or complements, with the value of screening growing monotonically with population-level aleatoric uncertainty. In particular, real-world risk distributions with intermediate aleatoric uncertainty fall precisely in the steep region of the curve, where screening yields the largest marginal gains.

6.2 Applications to Real-World Resource Allocation

We apply our optimal screening framework to two real-world settings: humanitarian demining operations in Colombia and identifying low-income individuals for social protection programs in the US. In both settings, we train off-the-shelf machine learning models (logistic regression, random forest, and gradient boosting machine) on available predictive features and obtain out-of-sample risk estimates used to compute the optimal targeting policy for varying screening budgets α. We compare against a random screening baseline that randomly observes outcomes and allocates greedily (Proposition 1) using predictions from logistic regression. Reproducibility details are provided in Appendix F.

Humanitarian Demining. Humanitarian demining is the process of detecting and clearing landmines and other explosive remnants of war from conflict-affected areas so communities can safely reclaim land for economic and social development [21]. A decision-maker must allocate limited specialized demining teams across a territory, but first needs to identify which areas are actually contaminated with explosive ordnance (EO). Traditional mine action operations do this by screening the entire territory, a costly and time-consuming procedure [11]. This has motivated recent proposals to estimate EO contamination probabilities using machine learning, and use these estimates to prioritize field screening. Such systems have shown promising results in guiding team allocation in Colombia [6], Afghanistan [4], and Cambodia [3]. We apply our framework to study algorithmic allocation of demining teams in rural Colombia, using publicly available data from [6]. The dataset includes labeled data from finalized demining operations, with 63 predictive features encompassing geographic, socio-demographic, and conflict-related characteristics. We obtain out-of-sample risk estimates for Granada municipality, where approximately 10% of spatial units are truly contaminated. Accordingly, we set the allocation budget to $\beta = 10\%$, matching the underlying contamination prevalence.[3]

Income-Based Social Programs. Income verification is commonly used as a gatekeeping mechanism for determining eligibility in social programs. However, collecting and verifying income information can be costly, creating a natural setting in which screening decisions must trade off verification costs against allocation accuracy. We deploy our framework to predict whether an individual's total personal income exceeds \$50,000 per year, using the 2018 American Community Survey (ACS) Public Use Microdata Sample, accessed via the `folktables` benchmark [5]. This threshold follows the standard `ACSIncome` prediction task in this benchmark, rather than the eligibility criteria of any particular social program. The dataset contains approximately 1.6 million individual records drawn

[3]This implies that an allocation procedure with access to true contamination status would achieve perfect recall. Moreover, since $\beta = \mathbb{E}[Y]$, precision and recall coincide in this setting.

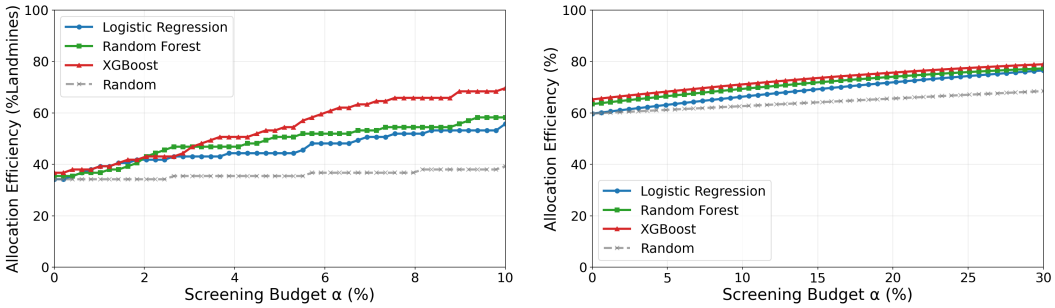

Figure 4: Precision (True Positives / β) as a function of the screening budget α under optimal screening allocation. Left: landmine contamination prediction in Colombia ($\beta = 10\%$). Right: ACS 2018 low-income prediction ($\beta = 30\%$). Across all models, optimal screening consistently improves allocation precision over both the no-screening baseline ($\alpha = 0$) and random screening.

from the U.S. Census Bureau's annual survey of socioeconomic conditions, with ten demographic and labor-market features per record. We split the data evenly into training and test sets, stratified on the outcome, and fix the allocation budget at $\beta = 30\%$.

Results. Figure 4 shows the percentage of landmines detected and the percentage of individuals correctly identified as low-income by the two-stage allocation procedure as a function of α. Our results demonstrate that optimal screening consistently improves allocation efficiency in real-world allocation settings. In particular, XGBoost achieves the largest gains in both experiments, raising precision from 39% to 70% and from 65% to 80%, respectively, at $\alpha = \beta$. Nevertheless, the monotone improvement with α is consistent across models regardless of their predictive performance, confirming that the framework is model-agnostic. Crucially, random screening yields small gains (35% to 39% and 60% to 65%, respectively), while optimal screening nearly doubles precision for the best-performing model. The comparison against random screening shows that the benefits of screening are not simply a function of the budget α, but of the screening criteria used. Unlike the synthetic experiments, here the risk scores are estimated from data, so results reflect both aleatoric uncertainty and estimation error. Hence, our experiments suggest that our framework can still deliver substantial gains in net benefit when using realistically estimated (rather than Bayes-optimal) risk scores.

7 Discussion and Limitations

Our work characterizes the fundamental limits of AI-driven allocation under aleatoric uncertainty and provides a principled framework for optimally combining algorithmic targeting and screening. We provide theoretical and empirical evidence that these two instruments are complements rather than substitutes: algorithmic targeting identifies the highest-risk units for direct allocation, while screening resolves aleatoric uncertainty for uncertain candidates. Operating under realistic budgets, our results identify where screening is most valuable and quantify the efficiency gains it yields. The concavity of the value of screening provides decision-makers with a principled basis for calibrating screening budgets, justifying targeted rather than universal screening, and determining when and where human verification is worth its cost. Taken together, these findings imply that improving algorithmic targeting should not come at the expense of eliminating screening, and that realizing the full potential of prediction-based allocation requires their principled integration.

Our framework is most useful when four conditions hold: (i) screening is cheaper per unit than allocation, $c_s < c_a$ (Corollary 2); (ii) universal screening is infeasible or prohibitively expensive; (iii) a risk model already exists or is feasible to obtain; and (iv) aleatoric uncertainty near the allocation cutoff is non-trivial, since the gains from screening grow with it, as we demonstrated empirically.

While our work assumes access to Bayes-optimal risk estimates, the practical value of screening is sensitive to estimation quality, as shown in our real-world applications. Epistemic uncertainty in the risk model may shift the screening set away from the optimal band, and understanding how estimation error propagates into screening decisions is a promising direction for future work. Extending the framework to imperfect screening, where noisy outcomes are revealed, is also a valuable research avenue.

References

[1] 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. 2, 3

[2] William Cai, Johann Gaebler, Nikhil Garg, and Sharad Goel. Fair allocation through selective information acquisition. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 22–28, 2020. 3, 5

[3] Flavio Cirillo, Gürkan Solmaz, Yi-Hsuan Peng, Christian Bizer, and Martin Jebens. Desk-AId: Humanitarian Aid Desk Assessment with Geospatial AI for Predicting Landmine Areas. *arXiv preprint arXiv:2405.09444*, 2024. 2, 3, 9

[4] Rory Collins, Lionel Fragniere, and Mateo Dulce Rubio. Advancements In Mine Action: Enhancing Remote Reporting and Analysis Through Innovative Technologies. *The Journal of Conventional Weapons Destruction*, 28(3):7, 2024. 9

[5] Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. Retiring Adult: New Datasets for Fair Machine Learning. *Advances in neural information processing systems*, 34:6478–6490, 2021. 9, 22

[6] Mateo Dulce Rubio, Siqi Zeng, Qi Wang, Didier Alvarado, Francisco Moreno Rivera, Hoda Heidari, and Fei Fang. RELand: Risk Estimation of Landmines via Interpretable Invariant Risk Minimization. *ACM Journal on Computing and Sustainable Societies*, 2(2):1–29, 2024. 2, 3, 9

[7] Unai Fischer-Abaigar, Christoph Kern, and Juan Carlos Perdomo. The Value of Prediction in Identifying the Worst-Off. In *International Conference on Machine Learning*, pages 17239–17261. PMLR, 2025. 3, 5

[8] Julien Grand-Clément, You Hui Goh, Carri Chan, Vineet Goyal, and Elizabeth Chuang. Interpretable Machine Learning for Resource Allocation with Application to Ventilator Triage. *arXiv preprint arXiv:2110.10994*, 2021. 3

[9] Margaret Grosh and Judy L Baker. Proxy Means Tests for Targeting Social Programs. *Living standards measurement study working paper*, 118:1–49, 1995. 2, 3

[10] Cornelia Gruber, Patrick Oliver Schenk, Malte Schierholz, Frauke Kreuter, and Göran Kauermann. Sources of uncertainty in supervised machine learning—a statisticians' view. *Statistical Science*, 41(3):671–696, 2026. 3, 4

[11] Armen Harutyunyan, Danielle Payne, David Hewitson, and Raphaela Lark. Operational Efficiency in Mine Action. Technical report, Geneva International Centre for Humanitarian Demining, Geneva, 2023. 1, 9

[12] Peter Hill. *Examination Systems*. Asia-Pacific Secondary Education System Review Series No. 1. UNESCO Bangkok, Asia and Pacific Regional Bureau for Education, Bangkok, 2010. 1

[13] Eyke Hüllermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods. *Machine Learning*, 110(3):457–506, March 2021. 3, 4

[14] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30, 2017. 3, 4

[15] Toru Kitagawa and Aleksey Tetenov. Who Should be Treated? Empirical Welfare Maximization Methods for Treatment Choice. *Econometrica*, 86(2):591–616, 2018. 5

[16] Tuo Liu, Cosima Schenk, Stephan Braun, and Andreas Frey. A Machine-Learning-Based Approach to Informing Student Admission Decisions. *Behavioral Sciences*, 15:330, 03 2025. 3

[17] Alexander R Luedtke and Mark J Van Der Laan. Optimal Individualized Treatments in Resource-Limited Settings. *The International Journal of Biostatistics*, 12(1):283–303, 2016. 5

[18] Juan Carlos Perdomo. The relative value of prediction in algorithmic decision making. In *Proceedings of the 41st International Conference on Machine Learning*, pages 40439–40460, 2024. 3

[19] Juan Carlos Perdomo, Tolani Britton, Moritz Hardt, and Rediet Abebe. Difficult lessons on social prediction from wisconsin public schools. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, pages 2682–2704, 2025. 3

[20] Freddie Bickford Smith, Jannik Kossen, Eleanor Trollope, Mark Van Der Wilk, Adam Foster, and Tom Rainforth. Rethinking Aleatoric and Epistemic Uncertainty. In *International Conference on Machine Learning*, pages 4345–4359. PMLR, 2025. 3

[21] UNMAS. The deadly legacy of landmines. `https://news.un.org/en/story/2023/04/1135252`, 2023. Accesed: 2024-07-13. 9

[22] World Health Organization. Regional Office for Europe. *Screening Programmes: A Short Guide. Increase Effectiveness, Maximize Benefits and Minimize Harm*. World Health Organization Regional Office for Europe, Copenhagen, 2020. 1, 3

[23] Hao Wu, Xiaoyu Lu, and Hanyu Wang. The Application of Artificial Intelligence in Health Care Resource Allocation Before and During the COVID-19 Pandemic: Scoping Review. *JMIR AI*, 2:e38397, 01 2023. 3

A Notation

Symbol	Definition
$X \in \mathcal{X}$	Covariates and feature space.
$Y \in \{0,1\}$	Vulnerability status indicating benefit from resource.
$\mu(X) = \mathbb{P}(Y = 1 \mid X)$	Bayes-optimal risk score.
$\beta \in (0,1)$	Fixed allocation budget as fraction of population.
$\alpha \in (0,\beta)$	Screening budget as fraction of population.
$S : \mathcal{X} \to \{0,1\}$	Screening rule.
$T : \mathcal{X} \to \{0,1\}$	Allocation rule for unscreened units based on covariates.
$\pi(X,Y)$	Two-stage allocation rule.
$\widetilde{q}_\beta$	$1-\beta$ quantile of $\mu(X)$; purely AI allocation threshold.
q_α, q_β	Lower and upper boundaries of the optimal screening band.
$[q_\alpha, q_\beta]$	Optimal screening band.
$V(\alpha)$	Allocation value with screening budget α (for fixed β).
$V^*(\alpha)$	Allocation value under optimal α-screening rule (for fixed β).
$\mathbb{P}_{X,Y}$	Joint distribution of covariates and outcome.
F	Cumulative distribution function (CDF) of μ.

Table 1: Summary of notation.

B Proofs

Proof of Proposition 1 (Optimal Allocation Rule without Screening). We proceed by contradiction. Let $q^* = q_{1-\beta}$ denote the $(1-\beta)$-quantile of $\mu(X)$, and suppose T is an optimal allocation rule that differs from the threshold rule $T^*(X) = \mathbf{1}\{\mu(X) > q^*\}$. Since $T \neq T^*$ and both exhaust exactly budget β, they must differ on a set of positive measure. In particular, there exists a unit x_0 allocated by T^* but not by T, i.e., $\mu(x_0) > q^*$ and $T(x_0) = 0$. Conversely, since both rules exhaust the same budget β but T does not allocate x_0, it must allocate some unit x_1 that T^* does not, i.e., $\mu(x_1) \leqslant q^*$ and $T(x_1) = 1$. In particular it follows that $\mu(x_0) > q^* \geqslant \mu(x_1)$, so $\mu(x_0) > \mu(x_1)$.

Define $\tilde{T}$ to agree with T everywhere except at x_0 and x_1, specifically:

$$\tilde{T}(x_0) = 1, \qquad \tilde{T}(x_1) = 0.$$

This swap preserves feasibility: $\mathbb{E}[\tilde{T}(X)] = \mathbb{E}[T(X)] = \beta$. The change in objective is:

$$\mathbb{E}[\tilde{T}(X)\,\mu(X)] - \mathbb{E}[T(X)\,\mu(X)] = \mu(x_0) - \mu(x_1) > 0,$$

strictly improving the value, contradicting the optimality of T. Hence no such T can exist, and the optimal rule must be the quantile threshold $T^*(X) = \mathbf{1}\{\mu(X) > q^*\}$. ■

Proof of Theorem 1 (Optimal Screening in Two-Stage Allocation). To prove the result, we first show that the optimal screening set $S^*(X)$ lies entirely below the optimal direct allocation set $T^*(X)$. With this, the second-stage allocation for unscreened units reduces to a no-screening optimization problem with a threshold rule solution. Finally, we show that the two optimal sets are contiguous, with the screening band sitting immediately below the direct allocation band, and that the boundaries are uniquely determined by the budget constraints.

Step 1: No directly allocated unit lies below a screened one. Suppose for contradiction that under the optimal policy, unit z with $\mu(z) < \mu(z')$ is directly allocated while unit z' is screened. Consider the alternative policy that swaps their roles: screen z and directly allocate z'. This swap preserves the screening budget constraint since $\mathbb{E}[S(X)]$ is unchanged. The direct contribution to the objective from these two units is unchanged since both screened and directly allocated units contribute $\mu(X)$ in expectation. However, the allocation budget used by the screened unit changes from $\mu(z')$ to $\mu(z) < \mu(z')$, saving $\mu(z') - \mu(z) > 0$ of allocation budget. This extra budget can be used to extend the direct allocation band to additional units with positive risk, strictly increasing the objective. This contradicts optimality of the original policy. Hence at optimality, no directly allocated unit has lower risk than any screened unit, i.e., the screening set lies entirely below the direct allocation set.

Step 2: Given S, the optimal allocation rule is a threshold in $\mu(X)$. Fix any feasible screening rule S with $\mathbb{E}[S(X)] \leqslant \alpha$. For screened units, the allocation rule is fixed at $\pi(X, Y) = Y$ by definition. For unscreened units, the remaining allocation problem is

$$\max_{T:\mathcal{X}\to\{0,1\}} \mathbb{E}[(1 - S(X))T(X)\mu(X)] \quad \text{s.t.} \quad \mathbb{E}[S(X)Y] + \mathbb{E}[(1 - S(X))T(X)] \leqslant \beta. \tag{11}$$

This is identical in structure to the no-screening problem of Proposition 1 with effective budget $\beta - \mathbb{E}[S(X)Y]$ on the subpopulation with $S(X) = 0$. By Proposition 1, the optimal unscreened allocation rule is a threshold rule in $\mu(X)$ within this subpopulation. From Step 1, we have that the screening set lies entirely below the direct allocation set, and the subpopulation threshold translates directly to a threshold in the original risk score space, $T^*(X) = \mathbf{1}\{\mu(X) > q_\beta\}$ for some q_β.

Step 3: The optimal screening interval is $[q_\alpha, q_\beta]$. By Steps 1 and 2, and Proposition 2 the optimal policy has the following form: screen some interval $[q_a, q_b]$ with $q_b \leqslant q_\beta$ and directly allocate $[q_\beta, 1]$, where q_β is determined by the allocation budget constraint. We now show the optimal screening interval satisfies $q_b = q_\beta$.

Consider the objective as a function of q_b, with $q_a = F^{-1}(F(q_b) - \alpha)$ determined by the screening budget constraint, and $q_\beta = q_\beta(q_b)$ determined by the allocation budget constraint:

$$\int_{q_a}^{q_b} \mu \, dF(\mu) + 1 - F(q_\beta(q_b)) = \beta. \tag{12}$$

The objective is:

$$V(q_b) = \int_{q_a}^{q_b} \mu \, dF(\mu) + \int_{q_\beta(q_b)}^{1} \mu \, dF(\mu). \tag{13}$$

Differentiating with respect to q_b, using $\frac{dq_a}{dq_b} = \frac{f(q_b)}{f(q_a)}$ from the screening budget constraint and $\frac{dq_\beta}{dq_b} = \frac{f(q_b)(q_b - q_a)}{f(q_\beta)}$ from the allocation budget constraint:

$$\frac{dV}{dq_b} = q_b f(q_b) - q_a f(q_b) - q_\beta f(q_b)(q_b - q_a) = f(q_b)(q_b - q_a)(1 - q_\beta) > 0, \tag{14}$$

since $q_b > q_a$ and $q_\beta < 1$, for $\alpha, \beta > 0$, and $f > 0$ from having full support on $[0, 1]$. Hence V is strictly increasing in q_b and the optimal screening interval is pushed as high as possible. An alternative approach is presented in the proof of Proposition 2 which directly shows that the gains from screening higher-risk units are larger than the reduction in direct allocation value from a smaller available budget.

The upper boundary is therefore $q_b = q_\beta$, and $q_\alpha = F^{-1}(F(q_\beta) - \alpha)$ from the screening budget constraint. ■

Proposition 2 (Interval Structure of Optimal Screening)**.** *Let $\alpha \in (0, \beta)$ and $\beta \in (0, 1)$ with $\alpha + \beta < 1$. The optimal screening rule S^* has an interval structure in the risk score space of the form $S^*(X) = \mathbb{1}\{\mu(X) \in [q_a, q_b]\}$ for some quantiles $0 < q_a \leqslant q_b < 1$.*

Proof of Proposition 2 (Interval Structure of Optimal Screening). Suppose by contradiction that there is a gap in the optimal screening set S^*. That is, there exist units x_a, x_b with $\mu(x_a) < \mu(x_b)$, both screened, and a unit x_m with $\mu(x_a) < \mu(x_m) < \mu(x_b)$ that is *not* screened. Note that x_m is also not directly allocated from Step 1 in the proof of Theorem 1.

Consider swapping x_a out of the screening set and replacing it with x_m, keeping all else equal, and define this rule as $\widetilde{S}$, satisfying the screening budget $\mathbb{E}[S^*] = \mathbb{E}[\widetilde{S}] = \alpha$. Since $\mu(x_m) > \mu(x_a)$, $\widetilde{S}$ now covers a unit with strictly higher risk, yielding strictly more expected positives from screening:

$$\mathbb{E}[\widetilde{S}Y] - \mathbb{E}[S^*Y] = \mu(x_m) - \mu(x_a) > 0.$$

However, the expected saved budget decreases by $\mu(x_m) - \mu(x_a)$, reducing the budget available for direct allocation. Define $q^* = q_{1-\beta+\mathbb{E}[S^*Y]}$ and $\widetilde{q} = q_{1-\beta+\mathbb{E}[\widetilde{S}Y]}$ as the optimal threshold rules for the direct allocation set from Proposition 1, from screening sets S^* and $\widetilde{S}$, respectively. We have that the change in contribution from direct allocation is then

$$\int_{q^*}^{1} \mu dF(\mu) - \int_{\widetilde{q}}^{1} \mu dF(\mu) = \int_{q^*}^{\widetilde{q}} \mu dF(\mu) \leqslant \widetilde{q}(\mu(x_m) - \mu(x_a)) < (\mu(x_m) - \mu(x_a)),$$

where the first inequality follows from $F(\widetilde{q}) - F(q^*) = \mathbb{E}[\widetilde{S}Y] - \mathbb{E}[S^*Y] = \mu(x_m) - \mu(x_a)$, and the second from $\widetilde{q} < 1$ provided $\alpha + \beta < 1$.

Therefore, the gains from screening higher-risk units exceed the reduction in contribution from direct allocation given a smaller available budget. That is, the net change in allocation efficiency is strictly positive, contradicting the optimality of S^*. Hence the optimal screening set is an interval in the risk space.

■

Proof of Theorem 2 (Marginal Value of Screening). The proof proceeds in two steps: we first differentiate $V^*(\alpha)$ via Leibniz's rule, and then obtain $dq_\beta/d\alpha$ via implicit differentiation of the budget constraint.

By Leibniz's rule applied to the allocation value under optimal screening (9):

$$\frac{dV^*}{d\alpha} = -q_\alpha f(q_\alpha)\frac{dq_\alpha}{d\alpha} = -q_\alpha f(q_\alpha)\frac{1}{f(q_\alpha)}\left(f(q_\beta)\frac{dq_\beta}{d\alpha} - 1\right),$$

where $\frac{dq_\alpha}{d\alpha}$ is obtained from the relationship $q_\alpha = F^{-1}(F(q_\beta(\alpha)) - \alpha)$. Moreover, from this expression we have the partial derivatives for q_α:

$$\frac{\partial q_\alpha}{\partial q_\beta} = \frac{f(q_\beta)}{f(q_\alpha)}, \qquad \frac{\partial q_\alpha}{\partial \alpha} = -\frac{1}{f(q_\alpha)}.$$

To derive $\frac{dq_\beta}{d\alpha}$, define the budget function

$$G(q_\beta, \alpha) \equiv \int_{q_\alpha}^{q_\beta} \mu dF(\mu) + 1 - F(q_\beta) - \beta = 0.$$

Applying Leibniz's rule to differentiate G with respect to q_β:

$$\frac{\partial G}{\partial q_\beta} = q_\beta f(q_\beta) - q_\alpha f(q_\alpha)\frac{\partial q_\alpha}{\partial q_\beta} - f(q_\beta) = f(q_\beta)\big(q_\beta - q_\alpha - 1\big).$$

Differentiating G with respect to α:

$$\frac{\partial G}{\partial \alpha} = -q_\alpha f(q_\alpha)\frac{\partial q_\alpha}{\partial \alpha} = q_\alpha.$$

By the implicit function theorem:

$$\frac{dq_\beta}{d\alpha} = -\frac{\partial G/\partial \alpha}{\partial G/\partial q_\beta} = -\frac{q_\alpha}{f(q_\beta)(q_\beta - q_\alpha - 1)} = \frac{q_\alpha}{f(q_\beta)(1 - q_\beta + q_\alpha)}.$$

This derivative is always positive as $1 > q_\beta \geqslant q_\alpha > 0$.

Finally, replacing this result for $dq_\beta/d\alpha$ in $dV^*/d\alpha$:

$$\frac{dV^*}{d\alpha} = -q_\alpha f(q_\alpha)\frac{1}{f(q_\alpha)}\left(f(q_\beta)\frac{dq_\beta}{d\alpha} - 1\right) = -q_\alpha\left(\frac{q_\alpha}{1 - q_\beta + q_\alpha} - 1\right) = \frac{q_\alpha(1 - q_\beta)}{1 - q_\beta + q_\alpha}.$$

Positivity follows immediately again by $1 - q_\beta + q_\alpha > 0$, from $q_\alpha, q_\beta \in (0, 1)$.

To establish concavity, we differentiate the expression for $dV^*/d\alpha$ with respect to α. Substituting the expressions for $dq_\alpha/d\alpha$ and $dq_\beta/d\alpha$ established above:

$$\frac{d^2V^*}{d\alpha^2} = -\frac{(1 - q_\beta)^3 f(q_\beta) + q_\alpha^3 f(q_\alpha)}{f(q_\alpha)f(q_\beta)(1 - q_\beta + q_\alpha)^3} < 0,$$

where the inequality follows since $q_\alpha, q_\beta \in (0, 1)$. Hence $V^*(\alpha)$ is concave in α, formalizing a law of diminishing returns to screening.

■

Proof of Corollary 1. Following the proof of Theorem 2, we have that

$$\frac{dq_\beta}{d\alpha} = \frac{q_\alpha}{f(q_\beta)(1 - q_\beta + q_\alpha)},$$

which is always positive given $0 < q_\alpha \leqslant q_\beta < 1$. That is, as the screening budget increases, the upper limit of the screening interval also increases.

Moreover, from the budget constraint $q_\alpha = F^{-1}(F(q_\beta(\alpha)) - \alpha)$ we have

$$\frac{dq_\alpha}{d\alpha} = \frac{1}{f(q_\alpha)}\left(f(q_\beta)\frac{dq_\beta}{d\alpha} - 1\right) = \frac{1}{f(q_\alpha)}\left(\frac{q_\alpha}{1 - q_\beta + q_\alpha} - 1\right) = \frac{1}{f(q_\alpha)}\left(\frac{q_\beta - 1}{1 - q_\beta + q_\alpha}\right),$$

which is always negative for $q_\beta < 1$. That is, as the screening budget grows, the lower limit of the screening interval decreases.

To establish that the purely algorithmic allocation threshold $\widetilde{q}_\beta$ falls within the optimal screening band, note that when $\alpha = 0$ the screening set collapses to $q_\alpha = q_\beta = \widetilde{q}_\beta$. Then, for $\alpha > 0$, q_α decreases and q_β increases, so $\widetilde{q}_\beta \in [q_\alpha, q_\beta]$ and the screening set selects units at the margin of algorithmic allocation around $\widetilde{q}_\beta$. ■

Corollary 3. *Assume the risk is uniformly distributed, $\mu \sim$ Unif$[0, 1]$. Then the optimal screening thresholds admit the closed-form expressions:*

$$q_\beta = \frac{1 - \beta - \alpha^2/2}{1 - \alpha}, \qquad q_\alpha = q_\beta - \alpha.$$

Proof. Under uniform risk $\mu \sim$ Unif$[0, 1]$, the screening constraint implies

$$F(q_\alpha) = F(q_\beta) - \alpha \quad \Longrightarrow \quad q_\alpha = q_\beta - \alpha.$$

Therefore, the allocation constraint (8) becomes

$$\beta = \int_{q_\alpha}^{q_\beta} \mu dF(\mu) + 1 - F(q_\beta) = \frac{1}{2}q_\beta^2 - \frac{1}{2}(q_\beta - \alpha)^2 + 1 - q_\beta.$$

Solving for q_β yields the result. ■

Proof of Proposition 1. Fix α and differentiate the budget constraints (**??**)–(**??**) with respect to β. Writing f for the density of F, the screening constraint gives $f(q_\beta)\, dq_\beta = f(q_\alpha)\, dq_\alpha =: A$, and the allocation constraint gives $q_\beta A - q_\alpha A - A = d\beta$, so that $A = -d\beta/(1 - q_\beta + q_\alpha)$. Since $V^*(\alpha, \beta) = \int_{q_\alpha}^1 \mu\, dF(\mu)$ (Eq. (9)), $dV^* = -q_\alpha f(q_\alpha)\, dq_\alpha = -q_\alpha A$, which yields the expression for $dV^*/d\beta$. Substituting it and $dV^*/d\alpha = q_\alpha(1 - q_\beta)/(1 - q_\beta + q_\alpha)$ (Theorem 2) into (**??**), the common factor $q_\alpha/(1 - q_\beta + q_\alpha)$ cancels and we obtain $1 - q_\beta^* = c_s/c_a$. ■

Proof of Theorem 3 (Marginal Rate of Substitution between Screening and Allocation). The proof follows the same strategy as Theorem 2, now differentiating with respect to β while holding α fixed. We first differentiate V^* via Leibniz's rule, then obtain $dq_\beta/d\beta$ via implicit differentiation of the budget constraint, and finally take the ratio with $dV^*/d\alpha$.

By Leibniz's rule applied to the allocation value under optimal screening (9):

$$\frac{dV^*}{d\beta} = -q_\alpha f(q_\alpha)\frac{dq_\alpha}{d\beta} = -q_\alpha f(q_\alpha)\frac{f(q_\beta)}{f(q_\alpha)}\frac{dq_\beta}{d\beta} = -q_\alpha f(q_\beta)\frac{dq_\beta}{d\beta},$$

where, since α is fixed, q_α depends on β only through q_β, with $\partial q_\alpha/\partial q_\beta = f(q_\beta)/f(q_\alpha)$ as in the proof of Theorem 2.

To derive $\frac{dq_\beta}{d\beta}$, consider the budget function G from the proof of Theorem 2, now as a function of (q_β, β) for fixed α:

$$G(q_\beta, \beta) \equiv \int_{q_\alpha}^{q_\beta} \mu dF(\mu) + 1 - F(q_\beta) - \beta = 0.$$

The partial derivative with respect to q_β is unchanged,

$$\frac{\partial G}{\partial q_\beta} = f(q_\beta)\big(q_\beta - q_\alpha - 1\big),$$

while differentiating with respect to β gives $\partial G/\partial\beta = -1$. By the implicit function theorem:

$$\frac{dq_\beta}{d\beta} = -\frac{\partial G/\partial\beta}{\partial G/\partial q_\beta} = \frac{1}{f(q_\beta)(q_\beta - q_\alpha - 1)} = -\frac{1}{f(q_\beta)(1 - q_\beta + q_\alpha)}.$$

This derivative is always negative as $1 > q_\beta \geqslant q_\alpha > 0$.

Replacing this result in $dV^*/d\beta$:

$$\frac{dV^*}{d\beta} = -q_\alpha f(q_\beta)\left(-\frac{1}{f(q_\beta)(1 - q_\beta + q_\alpha)}\right) = \frac{q_\alpha}{1 - q_\beta + q_\alpha},$$

which is positive since $q_\alpha > 0$ and $1 - q_\beta + q_\alpha > 0$.

Finally, dividing the expression for $dV^*/d\alpha$ in Theorem 2 by $dV^*/d\beta$:

$$\frac{dV^*/d\alpha}{dV^*/d\beta} = \frac{q_\alpha(1 - q_\beta)}{1 - q_\beta + q_\alpha} \cdot \frac{1 - q_\beta + q_\alpha}{q_\alpha} = 1 - q_\beta.$$

■

Proof of Corollary 2. Since $V^*(\alpha, \beta)$ is increasing in both arguments (Theorems 2 and 3), the budget constraint binds, and we can write $\beta(\alpha) = (C - c_s\alpha)/c_a$. Along this line, by Theorem 3,

$$\frac{dV^*(\alpha, \beta(\alpha))}{d\alpha} = \frac{dV^*}{d\alpha} - \frac{c_s}{c_a}\frac{dV^*}{d\beta} = \frac{dV^*}{d\beta}\left(1 - q_\beta - \frac{c_s}{c_a}\right),$$

whose sign is that of $1 - q_\beta - c_s/c_a$, since $dV^*/d\beta > 0$. Moreover, using the derivatives of q_β from the proofs of Theorems 2 and 3,

$$\frac{dq_\beta(\alpha, \beta(\alpha))}{d\alpha} = \frac{\partial q_\beta}{\partial \alpha} - \frac{c_s}{c_a}\frac{\partial q_\beta}{\partial \beta} = \frac{q_\alpha + c_s/c_a}{f(q_\beta)(1 - q_\beta + q_\alpha)} > 0,$$

so q_β is strictly increasing along the budget line, and the objective is increasing while $q_\beta < 1 - c_s/c_a$ and decreasing afterwards. Hence any interior maximum satisfies $q_\beta^* = 1 - c_s/c_a$, and it is unique.

At $\alpha = 0$, the entire budget is allocated directly, so $1 - F(q_\beta) = C/c_a$. If $C \leqslant c_a(1 - F(1 - c_s/c_a))$, then $q_\beta \geqslant 1 - c_s/c_a$ already at $\alpha = 0$, the objective is non-increasing along the whole budget line, and the optimum is $\alpha^* = 0$ and $\beta^* = C/c_a$. Otherwise, $q_\beta < 1 - c_s/c_a$ at $\alpha = 0$, and the maximum is attained where $q_\beta^* = 1 - c_s/c_a$. ■

C Algorithm for Computing the Optimal Screening Quantiles

Theorem 4 (Fixed Point: Existence, Uniqueness, and Convergence of Algorithm 1)**.** *Let $\alpha \in (0, \beta)$ and $\beta + \alpha < 1$. The update function $g(\rho) = \mathbb{E}[\mu \mid \mu \in [q_\alpha(\rho), q_\beta(\rho)]]$ of Algorithm 1 is a contraction on $[0, 1]$ with nonnegative Lipschitz constant*

$$c_F := F^{-1}(1 - \beta + \alpha) - F^{-1}(1 - \beta - \alpha) < 1.$$

By Banach's fixed point theorem, there exists a unique $\rho^* \in [0, 1]$ with $g(\rho^*) = \rho^*$, and Algorithm 1 converges for any $\rho^{(0)} \in [0, 1]$:

$$|\rho^{(k)} - \rho^*| \leqslant c_F^k\, |\rho^{(0)} - \rho^*|.$$

Proof of Theorem 4. Define the update function g as the average risk over the screening band:

$$g(\rho) = \frac{1}{\alpha}\int_{q_\alpha(\rho)}^{q_\beta(\rho)} \mu\, dF(\mu),$$

Algorithm 1 Fixed-Point Algorithm for Optimal Screening and Allocation Assignment

Require: Population sorted by increasing risk, allocation budget β, screening budget α.
Ensure: Band assignments that maximize allocation efficiency

$\rho_{\text{prev}} \leftarrow 0$ ▷ Initialize average risk of Screening Band
for $k = 1$ to `max_iterations` **do**
  Compute target mass for each band:
  $m_{alloc} \leftarrow \beta - \alpha$ ▷ Direct Allocation Band
  $m_{res} \leftarrow \alpha \cdot (1 - \rho_{\text{prev}})$ ▷ Residual Allocation Band from saved budget
  $m_{screen} \leftarrow \alpha$ ▷ Screening Band

  Assign units to bands by cumulative mass:
  $i_{alloc} \leftarrow \min\{i : \sum_{j=i}^{n} \frac{1}{n} \geqslant m_{alloc}\}$ ▷ Start of Direct Allocation Band
  $i_{res} \leftarrow \min\{i : \sum_{j=i}^{n} \frac{1}{n} \geqslant m_{alloc} + m_{res}\}$ ▷ Start of Residual Allocation Band
  $i_{screen} \leftarrow \min\{i : \sum_{j=i}^{n} \frac{1}{n} \geqslant m_{alloc} + m_{res} + m_{screen}\}$ ▷ Start of Screening Band

  Recompute average risk in Screening Band:
  $\rho_{\text{curr}} \leftarrow \frac{1}{i_{res} - i_{screen}} \sum_{j=i_{screen}}^{i_{res}-1} \mu_j$ ▷ μ_j = risk of unit j

  if $|\rho_{\text{curr}} - \rho_{\text{prev}}| < $ `tolerance` **then**
    break ▷ Convergence achieved
  end if
  $\rho_{\text{prev}} \leftarrow \rho_{\text{curr}}$ ▷ Update for next iteration
end for
return Band assignments $(i_{alloc}, i_{res}, i_{screen})$ with converged ρ_{prev}

where the band thresholds

$$q_\beta(\rho) = F^{-1}(1 - \beta + \alpha\rho),$$
$$q_\alpha(\rho) = F^{-1}(1 - \beta - \alpha + \alpha\rho),$$

are determined by the budget constraints of Theorem 1 and the average screening risk ρ, satisfying $q_\alpha < q_\beta$ and $F(q_\beta) - F(q_\alpha) = \alpha$.

First note that since $g(\rho)$ is an average of $\mu \in [0, 1]$, we have that $g(\rho) \in [0, 1]$ for all $\rho \in [0, 1]$, and g maps $[0, 1]$ into $[0, 1]$.

We now show that g is a contraction. Let $\rho_1 < \rho_2$, with corresponding screening bands $[q_\alpha(\rho_1), q_\beta(\rho_1)]$ and $[q_\alpha(\rho_2), q_\beta(\rho_2)]$ of F-mass α. Given that F^{-1} is nondecreasing and $\rho_2 - \rho_1 \leqslant 1$, we have $q_\alpha(\rho_2) = F^{-1}(1 - \beta - \alpha(1 - \rho_2)) \leqslant F^{-1}(1 - \beta + \alpha\rho_1) = q_\beta(\rho_1)$. That is, the two intervals always overlap (or touch at the boundary).

The overlapping region $[q_\alpha(\rho_2), q_\beta(\rho_1)]$ cancels in the difference, yielding

$$g(\rho_2) - g(\rho_1) = \frac{1}{\alpha}\left[\int_{q_\beta(\rho_1)}^{q_\beta(\rho_2)} \mu\, dF(\mu) - \int_{q_\alpha(\rho_1)}^{q_\alpha(\rho_2)} \mu\, dF(\mu)\right].$$

Both integrals are over intervals of F-mass $\alpha(\rho_2 - \rho_1)$. Bounding by the supremum and infimum of each region:

$$\int_{q_\beta(\rho_1)}^{q_\beta(\rho_2)} \mu\, dF(\mu) \leqslant q_\beta(\rho_2) \cdot \alpha(\rho_2 - \rho_1) \leqslant F^{-1}(1 - \beta + \alpha) \cdot \alpha(\rho_2 - \rho_1),$$
$$\int_{q_\alpha(\rho_1)}^{q_\alpha(\rho_2)} \mu\, dF(\mu) \geqslant q_\alpha(\rho_1) \cdot \alpha(\rho_2 - \rho_1) \geqslant F^{-1}(1 - \beta - \alpha) \cdot \alpha(\rho_2 - \rho_1),$$

where $F^{-1}(1 - \beta + \alpha)$ and $F^{-1}(1 - \beta - \alpha)$ correspond to $q_\beta(1)$ and $q_\alpha(0)$, respectively. Subtracting and dividing by α:

$$0 \leqslant g(\rho_2) - g(\rho_1) \leqslant c_F(\rho_2 - \rho_1),$$

where $c_F = F^{-1}(1-\beta+\alpha) - F^{-1}(1-\beta-\alpha)$. Since $\alpha < \beta$ implies $F^{-1}(1-\beta+\alpha) < 1$, and $\beta + \alpha < 1$ implies $F^{-1}(1-\beta-\alpha) > 0$, we have $c_F < 1$. The first inequality $g(\rho_2) - g(\rho_1) \geqslant 0$ follows because the first integral dominates the second, establishing that g is monotone increasing.

Hence $|g(\rho_1) - g(\rho_2)| \leqslant c_F|\rho_1 - \rho_2|$ and g is a contraction on the complete metric space $([0,1], |\cdot|)$. The result follows from Banach's fixed point theorem. ■

D Empirical Analysis on Synthetic Data

We consider four distributions governing the vulnerability risk in our synthetic population, depicted in Figure 5. All four distributions share the same expected fraction of vulnerable units ($\mathbb{E}[\mu] = 0.5$), so differences in targeting outcomes are driven purely by distributional shape. This allows us to study the limits of purely AI-driven allocation when no screening is employed, and to characterize the benefits of screening on allocation efficiency as aleatoric uncertainty increases. Specifically, we consider four risk distributions ordered by increasing aleatoric uncertainty. **Bimodal** (Beta$(0.1, 0.1)$): Mass concentrated near 0 and 1, with high separation between low- and high-risk units and minimal aleatoric uncertainty. **Uniform**: Flat over $[0, 1]$ representing moderate uncertainty with random risk level. **Unimodal bell-shaped** (Beta$(10, 10)$): Concentrated around 0.5, where most units have intermediate risk and the signal for targeting is weak. **Point mass** ($\delta_{0.5}$): All units have identical risk $\mu = 0.5$, representing the limiting case of maximum aleatoric uncertainty in which no algorithmic targeting is possible.

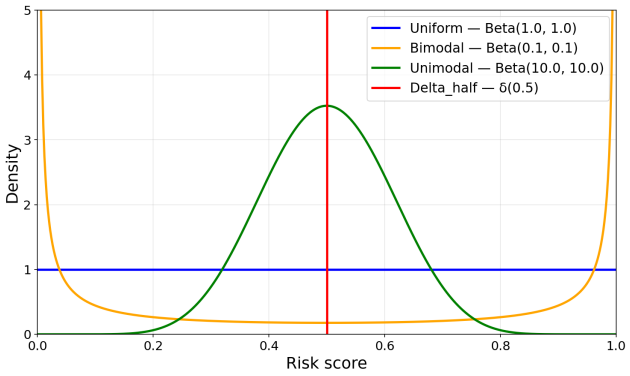

Figure 5: Risk distributions used in simulations. All four share the same mean ($\mathbb{E}[\mu] = 0.5$) but differ in aleatoric uncertainty, ranging from low (bimodal, with mass concentrated near 0 and 1) to maximum ($\delta_{0.5}$, with all mass at 0.5).

For each distribution, we sample $N = 100{,}000$ units with true conditional risks and realized outcomes $Y_i \mid \mu_i \sim \text{Bernoulli}(\mu_i)$. We fix the resource budget at $\beta = 35\%$ and run Algorithm 1 to obtain the optimal screening set for varying screening budgets $\alpha \in [0, \beta]$. We report allocation precision, defined as the fraction of allocated units that are truly vulnerable (True Positives / β). Since the total number of targeted units is fixed to $\beta = 35\%$, all improvements in allocation efficiency come from better targeting through observed outcomes via screening.[4]

We compare against the optimal no-screening policy (Proposition 1) in Figure 2. We further compare the optimal screening strategy against random screening, which selects units for physical verification uniformly at random (Figure 6, left), and a heuristic variant that places the screening band immediately below the purely algorithmic allocation threshold (Figure 6 right). These comparisons illustrate that the efficiency gains from screening depend not only on the total amount of screening capacity deployed, but critically on where in the score distribution screening is concentrated.

E Comparison with Fixed Heuristic Screening Rules

In practice, a natural way to deploy screening is to apply a simple rule of thumb: verify the units "near the cutoff," where the allocation decision is most uncertain. Such rules are practical because

[4] All computations were performed on a Lenovo laptop equipped with an Intel Core Ultra 9 258H CPU and 32 GB of RAM.

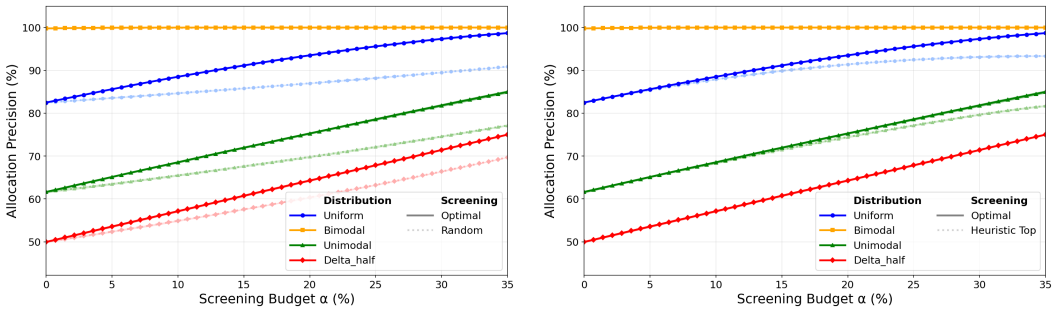

Figure 6: Allocation precision by screening budget α under optimal screening vs. random screening (left) and heuristic top-adjacent screening (right) across risk distributions ($\beta = 35\%$). Lines show averages over 10 simulations; shaded regions indicate $\pm$ one standard deviation, which is negligible at this scale for all α.

they require neither knowledge of the risk distribution F nor solving for the endpoints of the optimal band. This raises the question of whether computing the optimal policy is worth it, or whether a fixed heuristic performs comparably. In this section, we show that fixed heuristics can lose a substantial fraction of the value of screening, both theoretically, by characterizing their worst-case loss over F, and empirically.

We consider fixed screening intervals of the form $[a(u), u]$ parametrized by their upper endpoint u, with $F(u) - F(a(u)) = \alpha$. Let

$$\rho(u) := \int_{a(u)}^{u} \mu \, dF(\mu) \qquad \text{and} \qquad \rho^*(\alpha) := \int_{q_\alpha}^{q_\beta} \mu \, dF(\mu)$$

denote the expected number of screened positives under the heuristic and under the optimal band of Theorem 1, respectively. The residual budget $\beta - \rho(u)$ is allocated directly to the highest-risk unscreened units (Proposition 1), yielding value $V(u)$. Our analysis applies to any such interval, and we instantiate it with two natural heuristics:

- **Top-adjacent:** the screening band sits immediately below the pure algorithmic allocation threshold, $u = \widetilde{q}_\beta$.
- **Symmetric:** the screening band is centered, in quantile terms, at $\widetilde{q}_\beta$, that is, $u = q_s$ with $F(q_s) = F(\widetilde{q}_\beta) + \alpha/2$.

First, note that there is no nontrivial distribution-free *lower* bound on the gain of optimal screening over a fixed heuristic. For instance, as F approaches a point mass, any band of mass α selects essentially the same units as the optimal band, so the loss of a heuristic can be made arbitrarily small. We therefore characterize the worst-case loss of any fixed screening interval over risk distributions F, and show that the resulting bounds are sharp, including for the two heuristics above.

Proposition 3 (Worst-Case Loss of Fixed Screening Intervals)**.** *Fix $0 < \alpha < \beta < 1$ and a screening interval $[a(u), u]$ with $F(u) - F(a(u)) = \alpha$.*

(i) If $u < q_\beta$, let $r(u)$ solve $F(r(u)) = 1 - \beta + \rho(u)$. Then

$$V^*(\alpha) - V(u) \leqslant F(r(u)) - F(\max(q_\alpha, u)).$$

(ii) If $u > q_\beta$, then

$$V^*(\alpha) - V(u) \leqslant \rho(u) - \rho^*(\alpha).$$

Moreover, both bounds are sharp, in the sense that for each there exists a sequence of distributions F_n along which the loss $V^*(\alpha) - V(u)$ attains the upper bound in the limit.

Proof. **Case (i), $u < q_\beta$.** Screening lower-risk units yields fewer expected positives, $\rho(u) < \rho^*(\alpha)$, so the direct allocation cutoff $r(u)$ satisfies $u < r(u) < q_\beta$. This leaves a gap $(u, r(u))$ that is neither

screened nor allocated. Since $V(u) = \int_{a(u)}^{u} \mu\, dF(\mu) + \int_{r(u)}^{1} \mu\, dF(\mu)$ and $V^*(\alpha) = \int_{q_\alpha}^{1} \mu\, dF(\mu)$ (Eq. (9)), we have

$$V^*(\alpha) - V(u) = \int_{\max(q_\alpha,u)}^{r(u)} \mu\, dF(\mu) - \int_{a(u)}^{\min(q_\alpha,u)} \mu\, dF(\mu) \leqslant F(r(u)) - F(\max(q_\alpha, u)),$$

using $\mu \leqslant 1$ in the first integral and $\mu \geqslant 0$ in the second. That is, the loss is at most the mass of the uncovered gap. The bound is approached as F concentrates the mass of $[\max(q_\alpha, u), r(u)]$ near $\mu = 1$ and that of $[a(u), \min(q_\alpha, u)]$ near $\mu = 0$.

Case (ii), $u > q_\beta$**.** Screening higher-risk units yields $\rho(u) \geqslant \rho^*(\alpha)$ and frees budget $\Delta(u) := \beta - \rho(u) - (1 - F(u)) > 0$, which extends direct allocation *below* the screening window to a threshold $d(u)$ solving $F(a(u)) - F(d(u)) = \Delta(u)$. Coverage is then contiguous, so $V(u) = \int_{d(u)}^{1} \mu\, dF(\mu)$, with $F(d(u)) - F(q_\alpha) = \rho(u) - \rho^*(\alpha) \geqslant 0$. Hence

$$V^*(\alpha) - V(u) = \int_{q_\alpha}^{d(u)} \mu\, dF(\mu) \leqslant \rho(u) - \rho^*(\alpha),$$

which is approached as F concentrates the mass of $[q_\alpha, d(u)]$ near $\mu = 1$. ■

Proposition 3 holds for any fixed screening interval. Applying it to the two heuristics above gives their worst-case losses.

Top-adjacent heuristic. Since $q_\alpha \leqslant \widetilde{q}_\beta \leqslant q_\beta$ (Corollary 1), $u = \widetilde{q}_\beta$ falls in case (i) of Proposition 3 with $\max(q_\alpha, u) = \widetilde{q}_\beta$. Using $F(\widetilde{q}_\beta) = 1 - \beta$ and $F(r(\widetilde{q}_\beta)) = 1 - \beta + \rho(\widetilde{q}_\beta)$, its worst-case loss is

$$V^*(\alpha) - V(\widetilde{q}_\beta) \leqslant \rho(\widetilde{q}_\beta).$$

That is, in the worst case, the positives gained by the top-adjacent heuristic through screening are exactly offset by the uncovered gap it creates, so it performs no better than no screening at all.

Symmetric heuristic. Since $F(q_\beta) = 1 - \beta + \rho^*(\alpha)$ and $F(q_s) = 1 - \beta + \alpha/2$, we have $q_s \leqslant q_\beta$ if and only if $\alpha/2 \leqslant \rho^*(\alpha)$. In case (i), $\max(q_\alpha, q_s) = q_s$ and the bound is $F(r(q_s)) - F(q_s) = \rho(q_s) - \alpha/2$; in case (ii), the bound is $\rho(q_s) - \rho^*(\alpha)$. Combining both cases, its worst-case loss is

$$V^*(\alpha) - V(q_s) \leqslant \rho(q_s) - \min\Big(\frac{\alpha}{2}, \rho^*(\alpha)\Big).$$

That is, if its band sits below the optimal one, it leaves an interval of mass $\rho(q_s) - \alpha/2$ neither screened nor allocated; if it sits above, it consumes $\rho(q_s) - \rho^*(\alpha)$ more allocation budget than the optimal policy by screening higher-risk units.

Proposition 3 shows that the advantage of the fixed-point algorithm is not an artifact of a loose upper bound, as the worst-case loss of fixed heuristics is attained. For both heuristics, the optimal placement of the band depends on F, so no fixed rule can close the gap without knowledge of the risk distribution.

Empirical comparison. Table 2 reports allocation precision at $\alpha = \beta = 35\%$ for the synthetic distributions of Appendix D and for an *adversarial* distribution that places risk $\mu \approx 0$ on 58% of units and $\mu \approx 0.8$ on the remaining 42%. Unlike the other synthetic distributions, the adversarial distribution does not satisfy $\mathbb{E}[\mu] = 0.5$. The table also reports results for the humanitarian demining application, with XGBoost risk scores and $\alpha = \beta = 10\%$.

Because the Beta(t, t) family is symmetric around $\mu = 0.5$, the symmetric heuristic tracks the optimum closely on these distributions. The top-adjacent heuristic, nonetheless, shows a non-negligible gap in the realistic regime of moderate aleatoric uncertainty: 5.4 points on the uniform and 3.3 points on the unimodal distribution. On the adversarial distribution, the gap grows to 3.5 points for the symmetric and 12.2 points for the top-adjacent heuristic, the latter performing close to the no-screening baseline, as predicted by Proposition 3. On the demining task, where risk scores are estimated from data, the symmetric and top-adjacent heuristics fall 6.3 and 12.5 points short of the optimal policy, respectively. These results confirm that fixed heuristics can leave significant value on the table once risk departs from symmetric, idealized cases, whereas the fixed-point algorithm (Algorithm 1) adapts to how risk is actually distributed near the cutoff and attains $V^*(\alpha, \beta)$ for any F.

Table 2: Allocation precision (%) of optimal screening vs. fixed heuristics at $\alpha = \beta$ (35% for synthetic distributions, 10% for demining). Synthetic results report mean $\pm$ one standard deviation over 10 simulations. Bold indicates the best-performing screening rules for each distribution.

Distribution	No screening	Top-adjacent	Symmetric	Optimal
Uniform	82.49 ± 0.17	93.34 ± 0.17	97.84 ± 0.07	$\mathbf{98.70 \pm 0.19}$
Bimodal	99.78 ± 0.03	$\mathbf{99.99 \pm 0.00}$	$\mathbf{100.00 \pm 0.00}$	$\mathbf{99.99 \pm 0.02}$
Unimodal	61.63 ± 0.27	81.69 ± 0.21	$\mathbf{84.84 \pm 0.25}$	$\mathbf{84.97 \pm 0.35}$
Point mass ($\delta_{0.5}$)	49.97 ± 0.24	$\mathbf{75.03 \pm 0.13}$	$\mathbf{74.98 \pm 0.19}$	$\mathbf{75.01 \pm 0.22}$
Adversarial	81.08 ± 2.29	83.31 ± 1.54	91.97 ± 1.63	$\mathbf{95.51 \pm 1.51}$
Demining (XGBoost)	36.71	57.14	63.29	**69.62**

F Reproducibility

Datasets. *Humanitarian demining.* The dataset contains 26,761 grid cells of $500\,\text{m} \times 500\,\text{m}$ in Antioquia, Colombia, each described by 63 geographic, socio-demographic, and war-remnant features. The outcome is landmine contamination, with 1.6% of grid cells having a positive label.[5]

Income-based social programs. We use the 2018 American Community Survey (ACS) Public Use Microdata Sample, accessed via the `folktables` benchmark [5]. The dataset contains 1,611,572 individual records from the U.S. Census Bureau's annual survey of socioeconomic conditions. Each record includes ten demographic and labor-market features: age, class of worker, educational attainment, marital status, occupation code, place of birth, relationship to household head, usual hours worked per week, sex, and race. The binary outcome is whether an individual's total personal income exceeds $50,000 per year, which holds for approximately 37% of individuals.

Computing environment. All experiments use Python 3.11, with dependencies managed via `uv` and a committed `uv.lock` lockfile that pins the exact resolved versions. Running `uv sync` reproduces the environment.

Models. We use three tabular baselines (logistic regression, random forest, and XGBoost), each implemented as a `scikit-learn` pipeline with a `ColumnTransformer` preprocessor that applies `StandardScaler` to numeric columns and `OneHotEncoder` to categorical columns.

Data splitting. For humanitarian demining, we use a geographic split: the test set consists of all grid cells in the municipality of Granada (773 cells), and the training set of the remaining 14 municipalities. Since Granada is excluded before fitting, the model is evaluated on both unseen labels and an unseen feature distribution. For income-based social programs, we use a random 50/50 train–test split stratified on the outcome, with approximately 806,000 records in each split.

All precision curves, which report true positives recovered as a function of the screening budget α, are computed on the test split only, using the fitted model's predicted risk to rank units and the true outcomes to resolve screened units.

[5]The data is available at `https://huggingface.co/datasets/mdulcer10/landmine-detection`.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The list of contributions in the introduction are well explained in the main body (Sections 4 and 5) and experiments (Section 6).

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Yes, we explicitly discuss limitations in Section 7 and throughout the paper.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

Justification: All necessary assumptions are included in the theorem statements in the main body and in the appendices. All theoretical results are (correctly) proved in the appendix.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Section 6. Experiments - contains all details about the experimental setup, and Appendix F includes all details for reproducibility

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: In the Experiments section, we provide a link to a well-documented repository with instructions for accessing the data and executing the code to reproduce our results.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

Answer: [Yes]

Justification: Section 6. Experiments and and Appendix F contain thorough details about the experimental setup.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: All our results are averaged over several runs and reported with a one standard deviation confidence. We also give details on the initialization of each simulation in Section 6.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer:[Yes]

   Justification: We include the computational resources used in our experiments in the Appendix.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer:[Yes]

   Justification: Our work conforms, in every respect, with the NeurIPS Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: Section 7 thoroughly discusses broader impacts and limitations.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.

- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A].

    Justification:

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: Both datasets used are properly cited in Section 6 and Appendix F.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A].

    Justification:

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A].

    Justification:

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A].

    Justification:

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: The core methods in our work do not involve LLMs as any important, original, or non-standard components.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.