



THE LONDON SCHOOL
OF ECONOMICS AND
POLITICAL SCIENCE ■

Najafi-Zangeneh, Saeed, Shams-Gharneh, Naser & Gossner, Olivier (2025)
Two-sided matching with bounded rationality: a stochastic framework for personnel
selection. *Mathematics*, 13(19). <https://doi.org/10.3390/math13193173>

<https://researchonline.lse.ac.uk/id/eprint/129899/>

Version: Published Version

Licence: [Creative Commons: Attribution 4.0](#)

[LSE Research Online](#) is the repository for research produced by the London School
of Economics and Political Science. For more information, please refer to our [Policies](#)
page or contact lseresearchonline@lse.ac.uk

Article

Two-Sided Matching with Bounded Rationality: A Stochastic Framework for Personnel Selection

Saeed Najafi-Zangeneh ¹, Naser Shams-Gharneh ^{1,*} and Olivier Gossner ^{2,3}¹ Industrial Engineering Department, Amirkabir University of Technology, Tehran 1591634311, Iran; saeed.najafi@aut.ac.ir² CNRS-CREST, École Polytechnique, 91120 Palaiseau, France; olivier.gossner@polytechnique.edu³ Department of Mathematics, London School of Economics, London WC2A 2AE, UK

* Correspondence: nshams@aut.ac.ir

Abstract

Personnel selection represents a two-sided matching problem in which firms compete for qualified candidates by designing job-offer packages. While traditional models assume fully rational agents, real-world decision-makers often face bounded rationality due to limited information and cognitive constraints. This study develops a matching framework that incorporates bounded rationality through the Quantal Response Equilibrium, where firms and candidates act as probabilistic rather than perfect optimizers under uncertainty. Using Maximum Likelihood Estimation and organizational hiring data, we validate that both sides display bounded rational behavior and that rationality increases as the selection process advances. Building on these findings, we propose a two-stage stochastic optimization approach to determine optimal job-offer packages that balance organizational policies with candidate competencies. The optimization problem is solved using particle swarm optimization, which efficiently explores the solution space under uncertainty. Data analysis reveals that only 23.10% of low-level hiring decisions align with rational choice predictions, compared to 64.32% for high-level positions. In our case study, bounded rationality increases package costs by 26%, while modular compensation packages can reduce costs by up to 25%. These findings highlight the cost implications of bounded rationality, the advantages of flexible offers, and the systematic behavioral differences across job levels. The framework provides theoretical contributions to matching under bounded rationality and offers practical insights to help organizations refine their personnel selection strategies and attract suitable candidates more effectively.

Keywords: two-sided matching; quantal response equilibrium; bounded rationality; two-stage stochastic approach; particle swarm optimization

MSC: 90-10



Academic Editor: David Carfi

Received: 7 September 2025

Revised: 27 September 2025

Accepted: 1 October 2025

Published: 3 October 2025

Citation: Najafi-Zangeneh, S.;

Shams-Gharneh, N.; Gossner, O.

Two-Sided Matching with Bounded Rationality: A Stochastic Framework for Personnel Selection. *Mathematics* **2025**, *13*, 3173. <https://doi.org/10.3390/math13193173>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The core strength of any company lies in its human resources, which serve as the foundation and driving force behind its achievements. Managers devote significant time to recruiting skilled personnel, recognizing their pivotal role in the company's success [1]. In today's competitive employment landscape, it is essential for companies to optimize their job-offer packages, including salary, bonuses, remote work options, and other incentives, to attract prospective employees. At the same time, organizations must strike a balance in managing the costs of these packages to ensure long-term sustainability and profitability.

Designing offers that are both competitive and cost-effective enables organizations to stand out in attracting high-caliber candidates. Personnel selection is typically modeled as a two-sided matching market, where organizations and candidates are assumed to act as rational agents making decisions on opposite sides [2]. Organizations' job-offer items can affect their ranking among job seekers. However, the importance of these items may vary for each candidate. Similarly, organizations assess each candidate's competencies differently, depending on their strategies. Nevertheless, due to insufficient understanding of the decision environment and the physiological and psychological factors influencing agents, achieving fully rational conditions is challenging. To quantify agents' rationality in two-sided matching, a model of bounded rationality using the logit quantal response equilibrium (LQRE) can be employed. In this model, players are not perfect optimizers and uncertainty regarding the actions of others [3].

Personnel selection is commonly modeled as a two-sided matching market, where organizations and candidates are assumed to act as rational decision-makers. However, in practice, this assumption rarely holds. Real-world hiring involves uncertainty, incomplete information, and behavioral factors that prevent agents from being perfect optimizers. Existing research has primarily approached job-offer optimization from the perspective of post-hiring salary adjustments, leaving pre-hiring package design underexplored. Moreover, while the concept of bounded rationality has been studied extensively in economics and operations research, its integration into matching theory for recruitment remains limited. Prior work has often focused on laboratory experiments or behavioral explanations, rather than optimization frameworks that can guide practical decision-making in real hiring contexts.

This gap motivates the present study. We develop a two-stage stochastic optimization framework that incorporates bounded rationality, modeled through the Quantal Response Equilibrium (QRE), into a two-sided matching setting for personnel selection. By empirically estimating bounded rationality parameters from organizational data and embedding them into an optimization model, our approach provides actionable tools for designing job-offer packages that balance cost efficiency with candidate preferences. In doing so, the paper addresses the lack of integration between matching theory, bounded rationality modeling, and pre-hiring package optimization, offering both theoretical contributions and practical guidance for organizations.

2. Literature Review

Two-sided matching markets arise in many contexts, such as labor recruitment, marriage, and university admissions. In these settings, agents are divided into two disjoint sets, each ranking potential partners on the opposite side. A matching is deemed stable if no pair of agents would both prefer to deviate from their assigned partners to match with each other. Stability can be further refined to the side-optimal case for one side of the market, where every agent on that side weakly prefers the resulting assignment to any alternative stable matching.

Gale and Shapley [4] introduced a practical algorithm for finding optimal stable matchings, laying the foundation for a vast literature. Roth and Sotomayor [5] provide a comprehensive survey, and numerous studies have examined the problem from a computational perspective (Pu and Yuan [6]; Jiang and Guo [7]; Manlove, Irving, and Iwama [8]; Bansal and Gupta [9]; Ashlagi, Braverman, and Hassidim [10]; Hatfield and Kominers [11]; Manlove [12]). A central assumption in much of this literature is that participants behave as fully rational decision-makers [13]. While analytically convenient, this assumption relies on two strong premises: agents are perfect optimizers and possess complete knowledge of

their counterparts' decision models. In practice, particularly in personnel selection, these assumptions often do not hold.

To address this, bounded rationality can be incorporated through quantal choice [14], in which all potential matches are considered but alternatives yielding higher utility are chosen with higher probability [15]. The concept of bounded rationality has been studied for decades as a way to explain observed deviations from perfect rationality (Simon [16]; Nelson [17]; Rubinstein [18]). Among approaches to formalize bounded rationality, quantal-response functions are widely used, particularly for deriving analytical results [19]. McKelvey and Palfrey [20] introduced these functions and their equilibrium properties for normal-form games, and Goeree, Holt, and Palfrey [21] provide a detailed primer on Quantal Response Equilibrium (QRE). Subsequent studies have explored properties such as existence, uniqueness, and comparative statics for both discrete and continuous action spaces (Chen, Friedman, and Thisse [22]; Anderson, Goeree, and Holt [23]).

The predictive power of QRE has been demonstrated across various game-theoretic settings (Goeree and Holt [24]), including alternating-offer bargaining [25], coordination games [26], auctions [27], traveler's dilemma games [28], and pricing contracts (Lim and Ho [29]; Ho and Zhang [30]). Quantal-response functions have also found applications in operations management. For example, Su [31] applied them in a newsvendor setting; Liu, Methapattara, and Wynter [32] in IT service revenue management; Shang and Liu [33] in time-based capacity competition; Chen, Su, and Zhao [15] in capacity-allocation problems; Di and Liu [34] in route choice behavior; Canbolat [19] in clearing service systems; and Roemer, Müller, and Voigt [35] and Khanlarzade and Farughi [36] in supply chain contexts.

Despite this extensive literature, the integration of bounded rationality into matching mechanisms remains limited. Most studies focus on experimental validation rather than optimization for practical applications. Pais and Pintér [37] conducted one of the first experimental studies on matching with boundedly rational agents, analyzing Deferred Acceptance (DA), Boston, and Top Trading Cycles (TTC) mechanisms under varying informational conditions. Echenique, Wilson, and Yariv [38] estimated a logit-QRE in a two-sided matching market without deriving testable hypotheses. Dreyfuss, Heffetz, and Rabin [39] showed that loss aversion can explain boundedly rational choices, with the LQRE converging to truthful preference revelation in strategy-proof mechanisms. Alcalde-Unzu, Klijn, and Vorsatz [40] further demonstrated that QRE accurately describes individual behavior and resulting matchings in constrained settings.

While many studies address bounded rationality, relatively few examine the design of job-offer items, and most focus on post-hiring adjustments. Wallace Jr and Steuer [41] formulated a linear multi-objective model incorporating internal and external equity, requiring iterative adjustments to achieve satisfactory salary structures. Kwak, Allen, and Schniederjans [42] employed goal programming to analyze annual salary adjustments in large organizations. Garcia-Diaz and Hogg [43] developed mixed-integer linear programming models to assign executives to salary increase categories and schedule subsequent raises. Kassa [44] proposed a goal programming model to determine minimum and maximum wages for grades, minimizing average structure wages relative to external markets and costs. Tremblay, Piché-Meunier and Dubois [45] introduced a multi-objective a posteriori optimization model to maximize equity dimensions while minimizing costs. Building on Najafi-Zangeneh, Shams [46], who developed a two-stage MILP model for job-offer optimization, this study extends their approach by incorporating bounded rationality in decision-making.

Despite the breadth of this literature, key gaps remain. First, the integration of matching theory and bounded rationality is limited, with existing work (Pais and Pintér [37]; Echenique, Wilson and Yariv [38]) focusing on experiments rather than practical opti-

mization. Second, job-offer optimization research has largely addressed post-hiring salary adjustments and assumed rational agents, neglecting pre-hiring package design and behavioral realities. Third, no prior work unifies matching theory, bounded rationality, and job-offer optimization into a single framework, despite its relevance in real-world hiring markets. Fourth, empirical calibration of bounded rationality parameters for hiring contexts is lacking, limiting actionable guidance. Finally, although QRE has been applied successfully in operations management (Su [31]; Chen et al. [15]), its application to personnel selection and job-offer optimization remains unexplored.

These gaps motivate the development of a two-stage stochastic optimization framework that integrates QRE-based bounded rationality into matching theory, providing organizations with practical tools to design cost-efficient, behaviorally informed job packages.

Contributions of This Paper

To address the gaps highlighted earlier, the study presents six main contributions:

- **Integrating Bounded Rationality into Matching:** Extends classical two-sided matching by incorporating Quantal Response Equilibrium (QRE) to account for incomplete information in organizational hiring decisions.
- **Empirical Estimation of Parameters:** Provides the first empirical calibration of bounded rationality in recruitment using three years of data, showing learning effects vary by position level.
- **Computational Framework:** Introduces a two-stage stochastic optimization approach to handle non-convexities from QRE, with particle swarm optimization outperforming other heuristics.
- **Quantifying Costs of Bounded Rationality:** Demonstrates a measurable cost premium (~26%) in optimal job packages due to bounded rationality, with implications for HR budgeting.
- **Actionable Decision Support:** Offers insights on package design, showing single-item packages are costlier and lower-priority roles rely more on bounded rational behavior, sometimes reducing costs.
- **Dynamic Matching with Learning Effects:** Models evolving organizational expertise across repeated hiring cycles through time-varying bounded rationality parameters.

3. Materials and Methods

This section presents the proposed methodology, integrating both qualitative and quantitative analyses. We first define the problem, establish criteria and weights for candidates and organizations, and formulate the corresponding utility functions within a stepwise matching mechanism. Next, we analyze the case study data to formalize bounded rationality and conduct comparative analyses. Parameters are then calibrated using MLE, capturing decision-maker behavior over time, across job types, and under varying rationality levels. The core two-stage optimization framework models selection with bounded rational decision-makers via the Deferred Acceptance Algorithm (DAA) and then identifies optimal hiring packages through a scenario-based stochastic approach. Nonlinearities are addressed by generating multiple scenarios and applying PSO, chosen for its efficiency in handling non-convex, mixed-variable problems. Although a number of novel metaheuristics such as RFO, RSA, and COA have recently been introduced, the no free lunch theorem [47] demonstrates that no single optimization method is best for all problem types. We selected PSO because it is a proven and computationally efficient algorithm for solving non-convex optimization problems with mixed variable types, as in our model [48]. Moreover, PSO's population-based search allows for effective exploration and exploitation of the solution space [49]. To validate this choice, we compared PSO with

genetic algorithms and simulated annealing under the same computational budget, and the results showed that PSO achieved faster convergence and better solution quality in our case study. The final package composition is then determined via Mixed-Integer Linear Programming (MILP). The methodology is validated on a real case study, with results analyzed through sensitivity analyses. Figure 1 summarizes the overall workflow.




Figure 1. General structure and the process of the proposed model.

3.1. Problem Definition and Two-Sided Matching

This paper addresses a two-sided matching problem in the recruitment context, where a set of organizations $I = \{1, 2, \dots, m\}$ and a set of candidates $J = \{1, 2, \dots, n\}$ express preferences for one another. We specifically examine a one-to-one matching scenario where each candidate can only be paired with a single organization and vice versa, creating a matching function μ :

$$\mu : I \cup J \rightarrow I \cup J \text{ Where } \mu(i) = j \text{ if and only if } \mu(j) = i \quad (1)$$

Definition 1. *Preference Relations:* Each organization $i \in I$ has preferences over candidates based on their competencies, while each candidate $j \in J$ has preferences over organizations based on job offer packages. Under bounded rationality, these preferences are not deterministic orderings, but follow probabilistic QRE distributions.

Definition 2. *Stable Matching:* A matching μ is stable if there exists no blocking pair (i, j) where

$$\text{both } j \succ_i \mu(i) \text{ and } i \succ_j \mu(j). \quad (2)$$

In the final stage of hiring, organizations evaluate interested candidates using a Deferred Acceptance Algorithm adapted to bounded rationality. The key challenge is

to form stable matches under probabilistic choice behavior while optimizing job-offer packages within budget constraints. Candidate selection depends on both candidate utility and competing offers, with weights determined by role importance. Bounded rationality turns this process into a stochastic optimization problem, where successful matching probabilities depend on behavioral parameters estimated from real hiring data.

3.1.1. Deferred Acceptance Algorithm (DAA)

The Deferred Acceptance Algorithm (DAA) addresses the stable matching problem, where participants aim to secure mutually preferred partners. In its basic form, the problem involves two groups with ranked preferences, and a stable matching is achieved when no pair of individuals prefer to match with each other rather than with their assigned partners [4]. In the hiring context, the process begins with candidates applying to their most preferred organizations, which, in turn, place the highest-utility candidate on a waiting list while rejecting others. Rejected candidates then apply to their next preferred choice in subsequent rounds, and organizations compare new applicants with those on their waiting lists, retaining the candidate with the higher-utility score. This iterative process continues until no further changes occur, resulting in a stable matching.

3.1.2. Identification of Desired Criteria

A comprehensive literature review was conducted to identify potential criteria, which were then refined and prioritized using the Fuzzy Delphi Method through three rounds of consultation with ten university professors and senior recruitment managers. The final set of criteria reflects key organizational expectations of candidates: education level, ranging from no high school diploma to doctorate; general skills, from basic to advanced in areas such as interpersonal, organizational, and problem-solving abilities; English language proficiency, from low to native; work experience, from none to more than twenty years; and psychological test performance, from low to excellent. On the job-offer side, experts highlighted four decisive items: competitive salary, bonus plans, remote work options, and flexible working hours. While remote work and flexibility carry no direct financial cost, they involve intangible organizational trade-offs. In this study, these factors are treated as variables whose values are optimized for the selected organization.

To assess the relative importance of these criteria for both candidates and organizations, the Fuzzy Analytic Hierarchy Process (FAHP) was applied. This method employed pairwise comparisons with triangular fuzzy numbers, and fuzzy weights were derived by multiplying each criterion's value in the pairwise comparison matrix by the inverse of the sum of the other criteria's values. Based on the evaluations of hiring managers in the case study, this process yielded fuzzy weights that capture the importance of each criterion from both perspectives.

3.1.3. Payoff Formulations

To capture preferences, we apply the concept of equivalent utility, where job-offer package items generate utility for candidates and candidate competencies such as education, work experience, language proficiency, general skills, and psychological tests, generate utility for organizations, all measures are quantified and standardized on a common scale.

$$C_j = [C_{j1}, C_{j2}, \dots, C_{jn}] \quad (3)$$

Candidates' competencies generate utilities for organizations, calculated as the weighted sum of these competencies.

$$W^{(i)} = [W_1^{(i)}, W_2^{(i)}, \dots, W_n^{(i)}] \quad (4)$$

$$U_{ij} = \sum_{\forall f} W_f^{(i)} C_{jf} + \varepsilon \quad (5)$$

For organization i , the competency weight vector determines the utility U_{ij} of candidate j , with noise terms ε modeled as extreme-value distributed to capture unobserved preferences. The job-offer package of organization i is represented as a vector including salary, bonus plans, remote work options, and flexible hours.

$$D_i = [d_{i1}, d_{i2}, \dots, d_{if}] \quad (6)$$

$$d_{if_{min}} \leq d_{if} \leq d_{if_{max}} \quad (7)$$

$$Dn_i = [dn_{i1}, dn_{i2}, \dots, dn_{im}] \quad (8)$$

Each item of the job-offer package is bounded and normalized, and the package generates candidate utilities calculated as the inner product of the normalized offer vector and the candidate's weight vector.

$$W'^{(j)} = [W'_1{}^{(j)}, W'_2{}^{(j)}, \dots, W'_m{}^{(j)}] \quad (9)$$

$$U'_{ij} = \sum_{k \leq f} W'_k{}^{(j)} dn_{ik} + \varepsilon' \quad \forall i \in I - \{q\} \quad (10)$$

$W'^{(j)}$ represents the importance weights of job-offer items for candidate j , and U'_{ij} denotes the utility of organization i for candidate j , excluding the target organization's utility, with a noise term ε' analogous to ε .

3.2. Bounded Rationality Model

The Quantal Response Equilibrium (QRE) extends standard random utility models and generalizes Nash equilibrium to account for noisy, boundedly rational behavior, where choices are positively, but not perfectly, related to expected payoffs [50]. This framework is next applied to model bounded rationality.

3.2.1. Data

The study analyzes personnel selection data from a large Iranian holding company with three EPC subsidiaries (A, B, and C) over 36 months, covering 4400 final interviews. Company A, the primary focus, conducted 1390 final interviews and hired 730 candidates, with 660 chosen by hiring managers. Across all companies, 511 candidates were interviewed for high-level (185) and low-level (316) positions. Data sources included resumes, interview transcripts, evaluation scores, recruitment forms, offered packages, and hiring manager interviews. Table 1 summarizes the study's personnel selection data.

Table 1. Summary of data.

Company	Final Interviews	Hiring Manager Recruits	Common Interviewees	High-Level Jobs	Low-Level Jobs
A	1853	660	511	185	316
B	1358	539	427	142	282
C	1189	423	352	117	235

3.2.2. Bounded Rationality Based on QRE

The concept of bounded rationality, introduced by Simon [51], describes decision-making under cognitive limitations due to constraints in information, knowledge, and

computational capacity. While central to behavioral studies, perfect rationality is uncommon in practical personnel selection. As shown in Figure 2, only 23.10% of hiring decisions for low-level jobs align with rational choice predictions, compared to 64.32% for high-level positions, which illustrates the limitations of rational choice theory in real-world hiring.




Figure 2. Distribution of choices in the high-level jobs (**left**) and low-level jobs (**right**), compared to the rational choice prediction.

Our analysis highlights distinct patterns in personnel selection between job levels.

For high-level positions, the average chosen utility is 0.923, significantly higher than 0.750 for low-level roles. Moreover, 75.14% of high-level hires have utilities above 90, compared to 32.91% for low-level hires, showing that managers prioritize higher-utility candidates for senior roles. This behavior contrasts with rational choice theory, which predicts that the highest-utility candidate will always be selected, regardless of position.

This systematic gap reflects the higher bounded rationality (β) observed in low-level roles, where managers appear more willing to accept satisfactory candidates rather than optimize fully. The result highlights the importance of modeling bounded rationality explicitly, since standard rational models would fail to capture this discrepancy.

The data also reveals a learning effect over time. Over three years, managers' hiring decisions improve, as shown in Figure 3, which plots the average utilities chosen by Organization A's manager. Both low-level and high-level job graphs show an upward trend, indicating that managers progressively select higher-utility candidates as they gain experience, further deviating from static rational choice predictions. These trends are consistent with the decreasing bounded rationality parameter ($\delta > 0$) estimated in our model, confirming that experience reduces randomness in decisions over repeated hiring cycles. From a managerial perspective, this suggests that training and continuity of hiring personnel can improve recruitment efficiency over time.

McKelvey and Palfrey [20] applied QRE to model bounded rationality in game theory, where players choose based on relative expected utility. Using logit formulas, the probability of selecting a candidate or organization incorporates bounded rationality parameters (β), capturing the influence of biases on decisions. The probability of selecting candidate j by organization i is given by the following formula.

$$P_{ij} = \frac{e^{\frac{U_{ij}}{\beta_i}}}{\varepsilon + \sum_{k \in J} e^{\frac{U_{ik}}{\beta_i}}} \quad (11)$$

Here, β_i represents the bounded rationality parameter of organization i . Similarly, the probability of candidate j selecting organization i is given by a corresponding formula.

$$\hat{p}_{ij} = \frac{e^{\frac{u'_{ij}}{\beta_j}}}{\varepsilon + \sum_{f \in I} \left[e^{\frac{u'_{fj}}{\beta_j}} \right]} \quad (12)$$




Figure 3. Time trends in the data for the high-level jobs (**left**) and low-level jobs (**right**).

The bounded rationality parameter β_j represents candidate j 's level of bounded rationality, reflecting errors due to incomplete information about the other side's utility. Each agent's utility includes explicit and implicit components, and lack of awareness of the implicit part affects decision-making. β varies across agents to capture heterogeneity: as $\beta \rightarrow \infty$, choices are random; as $\beta \rightarrow 0$, choices become fully rational, consistent with QRE. Following Chen, Su, and Zhao [15], β can change over time, capturing the learning effect observed in Figure 3, where managers' choices increasingly align with rational predictions. We, therefore, propose a time-dependent $\beta_{(r)}$ that decays exponentially with accumulated decision-making experience.

$$\beta_{(r)} = \beta + (\alpha - \beta)e^{-\delta(r-1)} \quad (13)$$

The bounded rationality parameter $\beta_{(r)}$ decreases exponentially over round r , from an initial value α to a final value β , reflecting the manager's learning at rate δ .

3.3. Maximum Likelihood Estimation (MLE)

This section fits the model to real-world data using MLE, following Chen, Su, and Zhao [15]. It estimates key parameters, α , β , δ , and noise ε , and computes $f_{ij(r)}$, the probability that organization i selects candidate j at round r , based on the QRE model (Equations (11) and (12)).

$$f_{ij(r)} = \frac{e^{\frac{u_{ij(r)}}{\beta_{i(r)}}}}{\varepsilon + \sum_{k \in J} \left[e^{\frac{u_{ik(r)}}{\beta_{i(r)}}} \right]} \quad (14)$$

Applying the natural logarithm ($\ln$) to the likelihood function and simplifying yields a log-likelihood expression:

$$L((\alpha, \beta, \delta, \varepsilon | \hat{U})) = \prod_{r=1}^R \prod_{j=1}^J (1 - \varepsilon) \times f_{ij(r)} + \varepsilon \times \frac{1}{N_r} \quad \forall i \in I \quad (15)$$

A uniform error term, ε , is included to account for random selections by managers, occurring with probability ε and uniform likelihood $\frac{1}{N_r}$ across candidates. In a perfect fit, $\varepsilon = 0$, but practically it is expected to be small and positive. The model is fitted by maximizing the log-likelihood function over parameters α , β , δ , and ε .

$$Ln(L((\alpha, \beta, \delta, \varepsilon | \hat{N})) = \sum_{r=1}^R \sum_{j=1}^J Ln \left((1 - \varepsilon) \times \frac{e^{U_{ij(r)}/\beta_i(r)}}{\sum_{k \in J} [e^{U_{ik(r)}/\beta_i(r)}]} + \varepsilon \times \frac{1}{N_r} \right) \quad \forall i \in I \quad (16)$$

The model incorporates all four parameters and is fitted using data from both high- and low-level jobs. Numerical computations are performed with Python's SciPy library, and the resulting parameter estimates are detailed in Table 2.

Table 2. Estimation results for high-level and low-level jobs' parameters.

Parameters	α	β	δ	ε	Log-Likelihood
High-level job	8.2542	0.0432	0.0150	0.004	−166.20
Low-level job	10.248	0.1244	0.0047	0.005	−551.19

The bounded rationality parameters were estimated for both low- and high-level jobs (Table 2). δ values, representing learning, are positive but small: $\delta = 0.0150$ for high-level jobs and $\delta = 0.0047$ for low-level jobs, which indicates slower learning in low-level jobs. β values remain positive: $\beta = 0.0432$ for high-level jobs and $\beta = 0.1244$ for low-level jobs, indicating that some bounded rationality persists despite learning. Log-likelihood is higher in high-level jobs, suggesting a better fit for this mode. Overall, the analysis shows decision-makers deviate from perfect rationality, with stochastic behavior influenced by β and learning dynamics captured by δ . High-level jobs exhibit lower bounded rationality and faster learning, achieving 64.32% alignment with rational choice predictions, while low-level jobs show higher bounded rationality and slower learning, with 23.10% alignment, highlighting the importance of accounting for job-level differences in optimizing job offers. To mitigate sensitivity to data imperfections, the same framework can be estimated with (i) a robust log-likelihood that down-weights outliers, (ii) an ε —contaminated or heavy-tailed error for the logit shock to reduce leverage, and (iii) EM with multiple imputation for missing covariates. In practice, we also recommend nonparametric bootstrap confidence intervals and leave-one-group-out sensitivity checks to assess stability.

4. Modeling and Computational Study

Incorporating bounded rationality through the QRE transforms the Deferred Acceptance process from a deterministic mechanism into a probabilistic one. In the classical setting, agents are assumed to make perfectly rational choices, always selecting the highest-utility option available. In practice, however, hiring decisions are often influenced by incomplete information, uncertainty about competitors' actions, and cognitive limitations. QRE captures these realities by allowing both candidates and organizations to choose with probabilities that increase with utility but still permit occasional deviations from the optimal choice. As a result, the DAA retains its iterative structure, but the outcomes reflect the behavioral imperfections observed in real hiring environments. This adjustment preserves the theoretical logic of deferred acceptance while embedding a more realistic representation of decision-making under uncertainty.

The process of reaching final hiring decisions involves two main steps. Figure 4 provides a clear overview of the entire process. First, the Deferred Acceptance Algorithm (DAA) is applied to match organizations with candidates, incorporating bounded-rationality adjustments (Sections 3.2 and 3.3). Second, the optimal job offer for each matched

candidate is identified using particle swarm optimization (PSO) and refined with a linear update. Figure 4 summarizes the workflow: Phase 1—data preparation, Phase 2—DAA with QRE-based choice rules, and Phase 3—PSO and linear tuning to find the optimal package, showing stable matches and associated offer costs.




Figure 4. Algorithm Flowchart.

4.1. First Stage: Two-Sided Matching and Equivalent Utility Based on Bounded Rationality

In the first stage, the model applies the Deferred Acceptance Algorithm (DAA) to form stable matches between candidates and organizations, accounting for both preferences and bounded rationality. The algorithm ensures that no participant has an incentive to break an

existing match. Central to this stage is the concept of equivalent utility, which quantifies the value of each candidate for organizations and informs both the matching process and the creation of cost-efficient job-offer packages in the second stage.

$$\text{MIN } Z = U'_{eq_i} - rev_i * E(\alpha_{i_{last}}) \quad (17)$$

Subject to:

$$In_{ijt} = \begin{cases} 0 & \sum_{t' \leq t} ch'_{ijt's} \neq 1 \\ 1 & \sum_{t' \leq t} ch'_{ijt's} = 1 \end{cases}, \forall i \in I, \forall j \in J, \forall t \in T, \forall s \in S \quad (18)$$

$$\sum_{\forall t} In_{ijts} \leq 1 \quad (19)$$

$$P_{ijts} = \frac{e^{\frac{U_{ij(r)}}{\beta_i(r)}} * ch'_{ijts}}{\varepsilon + \sum_{\forall k \in J} \left[e^{\frac{U_{ik(r)}}{\beta_i(r)}} * ch'_{izts} \right]} \quad (20)$$

$$\hat{P}_{ijts} = \frac{e^{\frac{U'_{ij(r)}}{\beta_j(r)}} * \zeta_{ijts}}{\varepsilon + \sum_{\forall f \in I} \left[e^{\frac{U'_{fj(r)}}{\beta_j(r)}} * \zeta_{ijts} \right]} \quad (21)$$

$$\zeta_{ijts} = 1 - \left(\sum_{t' < t} In_{ijt's} \right) + ch_{ij(t-1)s} \quad (22)$$

$$r'_{jts} \sim U(0, 1) \quad (23)$$

$$\sum_{k=0}^{n-1} p'_{kjts} * ch'_{njts} \leq r'_{jts} \leq \sum_{k=0}^n p'_{kjts} * ch'_{njts} \quad (24)$$

$$\sum_{\forall i} ch'_{ijts} \leq 1 \quad (25)$$

$$\sum_{\forall i} ch'_{ijts} = \sum_{\forall i} p'_{ijts} \quad (26)$$

$$r_{its} \sim U(0, 1) \quad (27)$$

$$\sum_{q=0}^{m-1} p_{iqts} * ch_{iqts} \leq r_{its} \leq \sum_{q=0}^m p_{iqts} * ch_{iqts} \quad (28)$$

$$ch'_{ijts} \geq ch_{ijts} \quad (29)$$

$$\alpha_{i_{last}s} = \sum U_{ij} \cdot ch_{ij(t_{last})s} \quad (30)$$

$$\sum_{\forall i} ch_{ijts} \leq 1 \quad (31)$$

$$\sum_{\forall i} ch_{ijts} = \sum_{\forall i} p_{ijts} \quad (32)$$

$$U'_{eq_i} = \sum_{\forall j} ch_{ij t_{last}s} U'_{ij} \quad (33)$$

$$E(\alpha_{i_{last}}) = (\sum_{\forall s \in S} \alpha_{i_{last}s}) / n_S \quad (34)$$

$$In_{ij0s} = 0, \quad ch_{ij0s} = 0, \quad ch'_{ij0s} = 0 \quad (35)$$

The objective function in Equation (17) aims to minimize the gap between the utility required to attract candidates (U'_{eq_i}) and the utility the organization expects to gain, weighted by the position sensitivity parameter (rev_i). Equations (18) and (19) enforce interview constraints, ensuring that each candidate–organization pair can have at most one interview. Equations (20) and (21) introduce bounded-rational choice probabilities based on QRE, where β controls the degree of rationality: a smaller β implies choices driven more strongly by utility, while a larger β allows greater randomness. Equation (22) ensures candidates remain available only if they have not yet accepted another offer or remain on a waiting list.

Equations (23)–(26) govern the stochastic selection process for candidates, and Equations (27)–(32) extend this to organizations, enforcing two-sided consent where organizations may only select candidates who also chose them.

Finally, Equations (33)–(35) characterize the solution: (33) defines equivalent utility as the bridge to stage two, (34) computes expected values over all stochastic scenarios, and (35) specifies conditions so that the process begins with no prior matches or interviews.

4.2. Second Stage: Cost-Efficient Job Offer Design

In the second stage, the equivalent utility from stable matches is used to design cost-efficient job-offer packages, optimizing salaries and benefits to minimize organizational costs while keeping offers attractive to candidates.

$$in\ Cost_D = \sum_{\forall f} Cost_f d_f + \sum_{\forall f_{int}} Cost_{f_{int}} d_{f_{int}} + \sum_{\forall f_{bin}} Cost_{f_{bin}} d_{f_{bin}} \quad (36)$$

$$U'_{eq_i} = \sum_{\forall f} w'_f d_n f + \sum_{\forall f_{int}} w'_{f_{int}} d_n f_{int} + \sum_{\forall f_{bin}} w'_{f_{bin}} d_n f_{bin} \quad (37)$$

$$d_f = d_{min_f} + (d_{max_f} - d_{min_f}) d_n f, \quad \forall f \quad (38)$$

$$d_{f_{int}} = d_{min_{f_{int}}} + (d_{max_{f_{int}}} - d_{min_{f_{int}}}) d_n f_{int}, \quad \forall f_{int} \quad (39)$$

$$d_{f_{bin}} = d_n f_{bin}, \quad \forall f_{bin} \quad (40)$$

$$0 \leq d_n f \leq 1, \quad (41)$$

$$0 \leq d_n f_{int} \leq 1, \quad (42)$$

$$d_n f_{bin} \in \{0, 1\} \quad (43)$$

$$d_f, d_{f_{int}} \geq 0, d_{f_{int}} \in \mathbb{Z} \quad (44)$$

The second stage begins with the objective function in Equation (36), which seeks the lowest-cost job-offer package that meets the candidate's required desirability while maximizing organizational cost-effectiveness, as refined in Equation (37). Job-offer items are treated as decision variables: continuous (salary, bonus), integer (work hours), and binary (remote work). To normalize these items, Equations (38)–(40) are used according to the variable type, while Equations (41)–(44) constrain them within the $[0, 1]$ range. Table A1 (Appendix A) provides the full set of notations.

4.3. Particle Swarm Optimization

Since the first stage of the model is non-linear and probabilistic (Formulas (20) and (21)), the optimization problem is solved using the particle swarm optimization (PSO) algorithm. Originally introduced by Kennedy and Eberhart [52], PSO models the social behavior of animals such as bird flocking and fish schooling. Each candidate solution is treated as a particle with a velocity that updates iteratively, moving toward both its own best position (P_{best}) and the global best position (g_{best}) [53]. The overall search process is illustrated in Figure 5.




Figure 5. Flowchart of PSO Algorithm.

$$P_{best_l}^t = (x_l^* | f(x_l^*) = \min_{k=1,2,\dots,t} \left(\left\{ f(x_l^k) \right\} \right), \dots \text{where } l \in \{1, 2, \dots, N\}, \quad (45)$$

$$g_{best_l}^t = (x_*^t | f(x_*^t) = \min_{\substack{l=1,2,\dots,N \\ k=1,2,\dots,t}} \left(\left\{ f(x_l^k) \right\} \right), \quad (46)$$

Here, l denotes the particle index, t the current iteration, f the objective function to be minimized, x the position vector (candidate solution), and N the number of particles in the swarm. At each iteration $t + 1$, the velocity v and position x of particle l are updated using the following equations:

$$v_l^{t+1} = \omega v_l^t + c_1 r_1 (P_{best_l}^t - x_l^t) + c_2 r_2 (g_{best_l}^t - x_l^t), \quad (47)$$

$$x_l^{t+1} = x_l^t + v_l^{t+1} \quad (48)$$

Here, v is the velocity vector, ω the inertia weight controlling the balance between local exploitation and global exploration, r_1 and r_2 are random vectors uniformly distributed over $[0, 1]^D$ (with D denoting the problem's dimensionality), and c_1, c_2 are positive acceleration coefficients.

4.3.1. Algorithm Selection and Comparative Analysis

We conducted a comparative analysis of three metaheuristic algorithms, particle swarm optimization (PSO), genetic algorithm (GA), and simulated annealing (SA), to determine the most suitable method for our non-linear optimization problem with mixed decision variables. All three algorithms were implemented under comparable computational budgets and tested on our case study data. To ensure robustness, each algorithm was executed 30 times with different random seeds.

For the PSO, we used 50 particles with an inertia weight of $\omega = 0.5$ and acceleration coefficients $c_1 = c_2 = 1.5$. The GA was configured with a population size of 50, a crossover

rate of 0.8, a mutation rate of 0.1, and tournament selection with a size of 3. Finally, the SA employed an initial temperature of 100, a geometric cooling schedule with a rate of 0.98, and adaptive neighborhood sizing. The outcomes of this comparison are reported in Table 3.

Table 3. Comparative Performance of Metaheuristic Algorithms.

Performance Metrics	PSO	GA	SA
Convergence Iteration	15	28	35
Final Objective Value	−8.72	−8.68	−8.70
Initial Objective Value	−7.20	−7.20	−7.20
Improvement Rate (First 10 Iterations)	1.32	0.85	0.72
Stability After Convergence	High	Medium	Low
Oscillation During Search	Minimal	Moderate	High
Average Final Value	−8.71	−8.66	−8.67
Best Value Found	−8.72	−8.70	−8.71
Worst Value Found	−8.69	−8.58	−8.61
Standard Deviation	0.008	0.032	0.027

The empirical results highlight the superior performance of PSO. As shown in Figure 6, PSO reaches the optimal value of −8.72 within just 15 iterations, with major improvements occurring between iterations 2–8 (from −7.6 to −8.5). After iteration 15, it remains stable with no oscillations. GA, in contrast, shows stepped convergence with improvements at iterations 5, 12, 19, and 28, eventually reaching −8.68 but at a slower pace. SA exhibits the most erratic path due to its temperature-dependent acceptance, stabilizing only after 35 iterations at −8.70. The strong performance of PSO can be attributed to its velocity-based updates and continuous optimization nature, which enable efficient exploration early on and effective exploitation in later rounds.




Figure 6. Convergence Comparison of Metaheuristic Algorithms.

To further illustrate the performance consistency of the algorithms, Figure 7 presents box plots of the final objective values across 30 independent runs for PSO, GA, and SA.

These results confirm the summary statistics reported in Table 3. PSO shows the narrowest spread and highest stability, with almost no outliers, while GA and SA display greater dispersion and variability.




Figure 7. Distribution of Objective Values Across 30 Runs.

4.3.2. Model Limitations and Parameter Analysis

The computational complexity of the framework is driven by the matching algorithm with $O(n^2m)$ complexity, and the PSO with T iterations and P particles, leading to a total complexity of $O(TPn^2m)$. Empirical tests confirm that the framework performs efficiently for problems up to 100 candidates and 20 organizations (under 5 min). Larger cases may require decomposition or parallel computing. A sensitivity analysis of PSO parameters, presented in Table 4, examines their effects on solution quality and convergence speed.

Table 4. PSO parameters sensitivity analysis.

Parameter	Range Tested	Optimal Value	Impact on Performance
Inertia weight (ω)	[0.2, 0.9]	0.5	Lower values (<0.4) cause premature convergence; higher values (>0.7) slow convergence
Cognitive coefficient (c_1)	[0.5, 2.5]	1.5	Values outside [1.0, 2.0] significantly degrade performance
Social coefficient (c_2)	[0.5, 2.5]	1.5	Similar to c_1 ; balanced values ($c_1 \approx c_2$) work best
Population size	[20, 100]	50	<30 particles insufficient for exploration; >70 offers marginal improvement
Max iterations	[50, 200]	100	95% of runs converge by iteration 50; 100 ensures robustness

The robustness of the parameter settings was confirmed through 30 independent runs (Section 4.3.1), which demonstrated consistent performance and minimal variance in solution quality.

5. Discussion

5.1. Case Study

This section evaluates the model's efficiency using a real-world case involving three organizations and five candidates for high-level electrical engineering roles. Based on the data from the left side of Figure 2, the decentralized recruitment processes of Organizations A, B, and C are analyzed. The relevant profiles and importance weight vectors are summarized in Tables 5–7.

Table 5. Candidates' Profiles.

#	Education	Language Proficiency	General Skills	Work Experience	Psychological Test
Candidate 1	MSc	native	good	0–5	good
Candidate 2	PHD	medium	medium	10–15	low
Candidate 3	BSc	good	high	0–5	medium
Candidate 4	PHD	advanced	advanced	15–20	excellent
Candidate 5	BSc	advanced	high	0–5	good

Table 6. Competencies' weight importance vector of organizations.

#	Education	Language Proficiency	General Skills	Work Experience	Psychological Test
Organization A	0.40	0.10	0.10	0.25	0.15
Organization B	0.36	0.04	0.20	0.30	0.10
Organization C	0.31	0.14	0.16	0.14	0.25

Table 7. Items' weight importance vector of candidates.

#	Salary Level	Bonus Plans	Flexibility in Working Hours	Ability to Work Remotely	Salary Level
Candidate 1	0.50	0.20	0.20	0.10	0.50
Candidate 2	0.40	0.20	0.15	0.25	0.40
Candidate 3	0.35	0.10	0.25	0.30	0.35
Candidate 4	0.60	0.10	0.20	0.10	0.60
Candidate 5	0.55	0.15	0.15	0.15	0.55

The objective function focuses on minimizing Organization A's total package cost, while keeping the package items of Organizations B and C fixed. The cost details of these rival organizations are provided in Table 8.

Table 8. Organization B and C package items.

#	Salary (\$)	Bonus (\$)	Flexibility in Working Hours (h)	Working Remotely
Organization B	900	400	2	No
Organization C	800	500	1	Yes

We solve a case with three organizations and five candidates, where the decision variables are Organization A's package items. Two modes are tested: (i) both sides fully rational, and (ii) Organizations are boundedly rational while candidates remain rational (due to insufficient data to estimate candidate bounded rationality). Experiments ran on an Intel Core i7-11700K 2.4 GHz, 16 GB RAM, Windows 10 Pro, using Python 3.9.7 (with

NumPy 1.21.5), MIP and SciPy 1.7.3 for optimization, and Matplotlib 3.5.1 for visualization purposes. The case serves as a proof of concept, showing how bounded rationality affects matching outcomes and offer costs in a controlled setting; broader validation requires larger, more diverse datasets.

5.2. Results

5.2.1. Rational Mode

Table 9 presents the first-mode results, where both organizations and candidates are rational. With $rev_A = 1$ (average recruitment sensitivity), the optimal job-offer package for Organization A and its corresponding stable match are reported.

Table 9. The optimal job offer package in first mode.

#	Salary (\$)	Bonus (\$)	Flexibility in Working Hours (h)	Working Remotely	Stable Match
Organization A	1300	500	-	No	Candidate 4
Organization B	900	400	2	No	Candidate 2
Organization C	800	500	1	Yes	Candidate 1

The first-mode matching results show that Organization A successfully hires Candidate 4 with a package of \$1300 salary and \$500 bonus, excluding flexibility and remote work. This outcome highlights that direct income was the decisive factor for Candidate 4, who also ranked Organization A as a top preference, confirming a mutual match. Concurrently, Organization B matched with Candidate 2, and Organization C with Candidate 1, based on their respective packages and priority weights. The detailed candidate utilities (U'_{ij}) and organizational utilities (U_{ij}) are provided in Tables 10 and 11, respectively.

Table 10. Utility of Candidates (U'_{ij}).

	Organization A	Organization B	Organization C
Candidate 1	0.525	0.36	0.34
Candidate 2	0.46	0.30	0.54
Candidate 3	0.327	0.27	0.54
Candidate 4	0.42	0.30	0.339
Candidate 5	0.50	0.30	0.419

Table 11. Utility of Organizations (U_{ij}).

#	Candidate 1	Candidate 2	Candidate 3	Candidate 4	Candidate 5
Organization A	0.38	0.54	0.14	0.90	0.20
Organization B	0.32	0.54	0.18	0.90	0.21
Organization C	0.47	0.42	0.22	0.92	0.31

The robustness of the first mode was assessed by varying model parameters and assumptions to observe their impact on the outcomes. The results of these tests, specifically regarding the optimization of Organization A's job-offer package, are summarized in Table 12, providing insight into the model's sensitivity and stability.

Table 12. Sensitivity analysis for organization A's optimized package.

Scenarios		Optimal Package				
#		Salary (\$)	Bonus (\$)	Flexibility (Hours)	Remote Work	Match
1	$rev_A = 1$	1300	500	0	No	{A,4}, {B,2}, {C,1}
2	$rev_A = 0.5$	500	0	0	No	{A,2}, {B,4}, {C,1}
3	Reduced cost items 3 and 4 to $1/20$, $rev_A = 1$	967	500	0	Yes	{A,4}, {B,2}, {C,1}
4	Salary only, $rev_A = 1$	1634	0	0	No	{A,4}, {B,2}, {C,1}

When rev_A is reduced from 1 to 0.5, Organization A can still hire its preferred candidate, but the package costs only \$500, reflecting a lower willingness to invest in less critical roles. Reducing the costs of flexibility and remote work benefits to one-twentieth of their original values does not change the matching outcome, Organization A still chooses the same candidate, though the total cost decreases, bringing the salary down from \$1300 to \$967. On the other hand, if only salary is offered without any additional benefits, the organization must pay \$1634 to retain the same match, an increase of 25 percent, which shows the impact of limiting options in the job package. These scenarios serve as a rational benchmark and help illustrate how bounded rationality, discussed in the next section, can further increase costs. The results emphasize the importance of designing flexible offers and considering the overall value of the job when managing recruitment expenses.

5.2.2. Bounded Rational Mode

In this section, the objective is to minimize the total cost of package items for Organization A, which is treated as a bounded rational decision-maker. This is done while keeping the items for Organizations B and C fixed, as shown in Table 8. To capture Organization A's bounded rationality, we use the $\beta_{(r)}$ function and the data collected in Section 3.2 for high-level positions. Organizations B and C have their own $\beta_{(r)}$ functions, calculated as described in Section 3.3. Because candidate-side behavioral data are insufficient for identification, in the empirical implementation we model candidate choices as rational in this case study. Here, bounded-rational parameters are estimated for the organization side using maximum likelihood. Extending estimation to candidates would require richer, longitudinal microdata (e.g., sequential choice logs across rounds and offers, with timestamps on accept/reject decisions). The theoretical framework is two-sided; future data collection will enable full two-sided calibration.

Because the problem is non-linear and probabilistic, as defined by formulas 20 and 21, the PSO algorithm is employed to find optimal solutions. The implementation follows the setup and parameters listed in Table 13, as explained in Section 4.3.

Table 13. PSO Algorithm Parameters.

Parameters	Values
Population size	50
Number of iterations	100
Inertia weight	0.5
Cognitive coefficient	1.5
Social coefficient	1.5

The results of the second mode are presented in Table 14. They show the optimal job-offer package for Organization A when $rev_A = 1$, along with the stable match obtained for this organization.

Table 14. The optimal job offer package in second mode.

#	Salary (\$)	Bonus (\$)	Flexibility in Working Hours (h)	Working Remotely	Stable Match
Organization A	1636.8	500	-	No	Candidate 4
Organization B	900	400	2	No	Candidate 2
Organization C	800	500	1	Yes	Candidate 1

After running the matching algorithm and solving the optimization problem with PSO, while accounting for bounded rationality for the organizations, Organization A successfully hires Candidate 4 with a salary of \$1636.8 and a bonus of \$500, without offering flexibility or remote work. This indicates that Candidate 4 remains the most preferred based on their competencies and weights. The tables also indicate that Organization A was a high-priority choice for Candidate 4, reflecting a mutual preference. Organization B is matched to Candidate 2, and Organization C to Candidate 1, according to their respective packages and desired weights. The corresponding utilities, U'_{ij} for the candidates and U_{ij} for the organizations, are summarized in Table 15.

Table 15. The utilities of Organizations and Candidates in second mode.

U'_{ij}	Candidate	Organization	U_{ij}
0.39	Candidate 1	Organization C	0.47
0.35	Candidate 2	Organization B	0.52
0.40	Candidate 4	Organization A	0.88

Comparing the utility tables from both modes shows that Organization A's utility declines slightly from 0.90 to 0.88, while Candidate 4's utility drops from 0.42 to 0.40. This reduction reflects the impact of shifting Organization A's behavior from fully rational to bounded rational. The robustness of the bounded rational framework was confirmed through additional testing: the PSO algorithm consistently converged across 30 independent runs with different random seeds, producing a standard deviation of only 2.29% in final objective values, as illustrated in Figure 6. This confirms that the observed 26% increase in total cost when moving from rational to bounded rational matching is a genuine effect, not an artifact of the algorithm. A sensitivity analysis, similar to the one performed in the rational mode, was also conducted under bounded rationality, with results presented in Table 16.

Table 16. Sensitivity analysis for organization A under bounded rationality.

Scenarios		Optimal Package					Cost vs. Rational
#		Salary (\$)	Bonus (\$)	Flexibility (Hours)	Remote Work	Match	
1	$rev_A = 1$	1636.8	500	0	No	{A,4}, {B,2}, {C,1}	+26%
2	$rev_A = 0.5$	380	0	0	No	{A,2}, {B,4}, {C,1}	−24%
3	Reduced cost items 3 and 4 to 1/20, $rev_A = 1$	875	420	2	Yes	{A,2}, {B,4}, {C,1}	−10%
4	Salary only, $rev_A = 1$	1925	0	0	No	{A,4}, {B,2}, {C,1}	+18%

Under bounded rationality, the impact of parameters is stronger than in the rational model. When rev_A is reduced to 0.5, total costs decrease to \$380 (compared to \$500 in the rational case) because lower job importance increases bounded rationality, leading managers to accept a satisfactory option rather than the optimal one. Changes in intangible costs also affect the matching outcome: Organization A chooses Candidate 2, who values remote work more, as QRE probabilities shift toward candidates whose preferences align better with the cheaper benefits. Imposing a salary-only constraint leads to a sharp cost increase to \$1925 (versus \$1,634 in the rational case), since removing package flexibility reduces the ability to differentiate candidates, forcing organizations to offer higher salary offers to compensate for decision randomness. These results demonstrate that bounded rationality not only increases costs but also significantly alters optimization strategies depending on job importance and package design.

5.3. Further Discussion

We observed that the bounded rationality model based on QRE provides more nuanced and realistic predictions for the matching problem compared to a purely rational choice model. However, evaluating and interpreting the results in context is essential. The key insights are as follows:

1. As shown in Sections 3.2 and 3.3, the behavior of the hiring manager, who serves as the final decision layer after technical and behavioral interviews, follows the QRE function. The collected data aligns well with this function, reflecting the manager's incomplete information about the candidate's utility. A candidate's utility has both explicit and implicit components, and the manager's rationality is constrained by limited knowledge of the implicit component. Moreover, analysis of the estimated beta functions shows evidence of a learning process over time, gradually reducing bounded rationality, although it never disappears completely.
2. In organizations with bounded-rational decision-makers, perceived candidate utilities tend to decrease. Managers often select candidates with lower normalized utility rather than the optimal choice. This effect is more pronounced for lower-level or less critical positions. Because bounded rationality has a random component, there are occasional increases in perceived utility, but overall, a long-term decline is observed.
3. Comparing the matching outcomes under rational and bounded rational modes, the organization's utility fell slightly from 0.90 to 0.88. This minor change reflects the learning and improvement in the hiring manager's decisions over time. The final match itself remained unchanged, but larger differences between rationality and bounded rationality could lead to different matches.
4. The decline in utility translated into higher job-offer costs, rising from \$1300 to \$1636.8 to secure the best candidate. Essentially, the bounded rationality of decision-makers imposes additional costs on the organization.

These findings suggest that organizations should acknowledge the bounded rationality of their hiring managers and work to mitigate its effects through training and experience. Doing so can reduce recruitment costs and improve the effectiveness of job-offer strategies.

6. Conclusions

This study presents a two-stage stochastic model to help organizations design job-offer packages and match candidates while taking into account the bounded rationality of decision-makers. By combining the Gale-Shapley matching mechanism, Quantal Response Equilibrium for bounded rationality, Maximum Likelihood Estimation for parameter fitting, and particle swarm optimization for offer design, the model captures both the strategic and cognitive aspects of hiring decisions.

From the case study, we observed that accounting for bounded rationality can influence both the outcomes and costs of matching. In our example, even though the final matches remained the same, the total cost of the job-offer package increased under bounded rationality. This shows that decision-makers' limitations can subtly reduce efficiency and increase costs.

The results highlight the importance of recognizing and addressing bounded rationality in hiring. Organizations can potentially reduce costs and improve decision quality through training and experience, helping managers make more informed choices despite cognitive limitations.

6.1. Managerial Implications

This study provides several actionable insights for managers aiming to improve recruitment strategies under realistic constraints. First, the 26% cost premium we report is specific to our case study. Sensitivity analyses indicate that this premium is not fixed: when job importance is lower, the bounded-rational solution is actually cheaper than the rational benchmark; under a salary-only constraint, it becomes more expensive. These variations reflect how bounded rationality interacts with position value and package flexibility. We, therefore, view the 26% as an informative point estimate for this context, not a universal constant.

Second, the impact of bounded rationality varies across job levels. Senior roles benefit from thorough evaluation and stable hiring teams, where careful judgment and experience reduce randomness. For junior roles, simpler or streamlined selection processes are sufficient, as decisions are more affected by bounded rationality and rotation of hiring team members has less impact.

Third, flexible, modular compensation packages are cost-effective. Tailored offers, including remote work, schedule flexibility, or other valued non-monetary items, can reduce costs by up to 25 compared to salary-only packages, provided they match candidate preferences. For HR managers, this suggests that offering a menu of non-monetary benefits can serve as a substitute for high salaries, allowing budgets to be stretched further without undermining candidate attractiveness.

Fourth, the learning curve of hiring managers presents a valuable opportunity. MLE estimates show that managers approach near-optimal decision-making after roughly 20–25 recruitment cycles, highlighting the benefits of training programs and continuity in hiring teams, especially for strategically important roles. This indicates that organizations can reduce costs not only by designing better packages but also by investing in the decision-makers themselves through targeted training and experience accumulation.

Finally, in competitive labor markets, this approach offers a practical alternative to salary-driven competition. By accounting for bounded rationality, organizations can design job offers that remain competitive while aligning effectively with candidate preferences in a probabilistic decision environment. Practically, this means budgeting should include a contingency margin for bounded-rational costs and HR managers should anticipate variation across job levels and market conditions when planning recruitment strategies.

6.2. Limitations and Future Research

Despite its contributions, this study has several limitations. First, the model assumes complete and commonly known information about candidate preferences, organizational importance weights, and competitors' offers, which is rarely the case in practice. Decision-makers often face partial or uncertain information. A useful extension would be to relax this assumption by formulating the problem as an incomplete-information (Bayesian) game, where firms and candidates hold beliefs about counterpart utilities and rival offers. This

could be operationalized through a Bayesian-QRE with hierarchical priors on preference parameters and belief-based DA, estimated via hierarchical MLE or EM from observed choice data. These extensions would align the mechanism more closely with real recruitment settings, where information is partial, evolving, and asymmetric.

Second, the empirical evidence comes from a single holding company, with a limited number of organizations and candidates, within a specific sector and regulatory environment. This scope constrains external validity, as institutional, cultural, and industry factors may differ across settings. Cross-industry or cross-country applications would require recalibration of the estimated parameters (α , β , δ , ε) and reweighting of competencies and job-offer items to reflect sectoral and cultural conditions. Moreover, while the framework allows bounded rationality on both sides, the case study estimates parameters only for organizations; candidates are modeled as rational due to data limitations. These choices narrow the empirical claim and underscore that full two-sided calibration awaits richer candidate-side behavioral data. Taken together, these results underscore that measured cost effects of bounded rationality are context-dependent: institutional, cultural, and industry factors may shift magnitudes, and full two-sided calibration awaits richer candidate-side behavioral data. Future work should (i) assemble larger, multi-industry and multi-country datasets to test transferability, (ii) collect candidate-side behavioral traces to jointly estimate bounded-rational parameters for both sides, and (iii) conduct out-of-sample validations and sensitivity analyses across institutional regimes. To quantify candidates' bounded rationality specifically, subsequent studies should gather candidate-side choice sequences over multiple rounds and offers, implement discrete choice or conjoint designs that randomize job attributes, and, where feasible, use field experiments or A/B tests on application portals. These designs would identify candidate-side QRE parameters and allow full two-sided estimation within our matching framework.

Third, the model treats competitors' compensation packages as fixed, whereas in reality, organizations dynamically adjust their offers in response to rivals' strategies and candidate reactions. Future research should, therefore, explore multi-agent dynamic game formulations, where competing firms update their packages endogenously over repeated interactions. Embedding bounded rationality into such a dynamic setting would allow the simulation of realistic competitive adjustments and feedback effects. In addition, incorporating candidate-side bounded rationality in this multi-agent context would further improve realism and predictive power.

Fourth, our framework focuses on the final stage of hiring, after initial AI-based filtering and preliminary interviews, where managers select among a smaller pool of candidates. In larger-scale matching problems with thousands of candidates and hundreds of organizations, computational efficiency could be challenged. Although the complexity is $O(TPn^2m)$, techniques such as decomposition, parallel/distributed PSO implementations, or adaptive swarm structures could help maintain tractability in future applications.

Finally, the framework could extend beyond recruitment to other two-sided matching problems involving bounded rational agents, such as supplier selection, project bidding, or school choice. Future work could integrate learning algorithms or agent-based simulations to capture more dynamic interactions.

Author Contributions: Conceptualization, S.N.-Z. and N.S.-G.; methodology, S.N.-Z., N.S.-G. and O.G.; software, S.N.-Z.; validation, S.N.-Z., N.S.-G. and O.G.; formal analysis, S.N.-Z.; investigation, S.N.-Z.; resources, O.G.; data curation, S.N.-Z.; writing—original draft preparation, S.N.-Z. and N.S.-G.; writing—review and editing, N.S.-G. and O.G.; visualization, S.N.-Z.; supervision, N.S.-G. and O.G.; project administration, S.N.-Z. and N.S.-G.; funding acquisition, O.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research was financially supported by the French National Research Agency (ANR), “Investissements d’Avenir” (ANR-11-IDEX0003/LabEx Ecodec/ANR-11-LABX-0047), by the UK Research and Innovation (UKRI Frontier Research Grant EP/U537093/1), and by the European Union (SInfoNiA ERC-2023-ADG 101142530) for Olivier Gossner’s research.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Acknowledgments: We would like to thank the reviewers for their valuable comments and all the authors of the references cited in this article.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

QRE	Quantal Response Equilibrium
MLE	Maximum Likelihood Estimation
PSO	Particle swarm optimization
LQRE	logit quantal response equilibrium
DAA	Deferred Acceptance Algorithm
FAHP	Fuzzy Analytic Hierarchy Process
MILP	Mixed-Integer Linear Programming

Appendix A

Table A1. Main notations used in the Model (Section 4), organized by sets, parameters, and variables.

Title 1		Title 2	Title 3
Maximum utility of organization i at time t	α_{it}	Index of organizations	i
Utility of candidate j for organization i	U_{ij}	Index of candidates	j
Utility of organization i for candidate j	U'_{ij}	Index of packages’ item	k
Importance weights of competencies for organization i	$W^{(i)}$	Index of time	t
Importance weights of package items for candidate j	$W'^{(j)}$	Second index of time	t'
Equivalent utility of the target organization	U'_{eq_i}	Set of Time Indexes	T
Cost of organization i ’s package item f	$cost_{if}$	Index of job-offer items	f
f -th package item of organization i	d_{if}	Index variable for the organization under investigation	q
Index of package items of integer type	f_{int}	Index of package items of binary type	f_{bin}
Set of Candidates	J	Set of Organizations	I
Normalized value of the f ’th package item of organization i	dn_{if}	Competencies of candidate j	C_j
Upper bound of the f ’th item	d_{max_f}	Lower bound of the f -th item	d_{min_f}
Binary variable for whether organization i interviews candidate j at time t	In_{ijt}	maximum utility of candidate j at time t	α'_{jt}
Binary variable: organization i chooses candidate j at time t	ch_{ijt}	Binary variable: candidate j chooses organization i at time t	ch'_{ijt}
Binary variable: candidate j can choose organization i at time t	ζ_{ijt}	Sensitivity of organization i ’s position	rev_i

References

1. Najafi-Zangeneh, S.; Shams-Gharneh, N.; Arjomandi-Nezhad, A.; Zolfani, S.H. An improved machine learning-based employees attrition prediction framework with emphasis on feature selection. *Mathematics* **2021**, *9*, 1226. [CrossRef]
2. Nabeeh, N.A.; Smarandache, F.; Abdel-Basset, M.; El-Ghareeb, H.A.; Aboelfetouh, A. An integrated neutrosophic-topsis approach and its application to personnel selection: A new trend in brain processing and analysis. *IEEE Access* **2019**, *7*, 29734–29744. [CrossRef]

3. Wang, B.; Li, Z.; Wang, S.; Li, M.; Ji, A. Modeling bounded rationality in discretionary lane change with the quantal response equilibrium of game theory. *Transp. Res. Part B Methodol.* **2022**, *164*, 145–161. [\[CrossRef\]](#)
4. Gale, D.; Shapley, L.S. College admissions and the stability of marriage. *Am. Math. Mon.* **1962**, *69*, 9–15. [\[CrossRef\]](#)
5. Roth, A.E.; Sotomayor, M. Two-sided matching. In *Handbook of Game Theory with Economic Applications*; Elsevier: Amsterdam, The Netherlands, 1992; Volume 1, pp. 485–541.
6. Pu, D.; Yuan, G. Two-sided matching model considering multi-information fusion of stakeholders. *Expert Syst. Appl.* **2023**, *212*, 118784. [\[CrossRef\]](#)
7. Jiang, P.; Guo, S.; Du, B.; Guo, J. Two-sided matching decision-making model for complex product system based on life-cycle sustainability assessment. *Expert Syst. Appl.* **2022**, *208*, 118184. [\[CrossRef\]](#)
8. Manlove, D.F.; Irving, R.W.; Iwama, K. Guest editorial: Special issue on matching under preferences. *Algorithmica* **2010**, *58*, 1–4. [\[CrossRef\]](#)
9. Bansal, N.; Gupta, A.; Li, J.; Mestre, J.; Nagarajan, V.; Rudra, A. When LP is the cure for your matching woes: Improved bounds for stochastic matchings. *Algorithmica* **2012**, *63*, 733–762. [\[CrossRef\]](#)
10. Ashlagi, I.; Braverman, M.; Hassidim, A. Matching with couples revisited. In Proceedings of the 12th ACM Conference on Electronic Commerce, San Jose, CA, USA, 5–9 June 2011.
11. Hatfield, J.; Kominers, S.D. Multilateral matching. In Proceedings of the 12th ACM Conference on Electronic Commerce, San Jose, CA, USA, 5–9 June 2011.
12. Manlove, D. *Algorithmics of Matching Under Preferences*; World Scientific: Singapore, 2013; Volume 2.
13. Rastegari, B.; Condon, A.; Immorlica, N.; Leyton-Brown, K. Two-sided matching with partial information. In Proceedings of the Fourteenth ACM Conference on Electronic Commerce, Philadelphia, PA, USA, 16–20 June 2013.
14. Luce, R.D. *Individual Choice Behavior*; Wiley: New York, NY, USA, 1959; Volume 4.
15. Chen, Y.; Su, X.; Zhao, X. Modeling bounded rationality in capacity allocation games with the quantal response equilibrium. *Manag. Sci.* **2012**, *58*, 1952–1962. [\[CrossRef\]](#)
16. Simon, H.A. A behavioral model of rational choice. *Q. J. Econ.* **1955**, *69*, 99–118. [\[CrossRef\]](#)
17. Nelson, R.R. *An Evolutionary Theory of Economic Change*; Harvard University Press: Cambridge, MA, USA, 1985.
18. Rubinstein, A. *Modeling Bounded Rationality*; MIT Press: Cambridge, MA, USA, 1998.
19. Canbolat, P.G. Bounded rationality in clearing service systems. *Eur. J. Oper. Res.* **2020**, *282*, 614–626. [\[CrossRef\]](#)
20. McKelvey, R.D.; Palfrey, T.R. Quantal response equilibria for normal form games. *Games Econ. Behav.* **1995**, *10*, 6–38. [\[CrossRef\]](#)
21. Goeree, J.K.; Holt, C.A.; Palfrey, T.R. Stochastic game theory for social science: A primer on quantal response equilibrium. In *Handbook of Experimental Game Theory*; Edward Elgar Publishing: Cheltenham, UK, 2020; pp. 8–47.
22. Chen, H.-C.; Friedman, J.W.; Thisse, J.-F. Boundedly rational Nash equilibrium: A probabilistic choice approach. *Games Econ. Behav.* **1997**, *18*, 32–54. [\[CrossRef\]](#)
23. Anderson, S.P.; Goeree, J.K.; Holt, C.A. The logit equilibrium: A perspective on intuitive behavioral anomalies. *South. Econ. J.* **2002**, *69*, 21–47. [\[PubMed\]](#)
24. Goeree, J.K.; Holt, C.A. Ten little treasures of game theory and ten intuitive contradictions. *Am. Econ. Rev.* **2001**, *91*, 1402–1422. [\[CrossRef\]](#)
25. Goeree, J.K.; Holt, C.A. Asymmetric inequality aversion and noisy behavior in alternating-offer bargaining games. *Eur. Econ. Rev.* **2000**, *44*, 1079–1089. [\[CrossRef\]](#)
26. Anderson, S.P.; Goeree, J.K.; Holt, C.A. Minimum-effort coordination games: Stochastic potential and logit equilibrium. *Games Econ. Behav.* **2001**, *34*, 177–199. [\[CrossRef\]](#)
27. Goeree, J.K.; Holt, C.A.; Palfrey, T.R. Quantal response equilibrium and overbidding in private-value auctions. *J. Econ. Theory* **2002**, *104*, 247–272. [\[CrossRef\]](#)
28. Capra, C.M.; Goeree, J.K.; Gomez, R.; A Holt, C. Anomalous behavior in a traveler’s dilemma? *Am. Econ. Rev.* **1999**, *89*, 678–690. [\[CrossRef\]](#)
29. Lim, N.; Ho, T.-H. Designing price contracts for boundedly rational customers: Does the number of blocks matter? *Mark. Sci.* **2007**, *26*, 312–326. [\[CrossRef\]](#)
30. Ho, T.-H.; Zhang, J. Designing pricing contracts for boundedly rational customers: Does the framing of the fixed fee matter? *Manag. Sci.* **2008**, *54*, 686–700. [\[CrossRef\]](#)
31. Su, X. Bounded rationality in newsvendor models. *Manuf. Serv. Oper. Manag.* **2008**, *10*, 566–589. [\[CrossRef\]](#)
32. Liu, T.; Methapatara, C.; Wynter, L. Revenue management model for on-demand IT services. *Eur. J. Oper. Res.* **2010**, *207*, 401–408. [\[CrossRef\]](#)
33. Shang, W.; Liu, L. Promised delivery time and capacity games in time-based competition. *Manag. Sci.* **2011**, *57*, 599–610. [\[CrossRef\]](#)
34. Di, X.; Liu, H.X. Boundedly rational route choice behavior: A review of models and methodologies. *Transp. Res. Part B Methodol.* **2016**, *85*, 142–179. [\[CrossRef\]](#)

35. Roemer, N.; Müller, S.; Voigt, G. A choice-based optimization approach for contracting in supply chains. *Eur. J. Oper. Res.* **2023**, *305*, 271–286. [\[CrossRef\]](#)
36. Khanlarzade, N.; Farughi, H. Modeling the Stackelberg game with a boundedly rational follower in deterioration supply chain-based interaction with the leader's hybrid pricing strategy. *Expert Syst. Appl.* **2024**, *237*, 121302. [\[CrossRef\]](#)
37. Pais, J.; Pintér, Á. School choice and information: An experimental study on matching mechanisms. *Games Econ. Behav.* **2008**, *64*, 303–328. [\[CrossRef\]](#)
38. Echenique, F.; Wilson, A.J.; Yariv, L. Clearinghouses for two-sided matching: An experimental study. *Quant. Econ.* **2016**, *7*, 449–482. [\[CrossRef\]](#)
39. Dreyfuss, B.; Heffetz, O.; Rabin, M. Expectations-based loss aversion may help explain seemingly dominated choices in strategy-proof mechanisms. *Am. Econ. J. Microecon.* **2022**, *14*, 515–555. [\[CrossRef\]](#)
40. Alcalde-Unzu, J.; Klijn, F.; Vorsatz, M. Constrained school choice: An experimental QRE analysis. *Soc. Choice Welf.* **2023**, *61*, 587–624. [\[CrossRef\]](#)
41. Wallace, M.J.J.; Steuer, R.E. Multiple Objective Linear Programming in the Design of Internal Wage Structures. In *Academy of Management Proceedings*; Academy of Management Briarcliff Manor: New York, NY, USA, 1979; p. 10510.
42. Kwak, N.; Allen, T.; Schniederjans, M. A multilevel salary compensation model using goal programming. *RAIRO-Oper. Res.* **1982**, *16*, 21–31. [\[CrossRef\]](#)
43. Garcia-Diaz, A.; Hogg, G.L. A mathematical programming approach to salary administration. *Comput. Ind. Eng.* **1983**, *7*, 7–13. [\[CrossRef\]](#)
44. Kassa, B.A. A decision support model for salary structure design. *Compens. Benefits Rev.* **2020**, *52*, 109–120. [\[CrossRef\]](#)
45. Tremblay, F.-A.; Piché-Meunier, D.; Dubois, L. Multi-objective Optimization for the Design of Salary Structures. In *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*; Springer: Berlin/Heidelberg, Germany, 2023.
46. Najafi-Zangeneh, S.; Shams-Gharneh, N.; Gossner, O.; Abbasi, E. Optimizing Job Offer Packages: How Can Organizations Enhance Personnel Selection. In *Proceedings of the 33rd European Conference on Operational Research*, Copenhagen, Denmark, 30 June–3 July 2024.
47. Wolpert, D.H.; Macready, W.G. No free lunch theorems for search. In *Technical Report SFI-TR-95-02-010*; Santa Fe Institute: Santa Fe, NM, USA, 1995.
48. Bo, Y.; Guo, X.; Liu, Q.; Pan, Y.; Zhang, L.; Lu, Y. Prediction of tunnel deformation using PSO variant integrated with XGBoost and its TBM jamming application. *Tunn. Undergr. Space Technol.* **2024**, *150*, 105842. [\[CrossRef\]](#)
49. Jovanovic, L.; Bacanin, N.; Petrovic, A.; Zivkovic, M.; Antonijevic, M.; Gajic, V.; Elsayed, M.M.; Abouhawwash, M. Exploring artificial intelligence potential in solar energy production forecasting: Methodology based on modified PSO optimized attention augmented recurrent networks. *Sustain. Comput. Inform. Syst.* **2025**, *47*, 101174. [\[CrossRef\]](#)
50. Haile, P.A.; Hortaçsu, A.; Kosenok, G. On the empirical content of quantal response equilibrium. *Am. Econ. Rev.* **2008**, *98*, 180–200. [\[CrossRef\]](#)
51. Simon, H.A. Models of man: Social and rational. In *Mathematical Essays on Rational Human Behavior in Society Setting*; Wiley: New York, NY, USA, 1957.
52. Kennedy, J.; Eberhart, R. Particle swarm optimization. In *Proceedings of the ICNN'95-International Conference on Neural Networks*, Perth, WA, Australia, 27 November–1 December 1995; IEEE: New York, NY, USA.
53. Zhang, Y.; Balochian, S.; Agarwal, P.; Bhatnagar, V.; Houshia, O. *Artificial Intelligence and Its Applications*; Hindawi Publishing Corporation: Cairo, Egypt, 2014; p. 840491.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.