

Fair Allocation through Selective Information Acquisition

William Cai, Johann Gaebler, Nikhil Garg, Sharad Goel
Stanford University

Abstract

Public and private institutions must often allocate scarce resources under uncertainty. Banks, for example, extend credit to loan applicants based in part on their estimated likelihood of repaying a loan. But when the quality of information differs across candidates (e.g., if some applicants lack traditional credit histories), common lending strategies can lead to disparities across groups. Here we consider a setting in which decision makers—before allocating resources—can choose to spend some of their limited budget further screening select individuals. We present a computationally efficient algorithm for deciding whom to screen that maximizes a standard measure of social welfare. Intuitively, decision makers should screen candidates on the margin, for whom the additional information could plausibly alter the allocation. We formalize this idea by showing the problem can be reduced to solving a series of linear programs. Both on synthetic and real-world datasets, this strategy improves utility, illustrating the value of targeted information acquisition in such decisions. Further, when there is social value for distributing resources to groups for whom we have *a priori* poor information—like those without credit scores—our approach can substantially improve the allocation of limited assets.

1 Introduction

Approximately one in seven U.S. households have unmet demand for small-dollar loans, and are often unable to secure credit from traditional financial institutions as they have little or no formal credit history [1]. However, in the majority of these households, individuals receive regular income and typically pay their bills on time, which suggests many in fact would have low risk of default [1]. One barrier to providing loans to this low-risk yet underserved subpopulation is that it can be more expensive and time-consuming to screen individuals with non-traditional financial histories, limiting the inclusiveness of the banking system.

Motivated by this problem, we formalize and analyze a general setting in which one must allocate limited funds both to screen applicants and to distribute resources. The task for a budget-constrained decision maker is thus to first select a set of candidates to screen and then, given the results of that screening process, determine to whom to allocate the remaining resources. We assume there is a fixed

cost for screening each applicant, and that decision makers have prior knowledge of the distribution of information they would receive if they choose to screen an applicant. In practice, such prior knowledge could be obtained by screening a small random sample of applicants to learn the resulting information distributions.

We derive an efficient algorithm for computing an optimal, utility-maximizing strategy for the general screening and allocation problem. To do so, we first show that once a set of candidates has been selected to screen, it is optimal to allocate the remaining resources according to a threshold rule, with assets distributed to those candidates having post-screening expected utility above a fixed threshold. Further, for any fixed threshold policy, we show that one can find the optimal set of candidates to screen (while satisfying the budget constraint) via a linear program. Intuitively, one should screen candidates near the margin, for whom the screening process could reveal information that could push a candidate across the threshold. But to do this rigorously, one also needs to account for the precise structure of the prior information. Finally, we sweep over the possible thresholds, solving the corresponding linear program at each point. In this manner, we obtain both a rule to screen candidates and a specific threshold policy for distributing funds to candidates with sufficiently high post-screening value.

We further consider an extension of the above problem in which policymakers have explicit value for diversity. For example, instead of simply finding a max-utility policy, one might maximize utility subject to the constraint that a particular group—for example, those who traditionally have had limited access to credit markets—achieve at least a fixed minimum utility. In the United States, those with unmet demand for credit are disproportionately black and Hispanic [1], heightening the value of diversity considerations in allocation decisions. We show that this (and related) extensions can be incorporated into our general algorithmic approach in a straightforward manner.

To demonstrate the potential value of augmenting allocation decisions with a screening phase, we apply our methods to both synthetic datasets and one with real measures of creditworthiness. In particular, we examine the potential benefits of screening as a function of the cost and value of

information. Especially when we impose a diversity constraint, we find that screening strategies can significantly outperform a naive strategy that simply attempts to satisfy the constraint without screening any applicants.

For concreteness, we frame our discussion in terms of lending decisions, though our approach applies to many allocation settings. For example, it is often challenging to accurately assess household wealth—particularly in countries where informal and irregular work is more common—and, in turn, to appropriately target the distribution of government subsidies [22]. The simple strategy of distributing funds to those families clearly in need can systematically overlook populations with harder-to-verify financial status. As with lending, one can judiciously allocate some of the budget to more extensively screen certain applicants, ensuring funds are ultimately distributed to those who can benefit the most.

2 Related Work

Several papers address the problem of *active feature acquisition* [19, 20, 25], where one can selectively purchase missing data to improve the overall out-of-sample performance of a statistical model. We consider the related problem of acquiring features to identify specific, high-value individuals. While there is some shared intuition between the two settings—that one should seek information on individuals most likely to alter downstream decisions—the technical approach we take is different, in large part because our end goal is optimal allocation rather than statistical learning.

In a related, recent stream of research, Bakker et al. [2] and Noriega-Campero et al. [22] likewise consider a feature acquisition problem, but with fairness constraints. In their setting, the decision maker must acquire additional features for each individual to ensure classification decisions have similar errors rates across groups—including parity in false negative and false positive rates—a common measure of fairness in the machine learning community [8, 15, 17]. Our approach to the problem differs in three important respects. First, we adopt the perspective of constrained utility maximization. Past work has shown that directly equalizing error rates can lead to outcomes that, counterintuitively, may harm the very groups they were designed to protect [9, 10, 18]. We avoid such deleterious outcomes by instead framing the problem explicitly in terms of group-specific utilities: fairness is encoded into our requirement that the decision maker must allocate some minimum amount of utility to each group. Second, we focus on one-shot screening decisions, in which decision makers simply choose whether or not to acquire information on each individual, rather than sequentially deciding how much information to acquire based on the results of each past acquisition decision. Our one-shot formulation maps to the binary decision structure (i.e., to screen or not to screen) common in many institutions and leads to different optimization challenges. Third, we directly model the tradeoff between screening and allocation decisions by tying both to a common budget constraint (i.e., more screening means less funds are available to ultimately distribute to individuals).

Finally, our work touches on research from the fair division and allocation literature [5, 21, 26], which consid-

ers how to share resources while satisfying fairness properties defined between individuals. The former devises mechanisms wherein strategic agents self-divide the resources fairly, and the latter studies the existence of allocations that jointly satisfy various fairness notions. In contrast to our work, that line of research is particularly concerned with individual incentives, strategic action, and equilibrium effects.

3 A Model of Screening and Allocation

We model screening and allocation decisions as a sequential process in which a budget-constrained lender first selects a (possibly random) subset of candidates to further screen from a pool of n applicants, and then, based on the information revealed in that screening phase, selects a second (possibly random) subset of candidates to receive a loan.

We assume the value of lending to an applicant i is given by the random variable U_i . These utilities are intended to capture the full social value of providing loans, and we imagine the lender aims to optimize social welfare, as in the case of a government agency. In general, the lender has only partial information about U_i . More specifically, if the lender chooses not to screen an applicant, we assume the lender knows only the applicant’s conditional expectation $\mu_i = \mathbb{E}[U_i|X_i = x_i]$ given their pre-screening covariates x_i , such as credit score for those applicants who have traditional credit histories. On the other hand, if the lender opts to screen an applicant, they learn $\mathbb{E}[U_i|X_i = x_i, \tilde{X}_i = \tilde{x}_i]$, where \tilde{x}_i denotes the additional information one gains through screening. For example, \tilde{x}_i might encode applicant i ’s history of paying their electricity or phone bills—information that is often feasible to acquire with some extra effort and which is a good indicator of creditworthiness [1].

When deciding whom to screen, we assume the lender knows the distribution of $\mathcal{D} = (D_1, \dots, D_n)$, where $D_i = \mathbb{E}[U_i|X_i = x_i, \tilde{X}_i]$. That is, the lender knows how their estimate of utility could change if they decide to screen each applicant, where these distributions may depend on the available pre-screening covariates. A lender may, for example, thus choose only to screen applicants whose estimate is likely to substantially change given additional information. We further assume the lender must pay a fixed cost c_S for screening an applicant and a cost c_A for underwriting a loan. For simplicity we assume these costs do not vary across applicants, though it is straightforward to extend to the more general case; see Appendix A.

Based on knowledge of the above information and cost structure, the lender selects a randomized strategy to screen applicants. That is, the lender chooses a vector $\mathcal{P} = (p_1, \dots, p_n)$, meaning that each applicant i is selected to be screened independently with probability p_i . Let $\mathcal{S} = (S_1, \dots, S_n)$ indicate which applicants are ultimately screened under this policy; therefore, $S_i \in \{0, 1\}$ is a Bernoulli random variable with probability of success p_i .

Given this randomized screening policy, we can now write the information $\hat{\mathcal{U}} = (\hat{U}_1, \dots, \hat{U}_n)$ the lender has at the end of the screening phase as follows:

$$\hat{U}_i = \begin{cases} \mathbb{E}[U_i|X_i = x_i] & \text{if } S_i = 0, \\ \mathbb{E}[U_i|X_i = x_i, \tilde{X}_i] & \text{if } S_i = 1. \end{cases} \quad (1)$$

In other words, \hat{U}_i is the lender's post-screening estimated utility of giving a loan to applicant i . In particular, if $S_i = 1$ (i.e., the applicant is screened), the lender's estimate changes from $\mathbb{E}[U_i|X_i = x_i]$, the estimate based only on applicant i 's pre-screening covariates x_i , to $\mathbb{E}[U_i|X_i = x_i, \tilde{X}_i = \tilde{x}_i]$, which incorporates the post-screening information \tilde{x}_i .

Finally, the lender chooses an allocation policy $\mathcal{A}(\hat{\mathcal{U}}) = (A_1, \dots, A_n)$, where $A_i \in [0, 1]$ specifies the probability a loan is (independently) offered to each applicant. Importantly, \mathcal{A} is a function of the lender's post-screening utility estimates $\hat{\mathcal{U}}$. For example, the lender might give loans to the individuals with the highest post-screening utility estimates, up to the budget constraint.

Combining all of the above, the lender's optimization problem is to choose screening and allocation policies $(\mathcal{P}^*, \mathcal{A}^*)$ that maximize expected welfare,

$$(\mathcal{P}^*, \mathcal{A}^*) \in \arg \max_{\mathcal{P}, \mathcal{A}} \mathbb{E} \left[\sum_{i=1}^n U_i A_i \right], \quad (2)$$

subject to being budget-balanced in expectation,

$$\mathbb{E} \left[\sum_{i=1}^n c_S S_i + c_A A_i \right] \leq B, \quad (3)$$

where B is a fixed, non-negative constant.

In some settings, decision makers may value diversity in their allocations. For example, they may wish to ensure a certain minimum number of loans are provided to groups that historically have been excluded from credit markets. One can encode this policy preference directly into the utilities, in which case the resulting optimization problem would incorporate one's value for diversity. That approach, however, requires decision makers to agree upon these utilities to interpret the results, which can be challenging.

Here we take a complementary approach that explicitly allows value for diversity to differ across decision makers. Suppose $\{G_1, \dots, G_m\}$ partitions the applicant pool into m groups; for example, if $m = 2$, we might partition candidates into those who traditionally have had access to credit markets and those who have not. Then we require the selected policy $(\mathcal{P}^*, \mathcal{A}^*)$ to allocate at least $\Lambda_j \geq 0$ utility to group G_j , where these utilities do not themselves include any value for diversity:

$$\mathbb{E} \left[\sum_{i \in G_j} U_i A_i \right] \geq \Lambda_j. \quad (4)$$

In practice, as we discuss below, one would solve this optimization problem for a range of Λ , which traces out the Pareto frontier of possible policies across different group constraints, corresponding to different values for diversity. The diversity condition above is expressed in terms of utility, but we might, alternatively, simply lower bound the *number* of loans $\sum_{i \in G_j} A_i$ given to members of each group G_j . This alternative constraint can be handled in a straightforward manner by our algorithm detailed below.

A stylized example

We illustrate the above ideas in the context of a simple, stylized example. Suppose a lender must decide how best to allocate loans among an applicant pool of 13 people, with an overall budget of \$2,000. Providing a loan costs $c_A = \$400$, and additional screening of an applicant costs $c_S = \$50$.

Further suppose that five of the applicants are able to provide rich credit histories, and the expected utility of giving each of them a loan is $\mu_i = \$750$. For this group, additional screening would not provide any more information. Imagine that the other eight applicants do not have formal credit histories, and the utility of giving them a loan is accordingly lower due to the risk of default, with $\mu_i = \$500$. However, the lender knows that these applicants come in two types that could be disambiguated through additional screening. More specifically, the lender knows that after screening, individuals in this group can be divided into those with expected utility \$1,000 (with 50% chance) and those with expected utility \$0 (with 50% chance). The high-utility group could, for example, correspond to those who demonstrate a history of consistently paying their bills on time.

The naive strategy that does not screen any individuals would allocate five loans to hit the budget constraint, since $5 \times \$400 = \$2,000$. All five loans would go to applicants with rich credit histories ($\mu_i = \$750$) over those without ($\mu_i = \500). In this case, the total expected utility of the no-screening allocation is $5 \times \$750 = \$3,750$.

But in this scenario, one can improve overall utility by screening all eight applicants without formal credit histories and then granting loans to the high-utility applicants that are identified. Under that strategy, we expect four of the eight screened applicants to be identified as high utility (\$1,000), and so the expected utility of the allocation is $4 \times \$1,000 = \$4,000$, greater than the expected utility of \$3,750 under the no-screening strategy. Finally, the expected cost of the strategy is $8 \times \$50 = \400 for screening plus $4 \times \$400 = \$1,600$ for distributing the loans, totaling \$2,000 and satisfying the budget constraint. In this example, one can thus improve overall utility—and even allocate more loans to the group that *a priori* appears worse—by incorporating additional screening into the decision-making process.

4 Finding Optimal Policies

We now derive an efficient algorithm to find optimal screening and allocation policies $(\mathcal{P}^*, \mathcal{A}^*)$, subject to the budget and diversity constraints. We start by showing that over the full space of policies, it is optimal to allocate resources according to a *threshold policy*, with loans dispersed to individuals having post-screening expected utility \hat{U}_i above a group-specific threshold t_j for $j = 1, \dots, m$. Then, for each threshold policy, we show the optimal screening policy can be obtained by solving a linear program (LP), a type of optimization problem with linear objective function and linear constraints, for which there exist fast solution methods [3]. As a result, we can find an optimal combined screening and allocation policy by sweeping over threshold policies and solving the corresponding LP for each such policy.

We begin by formally defining threshold policies.

Definition 1 (Threshold Policy). A *threshold policy* is an allocation policy \mathcal{A} for some fixed $t_j \in \mathbb{R} \cup \{-\infty, \infty\}$ and $\alpha_j \in [0, 1]$, $j = 1, \dots, m$, such that

$$A_i = \begin{cases} 1 & \hat{U}_i > t_{g_i}, \\ \alpha_{g_i} & \hat{U}_i = t_{g_i}, \end{cases}$$

where g_i denotes the group membership of individual i .

Threshold policies deterministically allocate resources to those with post-screening expected utilities above a fixed, group-specific threshold t_j . Randomization (i.e., allocating resources with probability α_j) at the threshold t_j may be necessary to exactly satisfy the budget constraint, which is important because for an individual not screened, the distribution of \hat{U}_i is concentrated at a single point.

Theorem 2 below formally states that it is sufficient to restrict to the set of threshold policies when searching for a globally optimal screening and allocation policy.

Theorem 2. *Suppose the constrained optimization problem defined by Eqs. (2), (3), and (4) has a solution $(\mathcal{P}^*, \mathcal{A}^*)$. Then there is a threshold policy \mathcal{T}^* such that $(\mathcal{P}^*, \mathcal{T}^*)$ is also a solution.*

To see this, suppose that \mathcal{P}^* deterministically selects a subset of applicants to screen. Then, for each group G_j , it is clear one should allocate loans to the approximately k_j individuals in each group with the highest post-screening estimated utility, where k_j is the number of loans granted to each group under \mathcal{A}^* . Such a rank-based allocation can equivalently be written as a threshold rule with the same expected utility. The more general case, in which screening decisions are randomized, introduces some technical complications, but the spirit of the argument is similar. Full details can be found in Appendix A.

Now, given a threshold policy \mathcal{T} with thresholds $t_1, \dots, t_m \in \mathbb{R}$ and boundary randomization probabilities $\alpha_1, \dots, \alpha_m \in [0, 1]$, we turn to finding an optimal companion screening policy. We show that such an optimal screening policy \mathcal{P} can be found by solving an LP. In particular, the LP has n decision variables p_1, \dots, p_n , with $p_i \in [0, 1]$ specifying the probability that applicant i is screened; the objective equals the utility of the combined screening and allocation policy; and the constraints encode our budget and diversity conditions.

To construct the LP, we first define the following quantities that depend on both the threshold rule \mathcal{T} and the lender's prior knowledge on the value of screening:

$$\begin{aligned} q_i &= \Pr(D_i > t_{g_i}), \\ e_i &= \mathbb{E}[U_i \mid D_i > t_{g_i}], \\ o_i &= \mathbb{1}_{\mu_i > t_{g_i}} + \alpha_{g_i} \cdot \mathbb{1}_{\mu_i = t_{g_i}}, \end{aligned}$$

where g_i denotes applicant i 's group membership. Recall that μ_i denotes pre-screening expected utility and D_i denotes post-screening expected utility if applicant i is screened. Both μ_i and the distribution of D_i are known in advance.

For fixed screening probabilities $p = (p_1, \dots, p_n)$, the expected utility of the corresponding screening and allocation

policy can now be expressed as a linear function of p :

$$\sum_{i=1}^n q_i \cdot e_i \cdot p_i + o_i \cdot \mu_i \cdot (1 - p_i). \quad (5)$$

The first summand reflects the expected utility associated with applicant i if they were screened, and the second summand reflects the expected utility if they were not screened.

Our goal is to maximize the expression in Eq. (5) subject to the budget and diversity conditions, which we now show can also be expressed as linear constraints on p . In terms of the constants defined above, the budget constraint can be written as

$$\sum_{i=1}^n c_S \cdot p_i + c_A \cdot q_i \cdot p_i + c_A \cdot o_i \cdot (1 - p_i) \leq B. \quad (6)$$

Likewise, the j diversity constraints in Eq. (4) become

$$\sum_{i \in G_j} q_i \cdot e_i \cdot p_i + o_i \cdot \mu_i \cdot (1 - p_i) \geq \Lambda_j, \quad (7)$$

for $j = 1, \dots, m$.

Together, the objective given by Eq. (5), with constraints defined by Eqs. (6) and (7), define a linear program with decision variables $\{p_i\}$, the solution of which—if one exists—gives a screening policy \mathcal{P}^* that is optimal for the threshold policy \mathcal{T} . A jointly optimal screening and allocation policy $(\mathcal{P}^*, \mathcal{T}^*)$ can accordingly be found through a grid search over all threshold policies \mathcal{T} in a (discretized) space $\mathbb{R}^m \times [0, 1]^m$. For each choice of policy \mathcal{T} , defined by threshold and randomization parameters t_1, \dots, t_m and $\alpha_1, \dots, \alpha_m$, we find the corresponding optimal screening policy \mathcal{P} by solving the LP described above. Then, among the resulting screening-allocation pairs, the one maximizing utility is guaranteed to be globally optimal.

5 Experiments

We now investigate the value of our sequential screening and allocation approach through two simulation exercises, one based on synthetic data and another based on real loan data. First, with the synthetic data, we examine the optimal policies as we vary both the cost of screening and the value of the resulting information. Then, with the real-world data, we illustrate how a decision maker could, in practice, operationalize our screening and allocation approach to pursue equity when distributing limited resources.

In each experiment, we allocate loans to members of two groups. One of the groups—which we call the “targeted” group—can be screened for more information at some cost, and we imagine there is social value to distributing more resources to this group. For example, the targeted group may be comprised of those who traditionally have not had ready access to the banking system and accordingly have limited formal credit history. For simplicity, we further assume that those in the non-targeted group cannot be screened, perhaps because they already have complete credit histories.

We define the utility of lending to applicant i to be

$$U_i = \begin{cases} a & Y_i = 1, \\ b & Y_i = 0, \end{cases} \quad (8)$$




Figure 1: For an applicant with pre-screening likelihood of being creditworthy $x_i = 0.4$, the distribution of post-screening creditworthiness, $\Pr[Y_i = 1 | X_i = 0.4, \tilde{X}_i]$. Under a regime with high value of information (red line), the post-screening distribution has higher variance than in a setting with low value of information (blue line).

where $Y_i \in \{0, 1\}$ indicates whether they would pay back the loan, and a and b are fixed, known constants. In both of our experiments, we set $a = \$1,000$ and $b = -\$200$, meaning there is \$1,000 of social utility when an applicant receives and pays back a loan and $-\$200$ utility when an applicant defaults on a loan, for example because defaulting could trigger further financial distress. As above, we imagine the lender is a government institution or other agent attempting to maximize total social utility. We further assume that the pre-screening covariate $x_i \in [0, 1]$ specifies the lender’s pre-screening estimate of applicant i ’s likelihood to repay a loan; in other words, $x_i = \Pr(Y_i = 1 | X_i = x_i)$. The covariate x_i can thus be translated into (pre-screening) expected utility μ_i by Eq. (8):

$$\mathbb{E}[U_i | X_i = x_i] = ax_i + b(1 - x_i).$$

Synthetic data

We illustrate the value of screening in four regimes of low vs. high cost of information paired with low vs. high value of information. To do so, we created four synthetic datasets, each comprised of $n = 500$ individuals.

For all four datasets, we first evenly split the population into targeted and non-targeted groups. For each applicant in the targeted group, we generated their pre-screening probability of repayment x_i (or, equivalently, their observed pre-screening covariates) by independently drawing from a beta distribution with mean 0.5 and count parameter 50.¹ Similarly, for each applicant in the non-targeted group, we generated their pre-screening probability of repayment by independently drawing from a beta distribution with mean 0.70 and count parameter 50. The higher mean repayment probability for members of the non-targeted group corresponds to them being, on average, more creditworthy.

Now, for each applicant in the targeted group, the lender may elect to screen them. As a result of screening, the lender

¹In terms of the α and β shape parameters often used to parameterize beta distributions, the mean is $\alpha/(\alpha + \beta)$ and the count parameter is $\alpha + \beta$. A higher count corresponds to a lower variance.




Figure 2: Comparison of our optimal screening strategy (blue line) with one without screening (black line) in four regimes with different costs and values of information.

receives an improved estimate \tilde{x}_i of the applicant’s repayment probability, so that:

$$\mathbb{E}[U_i | X_i = x_i, \tilde{X}_i = \tilde{x}_i] = a\tilde{x}_i + b(1 - \tilde{x}_i).$$

We assume \tilde{x}_i is drawn from a beta distribution with mean x_i (i.e., the lender’s pre-screening estimate of the applicant’s repayment probability). In the high-information scenario, we set the count parameter for this beta distribution to be 5; and in the low-information scenario, we set it equal to 25.

Figure 1 shows these two post-screening information distributions for an applicant with $x_i = 0.4$. As illustrated in the plot, the high-information distribution (red line) has higher variance than the low-information distribution (blue line), and so screening is more likely to reveal very high risk and very low risk applicants in the high-information setting. Finally, we set the cost of screening c_S to be \$25 in the low-cost scenario and \$100 in the high-cost scenario, the cost of a loan to be $c_A = \$1,000$, and the total budget to be $B = \$50,000$.

Given these four datasets, we computed the optimal screening and allocation strategies using the algorithm described above.² In our original problem formulation, the diversity constraint in Eq. (4) specified only that we lower bound the utility of each group. To better understand the impact of diversity on utility, we modify this constraint to be a strict equality for the utility of the targeted group and set the lower bound on utility to be \$0 for the non-targeted group. Thus, across a range of exactly satisfied utilities for the targeted group, we find the strategy that maximizes overall utility.

Figure 2 shows the results of this analysis, with the blue lines tracing out the Pareto frontiers for each of the four sce-

²Because only one group can be screened in our experiments, we use a faster variant of our optimization algorithm, described in Appendix B. For any fixed set of parameters, this approach returns the optimal screening and allocation policy within a few seconds with the open-source LP solver SCS [23].

narios we consider. For comparison, the black lines show the corresponding result under a strategy that does not screen any applicants. Specifically, for any fixed utility constraint on the targeted group, the optimal no-screening policy first allocates loans to the k individuals in the targeted group most likely to repay based on their pre-screening estimates x_i , where k is chosen to satisfy the utility constraint; and then any remaining budget is used to allocate loans to those in the non-targeted group most likely to repay.

When either the value of information is high (left column) or the cost of screening is low (bottom row), we find that screening can be a valuable tool to improve utility. Notably, screening plays a more important role in these examples as we demand greater utility be allocated to the targeted group, since the no-screening strategy ends up dispersing loans to relatively high-risk applicants in the targeted group even though more creditworthy applicants in that group could be identified for little marginal cost. As one might expect, the gap between the screening and no-screening strategies is particularly large when both information is valuable and cheap. Indeed, in the high-value, low-cost setting (lower-left panel), one can achieve substantial diversity with little drop in overall utility.

Empirical credit data

We next apply our approach to the German Credit Dataset [16], which includes a variety of individual-level socioeconomic and financial characteristics (e.g., age, employment status, and credit history) on a sample of $n = 1,000$ people, of whom 700 are deemed creditworthy. We define the targeted group to be those who currently do not own their residence, a subpopulation that comprises 28% of the dataset. In this case, 60% of individuals in the targeted group are creditworthy compared to 74% in the non-targeted group. For members of the targeted group, we assume the lender, prior to screening, only knows the targeted group’s overall base rate of creditworthiness. If, however, the lender chooses to screen an applicant in the targeted group, they learn \hat{x}_i , the applicant’s likelihood of being creditworthy conditional on all the available features in the dataset. For members of the non-targeted group, we assume this full estimate of creditworthiness is available prior to screening, and that there is no opportunity to obtain additional information.³ Finally, we translate estimates of creditworthiness to estimates of utility via Eq. (8), in line with our simulations above.

Figure 3 shows the result of applying our screening and allocation algorithm to this dataset, where we assume the cost of screening c_S is \$100 (equal to our high cost regime in the synthetic datasets), the cost of a loan c_A is \$1,000, and the total budget is $B = \$150,000$. Like before, we compare the Pareto frontier of our approach (blue line) to that




Figure 3: For the German Credit Dataset, comparison of our optimal screening strategy (blue line) with one without screening (black line).

of a naive policy in which the lender does not screen applicants (black line). As with the synthetic datasets above, we find that the optimal policies with screening substantially outperform those without screening, particularly when we enforce a diversity constraint. For example, when we require \$50,000 of utility to come from allocating loans to the targeted group, the maximum total utility under the no-screening policy is \$102,000, compared to \$119,000 under the screening policy, an increase of 17%.

6 Discussion

Many creditworthy individuals often have difficulty gaining access to traditional credit markets due to lack of formal financial histories, an issue that can exacerbate existing socioeconomic disparities. To address this gap, we developed a simple and efficient algorithm for a budget-constrained decision maker to screen applicants and then allocate a limited resource, an approach that we find offers substantial benefits on both real and synthetic lending datasets. This joint screening-plus-allocation approach is especially useful in settings where a targeted subset of the population—those most in need of an intervention—are also those for whom the least information is available *a priori*, a common situation in many social welfare programs.

Past research has shown that a dearth of high-quality data for various subgroups of the population can lead to poor models in a variety of domains [14, 24], including text analysis [4, 7, 13], facial recognition [6], and automated hiring [12]. Looking forward, our combined data acquisition and decision-making approach provides one framework to address this challenge by jointly modeling the cost of data collection and the value for subsequent improvements in downstream decisions.

³More specifically, at the start of this exercise, we train a logistic regression model on the full dataset predicting creditworthiness as a function of the available covariates. Then, for members of the targeted group, \hat{x}_i is the model-estimated probability of creditworthiness for applicant i ; and for the members of the non-targeted group, that same model estimate is available prior to screening.

References

- [1] Apaam et al. 2018. 2017 FDIC National survey of unbanked and underbanked households. *US Federal Deposit Insurance Corporation*.
- [2] Bakker, M. A.; Noriega-Campero, A.; Tu, D. P.; Sattigeri, P.; and Varshney, K. R. 2019. On fairness in budget-constrained decision making. In *KDD 19: Explainable AI (XAI)*.
- [3] Bertsimas, D., and Tsitsiklis, J. N. 1997. *Introduction to linear optimization*, volume 6. Athena Scientific Belmont, MA.
- [4] Bolukbasi, T.; Chang, K.-W.; Zou, J. Y.; Saligrama, V.; and Kalai, A. T. 2016. Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In *Advances in Neural Information Processing Systems*, 4349–4357.
- [5] Brams, S. J., and Taylor, A. D. 1996. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press.
- [6] Buolamwini, J., and Gebru, T. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, 77–91.
- [7] Caliskan, A.; Bryson, J. J.; and Narayanan, A. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* 356(6334):183–186.
- [8] Chouldechova, A. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data* 5(2):153–163.
- [9] Corbett-Davies, S., and Goel, S. 2018. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023*.
- [10] Corbett-Davies, S.; Pierson, E.; Feller, A.; Goel, S.; and Huq, A. 2017. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 797–806. ACM.
- [11] Dantzig, G. B. 1957. Discrete-variable extremum problems. *Operations research* 5(2):266–288.
- [12] Dastin, J. 2018. Amazon scraps secret ai recruiting tool that showed bias against women.
- [13] Garg, N.; Schiebinger, L.; Jurafsky, D.; and Zou, J. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences* 115(16):E3635–E3644.
- [14] Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Daumeé III, H.; and Crawford, K. 2018. Datasheets for datasets. *arXiv preprint arXiv:1803.09010*.
- [15] Hardt, M.; Price, E.; and Srebro, N. 2016. Equality of opportunity in supervised learning. In *Advances In Neural Information Processing Systems*, 3315–3323.
- [16] Hofmann, H. 1994. Statlog (German Credit Data) data set. *UCI Repository of Machine Learning Databases*.
- [17] Kleinberg, J.; Mullainathan, S.; and Raghavan, M. 2017. Inherent trade-offs in the fair determination of risk scores. In *Proceedings of Innovations in Theoretical Computer Science (ITCS)*.
- [18] Liu, L. T.; Dean, S.; Rolf, E.; Simchowitz, M.; and Hardt, M. 2018. Delayed impact of fair machine learning. *arXiv preprint arXiv:1803.04383*.
- [19] Melville, P.; Saar-Tsechansky, M.; Provost, F.; and Mooney, R. 2004. Active feature-value acquisition for classifier induction. In *Fourth IEEE International Conference on Data Mining (ICDM'04)*, 483–486. IEEE.
- [20] Melville, P.; Saar-Tsechansky, M.; Provost, F.; and Mooney, R. 2005. An expected utility approach to active feature-value acquisition. In *Fifth IEEE International Conference on Data Mining (ICDM'05)*, 4–pp. IEEE.
- [21] Moulin, H. 2004. *Fair division and collective welfare*. MIT press.
- [22] Noriega-Campero, A.; Bakker, M. A.; Garcia-Bulle, B.; and Pentland, A. 2019. Active fairness in algorithmic decision making. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 77–83. ACM.
- [23] O'Donoghue, B.; Chu, E.; Parikh, N.; and Boyd, S. 2016. Conic optimization via operator splitting and homogeneous self-dual embedding. *Journal of Optimization Theory and Applications* 169(3):1042–1068.
- [24] Olteanu, A.; Castillo, C.; Diaz, F.; and Kiciman, E. 2019. Social data: Biases, methodological pitfalls, and ethical boundaries. *Frontiers in Big Data* 2:13.
- [25] Saar-Tsechansky, M.; Melville, P.; and Provost, F. 2009. Active feature-value acquisition. *Management Science* 55(4):664–684.
- [26] Thomson, W. 2011. Fair allocation rules. In *Handbook of social choice and welfare*, volume 2. Elsevier. 393–506.

A Proofs

We prove Theorem 2 in a more general setting in which we allow the cost of screening c_S and cost of allocation c_A to vary for each individual. In the special case of no screening (but with applicant-specific costs of allocation), we note that our optimization problem is equivalent to the fractional knapsack problem. For fractional knapsack, the greedy strategy—in which one packs items in descending order of their value per weight—yields an optimal solution [11]. In our case, we show that the same approach can be used to find an optimal allocation of loans once the set of applicants to screen has been appropriately selected.

To start, we generalize our definition of a threshold policy to account for the applicant-specific allocation costs.

Definition 3 (Cost-Aware Threshold Policy). A *cost-aware threshold policy* is an allocation policy \mathcal{A} for some fixed $t_j \in \mathbb{R} \cup \{-\infty, \infty\}$ and $\alpha_j \in [0, 1]$, $j = 1, \dots, m$, such that

$$A_i = \begin{cases} 1 & \hat{U}_i/c_i > t_{g_i}, \\ \alpha_{g_i} & \hat{U}_i/c_i = t_{g_i}, \end{cases}$$

where g_i denotes the group membership of individual i and c_i denotes the cost of allocating resources to individual i .

We first analyze the setting of a single group, in which case our threshold policy will have a single threshold, t . Lemma 4 shows that in this case cost-aware threshold policies are non-dominated: there cannot be an allocation policy with both higher expected utility and lower cost than a single-threshold policy. Lemma 5 is an existence theorem for cost-aware threshold policies, which shows that if a particular expected cost or utility can be achieved by an allocation policy, it can be achieved by a threshold policy.

Lemma 4. Let \mathcal{T} be a cost-aware threshold policy with a single threshold $t > 0$, and suppose \mathcal{T} has an expected cost C —i.e., $\sum_i \mathbb{E}[c_i \cdot T_i] = C$ —and expected utility Υ —i.e., $\sum_i \mathbb{E}[U_i \cdot T_i] = \Upsilon$. Let \mathcal{A} be any other allocation policy.

Case 1: If the expected cost of \mathcal{A} is C , then the expected utility of \mathcal{A} is less than or equal to Υ .

Case 2: If the expected utility of \mathcal{A} is Υ , then the expected cost of \mathcal{A} is greater than or equal to C .

Proof. Let $\Delta_i = T_i - A_i$, and set Δ_i^- be the negative part of Δ_i , i.e., $\Delta_i \cdot \mathbb{1}_{\Delta_i \leq 0}$. Similarly, let Δ_i^+ be the positive part.

The variables Δ_i^+ represent circumstances in which \mathcal{A} “saves money” in expectation over \mathcal{T} , and Δ_i^- represents how this portion of the budget is reallocated.

Consider the following events:

$$E_i^> = \{\hat{U}_i > c_i \cdot t\},$$

$$E_i^= = \{\hat{U}_i = c_i \cdot t\},$$

$$E_i^< = \{\hat{U}_i < c_i \cdot t\}.$$

Note that $\Delta_i^+ \neq 0$ only on some subset of $E_i^> \cup E_i^=$. likewise, $\Delta_i^- \neq 0$ only on a subset of $E_i^< \cup E_i^=$. Informally, \mathcal{A} can only save in expected cost over \mathcal{T} when $\hat{U}_i \geq t$, and can only reallocate savings to circumstances in which $\hat{U}_i \leq t$.

Because $\Delta_i^+ \geq 0$ and $\Delta_i^- \leq 0$, it consequently follows that

$$\Delta_i^+ \cdot \hat{U}_i \geq \Delta_i^+ \cdot c_i \cdot t, \quad (9)$$

and

$$\Delta_i^- \cdot \hat{U}_i \geq \Delta_i^- \cdot c_i \cdot t. \quad (10)$$

In particular,

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^n \Delta_i \hat{U}_i \right] &= \sum_{i=1}^n \mathbb{E}[\Delta_i \hat{U}_i] \\ &= \sum_{i=1}^n \left(\Pr(E_i^>) \cdot \mathbb{E}[\Delta_i^+ \cdot \hat{U}_i \mid E_i^>] \right. \\ &\quad \left. + \Pr(E_i^=) \cdot \mathbb{E}[\Delta_i \cdot \hat{U}_i \mid E_i^=] \right. \\ &\quad \left. + \Pr(E_i^<) \cdot \mathbb{E}[\Delta_i^- \cdot \hat{U}_i \mid E_i^<] \right) \\ &\geq \sum_{i=1}^n \left(\Pr(E_i^>) \cdot \mathbb{E}[\Delta_i^+ \cdot c_i \cdot t \mid E_i^>] \right. \\ &\quad \left. + \Pr(E_i^=) \cdot \mathbb{E}[\Delta_i \cdot c_i \cdot t \mid E_i^=] \right. \\ &\quad \left. + \Pr(E_i^<) \cdot \mathbb{E}[\Delta_i^- \cdot c_i \cdot t \mid E_i^<] \right) \\ &= t \cdot \mathbb{E} \left[\sum_{i=1}^n c_i \Delta_i \right]. \end{aligned}$$

Here the inequality follows from Eqs. (9) and (10).

Now, since any allocation policy is by definition conditionally independent of U_i given \hat{U}_i , it follows that $\mathbb{E}[A_i U_i] = \mathbb{E}[A_i \hat{U}_i]$ and $\mathbb{E}[T_i U_i] = \mathbb{E}[T_i \hat{U}_i]$. In consequence,

$$\mathbb{E} \left[\sum_{i=1}^n \Delta_i U_i \right] \geq t \cdot \mathbb{E} \left[\sum_{i=1}^n c_i \Delta_i \right]. \quad (11)$$

In Case 1, the expected costs of \mathcal{A} and \mathcal{T} are equal. It follows that $\mathbb{E}[\sum_i c_i \Delta_i] = 0$; consequently, by Eq. (11), $\mathbb{E}[\sum_i \Delta_i U_i] \geq 0$, and so the expected utility of \mathcal{T} is greater than or equal to that of \mathcal{A} .

In Case 2, the expected utilities of \mathcal{A} and \mathcal{T} are equal, so that $\mathbb{E}[\sum_i \Delta_i U_i] = 0$. Then, again by Eq. (11), since $t > 0$, $\mathbb{E}[\sum_i c_i \Delta_i] \leq 0$. Therefore the expected cost of \mathcal{T} is less than or equal to that of \mathcal{A} . \square

Lemma 5. Let $\Upsilon > 0$ be an achievable expected utility—that is, there is an allocation policy \mathcal{A} such that $\mathbb{E}[\sum_i A_i U_i] = \Upsilon$ —and let C be an achievable expected cost. Then,

- There exists a cost-aware threshold policy \mathcal{T} of expected utility Υ ; and,
- There exists a (possibly different) cost-aware threshold policy \mathcal{T} of expected cost C .

Proof. We first address the case of a given positive expected utility. Consider the function $\Upsilon(t) : [0, \infty) \rightarrow \mathbb{R}$ given by

$$\Upsilon(t) = \sum_{i=1}^n \mathbb{E} \left[U_i \cdot \mathbb{1}_{\hat{U}_i/c_i \geq t} \right]. \quad (12)$$

We note three facts about $\Upsilon(t)$. First, $\Upsilon(t)$ is monotonically non-increasing for $t \geq 0$. Second, $\Upsilon(0)$ is the maximum expected utility achievable by any allocation policy. Third, $\Upsilon(t) \rightarrow 0$ as $t \rightarrow \infty$.

Suppose $0 < v \leq \Upsilon(0)$. Then there exists some maximal t such that $\Upsilon(t) \geq v$. Let $\rho = \sum_i \Pr(\hat{U}_i = t)$ and $\delta = \Upsilon(t) - v$. Define $\alpha = \delta/(t \cdot \rho)$. Then the threshold policy defined by (t, α) will have expected utility v . The result follows as all achievable positive expected utilities lie in the half-open interval $(0, \Upsilon(0)]$.

The proof in the case of expected cost is similar. \square

We are now equipped to move to analyze settings in which there are multiple groups. We prove the following generalization of Theorem 2.

Theorem 6. *Suppose the constrained optimization problem defined by Eqs. (2), (3), and (4) has a solution $(\mathcal{P}^*, \mathcal{A}^*)$. Then there is a (not necessarily single-threshold) cost-aware threshold policy \mathcal{T}^* such that $(\mathcal{P}^*, \mathcal{T}^*)$ is also a solution.*

Proof of Theorem 2. Consider a solution screening and allocation policy pair $(\mathcal{P}^*, \mathcal{A}^*)$ satisfying the allocation diversity constraints.

Note that each of the collections $\mathcal{A}_j^* = \{A_i^*\}_{i \in G_j}$ defines an allocation policy on the subpopulation G_j . Applying Lemma 5 to each of them yields a threshold policy \mathcal{T}_j^* —defined by the pair (t_j, α_j) —of the same expected cost as \mathcal{A}_j^* . By Lemma 4, \mathcal{T}_j^* achieves the same or higher expected utility on G_j as \mathcal{A}_j^* .

Let \mathcal{T}^* be the threshold policy defined by t_1, \dots, t_m and $\alpha_1, \dots, \alpha_m$. The expected cost of \mathcal{T}^* is the same as \mathcal{A}^* . Moreover, \mathcal{T}^* satisfies the group utility constraints, since the utility it achieves on each group G_j is greater than or equal to \mathcal{A}^* . Lastly, the expected utility achieved by \mathcal{T}^* is greater than or equal to that of \mathcal{A}^* .

Since \mathcal{A}^* was assumed optimal, it must be that the expected utilities are equal. Therefore the pair $(\mathcal{P}^*, \mathcal{T}^*)$ is also a solution to the constrained optimization problem defined by Eqs. (2), (3), and (4). \square

Theorem 2 follows as an immediate corollary.

In some cases, it is useful to impose equality rather than inequality diversity allocation constraints. For instance, in some situations the decision maker may wish to obtain the Pareto frontier. Lemma 7 ensures that it is still sufficient when specific utilities must be achieved on subgroups to restrict to the set of threshold policies when searching for a globally optimal screening and allocation policy.

Lemma 7. *The conclusion of Theorem 2 still holds even if the inequalities in Eq. (4) are replaced with equalities, i.e., if the optimal policy pair $(\mathcal{P}^*, \mathcal{A}^*)$ satisfies*

$$\mathbb{E} \left[\sum_{i \in G_j} U_i A_i \right] = \Lambda_j. \quad (13)$$

Proof. The proof is identical to that of Theorem 2, except that Lemma 5 is used to construct a threshold policy \mathcal{T}^* achieving exactly utility Λ_j on any constrained group G_j . Case 2 of Lemma 4 then shows that this utility is achieved

at possibly less expected cost than \mathcal{A}^* . It follows immediately that the pair $(\mathcal{P}^*, \mathcal{T}^*)$ is a solution to the optimization problem.⁴ \square

B Optimized Procedure

In our experiments we consider cases in which screening provides no additional information about one group of individuals (i.e., D_i is concentrated at a single point for $i \in G_j$). In this case, it is possible to significantly reduce the number of LPs the decision maker must solve to find an optimal policy or calculate the Pareto frontier. For instance, a lender may already possess all available information relevant to the creditworthiness of those with traditional credit histories.

In this case, the cost of allocating resources to group G_j does not depend on the screening policy \mathcal{P} . Therefore, rather than sweeping over all possible pairs of thresholds and boundary randomization probabilities (t_j, α_j) , one can encode the allocation policy to G_j directly as a collection of allocation probabilities $\{a_i\}_{i \in G_j}$.

For notational simplicity, suppose without loss of generality that $m = 2$, that $G_1 = \{1, \dots, n_1\}$ and $G_2 = \{n_1 + 1, \dots, n_1 + n_2\}$, and that D_i is a pointmass for all $i \in G_2$.

Then, let $p = (p_1, \dots, p_{n_1})$ be a length n_1 vector of probabilities indicating the probability of screening the members of G_1 . Let $a = (a_{n_1+1}, \dots, a_{n_1+n_2})$ be a length n_2 vector of probabilities indicating the probability of allocating resources to members of G_2 . We keep the rest of the notation the same as in Section 4.

Fix threshold t_1 and boundary randomization probability α_1 for G_1 . Then, the expected utility of any pair of screening policies (for G_1) and allocation policies (for G_2) can be expressed as a linear function of p and a :

$$\sum_{i \in G_1} q_i \cdot e_i \cdot p_i + o_i \cdot \mu_i \cdot (1 - p_i) + \sum_{i \in G_2} \mu_i \cdot a_i. \quad (14)$$

Likewise, our budget constraint can be expressed as follows:

$$\sum_{i \in G_1} c_S \cdot p_i + c_A \cdot q_i \cdot p_i + c_A \cdot o_i \cdot (1 - p_i) + \sum_{i \in G_2} c_A \cdot a_i \leq B. \quad (15)$$

The allocation diversity constraint will take one of the following two forms. If the constrained group is G_1 ,

$$\sum_{i \in G_1} q_i \cdot e_i \cdot p_i + o_i \cdot \mu_i \cdot (1 - p_i) = \Lambda_1. \quad (16)$$

Otherwise, if the constrained group is G_2 , the constraint will be

$$\sum_{i \in G_2} \mu_i \cdot a_i = \Lambda_2. \quad (17)$$

Lastly, we ensure that a and p are probability vectors:

$$a_i \leq 1, \quad (18)$$

$$p_i \leq 1. \quad (19)$$

⁴In this case, it is actually possible that \mathcal{T}^* will result in expected cost savings over \mathcal{A}^* , even though \mathcal{A}^* is optimal. This can occur if there is no remaining positive utility to “spend” the savings on, since resources are already allocated at total expected cost less than B in all cases where \hat{U}_i is positive.

Together Eqs. (14), (15), (16), and (17)—given the initial data of the threshold policy \mathcal{T}_1 on G_1 defined by the pair (t_1, α_1) —define a linear program in the $n_1 + n_2$ decision variables $p = (p_1, \dots, p_n)$ and $a = (a_{n_1+1}, \dots, a_{n_1+n_2})$ that can be solved for an optimal screening policy \mathcal{P}^* (on G_1) and allocation policy \mathcal{T}_2^* (on G_2). By the results of Appendix A, \mathcal{T}_2^* can be assumed to have the form of a threshold policy. Sweeping over all such pairs in a (discretized) space $\mathbb{R} \times [0, 1]$, the resulting triple $(\mathcal{P}^*, \mathcal{T}_1^*, \mathcal{T}_2^*)$ that maximizes utility will represent a globally optimal screening and allocation policy.