

# Markovian Search with Socially Aware Constraints

---

Decisions that significantly impact lives, like hiring and admissions, have historically shown biases against certain societal groups, resulting in a lack of diversity and representation. The emergence of algorithmic hiring offered a chance to break this cycle, leading to the development of numerous hiring tools and platforms aiming to reduce biases and promote fairness, diversity, and inclusion. A significant challenge in algorithmic decision-making is the potential bias in the data used to make decisions. A school of thought to achieve socially desirable outcomes in the presence of such disparities is to influence the algorithm’s decisions by imposing certain socially aware constraints on its induced statistical behavior. These constraints, which are typically held on average and, therefore, are called ex-ante constraints, can capture various measures of fairness and diversity in the underlying decision processes. The sub-field of algorithmic fairness in computer science introduced and studied several such notions in machine learning problems such as classification and regression and provided means of satisfying those on-average constraints. However, search and selection processes can involve operational intricacies, which makes the underlying problem fundamentally different from classification or regression.

Motivated by the above, we study a general class of constrained sequential search problems for selecting multiple candidates from a pool that belongs to different societal groups. We focus on search processes under ex-ante constraints primarily motivated by inducing societally desirable outcomes such as attaining demographic parity among different groups, achieving diversity through quotas, or subsidizing disadvantaged groups within budget. We start with a canonical search model, known as the Pandora’s box model [Weitzman (1979)], under a single affine constraint on the probability of selection and inspection of each candidate. We show that the optimal policy for such a constrained problem retains the index-based structure of the optimal policy for the unconstrained one but potentially randomizes between two policies that are dual-based adjustments of the unconstrained problem; thus they are easy to compute and economically interpretable.

Building on these insights, we consider the richer class of search processes, such as search with rejection and multistage search, that can be modeled by joint Markov scheduling (JMS) [Dumitriu et al., 2003; Gittins, 1979]. Broadly, in our approach, we concentrate on the anticipated frequency of visits to each state within each category, and we apply arbitrary affine or convex restrictions on this frequency vector. This method allows us to accommodate different interpretations of fairness, both at the group level and for individuals, as outlined in existing literature, such as the 4/5 rule or ensuring a minimum level of welfare. Imposing general affine and convex ex-ante constraints, we give a primal-dual algorithm to find the near-feasible and near-optimal policy. This algorithm, too, randomizes over index-based policies; this time, over a polynomial number of policies whose indices are dual-based adjustments to the Gittins indices of the unconstrained JMS. Our algorithmic developments, while involving many intricacies, rely on a simple, yet powerful observation: There exists a relaxation to the Lagrange dual function of these constrained optimization problems that admit index-based policies akin to the original unconstrained ones.

We complement our theoretical framework with numerical simulations on synthetic data. Focusing on demographic parity, quota, and budget in the Pandora’s box model with two demographic groups, we numerically study both the *short-term* and the *long-term* impacts of imposing socially aware constraints. Concretely, we take the perspective that the prior available for the quality (or value) of each candidate and its realization upon inspection can be viewed as only a surrogate for the true quality of the candidate which is unobservable in the short run. As such, to measure the short-term impact, we rely on these observable signals. On the other hand, given the opportunity, the true qualities are revealed in the long run. The latter is inspired by the theory of downstream outcomes introduced in Becker (2010) and is further studied in the literature on the economics of labor discrimination [Canay et al., 2020].