

# Trading off Relevance and Revenue in the Jobs Marketplace: Estimation, Optimization and Auction Design

Farzad Pourbabae, Sophie Yanying Sheng, Peter McCrory, Luke Simon, Di Mo  
*LinkedIn Corporation*  
*Sunnyvale, California, USA*

## Abstract

We study the problem of position allocation in job marketplaces, where the platform determines the ranking of the jobs for each seeker. The design of ranking mechanisms is critical to marketplace efficiency, as it influences both short-term revenue from promoted job placements and long-term health through sustained seeker engagement. Our analysis focuses on the tradeoff between revenue and relevance, as well as the innovations in job auction design. We demonstrated two ways to improve relevance with minimal impact on revenue: incorporating the seekers preferences and applying position-aware auctions.

## 1. Introduction

In a jobs marketplace, when a seeker submits a job search request, a query of relevant jobs will be retrieved and shown to them according to a particular order. In economic terms, the posters of these jobs are the bidders and the list of all available positions/slots on the seeker’s page constitutes the goods. Since there is more than one position, the auction effectively involves multiple goods, each potentially having different values for both the bidders and the seeker.

Prevalent ranking mechanisms employ an “auction score” (see, e.g. [9]) that integrates a monetization component from the bidders’ side with a relevance component from the seekers’ side. Optimizing solely for revenue prioritizes maximizing the monetary component, often leading to less relevant jobs being ranked higher due to their higher submitted bids — a phenomenon known as bid dominance. On the other hand, optimizing for match quality focuses exclusively on the relevance component, potentially sacrificing revenue. This introduces a fundamental tension between the short-term gains of maximizing revenue and the long-term benefits of the users’ satisfaction in any ranking problem. Jobs with fewer suitable candidates are naturally incentivized to allocate higher budgets to secure visibility, whereas jobs with broader appeal face less pressure to do so.

At a high level, the tension between revenue and relevance involves balancing short-term immediate revenue generation with building long-term gains through improved user engagement and satisfaction. An optimal job matching auction should achieve this balance by adopting a dynamic approach that considers both short- and long-term objectives.

Formally, when a seeker  $i \in I$  submits a job search request, a query of  $j \in [n]$  relevant jobs is retrieved, ranked, and displayed to them.<sup>1</sup> To achieve this, we construct an auction score that incorporates features of both the seeker and the job posters.

A common approach to produce the auction score is to modulate the bid value submitted by job  $j$ , denoted by  $b_{ij}$ , with the estimated click-through probability  $\bar{\pi}_{ij}$ :

$$\bar{s}_{ij} = S(b_{ij} \cdot \bar{\pi}_{ij}).$$

Under the Generalized-First-Price (GFP) auction, jobs are ranked based on their auction score, that is if  $\bar{s}_{ij_1} \geq \bar{s}_{ij_2} \geq \dots \geq \bar{s}_{ij_n}$ , then the posting with index  $j_k$  will sit in the  $k$ -th place.

While this mechanism is effective due to its simplicity, it is not optimal for two reasons: (i) There is considerable dispersion in discounted-cumulative-Gain (DCG) between seekers due to the inherent sparsity in labor markets. A significant portion of (local) labor markets consists of a small number of hirers and seekers, highlighting the need to capture the heterogeneity in seekers’ experience within the auction score; (ii) Sequential ranking based on the position-unaware scores overlooks the fact that the click-through rate of the same job for the same seeker varies depending on its position.

To address the first drawback, we augment the auction score with a seeker component that captures the relevance of job  $j$  to seeker  $i$ , thereby incorporating the seeker’s preferences. This enhancement is detailed in Section 2. To mitigate the second drawback, we introduce *position-aware* auction scores. Specifically, in addition to the seeker index  $i$  and the job index  $j \in [n]$ , the auction score will also depend on the position  $k \in [n]$  in which the job is displayed to the seeker (see Section 3). These types of position-aware auctions are

<sup>1</sup>We adopt the convention that  $[n] := \{1, \dots, n\}$ , representing the set of  $n$  available jobs.

widely used in the advertising environment, where ad placement and performance depend heavily on the position within search results or web pages [11, 6, 4].

## 2. Relevance-Revenue Tradeoff

In this section, we study how we can augment the score function to account for the relevance of the job to the seeker. Formally, we introduce  $\bar{\mu}_{ij}$ , a model-based estimation of the relevance of job  $j$  to seeker  $i$ . We refer to this variable as **eRelevance**. It represents the quality of the match from the seeker’s perspective, capturing how well the job aligns with their preferences and qualifications. The monetary gain of adding one more unit of relevance to seeker  $i$  is denoted by  $w_i$ , and referred to as the **Seeker-Weight**. Hence, the augmented auction score now features two components:

$$\bar{s}_{ij} = S \left( \underbrace{b_{ij} \cdot \bar{\pi}_{ij}}_{\text{poster component}}, \underbrace{w_i \cdot \bar{\mu}_{ij}}_{\text{seeker component}} \right).$$

In the following, we present two approaches to calibrate the Seeker-Weight. In Section 2.1, we leverage experimental data to estimate how the relevance measurement responds to variations in the Seeker-Weight<sup>2</sup>, thereby quantifying its causal impact. In Section 2.2, we propose a dynamic framework to determine the optimal level of the Seeker-Weight that maximizes long-term gains. Overall, the first approach focuses on causal estimation, while the second approach adopts an optimization-based strategy.

### 2.1 Causal Estimation

A higher Seeker-Weight increases the influence of eRelevance on the final auction score, resulting in rankings that better reflect the seeker’s preferences. This improved alignment boosts seeker engagement, which ultimately encourages more active participation from both seekers and posters.

A good starting point is to increase Seeker-Weight for segments of the job market for whom relevance is undesirably low. A question arises: By how much should we increase the Seeker-Weight? To answer this question we need two things: a target relevance level and a parsimonious model that maps Seeker-Weight to a relevance metric.

**Parsimonious Model:** Consider the following model of relevance measurement as a function of the Seeker-Weight for a particular seeker segment  $g$  under the conditions that  $\alpha \in (0, 1)$ :

$$\text{Rel}_g = F_g(\text{Seeker-Weight}) = z_g \times \text{Seeker-Weight}^\alpha$$

<sup>2</sup>Relevance measurement refers to well-established metrics for evaluating ranking quality in recommender systems, such as Normalized Discounted Cumulative Gain (NDCG), Mean Average Precision (MAP), and Mean Reciprocal Rank (MRR).

The parameter  $z_g$  represents segment-specific efficiency in converting Seeker-Weight to relevance. Segments with higher  $z_g$  achieve higher relevance levels for any given Seeker-Weight, reflecting heterogeneity in baseline relevance efficiency. The parameter  $\alpha \in (0, 1)$  captures the elasticity of relevance with respect to the Seeker-Weight, featuring diminishing marginal returns in terms of relevance per unit increase in the Seeker-Weight.

We then used an experimental setting to capture the causal effect of the Seeker-Weight on the relevance metric. If the above model is a good approximation to how the marketplace works, then the observed lift in relevance across seeker segments caused by a higher Seeker-Weight should be unrelated to the baseline level of relevance (as measured in the control group). Moreover, the lift in relevance should only be a function of the change in Seeker-Weight (which is observed) and the unobserved elasticity term  $\alpha$ . Indeed, this is exactly what we find: Among seeker segments where the Seeker-Weight was increased from a *moderate* to a *high* level, relevance rose significantly on average. Conversely, among seeker segments where the weight was decreased from a *moderate* to a *low* level, relevance fell by a significant level.

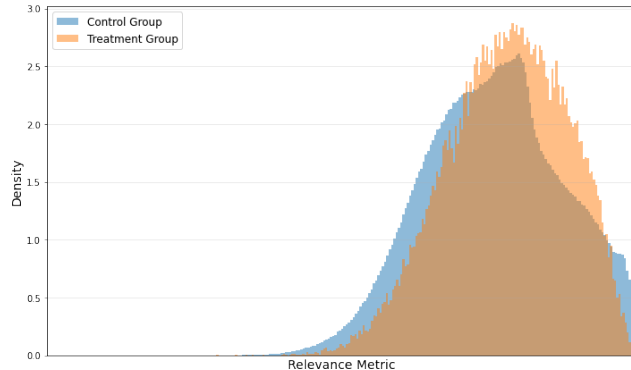


Figure 1: Dispersion in relevance metric in experiment

We calibrate the parsimonious model against experiment data to find elasticity of relevance with respect to the Seeker-Weight, namely,  $\alpha$ . We use this  $\alpha$  to tune personalized Seeker-Weights in jobs search auctions, and A/B test the new weights. Figure 1 shows that there is reduced relevance metric dispersion under new Seeker-Weights, as well as elevated relevance level for the mean and the median seeker. This improvement in relevance metric does not come at a cost of short-term revenue.

### 2.2 Dynamic Optimization

Having pinned down the causal impact of Seeker-Weight on the seeker’s relevance, the platform’s problem is now to choose an optimal Seeker-Weight that maximizes long-run discounted “gains”. Lower weights on relevance increase immediate gains but potentially

degrade future gains by reducing seeker engagement and poster participation. We model this dynamic problem by introducing a time dimension to the Seeker-Weight, i.e.,  $w_i(t)$ , and assume a discount factor  $\delta < 1$  across time periods. In each period  $t$ , the platform observes the submitted bids  $b_{ij}(t)$ 's, eRelevances  $\bar{\mu}_{ij}(t)$ 's, and sets the personalized Seeker-Weight  $w_i(t)$ . These jointly determine the auction score, which feeds into a particular mechanism of the slot assignment, and result in current gain. The platform's objective is to maximize the expected discounted gains:

$$\max_{\{w_i(t)\}} \mathbb{E} \left[ \sum_{t=0}^{\infty} \delta^t \text{Gain}(t) \right].$$

The dynamics are governed by a stochastic kernel  $Q$  that maps the current history of the states and Seeker-Weights to the distribution over future states. Formally, the state vector  $x_i(t)$  includes current bids and eRelevances, i.e.,  $x_i(t) = \{b_{ij}(t), \bar{\mu}_{ij}(t) : j \in [n]\}$ . The history  $h_i(t)$  includes all past states, that is  $h_i(t) = \{x_i(\tau) : \tau \leq t\}$ . Suppose  $\sigma$  is a policy function that maps the current history  $h_i(t)$  to the imputed Seeker-Weight. We can write the value of any such policy as:

$$V_{\sigma}(h_i(t)) = \text{Gain}(x_i(t); \sigma(h_i(t))) + \delta \mathbb{E}^Q[V_{\sigma}(h_i(t+1)) | h_i(t), \sigma(h_i(t))].$$

However, this formulation faces two key challenges. First, the transition kernel  $Q$  is unknown and must be learned. Second, the increasing length of the history makes optimization over the space of all policy functions computationally infeasible. To address the second challenge, we impose a Markovian structure: future states depend only on the current state and Seeker-Weight, rather than the entire history. This simplification enables us to focus on Markov strategies that map current states directly to weights. The simplified Bellman equation becomes:

$$V(x_i(t)) = \max_{\sigma} \{ \text{Gain}(x_i(t); \sigma(x_i(t))) + \delta \mathbb{E}^Q[V(x_i(t+1)) | x_i(t), \sigma(x_i(t))] \}. \quad (1)$$

We solve this recursively through value function iteration, starting with  $V^{(0)}(\cdot) \equiv 0$  and updating according to:

$$\sigma^{(m+1)}(x) = \operatorname{argmax}_{\sigma(x)} \left\{ \text{Gain}(x; \sigma(x)) + \delta \mathbb{E}^Q[V^{(m)}(x') | x, \sigma(x)] \right\},$$

where  $m$  is the index of the iteration.

Standard dynamic programming guarantees convergence to optimal values and policies as  $m \rightarrow \infty$ , provided  $\delta < 1$ . The key empirical challenge in implementing this framework is that the transition kernel  $Q$  must be learned from data. We address this through reinforcement learning methods that adaptively update beliefs about  $Q$  while optimizing the Seeker-Weights.

### 3. Auction Design

In the jobs marketplace, when a seeker submits a search request, a query of  $n$  relevant jobs is retrieved and shown according to a particular order. In economic terms, the posters of these jobs are the "bidders", and the list of all available positions on the seeker's page constitutes the "goods". Since there are more than one position, we are effectively selling multiple goods, with potentially different values to the bidders, as well as the seeker. That means we are in the world of *multiple/heterogeneous-goods* auction design. Similar to every other mechanism design problem, here we need to determine two functions, that jointly define the overall mechanism:

*Allocation rule*: also known as permutation, that prescribes the order of the postings shown to the seeker.

*Transfer rule*: also known as pricing, prescribes the price per slot that is charged to each bidder.

Suppose we want to assign multiple slots at the same time to a group of bidders. The well-developed theory of selling a single item to a group of buyers [7, 10] is not directly applicable to the case of allocating multiple goods. A common feature of single-good auctions is that the item should be allocated to the buyer who values it the most — the allocation is greedy. This allocation is both revenue maximizing for the seller and welfare maximizing for the buyers. However, this intuition does not immediately carry over to multiple goods, mainly because of two reasons: capacity constraints (agents cannot own multiple goods), and the possibility of bundling. We highlight these issues in the following examples, where there are two buyers Alice and Bob, and two items  $x, y$  to be allocated.

**Example 1** (Capacity constraints): Let  $v_a = (0, 1)$  and  $v_b = (2, 3)$  be the valuations of Alice and Bob respectively for goods  $x$  and  $y$ . If we allocate each good to the buyer who values it the most, then both items should be sold to Bob. In case that he has capacity constraints, then selling  $y$  to Bob and  $x$  to Alice is the optimal plan. Here is an example where allocating each good to the buyer who values it the most fails.

**Example 2** (Bundling): We add a third coordinate into the buyers' valuation vectors that represents their joint valuation of getting the bundle  $x$  and  $y$ . Let  $v_a = (2, 1, 7)$  and  $v_b = (0, 3, 4)$ . Allocating each good to the person who values it the most means that item  $x$  should be sold to Alice, and  $y$  should go to Bob, yielding the overall surplus of  $2 + 3 = 5$ . But if we sell both items to Alice the net surplus would be 7. Thus, once again greedily allocating each object according to the *maximum-demand* schedule (or first-price principle) fails.

To account for the valuation heterogeneity of different slots, we adopt position-aware scores. Specifically, a job poster  $j \in [n]$  submits a bid value  $b_{ij}$  for the available slots  $k \in [n]$  in the search results of seeker  $i \in I$ . The click-through probabilities are position-aware, denoted as  $\pi_{ijk}$ , representing the likelihood that seeker  $i$  clicks on job  $j$  displayed in position  $k$ . Additionally, the rele-

vance of job  $j$  in position  $k$  for seeker  $i$  is captured by the position-aware eRelevance,  $\mu_{ijk}$ . These position-aware parameters can be used to generate position-aware auction scores:

$$s_{ijk} = S(b_{ij} \cdot \pi_{ijk}, w_i \cdot \mu_{ijk}).$$

Below, in Section 3.1, we briefly review the application of GFP, which is inherently a position-unaware assignment method, to the above scores. Then, in Section 3.2, we extend our analysis to the assignment problem with position-aware scores, exploring the full-fledged bipartite matching approach, also known as VCG. Lastly, in Section 3.3 we present some simulations to compare these two assignment algorithms.

### 3.1 Position-Unaware and Generalized First-Price Ranking

Suppose instead of position-aware scores  $\{\pi_{ijk}, \mu_{ijk}, s_{ijk}\}$ , we only have their position-unaware estimates denoted by  $\{\bar{\pi}_{ij}, \bar{\mu}_{ij}, \bar{s}_{ij}\}$ , which can be found by taking the average over all slots  $k \in [n]$ . GFP prescription is to rank jobs based on their slot-averaged scores. Specifically, if  $\bar{s}_{ij_1} \geq \bar{s}_{ij_2} \geq \dots \geq \bar{s}_{ij_n}$ , then the job with index  $j_k$  is assigned to the  $k$ -th position. The running time is typical of any sorting algorithm,  $\mathcal{O}(n \log n)$ . As before, the transfer rule typically follows a ‘‘pay-for-performance model,’’ where performance is defined by metrics such as impressions, click-throughs, or job applications, which ultimately determine the revenue.

### 3.2 Position-Aware and Bipartite Matching

In this section, we study how the jobs should be ranked on the seeker’s query when we have position-aware scores as input. Given that the seeker is fixed, and to reduce clutter, we drop the seeker index  $i$  from all forthcoming variables, but it should be understood from the context that all of them depend on the index of the seeker.

Recall that  $j \in [n]$  refers to the job posting identifier, and there are  $n$  available positions indexed by  $k \in [n]$ . A matching function (or an allocation rule) is represented by an one-to-one function  $M : [n] \rightarrow [n]$ , mapping  $n$  job postings to  $n$  available slots. In particular,  $M(j) = k$  means that posting  $j$  will take the  $k$ -th position on the list, and vice-versa  $M^{-1}(k) = j$ , means that the position  $k$  is filled by the job posting  $j$ .

Our goal is to find a matching  $M$  that maximizes the sum of (position-aware) auction scores, that is, to solve the following combinatorial maximization problem:

$$\max_M \sum_{j \in [n]} s_{jM(j)}.$$

This problem falls under the classic bipartite matching for which various computational algorithms were proposed. Hungarian approach (a.k.a Kuhn-Munkres algorithm [8]) provides an exact solution in  $\mathcal{O}(n^3)$  time, and  $\varepsilon$ -approximate algorithms are known to achieve lower

computational complexity [1, 2]. The theory of selling heterogeneous goods is first studied in economics literature by [5]. The authors showed that repeated application of the  $\varepsilon$ -approximate algorithm, by gradually sending  $\varepsilon \rightarrow 0$ , results in the implementation of the optimal matching function with the *minimum Walrasian prices*, known as VCG transfer rule [12, 3]. More specifically, their algorithm not only outputs the optimal matching but also, as a byproduct, determines the prices for each item in the assignment problem. We only use the allocation rule of VCG (that is the same as the Hungarian output) to determine the optimal position-aware permutation of jobs, and similar to GFP, we use a pay-for-performance model to charge the bidders.

### 3.3 Simulations

In this section, we present simulations to compare the performance of GFP and VCG rankings based on two criteria: revenue and relevance. Specifically, we repeat the simulation over a fixed number of seekers  $|I|$ , and then take the sample average. Below, we explain the simulation steps for each  $i \in I$ :

- (i) Draw  $n$  random samples from the underlying distribution of bid values:  $\{b_{i,j} : j \in [n]\}$ .
- (ii) Draw  $n^2$  position-aware random samples from the underlying distribution of pCTR:  $\{\pi_{ijk} : j \in [n], k \in [n]\}$ .
- (iii) Draw a Seeker-Weight  $w_i \in \mathbb{R}_+$ .
- (iv) Draw  $n^2$  position-aware random samples from the underlying distribution of eRelevance:  $\{\mu_{ijk} : j \in [n], k \in [n]\}$ .
- (v) Compute the position aware scores  $s_{ijk} = S(b_{ij} \cdot \pi_{ijk}, w_i \cdot \mu_{ijk})$ .
- (vi) Compute the average values over  $n$  slots:  $\bar{\pi}_{ij}$ ,  $\bar{\mu}_{ij}$ , and  $\bar{s}_{ij}$ .
- (vii) Find the GFP ranking based on the average scores, and compute its resultant revenue  $\text{Rev-GFP}_i$  and relevance  $\text{Rel-GFP}_i$ .
- (viii) Find the VCG ranking based on the position-aware scores, and compute its resultant revenue  $\text{Rev-VCG}_i$  and relevance  $\text{Rel-VCG}_i$ .

Once we have the revenue and relevance per seeker  $i$ , we take their averages over the population of seekers  $I$  and compare them across the two matching algorithms (GFP and VCG):

$$\begin{aligned} \text{Rev-GFP} &:= \frac{1}{|I|} \sum_{i \in I} \text{Rev-GFP}_i, & \text{Rev-VCG} &:= \frac{1}{|I|} \sum_{i \in I} \text{Rev-VCG}_i \\ \text{Rel-GFP} &:= \frac{1}{|I|} \sum_{i \in I} \text{Rel-GFP}_i, & \text{Rel-VCG} &:= \frac{1}{|I|} \sum_{i \in I} \text{Rel-VCG}_i \end{aligned}$$

In Figure 2, the average revenue for both assignment algorithms are plotted as a function of the auction depth, i.e., the number of bidders or participating jobs. Both graphs show an increasing trend with  $n$ , which aligns with the intuition that more bidders lead




Figure 2: Average revenue per seeker vs. # of jobs




Figure 3: Average relevance per seeker vs. # of jobs

to greater competition and, consequently, higher revenue. Moreover, GFP consistently extracts higher average revenue than VCG. In Figure 3, the average relevance of both algorithms are plotted as a function of the number of participating jobs. VCG, by design, ranks jobs based on position-aware scores, whereas GFP relies on average auction scores across all positions. As a result, VCG outperforms GFP in terms of relevance.

These two graphs underscore the tradeoff between revenue and relevance. Despite incorporating seeker eRelevance into the auction score, the greedy nature of GFP continues to extract more revenue than the position-sensitive approach of VCG. However, on the relevance front, VCG demonstrates superior performance, highlighting its strength in aligning with seeker preferences.

## 4. Concluding Remarks

This paper makes two contributions to the design of job marketplace auctions. First, we capture the heterogeneity inherent in the jobs marketplace that leads to a wide dispersion of seekers' experiences under a uniform auction score formulation. We address this dispersion by augmenting the auction score with a parameter that allows for seeker personalization. We calibrate the parameter using (i) a parsimonious model of relevance as a function of the new parameter, and (ii) a reinforcement learning framework to optimally balance the tension between short- and long-term gains, incorporating dynamic behavioral responses of the seekers and posters into the optimization problem. Second, as an alternative to GFP, we explore a position-aware auction for multiple slots, demonstrating improvements in relevance even without altering the auction score formulation. We expect our approach to be widely applicable in other two-sided marketplaces where similar tensions between immediate monetization and long-term platform health exist.

## References

- [1] Dimitri Bertsekas. *Network optimization: continuous and discrete models*, volume 8. Athena Scientific, 1998.
- [2] Dimitri Bertsekas. New auction algorithms for the assignment problem and extensions. *Results in control and optimization*, 14:100383, 2024.
- [3] Edward H. Clarke. Multipart pricing of public goods. *Public choice*, pages 17–33, 1971.
- [4] Riccardo Colini-Baldeschi, Julián Mestre, Okke Schrijvers, and Christopher A. Wilkens. The ad types problem. In *International Conference on Web and Internet Economics*, pages 45–58. Springer, 2020.
- [5] Gabrielle Demange, David Gale, and Marilda Sotomayor. Multi-item auctions. *Journal of political economy*, 94(4):863–872, 1986.
- [6] Patrick Hummel. Position auctions with dynamic resizing. *International journal of industrial organization*, 45:38–46, 2016.
- [7] Vijay Krishna. *Auction theory*. Academic press, 2009.
- [8] Harold W. Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- [9] Ning Ma, Mustafa Ispir, Yuan Li, Yongpeng Yang, Zhe Chen, Derek Zhiyuan Cheng, Lan Nie, and Kishor Barman. An online multi-task learning framework for google feed ads auction models. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3477–3485, 2022.
- [10] Roger B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

- [11] Hal R. Varian. Position auctions. *international Journal of industrial Organization*, 25(6):1163–1178, 2007.
- [12] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.