Multi Stage Screening: Enforcing Fairness and Maximizing Efficiency in a Pre-Existing Pipeline

Avrim Blum
Toyota Technological Institute at Chicago
Chicago, Illinois, USA
avrim@ttic.edu

Kevin Stangl
Toyota Technological Institute at Chicago
Chicago, Illinois, USA
kevin@ttic.edu

Ali Vakilian
Toyota Technological Institute at Chicago
Chicago, Illinois, USA
vakilian@ttic.edu

ABSTRACT

Consider an actor making selection decisions (e.g., hiring) using a series of classifiers, which we term a *sequential screening process*. The early stages (e.g. resume screen, coding screen, phone interview) filter out some of the applicants, and in the final stage an expensive but accurate test (e.g. a full interview) is applied to those individuals that make it to the final stage. Since the final stage is expensive, if there are multiple groups with different fractions of positives in them at the penultimate stage (even if a slight gap), then the firm may naturally only choose to apply the final (interview) stage solely to the highest precision group which would be clearly unfair to the other groups. Even if the firm is required to interview all those who pass to the final round, the tests themselves could have the property that qualified individuals from some groups pass more easily than qualified individuals from others.

Accordingly, we consider requiring Equality of Opportunity (qualified members of each group have the same chance of reaching the final stage and being interviewed). We then examine the goal of maximizing quantities of interest to the decision maker subject to this constraint, via modification of the probabilities of promotion through the screening process at each stage based on performance at the previous stage.

We exhibit algorithms for satisfying Equal Opportunity over the selection process and maximizing precision (the fraction of interviews that yield qualified candidates) as well as linear combinations of precision and recall (recall determines the number of applicants needed per hire) at the end of the final stage. We also present examples showing that the solution space is non-convex, which motivate our combinatorial exact and (FPTAS) approximation algorithms for maximizing the linear combination of precision and recall. Finally, we discuss the 'price of' adding additional restrictions, such as not allowing the decision-maker to use group membership in its decision process.

ACM Reference Format:
Avrim Blum, Kevin Stangl, and Ali Vakilian. 2022. Multi Stage Screening: Enforcing Fairness and Maximizing Efficiency in a Pre-Existing Pipeline. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22), June 21–24, 2022, Seoul, Republic of Korea.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3531146.3533178

FAccT '22, June 21–24, 2022, Seoul, Republic of Korea

ACM ISBN 978-1-4503-9352-2/22/06…$15.00
https://doi.org/10.1145/3531146.3533178

1 INTRODUCTION

Consider what we will term *sequential screening processes*. In this setting a decision maker (e.g. a company seeking to hire applicants) makes a decision, like hiring, by using a sequence of intermediate decision-making steps that each filter out some candidates, in order to ideally produce a pool of mostly qualified candidates at the final step.

We assume some people are truly qualified for the position being filled, and we call them positive examples, and others are truly unqualified and we call them negative examples. And then the various intermediate steps have different probabilities of qualified/unqualified applicants passing each step, which could be different for different demographic groups. We also assume that the final (interview) stage of the process is particularly expensive for the decision-maker, and reveals the true label of the applicant.

To illustrate a concern that could arise in this setting, suppose there are two demographic groups A and B, and just one test t in the screening process prior to the final stage. Suppose that test t and the underlying base rates of the two groups have the property that $\Pr[y = 1|t(x) = 1, x \in A] \geq \Pr[y = 1|t(x) = 1, x \in B] + \epsilon$ for some $\epsilon > 0$. That is, the pool of group-A applicants who pass the test has a higher fraction of positive examples than the pool of group-B applicants who pass the test. Since the cost of final interviews is assumed to be high, in this case a rational decision maker would be sensitive to even a small ϵ gap, in order to minimize the expected number of interviews made per hire. In particular, small gaps between these groups in the population would be amplified in that the rational decision-maker would then choose not to promote *any* individuals from group B to the final interview round, which clearly violates common sense fairness norms. There is empirical evidence that similar phenomenon occurs in real world settings, when employers have limited information [3].

A second concern is that even if the decision-maker interviews all individuals who make it to the final round (and more generally, at each level promotes all individuals who pass the test to the next round), the tests themselves could have the property that qualified individuals from some groups pass them more easily than qualified individuals from others. So, in the end, a qualified individual from one group might have a much lower chance of making it to the final interview round than a qualified individual from another.

Because of fairness violations of this kind, we consider a regulator that requires the screening process to satisfy Equal Opportunity [9], that is, qualified individuals of each group have the same

chance of receiving an interview. This requirement motivates the problem of how to satisfy such a condition in the most efficient way, minimizing the number of interviews needed per successful hire as well as the number of overall applicants needed to enter the screening process per hire. This is the question we address in our paper.

We assume that the tests themselves and their order in the process are fixed beforehand and the action space of the firm (of our algorithm) is solely modifying how individuals move through the pipeline in response to their test outcomes (the promotion policy). More specifically, for each test, we need to decide the probability that an individual from a given group who passes or fails the test should continue on to the next stage. One can satisfy the fairness requirement with simple promotion policies (such as promoting all individuals regardless of whether they pass or fail each test), but the tension is how to do so in a way that results in a useful process.

This captures the scenario of performing modifications to pre-existing screening systems (the test themselves are fixed) in order to respond to fairness issues. We assume we are given, for each test, its statistical properties for each group (the probability that a random qualified or unqualified individual will pass the test).[1]

1.1 Our Results

We study how to implement the fairness requirement of Equal Opportunity in this sequential screening setting and what method of implementing it would achieve a high efficiency. One core result in our paper is that there is a solution that maximizes precision (minimizes the number of interviews needed per successful hire) subject to maintaining Equal Opportunity, that is given by promoting individuals from each group according to what we call the *opportunity ratio*. Moreover, it is possible to maximize overall precision subject to satisfying Equal Opportunity by a policy in which each level in the process satisfies Equal Opportunity individually (this property will not hold for the more general objective below).

Then we consider the more general case of satisfying Equal Opportunity while maximizing a linear combination of precision and recall (1/precision is the expected number of interviews needed per successful hire, and 1/recall is proportional to the number of overall applicants needed to enter the screening process per hire). This problem is challenging because, as we show, the space of Equal-Opportunity solutions is non-convex. Moreover, the optimal way to use one test to optimize a linear combination of precision and recall may depend on all other available tests.

Nonetheless, we are able to achieve an FPTAS for maximizing any linear combination of precision and recall, as well as an exact algorithm with running time that is 'only' exponential in the number of levels k and the number of the groups. This latter result relies on certain structural properties of optimal solutions that we develop in our analysis. Finally, we discuss extensions to our model such as requiring the screening process to be group-blind, and considering the requirement of satisfying Equalized Odds. Unfortunately, the optimal fair group-blind policy may be much worse than the optimal fair group-aware policy. For example, in some cases it may require a policy that completely bypasses all the tests.

[1] If we were to design a socio-technical system from first principles using the insights of machine learning research, we might seek to design tests that are ideally more robust to group difference and still predictive, however such a re-design process could be costly and slow. In a world of limited resources, re-purposing pre-existing tests to be more fairness aware in a timely manner and still maintaining effectiveness is necessary.

1.2 Related Work

Fairness in pipelines was initiated by Bower et al. [4] and follow up work by Dwork and Ilvento [7], Dwork et al. [8]. This paper differs from [8] in several keys ways. We both use the word 'pipelines' but our work is more focused on the specific case of hiring pipelines in which we are looking at the fairness of the final outcome for a given individual, drawn from the population, rather than considering the individual fairness [6] of the cohort context to which one is assigned. We do not consider cohort based scoring rules.

The structure of our model is very close to that of Kannan et al. [11], but the objective in that work is jointly designing college admission and grading schemes that satisfy Equal Opportunity over the admissions/college process *and in particular* incentivize a rational employer to use a group blind hiring policy. In contrast, our work considers maximizing precision or a linear combination of recall and precision while satisfying Equal Opportunity.

Another related work by Arunachaleswaran et al. [2] is the idea of pipeline interventions. In that paper there is a wide pipeline with a finite number of states at time t and the goal of the algorithm designer is to modify the transition probabilities from state to state in order to maximize a reward at the final step. This corresponds to efficiently allocating a government subsidy to aid dis-advantaged individuals, from the perspective of maximizing social welfare.

Intriguingly, the paper by [12] argues that Equal Opportunity is misaligned with fairness in screening allocation problems with a finite number of available items (think hiring a small number of engineers at a start-up vs accepting applicants for a credit card). In our work, we do not focus on modeling a finite number of available positions (e.g., we are in the case with a larger number of available items).

Most closely related to our work is Cohen et al. [5], in which there is noisy Bernoulli feedback in a hiring setting with sequential tests. In contrast to our scenario, they assume both underlying candidate skill levels and test results are sampled independently from Bernoulli distributions. Furthermore, they allow hiring an applicant before the end of the pipeline (e.g., if you pass the first three of five tests and those tests have high signal, you may skip the next two tests). In our model, we assume each stage of the process is memoryless (the probability of making it to stage 3 from stage 2 depends only on the result of the stage-2 test and group membership, and not the result of the stage-1 test) and we allow tests to be asymmetric (e.g., it could be that positive examples from a given group pass with probability 0.75 and negative examples pass with probability 0.5). In our motivation, we model the initial tests as cheap while the ultimate interview is expensive and accurate, while in Cohen et al. [5], each test is equally accurate and costly and additionally they want to minimize the expected number of tests to hire a candidate. Consistent with our perspective, the authors exhibit an impossibility result arguing that satisfying Equal Opportunity requires group dependent thresholds if the tests have different noise rates.

Additionally, there are connections between our work and classical economic discussions of statistical discrimination [1, 13] in that both perspectives model disparities in outcomes that derive from strategic actors making decisions to allocate goods differently based on perceived differences in predicted outcomes (termed statistical discrimination). Our models do not capture taste based discrimination.

1.3 Roadmap

In Section 2 we formally describe our model and present some examples that show key phenomena. In Section 3 we prove and discuss our first main theorem, about how to maximize precision (at the end of the screening process) subject to Equal Opportunity.

Then we consider the more general case of satisfying Equal Opportunity while maximizing a linear combination of precision and recall. This problem is challenging because, as we show in Section 2.2, the space of Equal-Opportunity solutions is non-convex. Moreover, how to effectively utilize a test may depend on all other available tests (Section 4.1.2). On the other hand, as we show in Section 4.2, the solution space does satisfy certain useful structural properties. We then use these structural results to to achieve an exact optimal algorithm, and in Section 4.3 to achieve an FPTAS for maximizing linear combination of precision and recall, as well as other functions of precision and recall.

Finally, in Section 5 we discuss extensions to our model such as requiring the screening process to be group-blind, and considering the requirement of satisfying Equalized Odds.

2 PRELIMINARIES

Now we formally define our model and introduce some informative examples. As mentioned above, the scenario to keep in mind is a stylized hiring process, consisting of a sequence of tests or interviews. Each candidate takes a test, and depending on their outcome on that test at that stage, is possibly promoted to the next stage of the screening process. We focus on modifying this promotion policy in response to satisfying the fairness constraints and achieving a high objective value or a low cost value. This is a constrained optimization problem, with structure.

2.1 Definitions

We use $\mathcal{X}$ to denote the set of demographic groups, and $X \in \mathcal{X}$ to denote a specific group. We assume group membership is known to the algorithm, groups are disjoint, and an individual from group X is promoted based on both their test performance and a promotion policy (defined below) for that corresponding group. We assume individuals are either truly qualified or truly unqualified, and use label $y = 1$ to denote a truly-qualified individual and label $y = 0$ to denote a truly-unqualified individual. For each group X, let q_X denote the base rate for that group, namely $\Pr[y = 1 | x \in X]$.

Definition 2.1 (**Test Statistics**). *For each test t and each group $X \in \mathcal{X}$, we define $\tau_{X1} := \Pr[t(x, y) = 1 | y = 1, x \in X]$ to be the probability a qualified candidate from group X passes the test, and $\tau_{X0} := \Pr[t(x, y) = 1 | y = 0, x \in X]$ to be the probability an unqualified candidate from group X passes the test. We assume all tests are* minimally effective *for all groups in that positive examples are more likely to pass than negative examples. More precisely,*

$$\tau_{X1} > \tau_{X0} \geq 0 \quad \forall X \in \mathcal{X} \quad \textit{(Minimal Effectiveness Property)} \quad (1)$$

Note that we assume that the probability of an individual passing a given test depends only on their true qualification y and their group membership X. We also assume test statistics are given and known to our algorithm.

We use τ^j_{X1}, τ^j_{X0} to denote the test statistics at *stage j* of the interview process. For convenience, we define $T^j_X = (\tau^j_{X1}, \tau^j_{X0})$ as useful shorthand to capture the test statistics at stage j for group X. Note that the same test may have different effectiveness per group.

Definition 2.2 (**Post-Processing Modification**). *We would like to modify the outcomes of the tests in the screening process so that some fairness goal (to be specified later) is achieved at the end of the screening (i.e., in the final interview stage). Further, we assume as part of the problem setting that the only 'allowed' correction is to modify how candidates are promoted to the next stage. The promotion probability of each candidate only depends on their group membership and performance at the* current *test (whether they passed or failed the test). Formally, for each group $X \in \mathcal{X}$, let π^j_{X1} denote the probability a candidate $x \in X$ who passes the test at stage j is promoted to stage $j + 1$, and π^j_{X0} the probability that a candidate who fails the test at stage j is promoted to stage $j + 1$.*[2] *We describe a policy for a given stage j as $\{(\pi^j_{X1}, \pi^j_{X0})\}_{X \in \mathcal{X}}$.*

For instance, a naive fairness respecting solution is to simply ignore the tests and promote all examples to the end of the pipeline, i.e., $\{(\pi^j_{X1} = 1, \pi^j_{X0} = 1)\}_{X \in \mathcal{X}, j \in [k]}$ where k is the number of tests in this screening process. However, this would result in a useless process from the perspective of the decision maker. The most straightforward use of tests is to promote all who pass and none who fail, i.e., $\{(\pi^j_{X1} = 1, \pi^j_{X0} = 0)\}_{X \in \mathcal{X}, j \in [k]}$. However, this might not satisfy required fairness properties. We now formally describe the fairness properties we consider.

Definition 2.3 (**Equal Opportunity and Equalized Odds [9]**). *Our paper primarily discusses two fairness notions, specifically* Equal Opportunity *and* Equalized Odds. *The first notion, Equal Opportunity requires that the classifier have equal* True Positive Rates *for each group in the population. Equivalently, for a classifier h and true labels y, $\Pr[h(x) = 1 | y(x) = 1, x \in A] = \Pr[h(x) = 1 | y(x) = 1, x \in B]$. Equalized Odds is similar but it also requires that the* False Positive Rates *are equal; formally, $\Pr[h(x) = 1 | y(x) = 0, x \in A] = \Pr[h(x) = 1 | y(x) = 0, x \in B]$.*

In our problem, Equal Opportunity is motivated by a desire that qualified individuals should have the same shot at an interview regardless of their group membership. In our problem, there is additionally a critical distinction between the fairness criteria (e.g. Equal Opportunity or Equalized Odds) being satisfied at the end pipeline and alternatively that requiring these criteria hold for every transition between stages as individuals move through the pipeline, a stronger notion.

Now that we have described the terms that characterize a problem instance and the action space of the algorithm, we describe

[2] Note, in general randomized promotion policies will be necessary to satisfy the fairness criteria.

the objective value that captures the usefulness of a screening process. We term these multiple different objective functions 'pipeline efficiency'.

Definition 2.4 (**Pipeline Efficiency**). *In our work we focus on two core notions of efficacy from the perspective of the firm deploying the screening process.* Interview efficiency *(equivalently,* precision*) is the fraction of candidates in the last round who are qualified, i.e., the fraction of interviews that lead to hires (or at least to job offers).* Throughput efficiency *(equivalently,* recall*) is fraction of qualified candidates who make it to the final round, and determines the expected number of applicants needed to enter the pipeline to hire one candidate. In this paper, we study cost functions that are functions of these two quantities only.*

We model the last available test as highly discriminative but extremely expensive per each test utilization and this is what motivates the interview efficiency. In particular, if we assume that the k stages prior to the interview round have zero or negligible cost per test, and there are many available candidates, then we presume that the goal of the firm is to maximize the interview efficiency (precision, at the final round).

2.2 Formal Problem Statement and Illustrative Examples

Now, we combine the above into a formal statement. Given a screening process/pipeline $\mathcal{P}$ with k stages, this pipeline consists of a collection of disjoint groups $\mathcal{X}$ and tests statistics $T_X = (T_X^1, T_X^2, \dots T_X^k)$ for every group $X \in \mathcal{X}$.

The goal of the algorithm designer is to exhibit a method to find promotion policies $\{(\pi_{X1}^j, \pi_{X0}^j)\}_{X\in\mathcal{X}, j\in[k]}$ denoted as π such that the overall policy satisfies the relevant fairness notion (either at the end of the screening process or at the end of each stage) and maximizes the given pipeline efficiency. Now we move into illustrative examples.

An illustrative one-stage example: Consider a one-stage pipeline with test parameters

$$((\tau_{A1}, \tau_{A0}), (\tau_{B1}, \tau_{B0})) = ((1, 0.5), (0.8, 0.5)).$$

Observe that the policy of promoting individuals if and only if they pass the test does not satisfy Equal Opportunity. Instead, two policies that satisfy Equal Opportunity are $P = ((\pi_{A1}, \pi_{A0}), (\pi_{B1}, \pi_{B0})) = ((0.8, 0), (1, 0))$ and policy $Q = ((1, 0), (1, 1))$. In words, the policy P would promote all individuals who passed the test from group B, but would only promote 80% of those from group A. This down-weighting of group A would suffice to satisfy Equal Opportunity. In contrast, policy Q promotes all individuals from group A who pass the test and promotes everyone from group B, regardless of their test score. In this example, P is the optimal Equal Opportunity policy with respect to precision.

The set of policies satisfying Equal Opportunity is not convex: Interestingly, for a two stage pipeline with two groups, the set of policies satisfying Equal Opportunity is not convex. Consider a pipeline with first level $T_A^1 = (3/4, 0)$ and $T_B^1 = (1/2, 1/2)$ and with second level $T_A^2 = (1/2, 1/2)$. and $T_B^2 = (3/4, 0)$. Consider policy P with $(P_A^1 = (1, 0), P_B^1 = (1, 1))$ and $(P_A^2 = (1, 1), P_B^2 = (1, 0))$. This policy has recall 3/4 for each group and therefore satisfies Equal Opportunity. Consider policy Q with parameters $(Q_A^1 = (1, 0), Q_B^1 = (1, 1/2))$ and $(Q_A^2 = (1, 1), Q_B^2 = (1, 1))$. This policy also has the recall of 3/4 for each group and therefore also satisfies Equal Opportunity. However, the average of these two policies denoted as π is $(\pi_A^1 = (1, 0), \pi_B^1 = (1, 3/4))$, while $(\pi_A^2 = (1, 1)\ , \pi_B^2 = (1, 1/2))$. The recall for group A is still $\frac{3}{4}$, while the recall for group B is $(\frac{1}{2} + \frac{1}{2} \cdot \frac{3}{4})(\frac{3}{4} + \frac{1}{4} \cdot \frac{1}{2}) = \frac{49}{64} \neq \frac{3}{4}$.

Thus this convex combination of policies does not satisfy Equal Opportunity and therefore the set of Equal Opportunity promotion policies is not convex.

Requiring Equalized Odds at each level can significantly harm performance: The above example also shows that requiring Equalized Odds at each level can significantly harm performance. Notice that policy P above satisfies Equalized Odds overall and has perfect precision and fairly high recall. However, the only way to satisfy Equalized Odds at each level is to completely bypass both tests, which would be much worse for precision.

Interestingly, as we show below, requiring *Equal Opportunity* at each level does *not* harm precision relative to requiring it for the pipeline as a whole (though it can hurt recall).

3 MAXIMIZING PRECISION SUBJECT TO EQUAL OPPORTUNITY

In this section, we exhibit a policy π that maximizes precision at the end of the screening process while satisfying Equal Opportunity over the entire process. To do this, we prove that the optimal method for this objective is given by promoting individuals from each group according to the *Opportunity Ratio* (which we will define shortly).

Definition 3.1. *For a test τ and associated promotion policy $\{(\pi_{X1}, \pi_{X0})\}_{X\in\mathcal{X}}$, define $M_{X,\tau,\pi} := (\tau_{X1}\pi_{X1} + (1 - \tau_{X1})\pi_{X0})$ and $N_{X,\tau,\pi} := (\tau_{X0}\pi_{X1} + (1 - \tau_{X0})\pi_{X0})$. Note that $M_{X,\tau,\pi}$ and $N_{X,\tau,\pi}$ are the probabilities that a positive and respectively a negative example from group X is promoted to the next level, and so will be important quantities for our analysis.*

Observation 3.1. *For any single-stage policy $\{(\pi_{X1}, \pi_{X0})\}_{X\in\mathcal{X}}$ satisfying Equal Opportunity for a test $\{(\tau_{X1}, \tau_{X0})\}_{X\in\mathcal{X}}$, there exists M such that $M_{X,\tau,\pi} = M$ for every $X \in \mathcal{X}$.*

Furthermore, for a k-stage screening process $\{\tau^i\}_{i\in[k]}$, a policy $\{(\pi_{X0}, \pi_{X1})\}_{X\in\mathcal{X}}$ is Equal Opportunity if there exists M such that $\Pi_{i=1}^k M_{X,\tau^i,\pi^i} = M$ for every group $X \in \mathcal{X}$.

Observation 3.2. *Recall that q_X denotes the base rate for group X, and let $u_X = 1 - q_X$. For a single-stage pipeline with test τ and promotion policy π, the interview efficiency (i.e., precision) is equal to*

$$\mathrm{IE}(q, u, \tau, \pi) := \frac{\sum_{X\in\mathcal{X}} q_X M_{X,\tau,\pi}}{\sum_{X\in\mathcal{X}} q_X M_{X,\tau,\pi} + u_X N_{X,\tau,\pi}}. \tag{2}$$

Similarly, when we consider the extension to a k-stage pipeline, the interview efficiency is equal to

$$\mathrm{IE}(q, u, \tau, \pi) := \frac{\sum_{X\in\mathcal{X}} q_X \prod_{i=1}^k M_{X,\tau^i,\pi^i}}{\sum_{X\in\mathcal{X}} q_X \prod_{i=1}^k M_{X,\tau^i,\pi^i} + u_X \prod_{i=1}^k N_{X,\tau^i,\pi^i}}. \tag{3}$$

Now, we formally define the policy given by the opportunity ratio as follows.

Definition 3.2 (**Opportunity Ratio Policy**). *Consider a screening process with k stages.*

For each $X \in \mathcal{X}$, let $\rho_X := \Pi_{j\in[k]}(\tau^j_{X^*1}/\tau^j_{X1})$, where $X^* = \text{argmin}_{X\in\mathcal{X}}\Pi_{j\in[k]}\tau^j_{X1}$. The Opportunity Ratio *policy, at the first stage for each $X \in \mathcal{X}$, promotes ρ_X fraction of those who pass the test and none of those who fail the test. For the remaining stages $(i = 2, 3, ..., k)$, the Opportunity Ratio policy fully trusts the result of the tests; a candidate is promoted to the next stage iff they pass the test at the current stage. Formally, for every $X \in \mathcal{X}, \pi^1_{X1} = \rho_X, \pi^1_{X0} = 0$ and $\pi^i_{X1} = 1, \pi^i_{X0} = 0, \forall i \geq 2$.*

In the rest of this section, we study the task of maximizing interview efficiency under different settings and fairness requirements.

3.1 Maximizing Interview Efficiency subject to Equal Opportunity at the Final Stage

As a warm-up, we start with the simplest setting where the screening process has only one test before the interview stage.

Theorem 3.3 (**Opportunity Ratio Policy Maximizes Precision for Single-Stage Process**). *Let $t = ((\tau_{A1}, \tau_{A0}), (\tau_{B1}, \tau_{B0}))$ be a test satisfying the Minimal Effectiveness Property. The maximum precision policy satisfying Equal Opportunity is the opportunity ratio policy. Moreover, for any group $X \in \mathcal{X}$, it is always sub-optimal to promote any candidates who failed the test (i.e., in any optimal policy, $\pi_{X0} = 0, \forall X \in \mathcal{X}$).*

Proof. First, for any policy π, we upper-bound the interview efficiency (i.e., precision) for a screening process with parameters q, u, τ. To bound the interview efficiency, for each $X \in \mathcal{X}$, we lower-bound the False Positive Rate $N_{X,\tau,\pi}$ in terms of the True Positive Rate $M_{X,\tau,\pi}$.

$$\begin{aligned} N_{X,\tau,\pi} &= \tau_{X0}\pi_{X1} + (1-\tau_{X0})\pi_{X0} \\ &= \tau_{X0}(\pi_{X1} - \pi_{X0}) + \pi_{X0} \\ &\geq \frac{\tau_{X0}}{\tau_{X1}}\big(\tau_{X1}(\pi_{X1} - \pi_{X0}) + \pi_{X0}\big) \qquad \triangleright \text{ by Eq. (1)} \\ &= \frac{\tau_{X0}}{\tau_{X1}} \cdot M_{X,\tau,\pi} \end{aligned} \tag{4}$$

By Equal Opportunity of π and employing Eq. (4) in the formula for the interview efficiency, Eq. (2),

$$\begin{aligned} \text{IE}(q,u,\tau,\pi) &= \frac{\sum_{X\in\mathcal{X}} q_X M_{X,\tau,\pi}}{\sum_{X\in\mathcal{X}} q_X M_{X,\tau,\pi} + u_X N_{X,\tau,\pi}} \\ &\leq \frac{\sum_{X\in\mathcal{X}} q_X M_{X,\tau,\pi}}{\sum_{X\in\mathcal{X}}(q_X + u_X \cdot \frac{\tau_{X0}}{\tau_{X1}}) M_{X,\tau,\pi}} \qquad \triangleright \text{ by Eq. (4)} \\ &= \frac{\sum_{X\in\mathcal{X}} q_X}{\sum_{X\in\mathcal{X}} q_X + u_X \cdot \frac{\tau_{X0}}{\tau_{X1}}} \quad \triangleright \forall X \in \mathcal{X}, M_{X,\tau,\pi} = M \end{aligned} \tag{5}$$

Note that the inequalities are tight when $\pi_{X0} = 0$ for all $X \in \mathcal{X}$.

Next, we show that the *opportunity ratio* policy satisfies Equal Opportunity and achieves the bound in Eq. (5). In the opportunity ratio policy π^*, only a $(\frac{\tau_{X^*1}}{\tau_{X1}})$-fraction of candidates in group X who pass the test t (picked uniformly at random) are promoted to the next stage. In other words, for any group $X \in \mathcal{X}$, we set $\pi^*_{X1} = \frac{\tau_{X^*1}}{\tau_{X1}}, \pi^*_{X0} = 0$. Then,

$$\begin{aligned} \text{IE}(q,u,\tau,\pi^*) &= \frac{\sum_{X\in\mathcal{X}} q_X M_{X,\tau,\pi^*}}{\sum_{X\in\mathcal{X}} q_X M_{X,\tau,\pi^*} + u_X N_{X,\tau,\pi^*}} \\ &= \frac{\sum_{X\in\mathcal{X}} q_X \tau_{X1}(\frac{\tau_{X^*1}}{\tau_{X1}})}{\sum_{X\in\mathcal{X}} q_X \tau_{X1}(\frac{\tau_{X^*1}}{\tau_{X1}}) + u_X \tau_{X0}(\frac{\tau_{X^*1}}{\tau_{X1}})} \\ &= \frac{\sum_{X\in\mathcal{X}} q_X}{\sum_{X\in\mathcal{X}}(q_X + u_X \cdot \frac{\tau_{X0}}{\tau_{X1}})} \end{aligned}$$

Hence, π^* is an equal opportunity policy with the maximum interview efficiency for any screening process with parameters q, u, τ, π. □

Remark 1. *Note that any policy π where for each $X \in \mathcal{X}$, $\pi_{X1} = \eta \cdot \pi^*_{X1}, \pi_{X0} = 0$ for a constant $\eta < 1$ also satisfies the Equal Opportunity and maximizes the interview efficiency objective (i.e., precision). However, π^* has a strictly higher* recall.

Next, we state our result for the general setting in which there are multiple stages and multiple groups in the screening process. The proof of the theorem is similar to the single test version and is deferred to Appendix A.

Theorem 3.4 (**Multi-Stage Screening Process**). *Consider a k-stage screening process whose all tests are minimally effective. The maximum interview efficiency policy satisfying Equal Opportunity is the Opportunity Ratio policy and has interview efficiency equal to $\frac{\|q\|_1}{\|q\|_1 + \sum_{X\in\mathcal{X}} u_X \Pi_{i=1}^k (\tau^i_{X0}/\tau^i_{X1})}$.*

3.2 Maximizing Interview Efficiency Subject to Equal Opportunity at the End of Each Stage

Here, we consider the setting in which the goal is find a policy that maximizes interview efficiency and satisfy Equal Opportunity at the end of each stage—*not only at the interview stage*. Following Theorem 3.4, the maximum interview efficiency in this setting is at most $\|q\|_1/(\|q\|_1 + \sum_{X\in\mathcal{X}} u_X \Pi_{i=1}^k \frac{\tau^i_{X1}}{\tau^i_{X0}})$. Next, we show that the following slightly modified opportunity ratio policy π that satisfies Equal Opportunity at the end of each stage maximizes the interview efficiency. The policy π applies the opportunity ratio at each stage of the pipeline.

$$\pi^i_{X0} = 0, \pi^i_{X1} = \frac{\tau^i_{X^*_i 1}}{\tau^i_{X1}} \quad \forall i \in [k], X \in \mathcal{X}, \text{ where } X^*_i := \text{argmin}_{X\in\mathcal{X}} \tau^i_{X1}$$

Again, it is straightforward to verify that π satisfies the Equality of Opportunity. Moreover,

$$\begin{aligned} \text{IE}(q,u,\tau,\pi) &= \frac{\sum_{X\in\mathcal{X}} q_X M_{X,\tau,\pi}}{\sum_{X\in\mathcal{X}} q_X M_{X,\tau,\pi} + u_X N_{X,\tau,\pi}} \\ &= \frac{\sum_{X\in\mathcal{X}} q_X \Pi_{i\in[k]} \tau^i_{X^*_i 1}}{\sum_{X\in\mathcal{X}} q_X \Pi_{i\in[k]} \tau^i_{X^*_i 1} + \sum_{X\in\mathcal{X}} u_X \frac{\tau^i_{X^*_i 1}\tau^i_{X0}}{\tau^i_{X1}}} \\ &= \frac{\|q\|_1}{\|q\|_1 + \sum_{X\in\mathcal{X}} u_X \Pi_{i=1}^k \frac{\tau^i_{X0}}{\tau^i_{X1}}} \end{aligned}$$

The only difference compared to the policy of Theorem 3.4 is that in the former policy the recall can be higher.

Remark 2. *Adding the condition to satisfy the Equality of Opportunity at the end of each stage does not harm interview efficiency. However, this condition may decrease the recall of the optimal policy.*

4 PIPELINE EFFICIENCY: MAXIMIZING LINEAR COMBINATIONS OF PRECISION AND RECALL

Now we shift our focus to exhibiting a promotion policy that satisfies Equal Opportunity and maximizes a linear combination of precision and recall given by the positive weight $\alpha \in \mathbb{R}_{\geq 0}$; $f_\alpha(\pi) := (1-\alpha) \cdot \text{recall}(\pi) + \alpha \cdot \text{precision}(\pi)$. As in Definition 2.4, higher precision corresponds to higher interview efficiency, and higher recall corresponds to higher throughput efficiency.

We start with a simple 2-approximation algorithm for maximizing any given linear of precision and recall.

Theorem 4.1 (**Approximation Algorithm for Linear Combination of Precision and Recall**). *There exists a polynomial time 2-approximation algorithm for maximizing any linear combination of precision and recall.*

Proof. Note that the policy that bypasses all tests is an Equal Opportunity policy and maximizes recall—it achieves recall equal to one. Moreover, by Theorem 3.4, the Opportunity Ratio is an Equal Opportunity policy maximizing precision. Hence, the better of the "bypassing all tests" policy and the Opportunity Ratio policy is a 2-approximation of any given linear combination of precision and recall. □

In order to obtain better performance for maximizing linear combinations of precision and recall, we develop structural properties of optimal solutions, and then use them to get an exact algorithm with running time that is exponential only in k and the number of groups. Additionally, by a dynamic programming approach we exhibit a *fully polynomial time approximation scheme (FPTAS).*

One challenge is that as shown in Section 2.2, the space of Equal Opportunity solutions is non-convex. Another is that as shown in Section 4.1.1 below, Opportunity Ratio is no longer optimal, and as shown in Section 4.1.2 below, there exists no function ranking the efficacy of tests solely based on their statistics.

We begin by presenting the examples mentioned above, and then developing the structural properties we will use.

4.1 Illustrative Examples

4.1.1 Opportunity Ratio not Optimal for Linear Combination of Precision and Recall. In the previous sections, our key algorithmic strategy is to use the Opportunity Ratio to re-weight the promotion policy. Since this policy satisfied Equal Opportunity and maximized precision (among Equal Opportunity policies), if our objective is to only maximize precision, then the Opportunity Ratio is sufficient. Now we exhibit an example where the Opportunity Ratio solution is not optimal when maximizing any linear combination of precision and recall when there is any nonzero weight on recall. Specifically, in this example there is an alternative policy with the same precision as the Opportunity Ratio solution but strictly higher recall.

Consider a pipeline with $T_A^1 = (3/4, 0)$ and $T_B^1 = (1/2, 1/4)$. In the second stage, $T_A^2 = (1/2, 1/4)$ and $T_B^2 = (3/4, 0)$. Consider policy P: ($P_A^1 = (1,0)$ and $P_B^1 = (1,1)$, while $P_A^2 = (1,1)$ and $P_B^2 = (1,0)$.

This policy has recall 3/4 and precision 1 for each group and therefore satisfies Equal Opportunity. Thus if our objective here is maximize the average of precision and recall, this policy has objective function value 7/8. In contrast, the Opportunity Ratio policy as given in Definition 3.2 is $P_A^1 = (1,0)$,$P_B^1 = (1,0)$ and $P_A^2 = (1,0)$, $P_B^2 = (1,0)$ which reduces our recall to $\frac{3}{4} \cdot \frac{1}{2} = \frac{3}{8}$ while to precision is still 1, for score of $\frac{11}{16}$. Clearly this is a lower objective function score than the first policy.

4.1.2 Optimal Policy Non-Locality for Linear Combination of Precision and Recall . Suppose we have one group in the population and want to optimize a linear combination of recall and precision. A baseline idea is whether we can solve this problem with a natural greedy algorithm that makes local decisions in a single pass of the test statistics [3].

We answer this question in the negative in by exhibiting an example pipeline with test statistics such that when two of three tests are available, using only the first test is strictly optimal, while when all three tests are available, the optimum is instead to use the other two tests and not the first test. This shows that an algorithm that maximizes a linear combination of precision and recall cannot simply assign separate scores to each test and then use only the highest-scoring tests. Our example is only for one group.

The counterexample is as follows. The base-rate in the population is $P(y=1) = 1/2$. Consider test $t_1 = (1/2, 0)$ and tests $t_2 = t_3 = (1-\delta, 1/2)$ where $\delta = \frac{1}{100}$. The objective function is $f(\pi) = \frac{1}{3} \cdot \text{recall}(\pi) + \frac{2}{3} \cdot \text{precision}(\pi)$. In the following, let $f(t_1)$ to denote the score of the policy that only promotes those who pass t_1 and bypasses all other tests while $f(t_2t_3)$ denotes bypassing t_1 and promoting individuals if and only if they pass tests t_2 and t_3. In the Appendix E we show while $f(t_1)$ is larger than any policy using t_1 and t_2 (possibly in fractions), $f(t_2t_3)$ is strictly larger than any policy using t_1, t_2 and t_3 (again, possibly in fractions).

4.2 An Exact Algorithm

In this section, we give an exact algorithm for maximizing any given linear combination of precision and recall subject to satisfying Equal Opportunity by the end of the screening process.

First we show that for any k-stage screening process over a population specified by a collection of groups $\mathcal{X}$, there exists a set of Equal Opportunity policies $\mathcal{P}_{k,\mathcal{X}}$ that *weakly Pareto dominate* (w.r.t. precision and recall) any policy satisfying Equal Opportunity. In particular, we show that each policy $\pi := (\pi^1, \cdots, \pi^k) \in \mathcal{P}_{k,\mathcal{X}}$ has the following structure, $(1 - \pi^i_{X1})\pi^i_{X0} = 0, \forall i \in [k], X \in \mathcal{X}$.

Definition 4.1 (**Pareto Dominant Policy**). *For a given screening process, a policy π* weakly *Pareto dominates a policy $\tilde{\pi}$ w.r.t. precision and recall iff,* $\text{recall}(\pi) \geq \text{recall}(\tilde{\pi})$ *and* $\text{precision}(\pi) \geq \text{precision}(\tilde{\pi})$. *Moreover, π* strictly *Pareto dominates $\tilde{\pi}$ if at least one of the above inequalities holds strictly.*

[3] In the related work by [5] the answer is in the affirmative, but their model is different and has uniform noise across true positives and true negatives.

Furthermore, a set of policies $\mathcal{P}$ weakly Pareto dominates a policy $\tilde{\pi}$ w.r.t. precision and recall iff there exists a policy $\pi \in \mathcal{P}$ that π weakly Pareto dominates $\tilde{\pi}$.

Lemma 4.2. *For any k-stage screening policy that satisfies the Minimal Effectiveness Property 1, the set of Equal Opportunity policies in $\mathcal{P} := \{\pi \in [0,1]^{2|\mathcal{X}|k} : (1-\pi^i_{X1})\pi^i_{X0} = 0, \forall X \in \mathcal{X}, i \in [k]\}$ weakly Pareto dominates all equal opportunity policies w.r.t. precision and recall.*

In other words, any equal opportunity policy violating $(1-\pi^i_{X1})\pi^i_{X0} = 0$ for a group $X \in \mathcal{X}$ and a stage $i \in [k]$ is weakly Pareto dominated by $\mathcal{P}$.

Proof. First, we show that in any policy π which is not strictly Pareto dominated (w.r.t. precision and recall), $\pi^i_{X1} > 0$ for every $X \in \mathcal{X}, i \in [k]$. Hence, we can only consider policies π where $\pi_{X1} > 0$ for all $X \in \mathcal{X}$. The proof of the following claim is deferred to Appendix B. □

Claim 4.3. *Consider a k-stage screening process whose tests satisfy the Minimal Effectiveness Property 1. In any optimal policy of this screening process that satisfies Equal Opportunity, for all $X \in \mathcal{X}$ and $i \in [k]$, $\pi^i_{X1} > 0$.*

Now, for the sake of contradiction, suppose that there exist a level $i \in [k]$ and a group $X \in \mathcal{X}$ such that $\pi^i_{X0} > 0$ and $\pi^i_{X1} < 1$. Note that w.l.o.g., we can assume that $\tau^i_{X1} < 1$; otherwise, by setting $\pi_{X0} = 0$, the recall of the policy does not decrease and the precision strictly increases. Hence, there exist $\epsilon_1, \epsilon_0 > 0$ such that $\tau^i_{X1}\epsilon_1 - (1-\tau^i_{X1})\epsilon_0 = 0$ where either $(\epsilon_1 = 1-\pi_{X1}, \epsilon_0 \le \pi_{X0})$ or $(\epsilon_1 \le 1-\pi_{X1}, \epsilon_0 = \pi_{X0})$.

We define a new policy $\tilde{\pi}$, which differs from π only in level i of group X, as follows: $\tilde{\pi}^i_{X1} = \pi^i_{X1} + \epsilon_1$ and $\tilde{\pi}^i_{X0} = \pi^i_{X0} - \epsilon_0$. Next, we show that $N_{X,\tau^i,\tilde{\pi}^i} < N_{X,\tau^i,\pi^i}$.

$$\begin{aligned}
N_{X,\tau^i,\tilde{\pi}^i} &= \tau^i_{X0}\tilde{\pi}^i_{X1} + (1-\tau^i_{X0})\tilde{\pi}^i_{X0} \\
&= \tau^i_{X0}(\pi^i_{X1}+\epsilon_1) + (1-\tau^i_{X0})(\pi^i_{X0}-\epsilon_0) \\
&= \tau^i_{X0}\pi^i_{X1} + (1-\tau^i_{X0})\pi^i_{X0} + (\tau^i_{X0}\epsilon_1 + \tau^i_{X0}\epsilon_0 - \epsilon_0) \\
&= \tau^i_{X0}\pi^i_{X1} + (1-\tau^i_{X0})\pi^i_{X0} \\
&\quad + (\tau^i_{X0}\epsilon_1 + \tau^i_{X0}\epsilon_0 - \tau^i_{X1}\epsilon_1 - \tau^i_{X1}\epsilon_0) \ \triangleright\ \epsilon_0 = \tau^i_{X1}(\epsilon_0+\epsilon_1) \\
&< \tau^i_{X0}\pi^i_{X1} + (1-\tau^i_{X0})\pi^i_{X0} \ \triangleright\ \tau^i_{X0} < \tau^i_{X1} \\
&= N_{X,\tau^i,\pi^i}
\end{aligned}$$

Further, since $\tau^i_{X1}\epsilon_1 - (1-\tau^i_{X1})\epsilon_0 = 0$, $\tilde{\pi}$ satisfies Equal Opportunity and has the same recall as π. Moreover, since $N_{X,\tau^i,\tilde{\pi}^i} < N_{X,\tau^i,\pi^i}$ and for all $j \in [k] \setminus \{i\}$, $N_{X,\tau^j,\pi^j} \ge 0$, $\Pi^k_{j=1} N_{X,\tau^j,\pi^j} \le \Pi^k_{j=1} N_{X,\tau^j,\tilde{\pi}^j}$. Hence the precision of $\tilde{\pi}$ is not less than the one of π. This contradicts the strict Pareto optimally of policy π. Thus the statement holds and for any level $i \in [k]$ and any group $X \in \mathcal{X}$, $(1-\pi^i_{X1})\pi^i_{X0} = 0$.

Next, we show additional structures of the set of Equal Opportunity policies $\mathcal{P}_{k,\mathcal{X}}$ that weakly Pareto dominates *all* Equal Opportunity policies.

Lemma 4.4. *Consider a k-stage screening process whose tests satisfy the Minimal Effectiveness Property 1. The set of Equal Opportunity policies $\mathcal{S} \subseteq \mathcal{P} = \{\pi \in [0,1]^{2|\mathcal{X}|k} : (1-\pi^i_{X1})\pi^i_{X0} = 0, \forall X \in \mathcal{X}, i \in [k]\}$ where for each group $X \in \mathcal{X}$, there exists at most one level $i \in [k]$ such that $0 < \pi^i_{X0} < 1$, weakly Pareto dominates all Equal Opportunity policies.*

In other words, any Equal Opportunity policy π of the screening process is weakly Pareto dominated by $\tilde{\pi} \in \mathcal{S}$ (in every policy $\tilde{\pi} \in \mathcal{S}$, for each group $X \in \mathcal{X}$, there exists at most one level i such that $0 < \tilde{\pi}^i_{X0} < 1$).

Proof. Suppose for contradiction that there exist a group $X \in \mathcal{X}$ and levels i, j such that $0 < \pi^i_{X0}, \pi^j_{X0} < 1$. Next, we show that we can modify π in levels i and j and replace π^i_{X0}, π^j_{X0} with $\tilde{\pi}^i_{X0}, \tilde{\pi}^j_{X0}$ such that

$$\begin{aligned}
M_{X,\tau^i,\pi^i} M_{X,\tau^j,\pi^j} &= (\tau^i_{X1} + \pi^i_{X0}(1-\tau^i_{X1}))(\tau^j_{X1} + \pi^j_{X0}(1-\tau^j_{X1})) \\
&= (\tau^i_{X1} + \tilde{\pi}^i_{X0}(1-\tau^i_{X1}))(\tau^j_{X1} + \tilde{\pi}^j_{X0}(1-\tau^j_{X1})) \\
&= M_{X,\tau^i,\tilde{\pi}^i} M_{X,\tau^j,\tilde{\pi}^j}, \qquad (6) \\
N_{X,\tau^i,\pi^i} N_{X,\tau^j,\pi^j} &= (\tau^i_{X0} + \pi^i_{X0}(1-\tau^i_{X0}))(\tau^j_{X0} + \pi^j_{X0}(1-\tau^j_{X0})) \\
&> (\tau^i_{X0} + \tilde{\pi}^i_{X0}(1-\tau^i_{X0}))(\tau^j_{X0} + \tilde{\pi}^j_{X0}(1-\tau^j_{X0})) \\
&= N_{X,\tau^i,\tilde{\pi}^i} N_{X,\tau^j,\tilde{\pi}^j} \qquad (7)
\end{aligned}$$

Note that Eq. (6) guarantees that the new policy $\tilde{\pi}$ satisfies Equal Opportunity and has the same recall as the policy π. Moreover, Eq. (7) shows that precision of the new policy is not less than than the precision of π. Next, we show that in the new policy, either $\tilde{\pi}^i_{X0} \in \{0,1\}$ or $\tilde{\pi}^j_{X0} \in \{0,1\}$.

Without loss of generality, we can assume that the feasible range of values for $\tilde{\pi}^i_{X0}$ to satisfy Equal Opportunity is $[\pi^i_{X0} - \epsilon^i, \pi^i_{X0} + \delta^i]$ which corresponds to $[\pi^j_{X0} - \delta^j, \pi^j_{X0} + \epsilon^j]$. Both intervals are sub-intervals of $[0,1]$ and since both $\tilde{\pi}^j_{X0}, \tilde{\pi}^i_{X0}$ belong to $[0,1]$, it is straightforward to verify that

$$(\pi^i_{X0} - \epsilon^i)(1 - (\pi^j_{X0} + \epsilon^j)) = (1 - (\pi^i_{X0} + \delta^i))(\pi^j_{X0} - \delta^j) = 0$$

Let $L = \frac{M_X}{\tau^i_{X1}\tau^j_{X1}}$ where $M_X = M_{X,\tau^i,\pi^i} M_{X,\tau^j,\pi^j} = M_{X,\tau^i,\tilde{\pi}^i} M_{X,\tau^j,\tilde{\pi}^j}$. By the Minimal Effectiveness Property, $1 < L < \frac{1}{\tau^i_{X1}\tau^j_{X1}}$. Then, satisfying Equal Opportunity is equivalent to satisfy the following constraint, $(1 + \tilde{\pi}^i_{X0}(\frac{1-\tau^i_{X1}}{\tau^i_{X1}}))(1 + \tilde{\pi}^j_{X0}(\frac{1-\tau^j_{X1}}{\tau^j_{X1}})) = L$. Hence, it implies that

$$\begin{aligned}
\tilde{\pi}^j_{X0} &= \left(\frac{L}{1 + \tilde{\pi}^i_{X0}(\frac{1-\tau^i_{X1}}{\tau^i_{X1}})} - 1\right) / \left(\frac{1-\tau^j_{X1}}{\tau^j_{X1}}\right) \\
&= \left(\frac{\tau^j_{X1}}{1-\tau^j_{X1}}\right)\left(\frac{L - 1 - \tilde{\pi}^i_{X0}(\frac{1-\tau^i_{X1}}{\tau^i_{X1}})}{1 + \tilde{\pi}^i_{X0}(\frac{1-\tau^i_{X1}}{\tau^i_{X1}})}\right)
\end{aligned}$$

Case 1: $\max(\tau^i_{X1}, \tau^j_{X1}) = 1$. Without loss of generality, suppose $\tau^i_{X1} = 1$. Then, we can simply set $\tilde{\pi}^i_{X0} = 0$ and the resulting policy $\tilde{\pi}$ will maintain Equal Opportunity. Moreover, since $1 - \tau^i_{X0} > 0$, $N_{X,\tau^i,\tilde{\pi}^i} \le N_{X,\tau^i,\pi^i}$. In the other case, we can similarly set $\tilde{\pi}^j_{X0} = 0$.

Case 2: $\tau^i_{X1}, \tau^j_{X1} < 1$. The task of finding $\tilde{\pi}^i_{X0}$ is as follows: $\tilde{\pi}^i_{X0} = \text{argmin}_{y \in [\pi^i_{X0}-\epsilon^i, \pi^i_{X0}+\delta^i]} f(y)$ which is equal to

$$(\tau^i_{X0} + y(1-\tau^i_{X0}))(\tau^j_{X0} + (\frac{\tau^j_{X1}}{1-\tau^j_{X1}})(\frac{L - 1 - y(\frac{1-\tau^i_{X1}}{\tau^i_{X1}})}{1 + y(\frac{1-\tau^i_{X1}}{\tau^i_{X1}})})(1-\tau^j_{X0})).$$

Next, we show that for any

$$y \in [0,1], f''(y) = -\frac{2L(\tau^j_{X0}-1)(\frac{\tau^j_{X1}}{1-\tau^j_{X1}})(\frac{1-\tau^i_{X1}}{\tau^i_{X1}})(\frac{\tau^i_{X0}}{\tau^i_{X1}}-1)}{(1+(\frac{1-\tau^i_{X1}}{\tau^i_{X1}})y)^3} < 0$$

To prove it note that the minimal "effectiveness property" of the tests $\{\tau^i\}_{i\in[k]}$ (i.e., $\tau^i_{X1} > \tau^i_{X0} \geq 0, \forall X \in \mathcal{X}, i \in [k]$) implies that $\frac{\tau^i_{X0}}{\tau^i_{X1}} - 1 < 0$. Moreover since by our assumption $\tau^j_{X1}, \tau^i_{X1} < 1$, $f''(y) < 0$ for all values of $y \in [0,1]$. Since f is a concave function in $[\pi^i_{X0}-\epsilon^i, \pi^i_{X0}+\delta^i]$, the minimum value of f in this interval obtained in one of its endpoints. In other words, the maximum precision corresponds to the case either $\tilde{\pi}^i_{X0} \in \{0,1\}$ or $\tilde{\pi}^j_{X0} \in \{0,1\}$. □

Finally, we show that each group can only have at most one level that partially uses its corresponding test.

Lemma 4.5. *Consider a k-stage screening process whose tests satisfy the "minimal effectiveness" property. The set of Equal Opportunity policies $\mathcal{P}_{k,\mathcal{X}} \subset \mathcal{P} = \{\pi \in [0,1]^{2|\mathcal{X}|k} : (1-\pi^i_{X1})\pi^i_{X0} = 0, \forall X \in \mathcal{X}, i \in [k]\}$ where for each group $X \in \mathcal{X}$, there exists at most one level $i \in [k]$ such that $\pi^i_{X1} < 1$ or $0 < \pi^i_{X0} < 1$, weakly Pareto dominates all Equal Opportunity policies.*

The proof is similar to the proof of Lemma 4.4 and we defer it to Appendix B. The above lemma enforces a very restricted structure on the set $\mathcal{P}_{k,\mathcal{X}}$ of Equal Opportunity policies that *weakly Pareto dominate* all Equal Opportunity policies. To summarize, in each policy $\pi \in \mathcal{P}_{k,\mathcal{X}}$, for each group $X \in \mathcal{X}$, the restriction of π on X has the following properties

(1) There is at most one level $i^* \in [k]$ such that π *partially uses* the test τ^{i^*}; i.e., *either* $0 < \pi^{i^*}_{X1} < 1$ and $\pi^{i^*}_{X0} = 0$, *or* $\pi^{i^*}_{X1} = 1$ and $0 < \pi^{i^*}_{X0} < 1$.
(2) In any remaining level i, π^i either *bypasses* τ^i (i.e., $\pi^i_{X1}\pi^i_{X0} = 1$), or *fully exploits* τ^i (i.e., $\pi^i_{X1} = 1, \pi^i_{X0} = 0$).

Theorem 4.6 (**Exact Algorithms for Linear Combination of Precision and Recall**). *Given any linear objective function of form $f_\alpha(\pi) := \alpha \cdot \text{precision}(\pi) + (1-\alpha) \cdot \text{recall}(\pi)$, There exists an exact algorithm that runs in time $O(k^{|\mathcal{X}|} \cdot 2^{k|\mathcal{X}|})$ and finds an Equal Opportunity policy of the screening process with parameters $(q, u, \tau, \mathcal{X})$ that maximizes f_α.*

Proof. Using the aforementioned set $\mathcal{P}_{k,\mathcal{X}}$ of weakly Pareto optimal policies (w.r.t. precision and recall) that satisfy the Equality of Opportunity, we enumerate over all policies in $\mathcal{P}_{k,\mathcal{X}}$ as follows.

- For each group $X \in \mathcal{X}$, pick a level $i_X \in [k]$ (*i.e., $k^{|\mathcal{X}|}$ possible configurations*).
- Fix an "integral" policy π for the rest of levels in each group $X \in \mathcal{X}$,
  - In each group $X \in \mathcal{X}$, for each level $i \neq i_X$, we decide whether to fully use the test ($\pi^i_{A1} = 1, \pi^i_{A0} = 0$) or to bypass the test ($\pi^i_{X1} = \pi^i_{X0} = 1$) (*i.e., $2^{(k-1)|\mathcal{X}|}$ possible configurations*).
- For each $X \in \mathcal{X}, i_X \in [k]$, we fix the policy π^{i_X} partially as follows,
  - $(1-\pi^{i_X}_{X1})\pi^{i_X}_{X0} = 0, \forall X \in \mathcal{X}$ (*i.e., $2^{|\mathcal{X}|}$ possible configurations*).

In each of the policies π as constructed above, we set the remaining π values (i.e., π^{i_X}) so that Equality of Opportunity is satisfied and the objective function f_α is maximized. Finally, we maintain the configuration π that maximizes f_α. Note that the whole process takes $O(k^{|\mathcal{X}|} \cdot 2^{k|\mathcal{X}|})$ time. □

Similarly, we can show the following.

Theorem 4.7 (**Exact Algorithms for Linear Combination of reciprocal of Precision and Recall**). *Given any objective function $g_\alpha(\pi) := \alpha/\text{precision}(\pi) + (1-\alpha)/\text{recall}(\pi)$, There exists an exact algorithm that runs in time $O(k^{|\mathcal{X}|} \cdot 2^{k|\mathcal{X}|})$ and finds an Equal Opportunity policy of the screening process with parameters $(q, u, \tau, \mathcal{X})$ that minimizes g_α.*

Remark 3 (**General Objective Functions**). *Our approach provides an exact algorithm for maximizing (resp., minimizing) a given pipeline efficiency objective f (resp., pipeline complexity cost g) over Equal Opportunity policies if f (resp., g) satisfies the following natural condition: for any pair of policies π_1, π_2 where π_1 weakly Pareto dominates π_2 w.r.t. precision and recall, $f(\pi_1) \geq f(\pi_2)$ (resp., $g(\pi_1) \leq g(\pi_2)$).*

4.3 An FPTAS Algorithm

In this section, we present FPTAS algorithms for maximizing a given pipeline efficiency objective (resp., minimizing a given pipeline cost function) while satisfying the Equal Opportunity requirement. We consider two regimes. In this section, as in previous sections, we consider the regime where we are allowed to treat individuals from different groups differently; more precisely, we can set $\pi^j_{Xi} \neq \pi^j_{Yi}$ for $j \in [k], i \in \{0,1\}$. Next, in Section 5.1, we consider a new regime where the goal is to achieve Equal Opportunity while treating individuals from both groups similarly; $\forall i \in [k], X \neq Y \in \mathcal{X}, \pi^i_{X1} = \pi^i_{Y1}, \pi^i_{X0} = \pi^i_{Y0}$.

To exploit our algorithm in different settings, we describe it for the most basic setting of the problem. Given a single group of applicants with parameters q, u and a pipeline $\mathcal{P} = \{\tau^i\}_{i\in[k]}$, the goal is find a policy π that maximizes a given *pipeline efficiency objective* $f(\text{recall}(\pi, \mathcal{P}), \text{precision}(\pi, q, u, \mathcal{P}))$. Our approach works for a quite general set of objective functions; more notably, as in the previous section, for two natural settings: maximizing a *linear combination of precision and recall* and minimizing a *linear combination of reciprocals of precision and recall.*

High-level Description of Algorithm. Now we write a dynamic program (DP) to optimize a given pipeline efficiency objective f up to a given accuracy parameter ϵ. We create a DP-table $M[i, \text{tpr}, \text{fpr}]$ where $i \in [k]$, $\text{tpr} \in [0, \ell_{\text{tpr}} := \log_{1-\epsilon} L_{\text{tpr}}]$ and $\text{fpr} \in [0, \ell_{\text{fpr}} := \log_{1-\epsilon} L_{\text{fpr}}]$ where $L_{\text{tpr}}, L_{\text{fpr}}$ are lower bounds on *True Positive Rate* and *False Positive Rate* respectively. For each set of parameters

$(i, \text{tpr}, \text{fpr})$, $M[i, \text{tpr}, \text{fpr}]$ will be a Boolean value indicating whether there exists a policy such that by the end of level i, the true positive rate becomes at least $(1-\epsilon)^{\text{tpr}}$ and the False Positive Rate becomes at most $(1-\epsilon)^{\text{fpr}}$. Without loss of generality and for the simplicity of the exposition, we assume L_{tpr} and L_{fpr} are powers of $(1-\epsilon)$; otherwise we can simply round the lower bounds to largest powers of $(1-\epsilon)$ smaller than actual bounds.

Solving the DP. We fill out the DP table starting from $i = 1$ as follows. First, for any $j_0 \in [0, \ell_{\text{fpr}}], j_1 \in [0, \ell_{\text{tpr}}]$, $M[1, j_1, j_0] = \text{true}$ iff the following system of linear inequalities has a feasible solution.

$$\tau_0^1 x + (1-\tau_0^1)y \le (1-\epsilon)^{j_0}, \quad \tau_1^1 x + (1-\tau_1^1)y \ge (1-\epsilon)^{j_1}. \tag{8}$$

Next, we describe the update rule for $i > 1$. For any $\text{tpr} \in [0, \ell_{\text{tpr}}], \text{fpr} \in [0, \ell_{\text{fpr}}]$, $M[i+1, \text{tpr}, \text{fpr}] = \bigvee_{(j_0, j_1) \in \mathcal{F}_{i+1}} M[i, \text{tpr} - j_1, \text{fpr} - j_0]$, where $\mathcal{F}_{i+1}$ is a set of $(j_0 \le \text{fpr}, j_1 \le \text{tpr})$ for which the following linear program has a feasible solution,

$$\tau_0^{i+1} x + (1-\tau_0^{i+1})y \le (1-\epsilon)^{j_0}, \tau_1^{i+1} x + (1-\tau_1^{i+1})y \ge (1-\epsilon)^{j_1}. \tag{9}$$

Note that x, y can be interpreted as π_1^{i+1}, π_0^{i+1}, respectively. Moreover, the system of linear inequalities of the update rule in level $i+1$ (Eq. (9)) is similar to the rules for the base case (Eq. (8)).

Lemma 4.8. *For any $i \in [k]$, if there exists a policy π with True Positive Rate $\mathsf{t}_i \ge L_{\text{tpr}}/(1-\epsilon)^{i-1}$ and False Positive Rate f_i by the end of level i, then for any $j_1 \in [0, \ell_{\text{tpr}}], j_0 \in [0, \ell_{\text{fpr}}]$ with $(1-\epsilon)^{j_1} \ge \mathsf{t}_i \cdot (1-\epsilon)^{i-1}$ and $(1-\epsilon)^{j_0} \le \min\{1, \max\{L_{\text{fpr}}, \mathsf{f}_i\}/(1-\epsilon)^{i-1}\}$, $M[i, j_1, j_0] = \text{true}$.*

In other words, if the policy π exists then the DP approach finds a policy with true positive rate at least $(1-\epsilon)^{j_1}$ and false positive rate at most $(1-\epsilon)^{j_0}$.

The proof is deferred to Section B.

Lemma 4.9 (**DP Main Lemma**). *For any group $X \in \mathcal{X}$, an accuracy parameter ϵ and lower bounds on the false positive rate, L_{fpr}, and the true positive rate, L_{tpr}, if there exists a policy π^* with true positive rate $\mathsf{t} \ge L_{\text{tpr}}/(1-\epsilon)^{k-1}$ and false positive rate $\mathsf{f} > 0$, then the DP algorithm runs in time $O(\frac{k \log^2(1/L_{\text{tpr}}) \log^2(1/L_{\text{fpr}})}{\epsilon^4})$ and finds a policy π with true positive rate at least $(1-\epsilon)^{k-1} \cdot \mathsf{t}$ and false positive rate at most $\min\{1, \max\{L_{\text{fpr}}, \mathsf{f}\}/(1-\epsilon)^{k-1}\}$.*

Proof. The size of table is $O(k\ell_{\text{tpr}}\ell_{\text{fpr}})$ and updating each entry in the table takes $O(\ell_{\text{tpr}}\ell_{\text{fpr}})$. Hence, the total runtime to compute all entries in the DP table is $O(k\ell_{\text{tpr}}^2\ell_{\text{fpr}}^2) = O(\frac{k \log^2(1/L_{\text{tpr}}) \log^2(1/L_{\text{fpr}})}{\epsilon^4})$.

Now we apply the DP approach and by Lemma 4.8, the solution returned by the algorithm has the true positive rate and the false positive rate satisfying the guarantee of the statement. □

Implications of DP. Here we present FPTAS algorithms using the described DP approach in different settings. We state the results formally and their proofs are deferred to Appendix B.

Theorem 4.10 (**FPTAS for Linear Combination of Precision and Recall**). *Consider a k-stage screening process with parameters $(u, q, \tau, \mathcal{X})$ and for any policy π, let $f_\alpha(\pi) = (1-\alpha) \cdot \text{recall}(\pi) + \alpha \cdot \text{precision}(\pi)$ where $\alpha > 0$. Given an accuracy parameter ϵ, there exists an FPTAS that runs in time $O(\frac{|\mathcal{X}|k^5 \log^4(1/\epsilon)}{\epsilon^4})$ and finds an Equal Opportunity policy π such that $f_\alpha(\pi) \ge (1-\epsilon) f_\alpha(\pi^*)$ where π^* maximizes f_α over Equal Opportunity policies.*

Theorem 4.11 (**FPTAS for Linear Combination of Reciprocals Precision and Recall**). *Consider a k-stage screening process with parameters $(u, q, \tau, \mathcal{X})$ and for any policy π, let $g_\alpha(\pi) = (1-\alpha)/\text{recall}(\pi) + \alpha/\text{precision}(\pi)$ where $\alpha > 0$. Given an accuracy parameter ϵ, there exists an FPTAS that runs in time $O(\frac{|\mathcal{X}|k^7(\log^2 \frac{1}{\epsilon} + k^2)}{\epsilon^4})$ and finds an Equal Opportunity policy π such that $g_\alpha(\pi) \le (1+\epsilon) g_\alpha(\pi^*)$ where π^* minimizes g_α over Equal Opportunity policies.*

Remark 4 (**General Objective Functions**). *In Theorem 4.10 and 4.11 we presented FPTAS for finding Equal Opportunity policies optimizing two standard pipeline efficiency objective functions. Here, we generalize the above theorems when the pipeline efficiency objective function $f : [0,1]^2 \to \mathbb{R}$ which maps precision and recall to efficiency scores have certain properties. Also, we define $g : [0,1]^2 \to \mathbb{R}$ such that for any $\mathsf{t}, \mathsf{f} \in [0,1]^2$, $g(\mathsf{t}, \mathsf{f}) := f(\text{recall}(\mathsf{t}), \text{precision}(\mathsf{t}, \mathsf{f}))$. We describe the properties when the goal is to maximize f*—the required conditions for the minimization version is similar.

- *f is non-decreasing w.r.t. both precision and recall*—equivalently, g is non-decreasing in t and non-increasing in f.
- *There exist $L_{\text{tpr}}, L_{\text{fpr}} > 0$ such that there exists a $(1-\alpha)$-approximate solution of f with $\mathsf{t} \in (L_{\text{tpr}}, 1], \mathsf{f} \in (L_{\text{fpr}}, 1]$.*
- *The function f is β-Lipschitz on $\{(x, y) | x \in (L_{\text{tpr}}, 1], y \in (L_{\text{fpr}}, 1]\}$.*

In particular, the above properties are sufficient to show that the DP approach finds a $(1-\epsilon)$-approximation of f in time
$\text{poly}(k, |\mathcal{X}|, \epsilon^{-1}, \log(1/L_{\text{tpr}}), \log(1/L_{\text{fpr}}))$.

Remark 5 (**Selecting from Available Tests**). *Suppose that in contrast to our previous approaches, we do allow for the design of the pipeline in that we allow the firm to select some tests to create a pipeline. For instance, imagine that there is a budget and the firm is allocating this budget to buy tests. The goal of the firm is the same, e.g. to exhibit a pipeline satisfying a fairness requirement. Our algorithms can be modified to handle to this case by adding a term in the DP table corresponding to the budget remaining, with a decision point of choosing to use a given test or not. Note that the ordering of tests in the pipeline does not matter for the objectives considered.*

5 ALTERNATE MODELS

In this section we describe some alternate settings, such as using a single promotion policy for both demographic groups (which might be required by regulation), or requiring Equalized Odds.

5.1 Screening Processes with Same Policy for All Groups

One alternate fairness model is to additionally require the same policy be used for all groups. While utilizing demographic features can aid in achieving fairness goals (e.g. [6, 9]), in some regulatory regimes, this fairness-through-awareness may be illegal or problematic, even when intended to ensure equitable treatment.

In our setting, if we are constrained to follow *group-blindness*, there be would only one set of tests and one ordering of the tests that all applicants are tested on. Analogously to the previous setting,

the action space of the algorithm remains modifying the promotion probabilities, but we now only have one set of policies to modify. We also exhibit a DP algorithm for this setting, which we defer to Section D. However, a simple example shows the inefficiencies in this regime. Suppose we have a single test with $T_A = (1, 0)$ and $T_B = (1/2, 0)$. Observe that since we are constrained to use group blindness and satisfy Equal Opportunity, there is no way to use the test without violating Equal Opportunity. Thus, the only way to satisfy Equal Opportunity is to completely bypass the test.

5.2 Equalized Odds

Next, recall that the requirement of Equalized Odds mandates equal True Positive and False Positive rates for all groups. In the appendix, we show structural properties of an optimal promotion policy that satisfies Equalized Odds. However, we also note the interview efficiency cost (precision) of requiring Equalized Odds. In particular, the gap between the interview efficiency of π_{EOdd} and π_{EOpp} can be as large as $\frac{1}{q} - \epsilon$ for any arbitrary $\epsilon > 0$. See Theorem C.3 for details.

5.3 Discussion Comparing Equalized Odds and Equal Opportunity

From the perspective of a decision maker in the wild, how to interpret and operationalize these results? A robust take-away is that requiring Equalized Odds and Equal Opportunity have substantially different efficiency consequences. Based on our examples, it seems unlikely that Equalized Odds is effective in this model, especially when requiring Equalized Odds at each stage. In contrast, the fact that requiring Equal Opportunity at each stage is equivalent to requiring Equal Opportunity of the overall process with respect to interview efficiency may have benefits in ensuring public confidence in the model.

5.4 Intersectionality

A natural question is how to think when the demographic groups may have an arbitrarily overlapping structure. This suggests several open questions in our model, e.g. if a person is in groups A and B, then which test parameter τ_A or τ_B corresponds to that person? Perhaps a direction is to assign to that person an interpolation between these values. A naive approach is when there are k groups, to create 2^k new groups and 2^k test parameters corresponding to every possible group intersection. If k is small, this may be computationally feasible, but is not responsive when the relevant sub-groups/intersections may not be known apriori. Perhaps our model could be merged with multi-calibration notions [10].

6 CONCLUSIONS

In contrast to some fairness in machine learning work, we focus on post-processing fairness modifications, rather than thinking about the fairness problem in screening processes where tests can be designed from scratch. While we believe that the more a priori design approach will have substantial benefits in practice, our approach of modifying pre-existing tests, combined with a concrete (and simple to evaluate) fairness notion, Equal Opportunity, is closely aligned with real world circumstances and models, especially in short term and iterative improvements to models. In some settings, the firm making hiring decisions will outsource some aspects of its pipeline to third party companies and the tests will be a black box, but possibly that come with statistics that can be used in our algorithms. This decoupling allows the effective implementation of fairness aware promotion policies in the short term.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation under grants CCF-1815011 and CCF-1934843, and by the Simons Foundation under the Simons Collaboration on the Theory of Algorithmic Fairness. The authors also greatly appreciate the comments of the reviewers and the work of the program committee and editors.

REFERENCES

[1] Kenneth J. Arrow. 2015. *The Theory Of Discrmination.* Princeton University Press, 1–33. https://doi.org/doi:10.1515/9781400867066-003
[2] Eshwar Ram Arunachaleswaran, Sampath Kannan, Aaron Roth, and Juba Ziani. 2021. Pipeline Interventions. In *12th Innovations in Theoretical Computer Science Conference (ITCS).*
[3] Marianne Bertrand and Sendhil Mullainathan. 2004. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review* 94, 4 (2004), 991–1013.
[4] Amanda Bower, Sarah N. Kitchen, Laura Niss, Martin J. Strauss, Alexander Vargas, and Suresh Venkatasubramanian. 2017. Fair Pipelines. arXiv:1707.00391 [cs.CY]
[5] Lee Cohen, Zachary C. Lipton, and Yishay Mansour. 2020. Efficient candidate screening under multiple tests and implications for fairness. In *1st Symposium on Foundations of Responsible Computing (FORC).*
[6] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *3rd Innovations in Theoretical Computer Science Conference (ITCS).* 214–226.
[7] Cynthia Dwork and Christina Ilvento. 2019. Fairness Under Composition. In *10th Innovations in Theoretical Computer Science Conference (ITCS).*
[8] Cynthia Dwork, Christina Ilvento, and Meena Jagadeesan. 2020. Individual Fairness in Pipelines. In *1st Symposium on Foundations of Responsible Computing (FORC).*
[9] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. *Advances in Neural Information Processing Systems (NeurIPS)* 29 (2016), 3315–3323.
[10] Ursula Hebert-Johnson, Michael Kim, Omer Reingold, and Guy Rothblum. 2018. Multicalibration: Calibration for the (Computationally-Identifiable) Masses. In *Proceedings of the 35th International Conference on Machine Learning*, Vol. 80. 1939–1948.
[11] Sampath Kannan, Aaron Roth, and Juba Ziani. 2019. Downstream Effects of Affirmative Action. *ACM FAT** (2019).
[12] Mohammad Mahdi Khalili, Xueru Zhang, and Mahed Abroshan. 2021. Fair Sequential Selection Using Supervised Learning Models. *Advances in Neural Information Processing Systems* 34 (2021).
[13] Edmund S. Phelps. 1972. The Statistical Theory of Racism and Sexism. *The American Economic Review* 62, 4 (1972), 659–661.

A PROOFS FROM SECTION 3

Proof of Theorem 3.4. First, we show that for any $M \in (0, 1]$, any Equal Opportunity policy π_M with recall M has interview efficiency at most

$$\begin{aligned}\text{IE}(q, u, \tau, \pi_M) &= \frac{\sum_{X \in \mathcal{X}} q_X M_{X,\tau,\pi_M}}{\sum_{X \in \mathcal{X}} q_X M_{X,\tau,\pi_M} + u_X N_{X,\tau,\pi_M}} \\ &= \frac{\|q\|_1}{\|q\|_1 + \sum_{X \in \mathcal{X}} u_X \frac{N_{X,\tau,\pi_M}}{M}} \\ &\leq \frac{\|q\|_1}{\|q\|_1 + \sum_{X \in \mathcal{X}} u_X \Pi_{i=1} \frac{\tau^i_{X0}}{\tau^i_{X1}}}, \end{aligned} \tag{10}$$

where the last inequality follows from the minimal effectiveness of tests in the screening process and an argument identical to Eq. (4). Note that the inequality holds no matter what the value of M is. Next, we show that opportunity ratio policy achieves the maximum possible interview efficiency as shown in Eq. (10). Let $X^* = \text{argmin}_{X\in\mathcal{X}}\Pi_{j\in[k]}\tau^j_{X1}$. Recall that the opportunity ratio policy π is defined as follow.

$$\pi^1_{X0} = 0 \text{ and } \pi^1_{X1} = \Pi_{i\in[k]}(\tau^i_{X^*1}/\tau^i_{X1}) \qquad \forall X \in \mathcal{X}$$
$$\pi^i_{X0} = 0 \text{ and } \pi^i_{X1} = 1 \qquad \forall X \in \mathcal{X}, i \geq 2$$

It is straightforward to check that π is an Equal Opportunity policy with recall $\Pi_{i\in[k]}\tau^i_{X^*1}$. Moreover, the interview efficiency of π is

$$\begin{aligned}\text{IE}(q,u,\tau,\pi) &= \frac{\sum_{X\in\mathcal{X}} q_X M_{X,\tau,\pi}}{\sum_{X\in\mathcal{X}} q_X M_{X,\tau,\pi} + u_X N^X_{\tau,\pi}} \\ &= \frac{\sum_{x\in\mathcal{X}} q_X \Pi_{i\in[k]}\tau^i_{X^*1}}{\sum_{x\in\mathcal{X}} q_X \Pi_{i\in[k]}\tau^i_{X^*1} + \sum_{X\in\mathcal{X}} u_X \Pi_{i\in[k]} \frac{\tau^i_{X^*1}\tau^i_{X0}}{\tau^i_{X1}}} \\ &= \frac{\|q\|_1}{\|q\|_1 + \sum_{X\in\mathcal{X}} u_X \Pi^k_{i=1} \frac{\tau^i_{X0}}{\tau^i_{X1}}}\end{aligned}$$

Hence, π is the Equal Opportunity policy maximizing the interview efficiency. □

B PROOFS FROM SECTION 4

Proof of Claim 4.3. Suppose for contradiction that there exists a group $X \in \mathcal{X}$ and a level $i \in [k]$ such that $\pi^i_{X1} = 0$. First note that $(1-\tau^i_{X1})\pi^i_{X0} > 0$; otherwise, the policy is useless because it prevents candidates of group X, in particular the qualified ones, from reaching the interview stage. Hence, by the Equal Opportunity requirement, no qualified candidate will reach the interview stage.

Next, we show that there exists a policy $\tilde{\pi}$ (which only differs from π in level i of group X) that satisfies Equal Opportunity for the given screening process and strictly Pareto dominates π; $\tilde{\pi}^i_{X1} = (\frac{1-\tau^i_{X1}}{\tau^i_{X1}})\pi^i_{X0}$ and $\tilde{\pi}^i_{X0} = 0$.

Since $M_{X,\tau^i,\tilde{\pi}^i} = \tau^i_{X1}\tilde{\pi}^i_{X1} + (1-\tau^i_{X1})\tilde{\pi}^i_{X0} = \tau^i_{X1}\tilde{\pi}^i_{X1} = (1-\tau^i_{X1})\pi^i_{X0} = \tau^i_{X1}\pi^i_{X1} + (1-\tau^i_{X1})\pi^i_{X0} = M_{X,\tau^i,\pi^i}$ and π satisfies the Equal Opportunity, $\tilde{\pi}$ also satisfies Equal Opportunity and has the same recall as π. Moreover, since

$$N_{X,\tau^i,\tilde{\pi}^i} = \tilde{\pi}^i_{X1}\tau^i_{X0} = (\frac{1-\tau^i_{X1}}{\tau^i_{X1}})\pi^i_{X0}\tau^i_{X0} < (1-\tau^i_{X0})\pi^i_{X0} = N_{X,\tau^i,\pi^i},$$

precision($\tilde{\pi}$) > precision(π). Note that $\frac{1-\tau^i_{X1}}{\tau^i_{X1}} < \frac{1-\tau^i_{X0}}{\tau^i_{X0}}$ holds by the minimal effectiveness property of tests. □

Lemma B.1. *Consider a k-stage screening process whose tests satisfy the Minimal Effectiveness Property. The set of Equal Opportunity policies $\mathcal{S} \subseteq \mathcal{P} = \{\pi \in [0,1]^{2|\mathcal{X}|k} | (1-\pi^i_{X1})\pi^i_{X0} = 0, \forall X \in \mathcal{X}, i \in [k]\}$, where for each group $X \in \mathcal{X}$, there exists at most one level $i \in [k]$ such that $\pi^i_{X1} < 1$, weakly Pareto dominates all Equal Opportunity policies.*

Proof. Suppose for contradiction that there are two levels i, j such that $\pi^i_{X1}, \pi^j_{X1} < 1$. First note that by Claim 4.3, $\pi^i_{X1}, \pi^j_{X1} > 0$. Moreover, by Lemma 4.2, since $\pi^i_{X1}, \pi^j_{X1} < 1$, $\pi^i_{X0} = \pi^j_{X0} = 0$.

Next, we show that we can modify π in levels i and j and replace π^i_{X0}, π^j_{X0} with $\tilde{\pi}^i_{X0}, \tilde{\pi}^j_{X0}$ as follows: $\tilde{\pi}^i_{X1} = \pi^i_{X1}\pi^j_{X1}$ and $\tilde{\pi}^j_{X1} = 1$. Then, $M_{X,\tau^i,\tilde{\pi}^i}M_{X,\tau^j,\tilde{\pi}^j} = (\tilde{\pi}^i_{X1}\tau^i_{X1})(\tilde{\pi}^j_{X1}\tau^j_{X1}) = (\pi^i_{X1}\tau^i_{X1})(\pi^j_{X1}\tau^j_{X1}) = M_{X,\tau^i,\pi^i}M_{X,\tau^j,\pi^j}$. In other words, the policy $\tilde{\pi}$ satisfies Equal Opportunity and has the same recall as π. Similarly, this modification does not decrease precision. Formally, $N_{X,\tau^i,\tilde{\pi}^i}N_{X,\tau^j,\tilde{\pi}^j} = (\tilde{\pi}^i_{X1}\tau^i_{X0})(\tilde{\pi}^j_{X1}\tau^j_{X0}) = (\pi^i_{X1}\tau^i_{X0})(\pi^j_{X1}\tau^j_{X0}) = N_{X,\tau^i,\pi^i}N_{X,\tau^j,\pi^j}$. Hence, for each policy π, there exists another policy with at most one level $i \in [k]$ such that $\pi^i_{X0} < 1$ and weakly Pareto dominates π. □

Proof of Lemma 4.5. We follow a similar arguments as in the proof of Lemma 4.4. Note that by Lemma 4.4 and Lemma B.1 there is at most one level $i_1 \in [k]$ such that $0 < \pi^{i_1}_{X1} < 1$ and $\pi^{i_1}_{X0} = 0$, and there is at most one level $i_0 \in [k]$ such that $\pi^{i_0}_{X1} = 1$ and $0 < \pi^{i_0}_{X0} < 1$. Next, we show that we can modify the policy π in levels i_0 and i_1 and replace $\pi^{i_0}_{X0}, \pi^{i_1}_{X1}$ with $\tilde{\pi}^{i_0}_{X0}, \tilde{\pi}^{i_1}_{X1}$ such that

$$\begin{aligned}M_{X,\tau^{i_0},\pi^{i_0}}M_{X,\tau^{i_1},\pi^{i_1}} &= (\tau^{i_0}_{X1} + \pi^{i_0}_{X0}(1-\tau^{i_0}_{X1}))(\pi^{i_1}_{X1}\tau^{i_1}_{X1}) \\ &= (\tau^{i_0}_{X1} + \tilde{\pi}^{i_0}_{X0}(1-\tau^{i_0}_{X1}))(\tilde{\pi}^{i_1}_{X1}\tau^{i_1}_{X1}) \\ &= M_{X,\tau^{i_0},\tilde{\pi}^{i_0}}M_{X,\tau^{i_1},\tilde{\pi}^{i_1}}, \\ N_{X,\tau^{i_0},\pi^{i_0}}N_{X,\tau^{i_1},\pi^{i_1}} &= (\tau^{i_0}_{X0} + \pi^{i_0}_{X0}(1-\tau^{i_0}_{X0}))(\pi^{i_1}_{X1}\tau^{i_1}_{X0}) \\ &< (\tau^{i_0}_{X0} + \tilde{\pi}^{i_0}_{X0}(1-\tau^{i_0}_{X0}))(\tilde{\pi}^{i_1}_{X1}\tau^{i_1}_{X0}) \\ &= N_{X,\tau^{i_0},\tilde{\pi}^{i_0}}N_{X,\tau^{i_1},\tilde{\pi}^{i_1}}\end{aligned}$$

Now, we show that in the new solution, either $\tilde{\pi}^{i_0}_{X0} \in \{0,1\}$ or $\tilde{\pi}^{i_1}_{X1} = 1$.

Without loss of generality, we can assume that the feasible range of values for $\tilde{\pi}^{i_0}_{X0}$ to satisfy Equal Opportunity is $[\pi^{i_0}_{X0} - \epsilon^{i_0}, \pi^{i_0}_{X0} + \delta^{i_0}]$ which corresponds to $[\pi^{i_1}_{X0} - \delta^{i_1}, \pi^{i_1}_{X0} + \epsilon^{i_1}]$. Both intervals are sub-intervals of $[0,1]$ and it is straightforward to verify that $(\pi^{i_0}_{X0} - \epsilon^{i_0})(1-(\pi^{i_1}_{X0} + \epsilon^{i_1})) = (1-(\pi^{i_1}_{X0} + \delta^{i_1})) = 0$.

Let $L = M_X/(\tau^{i_0}_{X1}\tau^{i_1}_{X1})$ where

$$M_X = M_{X,\tau^{i_0},\pi^{i_0}}M_{X,\tau^{i_1},\pi^{i_1}} = M_{X,\tau^{i_0},\tilde{\pi}^{i_0}}M_{X,\tau^{i_1},\tilde{\pi}^{i_1}}$$

. By the Minimal Effectiveness Property, $0 < L < 3$. Then, satisfying Equal Opportunity is equivalent to satisfy $(1+\tilde{\pi}^{i_0}_{X0}(\frac{1-\tau^{i_0}_{X1}}{\tau^{i_0}_{X1}}))\tilde{\pi}^{i_1}_{X1} = L$, which implies that $\tilde{\pi}^{i_1}_{X1} = L/(1+\tilde{\pi}^{i_0}_{X0}(\frac{1-\tau^{i_0}_{X1}}{\tau^{i_0}_{X1}}))$. The task of finding $\tilde{\pi}^i_{X0}$ is as follows:

$$\begin{aligned}\tilde{\pi}^i_{X0} &= \text{argmin}_{y\in[\pi^{i_0}_{X0}-\epsilon^{i_0}, \pi^{i_0}_{X0}+\delta^{i_0}]} f(y) \\ &:= (\tau^{i_0}_{X0} + y(1-\tau^{i_0}_{X0}))(\tau^{i_1}_{X0} \cdot \frac{L}{1+y(\frac{1-\tau^{i_0}_{X1}}{\tau^{i_0}_{X1}})}).\end{aligned}$$

Next, we show that for any $y \in [0,1]$,

$$f''(y) = \frac{2L\tau^{i_1}_{X0}(\frac{1-\tau^{i_0}_{X1}}{\tau^{i_0}_{X1}})(\tau^{i_0}_{X0}(\frac{1-\tau^{i_0}_{X1}}{\tau^{i_0}_{X1}}) + \tau^{i_0}_{X0} - 1)}{(1+\tau^{i_0}_{X0}y)^3} < 0$$

To prove it note that the Minimal Effectiveness Property of the tests $\{\tau^i\}_{i\in[k]}$ (i.e., $\tau^i_{X1} > \tau^i_{X0} \geq 0, \forall X \in \mathcal{X}, i \in [k]$) implies that $\frac{\tau^{i_0}_{X0}}{\tau^{i_0}_{X1}} - 1 < 0$. Since f is a concave function in $[\pi^{i_0}_{X0} - \epsilon^{i_0}, \pi^{i_0}_{X0} + \delta^{i_0}]$, the minimum value of f in this interval obtained in one of its endpoints. In other words, the maximum precision corresponds to the case either $\tilde{\pi}^{i_0}_{X0} \in \{0, 1\}$ or $\tilde{\pi}^{i_1}_{X1} = 1$. □

Proof of Lemma 4.8. The proof is by induction. For the base case ($i = 1$), let t_1 and f_1 denote the true positive rate and the false positive rate of π by the end of level 1. The existence of π guarantees that the system of inequalities Eq. (8) with ($j_0 = \lfloor \log_{1-\epsilon} \mathrm{f}_1 \rfloor, j_1 = \lceil \log_{1-\epsilon} \mathrm{t}_1 \rceil \leq \ell_{\mathrm{tpr}}$) has a feasible solution. More precisely, by setting ($x = \pi_1, y = \pi_0$),

$$\begin{aligned}
&\tau^1_0 x + (1-\tau^1_0)y = \mathrm{f}_1 \leq (1-\epsilon)^{\lfloor \log_{1-\epsilon} \mathrm{f}_1 \rfloor} \\
&= (1-\epsilon)^{j_0}, \quad \tau^1_1 x + (1-\tau^1_1)y \\
&= \mathrm{t}_1 \geq (1-\epsilon)^{\lceil \log_{1-\epsilon} \mathrm{t}_1 \rceil} = (1-\epsilon)^{j_1}
\end{aligned}$$

Next, we consider $i > 1$ and we assume that the claim holds for all values $i' < i$. Let $M_i := \tau^i_1\pi^i_1 + (1-\tau^i_1)\pi^i_0$ and $N_i := \tau^i_0\pi^i_1 + (1-\tau^i_0)\pi^i_0$. Note that $\mathrm{t}_i = \mathrm{t}_{i-1} \cdot M_i$ and $\mathrm{f}_i = \mathrm{f}_{i-1} \cdot N_i$.

By the induction hypothesis and considering the first $i-1$ levels in the pipeline, since $\mathrm{t}_{i-1} \geq \mathrm{t}_i \geq L_{\mathrm{tpr}}/(1-\epsilon)^{i-1} > L_{\mathrm{tpr}}/(1-\epsilon)^{i-2}$ and $\mathrm{f}_{i-1} \geq \mathrm{f}_i$, there exist $j'_1 \in [0, L_{\mathrm{tpr}}]$ and $j'_0 \in [0, L_{\mathrm{fpr}}] \cup \{\infty\}$ such that $M[i-1, j'_1, j'_0] = \text{true}$ and $(1-\epsilon)^{j'_1} \geq \mathrm{t}_{i-1} \cdot (1-\epsilon)^{i-2}$ and $(1-\epsilon)^{j'_0} \leq \min\{1, \max\{L_{\mathrm{fpr}}, \mathrm{f}_{i-1}\}/(1-\epsilon)^{i-2}\}$. More precisely, the algorithm finds a policy $\bar{\pi}$ with true positive rate at least $(1-\epsilon)^{j'_1}$ and false positive rate at most $(1-\epsilon)^{j'_0}$.

Next, by setting ($\bar{\pi}^i_1 = \pi^i_1, \bar{\pi}^i_0 = \pi^i_0$) and ($j_1 := \mathrm{argmin}_j\{(1-\epsilon)^j \leq \mathrm{t}_i(\bar{\pi})\}, j_0 := \mathrm{argmax}_j\{(1-\epsilon)^j \geq \mathrm{f}_i(\bar{\pi})\}$),

$$\begin{aligned}
(1-\epsilon)^{j_1} &> (1-\epsilon) \cdot \mathrm{t}_i(\bar{\pi}) \\
&= (1-\epsilon) \cdot \mathrm{t}_{i-1}(\bar{\pi}) \cdot M_i && \triangleright \text{ by definition of } j_1 \\
&\geq (1-\epsilon) \cdot (1-\epsilon)^{j'_1} \cdot M_i && \triangleright \text{ by } \mathrm{t}_{i-1}(\bar{\pi}) \geq (1-\epsilon)^{j'_1} \\
&\geq \mathrm{t}_{i-1} \cdot (1-\epsilon)^{i-1} \cdot M_i && \triangleright \text{ by induction hypothesis} \\
&= \mathrm{t}_i \cdot (1-\epsilon)^{i-1}.
\end{aligned}$$

Similarly,

$$\begin{aligned}
(1-\epsilon)^{j_0} &< \min\{1, \frac{\mathrm{f}_i(\bar{\pi})}{1-\epsilon}\} \\
&= \min\{1, N_i \cdot \frac{\mathrm{f}_{i-1}(\bar{\pi})}{1-\epsilon}\} \ \triangleright \text{ by definition of } j_0 \\
&\leq \min\{1, (1-\epsilon)^{j'_0} \cdot \frac{N_i}{1-\epsilon}\} \ \triangleright \ \mathrm{f}_{i-1}(\bar{\pi}) \leq (1-\epsilon)^{j'_0} \\
&\leq \min\{1, \frac{\max\{L_{\mathrm{fpr}}, \mathrm{f}_{i-1}\}}{(1-\epsilon)^{i-2}} \cdot \frac{N_i}{1-\epsilon}\} \ \triangleright \text{ induction hypoth.} \\
&\leq \min\{1, \frac{\max\{L_{\mathrm{fpr}}, \mathrm{f}_i\}}{(1-\epsilon)^{i-1}}\}
\end{aligned}$$

which completes the proof. □

Proof of Theorem 4.10. First, as we are aiming for a $(1-\epsilon)$-approximation, we only need to consider $\alpha \in (\epsilon, 1-\epsilon)$. Otherwise, either the policy maximizing recall (i.e. bypassing all tests) or the policy maximizing precision (Opportunity Ratio policy) is a $(1-\epsilon)$-approximation for f_α.

Next we show in order to guarantee $(1-\epsilon)$-approximations of recall and precision of the policy maximizing f_α, it suffices to run the described DP and consider estimates of t (true positive rate) and f (false positive rate) of form $(1-\bar{\epsilon})^i$ for $i \in \mathbb{N}$ in intervals $[L_{\mathrm{tpr}}, 1]$ and $[L_{\mathrm{fpr}}, 1]$ respectively, where $\bar{\epsilon} \leq \epsilon/(2k)$. We provide tight bounds for L_{tpr} and L_{fpr}. Note that since for any policy π, the true positive rate (t_i) and the false positive rate (f_i) are non-decreasing in i, it suffices to provide "large enough" lowerbounds L_{tpr} and L_{fpr} for t and f in the final stage respectively.

Bounding L_{tpr}. Consider the policy π_{bypass}, which bypasses all the tests in both groups, i.e., $\pi^i_{X0} = \pi^i_{X1} = 1$ for all $i \in [k], X \in \mathcal{X}$. Since π_{bypass} is an Equal Opportunity policy for the pipeline and $f_\alpha(\pi_{\mathrm{bypass}}) = (1-\alpha) + \alpha\|q\|_1$, any optimal Equal Opportunity policy π^* for f_α has recall at least $(1 - 2\alpha + \alpha\|q\|_1)/(1-\alpha)$. Thus, since $\alpha \in (\epsilon, 1-\epsilon)$, $\mathrm{t} \geq (1 - 2\alpha + \alpha\|q\|_1)/(1-\alpha) \geq \epsilon/(1-\epsilon)$ which implies that in our DP with accuracy parameter $\bar{\epsilon}$ it suffices to set $L_{\mathrm{tpr}} = (\frac{\epsilon}{1-\epsilon}) \cdot (1-\bar{\epsilon})^{k-1} \geq (\frac{\epsilon}{1-\epsilon}) \cdot \exp(-\epsilon)$.

Bounding L_{fpr}. For each $X \in \mathcal{X}$, let f_X denote the false positive rate of the optimal Equal Opportunity policy for group X. Similarly, let t_X denote the positive rate of (i.e., recall) the optimal policy π^* for group $X \in \mathcal{X}$. By Equality of Opportunity property of π^*, $\mathrm{t}_X = \mathrm{t}$ for each $X \in \mathcal{X}$. Next, we consider the following cases.

For any sufficiently small $\epsilon > 0$, we need to set L_{fpr} so that by running the DP with accuracy parameter $\bar{\epsilon}$, we can approximate both true positive rate and false positive rate of the optimal Equal Opportunity policy within $(1-\epsilon)$-factor of their values. More precisely, we set L_{fpr} so that if for each group $X \in \mathcal{X}$ and any pair $(\mathrm{t}_X, \mathrm{f}_X)$ with $t_X \geq L_{\mathrm{tpr}}/(1-\epsilon/2)$, there exists a pair $(\bar{\mathrm{t}}_X, \bar{f}_X)$ such that $\bar{\mathrm{t}}_X \geq (1-\epsilon/2)\mathrm{t}_X, \bar{\mathrm{f}}_X \leq \min(1, \max(L_{\mathrm{fpr}}, \mathrm{f}_X)/(1-\epsilon/2))$. Finally, once the above property holds for all groups $X \in \mathcal{X}$, then for the corresponding policy π, $\mathrm{precision}(\pi) > (1-\epsilon) \cdot \mathrm{precision}(\pi^*)$.

Let $\mathcal{X}_1 := \{X \in \mathcal{X} | \mathrm{f}_X/(1-\frac{\epsilon}{2}) \geq L_{\mathrm{fpr}}\}$ and $\mathcal{X}_2 := \{X \in \mathcal{X} | \mathrm{f}_X/(1-\frac{\epsilon}{2}) < L_{\mathrm{fpr}}\}$. Then,

$$\begin{aligned}
\frac{\mathrm{precision}(\pi)}{\mathrm{precision}(\pi^*)} &= \frac{\left(\frac{\|q\|_1 \cdot \bar{\mathrm{t}}}{\|q\|_1 \cdot \bar{\mathrm{t}} + \sum_{X\in\mathcal{X}} u_X \cdot \bar{\mathrm{f}}_X}\right)}{\left(\frac{\|q\|_1 \cdot \mathrm{t}}{\|q\|_1 \cdot \mathrm{t} + \sum_{X\in\mathcal{X}} u_X \cdot \mathrm{f}_X}\right)} \\
&\geq \frac{\left(\frac{\|q\|_1 \cdot \bar{\mathrm{t}}}{\|q\|_1 \cdot \bar{\mathrm{t}} + \sum_{X\in\mathcal{X}_1} u_X \cdot \bar{\mathrm{f}}_X + \sum_{X\in\mathcal{X}_2} u_X \cdot \bar{\mathrm{f}}_X}\right)}{\left(\frac{\|q\|_1 \cdot \mathrm{t}}{\|q\|_1 \cdot \mathrm{t} + \sum_{X\in\mathcal{X}_1} u_X \cdot \mathrm{f}_X}\right)} \\
&\geq \frac{\left(\frac{\|q\|_1 \cdot (1-\epsilon/2)\mathrm{t}}{\|q\|_1 \cdot \mathrm{t} + \sum_{X\in\mathcal{X}_1} \frac{u_X \cdot \mathrm{f}_X}{1-\epsilon/2} + \sum_{X\in\mathcal{X}_2} u_X \cdot L_{\mathrm{fpr}}}\right)}{\left(\frac{\|q\|_1 \cdot \mathrm{t}}{\|q\|_1 \cdot \mathrm{t} + \sum_{X\in\mathcal{X}_1} u_X \cdot \mathrm{f}_X}\right)}
\end{aligned}$$

Next, we set L_{fpr} so that $\|q\|_1 \cdot \mathrm{t} + \sum_{X\in\mathcal{X}_2} u_X L_{\mathrm{fpr}} \leq \frac{\|q\|_1 \cdot \mathrm{t}}{1-\frac{\epsilon}{2}}$. Since $\mathrm{t} \geq \epsilon/(1-\epsilon)$, it suffices to set $L_{\mathrm{fpr}} = \frac{\epsilon^2\|q\|_1}{(2-\epsilon)(1-\epsilon)(1-\|q\|_1)} = \Omega(\epsilon^2)$. Hence,

$$\frac{\left(\frac{\|q\|_1 \cdot (1-\frac{\epsilon}{2})\mathrm{t}}{\|q\|_1 \cdot \mathrm{t} + \sum_{X\in\mathcal{X}_1} \frac{u_X \cdot \mathrm{f}_X}{1-\frac{\epsilon}{2}} + \sum_{X\in\mathcal{X}_2} u_X \cdot L_{\mathrm{fpr}}}\right)}{\left(\frac{\|q\|_1 \cdot \mathrm{t}}{\|q\|_1 \cdot \mathrm{t} + \sum_{X\in\mathcal{X}_1} u_X \cdot \mathrm{f}_X}\right)} \geq (1 - \frac{\epsilon}{2})^2 > (1-\epsilon).$$

Finally, for each $X \in \mathcal{X}$, we run the DP algorithm for each group with accuracy parameter $\bar{\epsilon}$. By Lemma 4.8, the DP algorithm finds a

set $\{\mathrm{t}_X = (1-\bar{\epsilon})^{i_X}, \mathrm{f}_x = (1-\bar{\epsilon})^{j_X}\}_{X\in\mathcal{X}}$ (and a policy π achieving these rates) where for each $X \in \mathcal{X}$, $\mathrm{t}_X \in [L_{\mathrm{tpr}}, 1], \mathrm{f}_X \in [L_{\mathrm{fpr}}, 1]$ such that

$$\mathrm{t}_X = \mathrm{t} \geq (1-\frac{\epsilon}{2})\cdot \mathrm{t}(\pi^*),\ \mathrm{f}_X \leq \min\{1, \frac{\max\{L_{\mathrm{fpr}}, \mathrm{f}_X(\pi^*)\}}{1-\epsilon/2}\}\ \forall X \in \mathcal{X},$$

and for each $X \in \mathcal{X}$, $M_X[k, \mathrm{t}_X, \mathrm{f}_X] = \text{true}$. Thus, by the bounds we just showed for the precision of such a policy, $\mathrm{precision}(\pi) \geq (1-\epsilon)\cdot\mathrm{precision}(\pi^*)$. Thus, $f_\alpha(\pi) \geq (1-\epsilon)\cdot f_\alpha(\pi^*)$.

As we need to run the DP algorithm for any of the $|\mathcal{X}|$ groups separately with the specified parameters $L_{\mathrm{tpr}}, L_{\mathrm{fpr}}$ and $\bar{\epsilon} = O(\epsilon/k)$, by Lemma 4.9, the total time of the DP approach is

$$O(\frac{|\mathcal{X}|k\log^2\frac{1}{L_{\mathrm{tpr}}}\log^2\frac{1}{L_{\mathrm{fpr}}}}{\bar{\epsilon}^4}) = O(\frac{|\mathcal{X}|k^5(\epsilon^2+\log^2\frac{1}{\epsilon})\log^2\frac{1}{\epsilon}}{\epsilon^4}) = O(\frac{|\mathcal{X}|k^5\log^4\frac{1}{\epsilon}}{\epsilon^4})$$

□

Proof of Theorem 4.11. First we show that in our setting, in order to guarantee $(1+\epsilon)$-approximations of recall and precision, it suffices to run the described DP and consider estimates of t (true positive rate) and f (false positive rate) of form $(1-\bar{\epsilon})^i$ for $i \in \mathbb{N}$ in intervals $[L_{\mathrm{tpr}}, 1]$ and $[L_{\mathrm{fpr}}, 1]$ respectively, where $\bar{\epsilon} \leq \epsilon/(2k)$. We provide tight bounds for L_{tpr} and L_{fpr}.

Note that since for any policy π, $\mathrm{tpr}_{i,\pi}, \mathrm{fpr}_{i,\pi}$ are non-decreasing in i, it suffices to provide "large enough" lowerbounds L_{tpr} and L_{fpr} for true positive rate and false positive rate in the final stage respectively (i.e., for t and f).

Bounding L_{tpr}. Consider the policy π_{bypass}, which bypasses all the tests in both groups, i.e., $\pi^i_{X0} = \pi^i_{X1} = 1$ for all $i \in [k], X \in \mathcal{X}$. Let $\tau_{\min} = \min_{X\in\mathcal{X}, j\in[k]} \tau^j_{X1}$. Then, by Theorem 3.4, Opportunity Ratio maximizes the precision and has recall at least $(\tau_{\min})^k$, in the optimal policy $\mathrm{t} \geq (\tau_{\min})^k$ which implies that in our DP with accuracy parameter $\bar{\epsilon}$ it suffices to set $L_{\mathrm{tpr}} = (\tau_{\min})^k\cdot(1-\bar{\epsilon})^{k-1} \geq \exp(-\epsilon - k\ln(1/\tau_{\min}))$.

Bounding L_{fpr}. For each $X \in \mathcal{X}$, let f_X denote the false positive rate of the optimal Equal Opportunity policy for group X. Similarly, let t_X denote the positive rate of (i.e., recall) the optimal policy π^* for group $X \in \mathcal{X}$. By Equality of Opportunity property of π^*, $\mathrm{t}_X = \mathrm{t}$ for each $X \in \mathcal{X}$. Next, we consider the following cases.

For any sufficiently small $\epsilon > 0$, we need to set L_{fpr} so that by running the DP with accuracy parameter $\bar{\epsilon}$, we can approximate both true positive rate and false positive rate of the optimal Equal Opportunity policy within $(1-\epsilon)$-factor of their values. More precisely, we set L_{fpr} so that if for each group $X \in \mathcal{X}$ and any pair $(\mathrm{t}_X, \mathrm{f}_X)$ with $t_X \geq L_{\mathrm{tpr}}$, there exists a pair $(\bar{\mathrm{t}}_X, \bar{f}_X)$ such that $\bar{\mathrm{t}}_X \geq (1-\epsilon/2)\mathrm{t}_X$, $\bar{\mathrm{f}}_X \leq \min(1, \max(L_{\mathrm{fpr}}, \mathrm{f}_X)/(1-\epsilon/2))$. Finally, once the above property holds for all groups $X \in \mathcal{X}$, then for the corresponding policy π, $\mathrm{precision}(\pi) > (1-\epsilon)\cdot\mathrm{precision}(\pi^*)$.

Let $\mathcal{X}_1 := \{X \in \mathcal{X} | \mathrm{f}_X/(1-\epsilon/2) \geq L_{\mathrm{fpr}}\}$ and let $\mathcal{X}_2 := \{X \in \mathcal{X} | \mathrm{f}_X/(1-\epsilon/2) < L_{\mathrm{fpr}}\}$. Note that $\mathcal{X} = \mathcal{X}_1 \dot\cup \mathcal{X}_2$. Then,

$$\begin{aligned}\frac{\mathrm{precision}(\pi)}{\mathrm{precision}(\pi^*)} &= \Big(\frac{\|q\|_1\cdot\bar{\mathrm{t}}}{\|q\|_1\cdot\bar{\mathrm{t}} + \sum_{X\in\mathcal{X}} u_X\cdot\bar{\mathrm{f}}_X}\Big) \Big/ \Big(\frac{\|q\|_1\cdot\mathrm{t}}{\|q\|_1\cdot\mathrm{t} + \sum_{X\in\mathcal{X}} u_X\cdot\mathrm{f}_X}\Big)\\ &\geq \Big(\frac{\|q\|_1\cdot\bar{\mathrm{t}}}{\|q\|_1\cdot\bar{\mathrm{t}} + \sum_{X\in\mathcal{X}_1} u_X\cdot\bar{\mathrm{f}}_X + \sum_{X\in\mathcal{X}_2} u_X\cdot\bar{\mathrm{f}}_X}\Big) \Big/ \Big(\frac{\|q\|_1\cdot\mathrm{t}}{\|q\|_1\cdot\mathrm{t} + \sum_{X\in\mathcal{X}_1} u_X\cdot\mathrm{f}_X}\Big)\\ &\geq \Big(\frac{\|q\|_1\cdot(1-\epsilon/2)\mathrm{t}}{\|q\|_1\cdot\mathrm{t} + \sum_{X\in\mathcal{X}_1}\frac{u_X\cdot\mathrm{f}_X}{1-\epsilon/2} + \sum_{X\in\mathcal{X}_2} u_X\cdot L_{\mathrm{fpr}}}\Big) \Big/ \Big(\frac{\|q\|_1\cdot\mathrm{t}}{\|q\|_1\cdot\mathrm{t} + \sum_{X\in\mathcal{X}_1} u_X\cdot\mathrm{f}_X}\Big)\end{aligned}$$

Next, we set L_{fpr} so that $\|q\|_1\cdot\mathrm{t} + \sum_{X\in\mathcal{X}_2} u_X L_{\mathrm{fpr}} \leq \frac{\|q\|_1\cdot\mathrm{t}}{1-\frac{\epsilon}{2}}$. Since $\mathrm{t} \geq (\tau_{\min})^k$, it suffices to set $L_{\mathrm{fpr}} = \frac{\epsilon\|q\|_1\cdot(\tau_{\min})^k}{(2-\epsilon)(1-\|q\|_1)} = \Omega(\epsilon\cdot(\tau_{\min})^k)$. Hence,

$$\Big(\frac{\|q\|_1\cdot(1-\epsilon/2)\mathrm{t}}{\|q\|_1\mathrm{t} + \sum_{X\in\mathcal{X}_1}\frac{u_X\mathrm{f}_X}{1-\epsilon/2} + \sum_{X\in\mathcal{X}_2} u_X L_{\mathrm{fpr}}}\Big) \Big/ \Big(\frac{\|q\|_1\mathrm{t}}{\|q\|_1\mathrm{t} + \sum_{X\in\mathcal{X}_1} u_X\mathrm{f}_X}\Big)$$
$$\geq (1-\frac{\epsilon}{2})^2 > (1-\epsilon).$$

Finally, for each $X \in \mathcal{X}$, we run the DP algorithm for each group with accuracy parameter $\bar{\epsilon}$. By Lemma 4.8, the DP algorithm finds a set $\{\mathrm{t}_X = (1-\bar{\epsilon})^{i_X}, \mathrm{f}_X = (1-\bar{\epsilon})^{j_X}\}_{X\in\mathcal{X}}$ (and a policy π corresponding to these values) where for each $X \in \mathcal{X}$, $\mathrm{t}_X \in [L_{\mathrm{tpr}}, 1], \mathrm{f}_X \in [L_{\mathrm{fpr}}, 1]$ such that $\forall X \in \mathcal{X}$

$$\mathrm{t}_X = \mathrm{t} \geq (1-\epsilon/2)\cdot\mathrm{t}(\pi^*) \quad \mathrm{f}_X \leq \min(1, \frac{\max(L_{\mathrm{fpr}}, \mathrm{f}_X(\pi^*))}{1-\epsilon/2})$$

and for each $X \in \mathcal{X}$, $M_X[k, \mathrm{t}_X, \mathrm{f}_X] = \text{true}$. Thus, by the bounds we just showed for the precision of such a policy, $1/\mathrm{precision}(\pi) \leq (1+\epsilon)/\mathrm{precision}(\pi^*)$. Thus, $g_\alpha(\pi) \leq (1+\epsilon)\cdot g_\alpha(\pi^*)$.

As we need to run the DP algorithm for any of the $|\mathcal{X}|$ groups separately with the specified parameters $L_{\mathrm{tpr}}, L_{\mathrm{fpr}}$ and $\bar{\epsilon} = O(\epsilon/k)$, by Lemma 4.9, the total runtime is $O(\frac{|\mathcal{X}|k\log^2(1/L_{\mathrm{tpr}})\log^2(1/L_{\mathrm{fpr}})}{\bar{\epsilon}^4}) = O(\frac{|\mathcal{X}|k^7(\log^2(1/\epsilon)+k^2)}{\epsilon^4})$. □

C MISSING PROOFS OF SECTION 5.2

Similarly to Observation 3.1, we can show the following observation for the policies that satisfies the Equalized Odds requirement.

Observation C.1. *For any policy π that satisfies the Equalized Odds for a k-stage screening process with parameters $(\{\tau^i\}_{i\in[k]}, \{q_X, u_X\}_{X\in\mathcal{X}})$, there exists M and N such that for each $X \in \mathcal{X}$,*

$$M := \Pi_{i=1}^k \tau^i_{X1}\pi^i_{X1} + (1-\tau^i_{X1})\pi^i_{X0}, \quad N := \Pi_{i=1}^k \tau^i_{X0}\pi^i_{X1} + (1-\tau^i_{X0})\pi^i_{X0}$$

Note that as computed in Observation 3.2, for policy any satisfying the Equalized Odds, the interview efficiency of a policy π for a k-stage process with parameters $(q, u, \{\tau^i\}_{i\in[k]})$ is $\frac{\|q\|_1 M}{\|q\|_1 M + \|u\|_1 N}$.

In the rest of the section and for the simplicity of the exposition, we assume there are exactly two groups in the population; $\mathcal{X} = \{A, B\}$. The result for the general setting can be derived similarly.

Theorem C.2. *The interview efficiency of any policy satisfying Equalized Odds for a single-stage screening process with parameters (q, u, τ) is at most $\frac{1}{1+\frac{u_A+u_B}{q_A+q_B}\cdot\max(\frac{\tau_{A0}}{\tau_{A1}}, \frac{\tau_{B0}}{\tau_{B1}})}$.*

Proof. Maximizing the interview efficiency, is equivalent to minimizing $N_{\tau,\pi}/M_{\tau,\pi}$; a minimizer of the inverse ratio is a maximizer of the interview efficiency and vice versa. Moreover, note that by the Minimal Effectiveness Property of the given test (i.e., Eq. (1)), $N_{\tau,\pi} < M_{\tau,\pi}$.

$$\frac{N_{\tau,\pi}}{M_{\tau,\pi}} = \frac{\tau_{A0}(\pi_{A1}-\pi_{A0})+\pi_{A0}}{\tau_{A1}(\pi_{A1}-\pi_{A0})+\pi_{A0}} \geq \frac{\tau_{A0}(\pi_{A1}-\pi_{A0})}{\tau_{A1}(\pi_{A1}-\pi_{A0})} = \frac{\tau_{A0}}{\tau_{A1}} \text{ and}$$
$$\frac{N_{\tau,\pi}}{M_{\tau,\pi}} = \frac{\tau_{B0}(\pi_{B1}-\pi_{B0})+\pi_{B0}}{\tau_{B1}(\pi_{B1}-\pi_{B0})+\pi_{B0}} \geq \frac{\tau_{B0}(\pi_{B1}-\pi_{B0})}{\tau_{B1}(\pi_{B1}-\pi_{B0})} = \frac{\tau_{B0}}{\tau_{B1}}.$$

In other words, $N_{\tau,\pi} \geq \max(\frac{\tau_{A0}}{\tau_{A1}}, \frac{\tau_{B0}}{\tau_{B1}}) \cdot M_{\tau,\pi}$. Hence,

$$\frac{(q_A+q_B)M_{\tau,\pi}}{(q_A+q_B)M_{\tau,\pi}+(u_A+u_B)N_{\tau,\pi}} \leq \frac{1}{1+\frac{u_A+u_B}{q_A+q_B}\cdot\max(\frac{\tau_{A0}}{\tau_{A1}}, \frac{\tau_{B0}}{\tau_{B1}})}$$

□

Remark 6. *Note that we can generalize the result of Lemma C.2 to a k-stage screening process with multiple groups $\mathcal{X}$. For any $j \in [k]$, let $\rho := \max_{X\in\mathcal{X}} \Pi_{j\in[k]} \frac{\tau^j_{X0}}{\tau^j_{X1}}$. Any policy that satisfies Equalized Odds requirement at the end of the process (i.e., before the interview stage) has interview efficiency at most $\frac{\|q\|_1}{\|q\|_1+\sum_{X\in\mathcal{X}}\rho u_X}$. To see this, note that similarly to the proof of Theorem C.2 we can show that for every group $X \in \mathcal{X}$, $N_X \geq \rho \cdot M_X$.*

Theorem C.3. *Consider a k-stage screening process (q, u, τ) with multiple groups $\mathcal{X}$ whose tests are minimally effective. Let $\pi_{\text{EOdd}}, \pi_{\text{EOpp}}$ denote the interview efficiency maximizing policy that satisfies Equalized Odds and Equal Opportunity at the end of the process respectively. If $\max_{X\in\mathcal{X}} \Pi_{i\in[k]} \frac{\tau^i_{X0}}{\tau^i_{X1}} > \min_{X\in\mathcal{X}} \Pi_{i\in[k]} \frac{\tau^i_{X0}}{\tau^i_{X1}}$, then $\text{IE}(q, u, \tau, \pi_{\text{EOdd}}) < \text{IE}(q, u, \tau, \pi_{\text{EOpp}})$.*

In particular, the gap between the interview efficiency of π_{EOdd} and π_{EOpp} can be as large as $\frac{1}{\|q\|_1} - \epsilon$ for any arbitrary $\epsilon > 0$.[4]

Proof. The proof of the first part directly follows from the interview efficiency of opportunity ratio policy (Theorem 3.4) and the upper bound for the interview efficiency of Equalized Odds policies (Theorem C.2)

For the second part, consider a pipeline in which there exists a $X^* \in \mathcal{X}$ such that for every $X \in \mathcal{X} \setminus X^*$, $\Pi_{i\in[k]} \frac{\tau^i_{X0}}{\tau^i_{X1}} = 0$ and $\Pi_{i\in[k]} \frac{\tau^i_{X^*0}}{\tau^i_{X^*1}} = (1-\delta)^k$. Further, for every $X \in \mathcal{X} \setminus X^*$, let $q_X = \frac{\gamma}{k}, u_X = \frac{1-\gamma-\mu}{k-1}$ and $q_{X^*} = \frac{\gamma}{k}, u_{X^*} = \mu$. Then, it is straightforward to check that $\text{IE}(\pi_{\text{EOpp}}) = \frac{\gamma}{\gamma+\mu\cdot(1-\delta)^k}$ and $\text{IE}(\pi_{\text{EOdd}}) = \frac{\gamma}{\gamma+(1-\gamma)\cdot(1-\delta)^k}$. As we set δ, μ to sufficiently small values, $\text{IE}(\pi_{\text{EOpp}})/\text{IE}(\pi_{\text{EOpp}}) = 1/\gamma - \epsilon = 1/\|q\|_1 - \epsilon$. □

Next, we show the following structure on a non-trivial optimal solution (i.e., one maximizing the interview efficiency). Note that $\pi = \mathbf{1}$ or $\pi = \mathbf{0}$ are the two trivial solutions satisfying the Equalized Odds for any given test.

Observation C.4. *For any pipeline (τ, q, u), in any* non-trivial *optimal policy π, $\min(\pi_{A1}, \pi_{A0}, \pi_{B1}, \pi_{B0}) = 0$. Moreover, there exists an optimal policy such that $\max(\pi_{A1}, \pi_{A0}, \pi_{B1}, \pi_{B0}) = 1$.*

Proof. First, note that by the Minimal Effectiveness Property of the given test (i.e., Eq. (1)), $N_{\tau,\pi} < M_{\tau,\pi}$.

Suppose that $\min(\pi_{A1}, \pi_{A0}, \pi_{B1}, \pi_{B0}) = \epsilon$. This implies that $M_{\tau,\pi} > N_{\tau,\pi} \geq \epsilon$ Then, by subtracting ϵ from all π values, the new policy still satisfies the Equalized Odds and it only increases the interview efficiency. Formally, for $\epsilon > 0$

$$\frac{\|q\|_1 \cdot M_{\tau,\pi}}{\|q\|_1 \cdot M_{\tau,\pi} + \|u\|_1 \cdot N_{\tau,\pi}} < \frac{\|q\|_1 \cdot (M_{\tau,\pi}-\epsilon)}{\|q\|_1 \cdot (M_{\tau,\pi}-\epsilon) + \|u\|_1 \cdot (N_{\tau,\pi}-\epsilon)}$$

The above inequality holds since

$$\begin{aligned}
&N_{\tau,\pi} < M_{\tau,\pi} \\
&\Rightarrow -\|u\|_1\epsilon N_{\tau,\pi} > -\|u\|_1\epsilon M_{\tau,\pi} \\
&\Rightarrow (\|q\|_1 M^2_{\tau,\pi} - \|q\|_1\epsilon M_{\tau,\pi} + \|u\|_1 M_{\tau,\pi} N_{\tau,\pi}) - \|u\|_1\epsilon N_{\tau,\pi} \\
&\quad > (\|q\|_1 M^2_{\tau,\pi} - \|q\|_1\epsilon M_{\tau,\pi} + \|u\|_1 M_{\tau,\pi} N_{\tau,\pi}) - \|u\|_1\epsilon M_{\tau,\pi} \\
&\Rightarrow M_{\tau,\pi}(\|q\|_1 M_{\tau,\pi} + \|u\|_1 N_{\tau,\pi}) - \epsilon(\|q\|_1 M_{\tau,\pi} + \|u\|_1 N_{\tau,\pi}) \\
&\quad > M_{\tau,\pi}(\|q\|_1 (M_{\tau,\pi}-\epsilon) + \|u\|_1 (N_{\tau,\pi}-\epsilon)) \\
&\Rightarrow \frac{M_{\tau,\pi}-\epsilon}{\|q\|_1 (M_{\tau,\pi}-\epsilon) + \|u\|_1 (N_{\tau,\pi}-\epsilon)} \\
&\quad > \frac{M_{\tau,\pi}}{\|q\|_1 M_{\tau,\pi} + \|u\|_1 N_{\tau,\pi}} \quad \triangleright\ \|q\|_1 (M_{\tau,\pi}-\epsilon) + \|u\|_1 (N_{\tau,\pi}-\epsilon) > 0
\end{aligned}$$

In particular, this implies that in any optimal policy,

$$\min(\pi_{A1}, \pi_{A0}, \pi_{B1}, \pi_{B0}) = 0.$$

The second part of the statement follows simply from the fact that if we multiply all π values by a constant $c > 1$ so that they remain feasible (i.e., none of π values goes above one), the interview efficiency of the policy $c\pi$ and the policy π are the same. □

Note that though it seems counter-intuitive, it might be the case

$$\pi_{A0} = \text{argmax}(\pi_{A1}, \pi_{A0}, \pi_{B1}, \pi_{B0})$$

and/or

$$\pi_{A1} = \text{argmin}(\pi_{A1}, \pi_{A0}, \pi_{B1}, \pi_{B0}).$$

D AN FPTAS ALGORITHM FOR SCREENING PROCESSES WITH SAME POLICY FOR ALL GROUPS

Here, we devise a slightly different DP algorithm. Instead of running the DP algorithm for each group separately (as in Section 4.3), we run a single DP algorithm for all groups simultaneously. Hence, all policies $\{\pi_X\}_{X\in\mathcal{X}}$ are the same. In our DP approach, we use the same discretization technique and only consider powers of $(1-\epsilon)$.

[4] Note that the interview efficiency is always at most 1 and the trivial Equalized Odds policy that bypasses all tests has interview efficiency q.

Solving the DP. Consider the first level, $i = 1$. For any given parameters $\{j_{X,0}, j_{X,1}\}_{X\in\mathcal{X}}$, where for each group $X \in \mathcal{X}$, $j_{X,0} \in [0, L_{\text{fpr}}]$ and $j_{X,1} \in [0, L_{\text{tpr}}]$, $M[1, \{j_{X0}, j_{X1}\}_{X\in\mathcal{X}}]$ = true iff the following has a feasible solution.

$$\tau^1_{X0}x + (1-\tau^1_{X0})y \le (1-\epsilon)^{j_{X0}} \text{ and,}$$
$$\tau^1_{X1}x + (1-\tau^1_{X1})y \ge (1-\epsilon)^{j_{X1}} \quad \forall X \in \mathcal{X} \tag{11}$$

Next, we describe the update rule for $i > 1$. For any $X \in \mathcal{X}$, $\text{fpr}_X \in [0, \ell_{\text{fpr}}]$ and $\text{tpr}_X \in [0, \ell_{\text{tpr}}]$,

$$M[i+1, \{\text{tpr}_X, \text{fpr}_X\}_{X\in\mathcal{X}}] = \bigvee_{\{j_{X1}, j_{X0}\}_{X\in\mathcal{X}} \in \mathcal{F}_{i+1}} M[i, \{\text{tpr}_X - j_{X1}, \text{fpr}_X - j_{X0}\}_{X\in\mathcal{X}}]$$

where $\mathcal{F}_{i+1}$ is the set of $\{j_{X1} \le \text{tpr}_X, j_{X0} \le \text{fpr}_X\}_{X\in\mathcal{X}}$ for which the following system of linear inequalities has a feasible solution $\forall X \in \mathcal{X}$

$$\tau^{i+1}_{X1}x + (1-\tau^{i+1}_{X1})y \ge (1-\epsilon)^{j_{X1}}, \tau^{i+1}_{X0}x + (1-\tau^{i+1}_{X0})y \le (1-\epsilon)^{j_{X0}} \tag{12}$$

Lemma D.1. *For any $i \in [k]$, if there exists an Equal Opportunity policy π treating all groups similarly, with true positive rate $\text{t}_{X,i} \ge L_{\text{tpr}}/(1-\epsilon)^{i-1}$, false positive rate $\text{f}_{X,i}$ for $X \in \mathcal{X}$, then there exist $\{j_{X1}, j_{X0}\}_{X\in\mathcal{X}}$ such that $M[i, \{j_{X1}, j_{X0}\}_{X\in\mathcal{X}}]$ = true, where for each $X \in \mathcal{X}$, $(1-\epsilon)^{j_{X1}} \ge \text{t}_{X,i} \cdot (1-\epsilon)^{i-1}$ and $(1-\epsilon)^{j_{X0}} \le \min\{1, \max\{L_{\text{fpr}}, \text{f}_{X,i}\}/(1-\epsilon)^{i-1}\}$.*

In other words, if the policy π exists then the DP approach finds a policy with true positive rate at least $(1-\epsilon)^{j_{X1}}$ and false positive rate at most $(1-\epsilon)^{j_{X0}}$ for each $X \in \mathcal{X}$.

Proof. The proof is by induction. For the base case ($i = 1$), let $\text{t}_{X,1}$ and $\text{f}_{X,1}$ denote the true positive rate and the false positive rate of π by the end of level 1 for each group $X \in \mathcal{X}$. The existence of π guarantees that the system of inequalities Eq. (11) with ($j_{X0} = \lfloor \log_{1-\epsilon} \text{f}_{X,1} \rfloor, j_{X1} = \lceil \log_{1-\epsilon} \text{t}_{X,1} \rceil \le \ell_{\text{tpr}}$) has a feasible solution. More precisely, by setting ($x_X = \pi_{X1}, y_X = \pi_{X0}$), $\forall X \in \mathcal{X}$,

$$\tau^1_{X0}x_X + (1-\tau^1_{X0})y_X = \text{f}_{X1} \le (1-\epsilon)^{\lfloor \log_{1-\epsilon} \text{f}_{X1} \rfloor} = (1-\epsilon)^{j_{X0}}$$
$$\tau^1_{X1}x_X + (1-\tau^1_{X1})y_X = \text{t}_{X1} \ge (1-\epsilon)^{\lceil \log_{1-\epsilon} \text{t}_{X1} \rceil} = (1-\epsilon)^{j_{X1}}$$

Next, we consider $i > 1$ and we assume that the claim holds for all values $i' < i$. For each $X \in \mathcal{X}$, let $M_{X,i} := \tau^i_{X1}\pi^i_{X1} + (1-\tau^i_{X1})\pi^i_{X0}$ and $N_{X,i} := \tau^i_{X0}\pi^i_{X1} + (1-\tau^i_{X0})\pi^i_{X0}$. Note that for each $X \in \mathcal{X}$, $\text{t}_{X,i} = \text{t}_{X,i-1} \cdot M_{X,i}$ and $\text{f}_{X,i} = \text{f}_{X,i-1} \cdot N_{X,i}$.

By the induction hypothesis and considering the first $i-1$ levels in the pipeline, since $\text{t}_{X,i-1} \ge \text{t}_{X,i} \ge L_{\text{tpr}}/(1-\epsilon)^{i-1} > L_{\text{tpr}}/(1-\epsilon)^{i-2}$ and $\text{f}_{X,i-1} \ge \text{f}_{X,i}$, there exist $j'_{X1} \in [0, L_{\text{tpr}}]$ and $j'_{X0} \in [0, L_{\text{fpr}}]$ such that $M[i-1, \{j'_{X1}, j'_{X0}\}_{X\in\mathcal{X}}]$ = true and $(1-\epsilon)^{j'_{X1}} \ge \text{t}_{X,i-1} \cdot (1-\epsilon)^{i-2}$ and $(1-\epsilon)^{j'_{X0}} \le \min\{1, \max\{L_{\text{fpr}}, \text{f}_{X,i-1}\}/(1-\epsilon)^{i-2}\}$. More precisely, the algorithm finds a policy $\bar{\pi}$ with true positive rate at least $(1-\epsilon)^{j'_{X1}}$ and false positive rate at most $(1-\epsilon)^{j'_{X0}}$ for each $X \in \mathcal{X}$.

Next, for each $X \in \mathcal{X}$, by setting ($\bar{\pi}^i_{X1} = \pi^i_{X1}, \bar{\pi}^i_{X0} = \pi^i_{X0}$) and ($j_{X1} := \text{argmin}_j\{(1-\epsilon)^j \le \text{t}_{X,i}(\bar{\pi})\}, j_{X0} := \text{argmax}_j\{(1-\epsilon)^j \ge \text{f}_{X,i}(\bar{\pi})\}$),

$$\begin{aligned}
(1-\epsilon)^{j_{X1}} &> (1-\epsilon)\cdot \text{t}_{X,i}(\bar{\pi}) \\
&= (1-\epsilon)\cdot \text{t}_{X,i-1}(\bar{\pi}) \cdot M_{X,i} && \triangleright \text{ by definition of } j_{X1} \\
&\ge (1-\epsilon)\cdot(1-\epsilon)^{j'_{X1}} \cdot M_{X,i} && \triangleright\ \text{t}_{X,i-1}(\bar{\pi}) \ge (1-\epsilon)^{j'_{X1}} \\
&\ge \text{t}_{X,i-1}\cdot(1-\epsilon)^{X,i-1}\cdot M_{X,i} && \triangleright \text{ induction hypothesis} \\
&= \text{t}_{X,i}\cdot(1-\epsilon)^{i-1}.
\end{aligned}$$

Similarly,

$$\begin{aligned}
(1-\epsilon)^{j_{X0}} &< \min\{1, \frac{\text{f}_{X,i}(\bar{\pi})}{1-\epsilon}\} \\
&= \min\{1, N_{X,i}\cdot\frac{\text{f}_{X,i-1}(\bar{\pi})}{1-\epsilon}\} && \triangleright \text{ by definition of } j_{X0} \\
&\le \min\{1, (1-\epsilon)^{j'_{X0}}\cdot\frac{N_{X,i}}{1-\epsilon}\} && \triangleright\ \text{f}_{X,i-1}(\bar{\pi}) \le (1-\epsilon)^{j'_{X0}} \\
&\le \min\{1, \frac{\max\{L_{\text{fpr}}, \text{f}_{X,i-1}\}}{(1-\epsilon)^{i-2}}\cdot\frac{N_{X,i}}{1-\epsilon}\} && \triangleright \text{ induction hyp.} \\
&\le \min\{1, \frac{\max\{L_{\text{fpr}}, \text{f}_{X,i}\}}{(1-\epsilon)^{i-1}}\}
\end{aligned}$$

which completes the proof. □

Lemma D.2. *For an accuracy parameter ϵ and lowerbounds on the false positive rate, L_{fpr}, and the true positive rate, L_{tpr}, the (single policy) DP algorithm runs in time $O(\frac{k\log^{2|\mathcal{X}|}(1/L_{\text{tpr}})\log^{2|\mathcal{X}|}(1/L_{\text{fpr}})}{\epsilon^{4|\mathcal{X}|}})$ and finds a policy π with true positive rate at least $(1-\epsilon)^{k-1}\cdot\text{t}_X$ and false positive rate at most $\min\{1, \max\{L_{\text{fpr}}, \text{f}_X\}/(1-\epsilon)^{k-1}\}$ for each $X \in \mathcal{X}$.*

Proof. The size of table is $O(k\ell^{|\mathcal{X}|}_{\text{tpr}}\ell^{|\mathcal{X}|}_{\text{fpr}})$ and updating each entry in the table takes $O(\ell^{|\mathcal{X}|}_{\text{tpr}}\ell^{|\mathcal{X}|}_{\text{fpr}})$. Hence, the total runtime to compute all entries in the DP table is

$$O(k\ell^{2|\mathcal{X}|}_{\text{tpr}}\ell^{2|\mathcal{X}|}_{\text{fpr}}) = O(\frac{k\log^{2|\mathcal{X}|}(1/L_{\text{tpr}})\log^{2|\mathcal{X}|}(1/L_{\text{fpr}})}{\epsilon^{4|\mathcal{X}|}}).$$

Now we apply the DP approach and by Lemma D.1, the solution returned by the algorithm has the true positive rate and the false positive rate satisfying the guarantee of the statement. □

Implications of DP. Here, similarly to Section 4.3, we present FPTAS algorithms for the single policy setting with various pipeline efficiency objective using the modified DP approach described above when the number of different protected groups in the population is a fixed constant; $|\mathcal{X}| = O(1)$.

Theorem D.3. *Consider a k-stage screening process with parameters $(u, q, \tau, \mathcal{X})$ where $|\mathcal{X}| = O(1)$. For any policy π, let $f_\alpha(\pi) = \text{recall}(\pi) + \alpha\cdot\text{precision}(\pi)$ where $\alpha > 0$. Given an accuracy parameter ϵ, there exists an FPTAS that runs in time $O(\frac{k^{4|\mathcal{X}|}\log^{2|\mathcal{X}|}(1/\epsilon)}{\epsilon^{4|\mathcal{X}|}})$ and finds an Equal Opportunity policy π treating all groups similarly such that $f_\alpha(\pi) \ge (1-\epsilon)f_\alpha(\pi^*)$ where π^* maximizes f_α over Equal Opportunity policies treating all groups similarly.*

Theorem D.4. *Consider a k-stage screening process with parameters $(u, q, \tau, \mathcal{X})$ where $|\mathcal{X}| = O(1)$. For any policy π, let $g_\alpha(\pi) = 1/\text{recall}(\pi) + \alpha/\text{precision}(\pi)$ where $\alpha > 0$. Given an accuracy parameter ϵ, there exists an FPTAS that runs in time $O(\frac{k^{4|\mathcal{X}|}\log^{2|\mathcal{X}|}(1/\epsilon)}{\epsilon^{4|\mathcal{X}|}})$ and finds an Equal Opportunity policy π treating all groups similarly such that $g_\alpha(\pi) \le (1+\epsilon)g_\alpha(\pi^*)$ where π^* minimizes g_α over Equal Opportunity policies treating all groups similarly.*

The proof of above theorems are identical to Theorem 4.10 and Theorem 4.11.

E ADDITIONAL DETAILS IN LINEAR COMBINATION COUNTER EXAMPLES

In this section, we show that one cannot "locally score" tests when determining the optimum policy (the policy that maximizes a linear combination of precision and recall). Specifically, we give a setting with three levels of tests t_1, t_2, t_3 such that if only the first two levels t_1 and t_2 are available, then the optimal solution is to use t_1 and bypass t_2, but if t_3 is also available then the optimal solution is to bypass t_1 and use t_2 and t_3. Therefore, the question of how to best use two tests may depend on what tests are available at other levels. Note that in this example there is only one group and we do not have fairness constraints.

First, we show the following useful property of optimal policies for a pipeline where the first level has test statistics $(1/2, 0)$ and all other levels have test statistics $(1-\delta, 1/2)$.

Lemma E.1. *In any k-stage pipeline where the first stage has test statistics $(1/2, 0)$ and the rest of the stages have tests with statistics $(1-\delta, 1/2)$, the optimal policy is of the form $(1, \pi_0^1), \cdots, (1, \pi_0^k)$.*

Proof. By Lemma 4.2, if the False Positive rate is non-zero, in the optimal policy, for every $i \in [k]$, $(1-\pi_1^i)\pi_0^i = 0$. Next, we show that in this setting with only one group, for every $i \in [k]$, $\pi_1^i = 1$. Suppose that there exists a level $i \in [k]$ such that $\pi_0^i = 0$. Then, if $\pi_1^i < 1$, by increasing π_1^i to 1, the True Positive rate and False Positive rate increase by the same factor. Therefore, the precision remains unchanged and the recall increases; hence, the pipeline efficiency strictly increases.

Next, we consider the case where the optimal policy has precision one (i.e., its False Positive is zero). In any such policy, $\pi_0^1 = 0$. Note that once the precision is 1, the optimal policy maximizes recall. Hence, the optimal policy is to fully use t_1 ($\pi_1^1 = 1, \pi_0^1 = 0$) and bypass the rest of tests (for every $1 < i \le k$, $\pi_1^i = \pi_0^i = 1$). □

Theorem E.2. *When the objective is to maximize a linear combination of precision and recall in a multi-stage screening process, there exist test parameters T and base rate p such that the maximal score policy switches when more tests become available. Specifically, when only tests t_1 and t_2 are available, the optimal policy is to use t_1 and bypass t_2 $((1,0),(1,1))$, but if test t_3 is also available, the optimal policy is to bypass t_1 and use t_2 and t_3 $((1,1),(1,0),(1,0))$.*

Proof. Consider base rate $p = P(x = 1) = 1/2$ and test $t_1 = (\tau_1, \tau_0) = (1/2, 0)$ and test $t_2 = t_3 = (1-\delta, 1/2)$. Let $\delta = \frac{1}{100}$. The linear objective function is $f(\pi) = \text{recall}(\pi) + 2 \cdot \text{precision}(\pi)$. Next, we consider two cases: (1) $k = 2$ and (2) $k = 3$.

Case 1: Two test ($k = 2$). By Lemma E.1, the optimal policy is of form $((1, \pi_0^1), (1, \pi_0^2))$. By numerical analysis[5], the local optimum policies (w.r.t. f) are $((1,0),(1,1))$ and $((1,1),(1,0))$. Next, we compute the score of these two policies: $f((1,0),(1,1)) = 2.5$ and $f((1,1),(1,0)) < 2.32$. Hence, in this case, the optimal policy is to fully use t_1 and bypass t_2, i.e., $((1,0),(1,1))$.

Case 2: Three tests ($k = 3$). Similarly to the previous case, the optimal policy for the given pipeline efficiency objective is of form $((1, \pi_0^1), (1, \pi_0^2), (1, \pi_0^3))$. By numerical analysis, the local optimum policies (w.r.t. f) are $((1,0),(1,1),(1,1))$ and $((1,1),(1,0),(1,0))$. Next, we compute the score of these two policies:

$$f((1,0),(1,1),(1,1)) = 2.5 \text{ and } ((1,1),(1,0),(1,0)) > 2.57.$$

This time, the optimal policy is to bypass t_1 and fully use t_2, t_3, i.e. $((1,1),(1,0),(1,0))$.

Therefore, while in the first setting (*only t_1 and t_2 are available*) the optimal policy is to fully use t_1 and bypass t_2, once t_3 becomes available, the optimal policy changes to bypass t_1 and fully use t_2 and t_3. □

[5] Using WolframAlpha.