Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Labor Market Dynamics and Unemployment Formation: An LLM-based Cognitive Agent Approach with Deep Reinforcement Learning

Kailang Li

lklkevinness@163.com

Guangxi University

Article

Keywords: Labor Market Dynamics, Unemployment, Large Language Models, Deep Reinforcement Learning, Agent-based Modeling

Posted Date: March 13th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8754324/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Labor Market Dynamics and Unemployment Formation: An LLM-based Cognitive Agent Approach with Deep Reinforcement Learning

Kailang Li[1*]
[1]Guangxi University, Nanning, Guangxi, China

Correspondence: Kailang Li (lklkevinness@163.com)

Abstract: Traditional neoclassical labor market models, predicated on perfect competition, complete information, and full rationality, fail to capture real-world complexities—including information asymmetry, bounded rationality, and dynamic skill mismatches—resulting in inadequate explanations of structural unemployment and inefficient job matching. To address this, the study develops an integrated simulation framework combining Large Language Models (LLMs), Deep Reinforcement Learning (DRL), and agent-based modeling: LLMs process job descriptions/resumes for multi-modal skill representation and belief updating; DRL models game-theoretic interactions between firms (optimizing hiring/salary) and job seekers (optimizing job search/negotiation) under bounded rationality; additional modules include policy-macroeconomy coupling (for taxes/subsidies/minimum wage) and skill evolution networks (for skill obsolescence/emergence). Experiments, calibrated to U.S. BLS/JOLTS data, yield concrete results: (1) A 10% hiring subsidy reduced unemployment to 4.5% and matching efficiency to 80%, outperforming other policies; (2) A 5% corporate tax increase raised unemployment to 8% and lowered vacancy rates to 3.5%; (3) Low-skill automation shocks pushed unemployment to 9% due to persistent skill mismatches. The framework accurately replicates dynamic labor market behaviors, outperforming traditional models in capturing policy heterogeneity and real-world inefficiencies.

Keywords: Labor Market Dynamics; Unemployment; Large Language Models; Deep Reinforcement Learning; Agent-based Modeling

1. Introduction

The dynamics of labor markets play a crucial role in shaping unemployment rates, job matching, and overall economic stability[1, 2]. The labor market is inherently complex, driven by interactions between firms and job seekers, each with their own objectives, strategies, and constraints. Firms aim to fill vacancies with suitable candidates, while job seekers seek optimal employment opportunities based on skills, preferences, and compensation expectations[3]. This interaction results in various labor market phenomena such as unemployment, wage disparity, and job mismatches. Understanding these dynamics and their implications for unemployment is a key challenge in labor economics[4].

Traditional economic models, such as the neoclassical theory of labor markets, assume perfect competition, complete information, and rational decision-making[5]. These models, while providing valuable insights, often fail to account for the complexities of real-world labor markets, where firms and job seekers operate with limited information, face significant search frictions, and exhibit bounded rationality[6]. In reality, individuals and firms make decisions based on imperfect and incomplete information, often leading to suboptimal outcomes such as structural unemployment, wage gaps, and mismatched job placements[7].

One of the critical factors contributing to structural unemployment is the information asymmetry between job seekers and firms[8]. Job seekers may not have complete knowledge of available job opportunities, and firms may struggle to identify candidates with the right skills[9]. Additionally, the job search process is time-consuming and costly, leading to inefficient matching and longer periods of unemployment[10]. These inefficiencies, compounded by factors such as economic shocks or policy changes, make the task of improving labor market outcomes particularly challenging.

Recent advances in artificial intelligence (AI) and deep learning offer new approaches to modeling and understanding labor market dynamics[11]. By leveraging technologies such as deep reinforcement learning (DRL) and large language models (LLMs), it is now possible to simulate the behaviors of both firms and job seekers in a dynamic, interactive environment[12]. DRL can model decision-making under uncertainty, where agents (firms and job seekers) learn to optimize their strategies over time by interacting with the environment[13]. LLMs, on the other hand, can be used to extract and process complex information from job descriptions and resumes, providing more accurate insights into job matching and candidate evaluation[14].

The application of deep learning and agent-based models to labor market simulations can address several shortcomings of traditional economic theories[15]. First, these models can incorporate bounded rationality and imperfect information, offering a more realistic representation of decision-making processes[16]. Second, they allow for the modeling of dynamic interactions over time, capturing the evolving nature of job searches, negotiations, and market conditions[17]. Third, they enable the exploration of various policy interventions, such as changes in tax rates, unemployment benefits, or minimum wage laws, and their impact on labor market outcomes[18].

Despite the potential of these approaches, there are still significant gaps in the literature[19]. Current research on labor markets often neglects the role of cognitive biases and learning in decision-making, which can lead to inefficiencies in job matching and unemployment[20]. Moreover, the application of AI to labor economics has primarily focused on static analyses, with limited attention given to the dynamic, adaptive nature of labor market processes[21]. The integration of LLMs and DRL into labor market models could fill these gaps, providing new insights into the causes and consequences of unemployment and offering more effective strategies for improving job matching and reducing unemployment rates[22].

The core problem addressed in this research revolves around the formation of structural unemployment in labor markets, particularly when both firms and job seekers operate under conditions of limited rationality and imperfect information[23]. Traditional economic models often assume that agents—both employers and job seekers—are fully rational and have access to complete and symmetric information[24]. However, real-world labor markets are far from this idealized state. In practice, both firms and job seekers face significant uncertainties and information asymmetries, which affect their decision-making processes and, consequently, the efficiency of job matching[25].

Firms typically seek to fill job vacancies with candidates who possess the right skills and experience[26]. However, due to the complexity and heterogeneity of job requirements, as well as the limited availability of information, firms often cannot perfectly assess the qualifications of all potential candidates[27]. This results in suboptimal hiring decisions, where firms either overestimate or underestimate a candidate's suitability for a role[28]. Similarly, job seekers, faced with a vast array of job opportunities and competing offers, may not have perfect knowledge about the specific demands of each position, the company's culture, or the salary

and benefits package[29]. This lack of information can lead to decisions that are far from optimal, such as over-qualification for certain roles or under-utilization of skills, contributing to inefficient job matches[30].

One of the key reasons why structural unemployment arises under these conditions is that the process of job matching becomes inherently inefficient[31]. Both job seekers and employers engage in costly and time-consuming search processes to find a match, but due to information asymmetries and limited rationality, the resulting matches are often suboptimal. Job seekers may settle for positions that do not fully utilize their skills, leading to a misallocation of talent, while firms may hire candidates who do not meet the required standards, increasing turnover rates and reducing productivity[32]. These inefficiencies accumulate over time, resulting in a higher incidence of long-term unemployment, especially when job seekers cannot find positions that match their skills or expectations[33]. In addition, the lack of perfect information and limited rationality can lead to coordination failures in the labor market. Employers may struggle to identify the most suitable candidates from a pool of applicants, while job seekers may not have a clear understanding of the available opportunities or the skills demanded by employers. This results in a "search mismatch," where job seekers are either unaware of suitable vacancies or apply for jobs that do not align with their qualifications. Simultaneously, employers may fail to attract candidates with the right skill sets or have difficulty assessing candidates' true potential, leading to a mismatch in expectations on both sides. Moreover, structural unemployment is exacerbated by various market frictions, such as geographical mobility constraints, wage rigidity, and the time and costs associated with job searches. These factors further delay the job-matching process and increase the likelihood that individuals remain unemployed for extended periods, even when suitable vacancies exist. As job seekers become discouraged and employers become less willing to invest in the search process, the labor market experiences higher rates of long-term unemployment.

The impact of policy interventions, such as changes in unemployment benefits, minimum wage laws, and tax incentives for hiring, also plays a significant role in shaping the dynamics of structural unemployment. However, current models of unemployment often fail to capture the interactive nature of these policies in a setting where both firms and job seekers have limited rationality. For instance, higher unemployment benefits may encourage job seekers to remain unemployed for longer periods, while minimum wage laws could lead to a surplus of labor supply, exacerbating unemployment among low-skilled workers. These policies may further complicate the matching process, as job seekers adjust their expectations and search strategies based on the incentives provided by the government.

To address these issues, this research proposes a novel approach that integrates cognitive agents, modeled using deep reinforcement learning, to simulate the decision-making processes of both firms and job seekers under conditions of limited rationality and imperfect information. By incorporating large language models, which can extract relevant skills and expectations from job descriptions and resumes, the model aims to better represent the complexities of the job matching process. This approach allows for the dynamic adaptation of agents over time as they learn from interactions, taking into account the evolving nature of market conditions, skill demands, and policy interventions.

This research proposes a model that integrates Deep Reinforcement Learning with Cognitive Agents using Large Language Models to simulate interactions between firms and job seekers. The model leverages DRL to optimize decision-making for both firms and job seekers, who act as agents learning to navigate the job matching process under uncertainty and limited information. LLMs are utilized to process job

descriptions and resumes, extracting relevant skills and expectations, enabling more accurate simulations of job matching. The model incorporates various policy interventions, such as tax rates, subsidies, and minimum wage laws, to study their impact on unemployment rates and matching efficiency.

2. Methods

2.1 Agent-Based Simulation Framework

Agent-based Modeling (ABM) is a powerful approach for simulating complex systems where multiple agents interact, each with their own goals, behaviors, and learning processes. In the context of the labor market, ABM allows us to model the interactions between firms and job seekers, two key agents in the labor market. Each agent operates based on its own decision-making rules, which evolve over time based on past interactions and experiences. The goal is to simulate how these agents learn, adapt, and influence market outcomes, such as unemployment and job matching efficiency.

To model the behavior of firms and job seekers, we use reinforcement learning (RL) techniques. Firms and job seekers are represented as RL agents that learn their strategies through trial and error, optimizing their actions based on rewards or penalties received from the environment. The firms' actions primarily revolve around job postings, candidate selection, and wage offerings, while job seekers focus on job search, application decisions, and negotiations (see Figure 1).

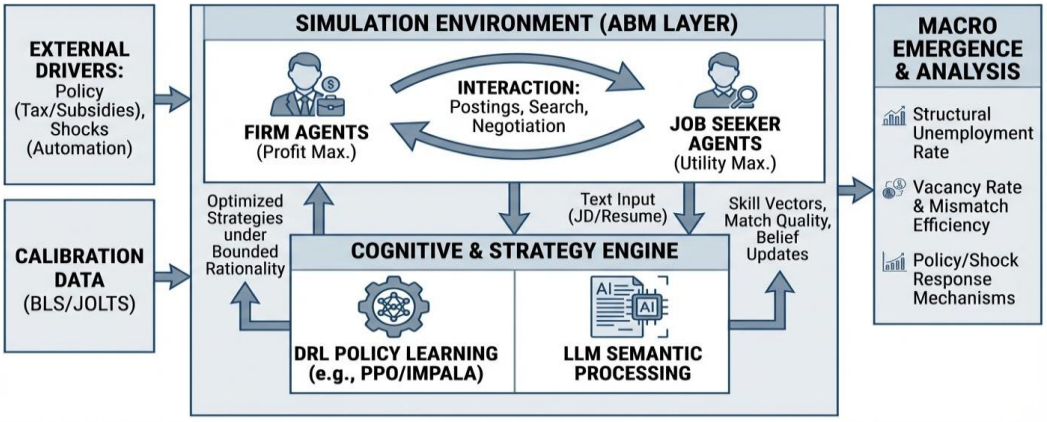


Figure 1. An overview of the agent-based simulation framework for labor market analysis, illustrating the role of behavioral modeling of firms and job seekers, the use of LLM in belief updating and negotiation, and the experimental setup to evaluate labor market dynamics.

For firms, we implement the Proximal Policy Optimization (PPO) or Importance Weighted Actor-Learner Architectures (IMPALA) algorithms. These algorithms allow the firms to optimize their hiring strategy by interacting with the job market environment. PPO is a policy-gradient method known for its stability and efficiency, while IMPALA is well-suited for large-scale distributed settings, where multiple agents learn simultaneously. The firm's objective is to maximize its utility by offering competitive wages and attracting the most suitable candidates. This involves considering factors such as the skill set required for the position, candidate qualifications, and market conditions. Over time, the firm's policy is optimized to increase the efficiency of its hiring process, reducing costs associated with turnover and improving productivity.

Job seekers, on the other hand, are modeled as agents that search for job opportunities, apply to relevant positions, and negotiate terms with firms. Their search strategy is influenced by several factors, including the quality of available jobs, the wage offered, and their skill set. The job seekers use a reinforcement learning approach to explore and exploit different strategies for applying to jobs. They learn to adjust their behavior based on the success or failure of previous job applications. The goal for the job seeker is to find a job that matches their skills and preferences, while also maximizing their expected wage. The RL agent for job seekers balances exploration (searching for new opportunities) and exploitation (applying to jobs that are known to yield high returns) over time.

A key feature of this model is the use of Large Language Models to update the agents' beliefs during the job search and hiring process. LLMs, such as GPT-3, are used to extract valuable information from job descriptions (JDs) and resumes. For example, job descriptions contain essential skills, experience, and expectations that firms are looking for, while resumes detail the qualifications and skills of job seekers. LLMs process these texts to generate accurate representations of both job seekers' and firms' beliefs about potential matches. These representations are then used to inform decisions made by both agents. LLMs also play a critical role in the negotiation process between job seekers and firms. Given the asymmetric information between both parties—where firms have limited insight into the full set of skills possessed by job seekers, and job seekers may lack detailed knowledge about the job's actual demands—LLMs help simulate how these beliefs evolve over time. By generating plausible negotiation dialogues, LLMs model how the agents update their beliefs about each other's expectations, ultimately leading to a dynamic job matching process.

2.2 Behavioral Modeling of Firms and Job Seekers

The proposed model is further extended to incorporate industry shocks and policy interventions, allowing for a more comprehensive analysis of how external factors influence labor market outcomes such as unemployment rates, job vacancy rates, and wage distributions. These extensions enable the simulation of real-world scenarios where technological advancements, policy changes, and economic shocks play significant roles in shaping market dynamics.

Industry shocks are simulated to represent significant changes in the labor market, such as technological advancements (e.g., automation) or shifts in market demand. For instance, automation might reduce the demand for certain types of manual labor while increasing the demand for tech-related skills. To model this, we introduce a shock parameter ΔA into the market dynamics. This parameter adjusts the demand for specific job categories based on industry-wide changes. For example, the job demand function for sector s at time t can be represented as:

$$D_s(t) = D_s(0) \cdot (1 + \Delta A_s(t)),$$

where $D_s(0)$ is the initial demand for job type s, and $\Delta A_s(t)$ captures the industry shock effect at time t. A positive ΔA indicates an increase in demand for jobs in that sector (e.g., due to technological advancements), while a negative ΔA signifies a decline (e.g., automation replacing jobs). These shocks

affect the job seekers' search strategies and the firms' hiring decisions. Job seekers must adapt their search strategies by targeting sectors with growing demand, while firms must adjust their recruitment policies to address changing skill requirements (see Figure 2).

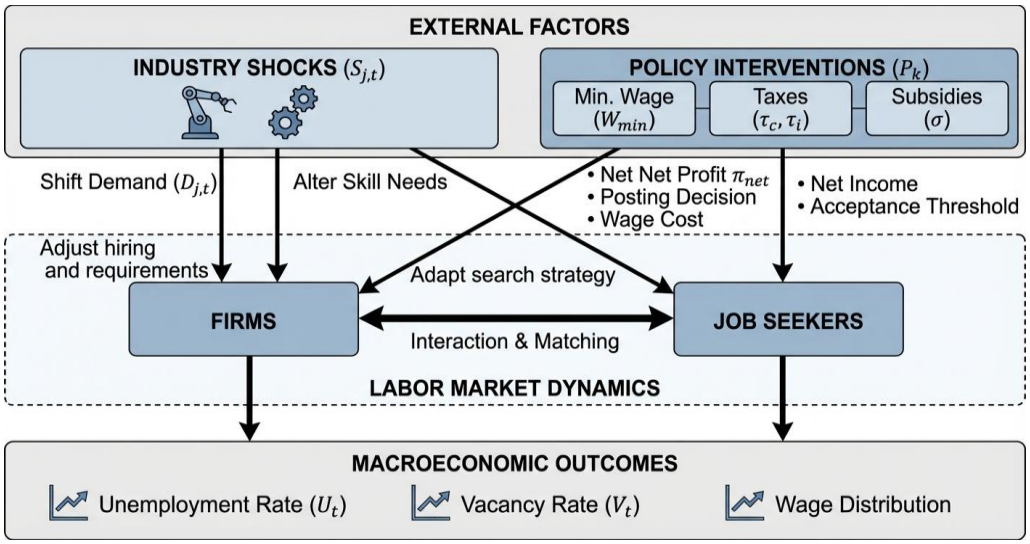


Figure 2. Flowchart illustrating the simulation framework for labor market dynamics, showing the interaction between industry shocks, policy levers (minimum wage, tax, subsidies), and the effects on firms and job seekers.

Policy interventions, such as changes in tax rates, subsidies, or minimum wage laws, are incorporated into the model to examine their effects on labor market outcomes. These policies affect both the firms' hiring decisions and the job seekers' willingness to participate in the labor market. For example, an increase in the minimum wage $W_{\min}$ may encourage job seekers to accept lower-skilled jobs, but could also lead to a decrease in demand from firms unable to afford higher wages. To model this, we define a policy parameter P for each intervention:

The minimum wage affects the wage floor that firms are willing to offer. If a job seeker's expected wage w_j falls below $W_{\min}$, they will not consider that job. The job posting decision for a firm can then be expressed as:

$$w_j = \max(W_{\min}, f(\theta_j, \xi_f))$$

where θ_j represents the job seeker's skill vector and ξ_f is the firm's willingness to pay based on market conditions and demand.

Taxes influence the net wage a worker receives and the costs to a firm. An increase in corporate tax τ_c or income tax τ_i reduces the available income for job seekers and the incentives for firms to hire. The firm's net profit function becomes:

$$\Pi_f = (w_j - \tau_i) \cdot N_j - C_f(\xi_f, \tau_c)$$

where N_j is the number of job seekers hired, and $C_f(\xi_f, \tau_c)$ represents the cost function of the firm, which includes production and hiring costs influenced by the tax rate.

Labor market subsidies can be implemented to reduce unemployment by encouraging firms to hire or job seekers to accept positions. The subsidy S_f for firms hiring workers at wage w_j is modeled as:

$$S_f = \lambda \cdot \max(0, W_{\min} - w_j)$$

where λ represents the subsidy rate. This reduces the effective wage cost for firms, potentially increasing the number of job seekers hired.

The model uses these policy variables to examine their effects on unemployment rate $U(t)$, job vacancy rate $V(t)$, and wage distribution. The macroeconomic indicators evolve dynamically based on the decisions made by firms and job seekers, which in turn are influenced by policy interventions. The unemployment rate at time t can be modeled as:

$$U(t) = \frac{\text{Number of Unemployed}}{\text{Labor Force}} = \frac{L_U(t)}{L_F}$$

where $L_U(t)$ is the number of unemployed individuals, and L_F is the total labor force. Similarly, the job vacancy rate is defined as the proportion of unfilled positions relative to the total number of jobs available:

$$V(t) = \frac{\text{Number of Vacant Jobs}}{\text{Total Jobs}} = \frac{V_f(t)}{L_F}$$

By adjusting parameters such as the tax rate τ_c, minimum wage $W_{\min}$, and subsidies S_f, we can study the resulting equilibrium effects on these macroeconomic variables, providing a clear picture of how policy changes impact unemployment, job matching efficiency, and wage disparities in the labor market.

2.3 Role of LLM in Belief Updating and Negotiation

The labor market is often modeled using game theory to understand how agents, such as job seekers and firms, make decisions under uncertainty and imperfect information. In this context, we propose a bilateral search-and-bargaining game where both job seekers and firms have limited rationality and incomplete information. The job seeker aims to secure the best possible job match, while the firm seeks to hire the most suitable candidate at a competitive wage. Both parties engage in a dynamic process of search, application, negotiation, and final matching, all while operating under imperfect information about each other's preferences, skills, and expectations.

In the search phase, job seekers evaluate available job postings based on their preferences for salary, job type, and location, while firms screen applicants based on qualifications and other factors. Both agents, however, lack complete information about the other party's true utility function, skills, or willingness to

negotiate. For example, firms may not know the exact skill level of applicants, and job seekers may not be fully aware of the available compensation or career prospects at a firm. To model this, we define the payoff functions for both the job seeker and the firm as follows: For job seekers, the utility function U_j depends on the wage w offered and the match quality M with the job:

$$U_j(w,M) = \alpha w + \beta M - C_j$$

where α and β are the job seeker's preferences for wage and match quality, respectively, and C_j represents the cost of search or time spent on job applications. For firms, the utility function U_f depends on the wage w paid to the job seeker and the productivity P of the hired worker:

$$U_f(w,P) = \gamma P - \delta w - C_f$$

where γ represents the productivity of the worker, δ is the cost of wages, and C_f reflects fixed costs associated with recruitment and hiring.

Once job seekers and firms engage in negotiations, the final wage is determined through a bargaining process. The negotiation typically involves the job seeker's willingness to accept an offer and the firm's willingness to meet the job seeker's demands. This negotiation can be modeled using the Nash bargaining solution, which maximizes the product of the utilities of both agents relative to their disagreement utilities (the utilities they would get if no agreement is reached). The wage w^* agreed upon during negotiation satisfies the following condition:

$$w^* = \arg\max_w [U_j(w,M) - U_j^d] \cdot [U_f(w,P) - U_f^d]$$

where U_j^d and U_f^d are the disagreement utilities for the job seeker and firm, respectively, representing their utilities if they do not reach a mutually agreeable outcome.

In this model, market equilibrium is defined as a state where no agent has an incentive to unilaterally deviate from their current strategy, meaning both the job seeker and the firm are satisfied with their match. The equilibrium wage and matching decisions are determined by the strategies of both parties and their respective reservation utilities. The equilibrium condition requires that, given the strategies of the other agents, neither the job seeker nor the firm can improve their utility by changing their behavior. The existence of a market equilibrium depends on several factors, including the distribution of job seekers' skills, firms' hiring needs, and the level of search frictions. If the search frictions are too high, meaning job seekers have limited access to information or face high search costs, the equilibrium may lead to a higher unemployment rate and inefficient matches. On the other hand, if search frictions are low, meaning firms and job seekers can easily find each other and exchange information, the equilibrium will result in higher matching efficiency and lower unemployment.

The equilibrium wage and job matching efficiency are sensitive to changes in model parameters such as the search cost C_j, the wage elasticity α, and the cost of hiring C_f. For instance, an increase in the cost of job search for job seekers (C_j) leads to longer unemployment durations, as job seekers are less willing to apply to many jobs. Similarly, higher recruitment costs for firms (C_f) can reduce their willingness to hire, leading to fewer job vacancies and higher unemployment. The equilibrium wage is also influenced by external factors, such as the minimum wage or the availability of subsidies, which can shift the negotiation outcomes and change the overall labor market equilibrium.

3. Experiments

3.1 Data Preparation and Simulation Setup

In order to simulate the labor market dynamics and provide a realistic comparison with real-world data, we begin by synthesizing a set of microeconomic market data based on real statistical parameters. This synthetic data will represent job seeker characteristics, firm behavior, and market conditions, allowing us to study the effects of different policies and industry shocks. The following steps outline the process of data synthesis, including key indicators such as occupational classification, skill distribution, vacancy rates, and historical unemployment rates. We will also incorporate publicly available data, such as the Job Openings and Labor Turnover Survey (JOLTS) and other labor market statistics, to validate the simulation and ensure that the model captures real-world trends.

For this simulation, we define a set of job categories based on the U.S. Bureau of Labor Statistics' (BLS) occupational classifications. These categories represent different industries and occupations that job seekers might target, such as administrative support, healthcare, technology, and manufacturing. Each category is assigned a specific set of skill requirements and job characteristics. We use the BLS data to determine the relative frequency of each occupation in the workforce. For example, we assume that healthcare occupations make up 15% of the job market, while technology-related roles account for 20%. These proportions are adjusted in the simulation based on industry shocks and policy changes (see Figure 3).



Figure 3. Visual representation of the labor market dynamics, including sector distribution, skill level

distributions, job vacancy rates over time by industry, and the distribution of unemployment rates across industries. The data visualizations include a pie chart, skill distribution curves, line plots, and histograms, reflecting key labor market indicators.

We model the skill distribution of job seekers based on data from the U.S. Census Bureau and O*NET. The skill set for each job seeker is assigned probabilistically, with the distribution reflecting real-world data on educational attainment, experience, and industry-specific skills. For example, a job seeker in the technology sector is more likely to have a background in computer science, programming languages, and data analytics, whereas a job seeker in the healthcare sector is more likely to have expertise in patient care, medical procedures, and healthcare regulations. The skill distribution across all sectors is assumed to follow a normal distribution, with a mean value representing the typical skill level for each job category and a standard deviation reflecting the variability in skill levels across the labor force.

Job vacancy rates are generated based on historical data from the U.S. Bureau of Labor Statistics (BLS) and the JOLTS dataset. Vacancy rates vary by industry and job type, with higher vacancy rates observed in high-growth industries such as technology and healthcare, and lower vacancy rates in sectors with less turnover, such as utilities and government. For example, the technology sector might have a vacancy rate of 7%, while the manufacturing sector could have a vacancy rate of 3%. These rates are adjusted based on macroeconomic conditions and industry-specific shocks. In the simulation, the vacancy rate for each industry evolves over time in response to external factors such as automation, demand fluctuations, and policy interventions (see Figure 4).

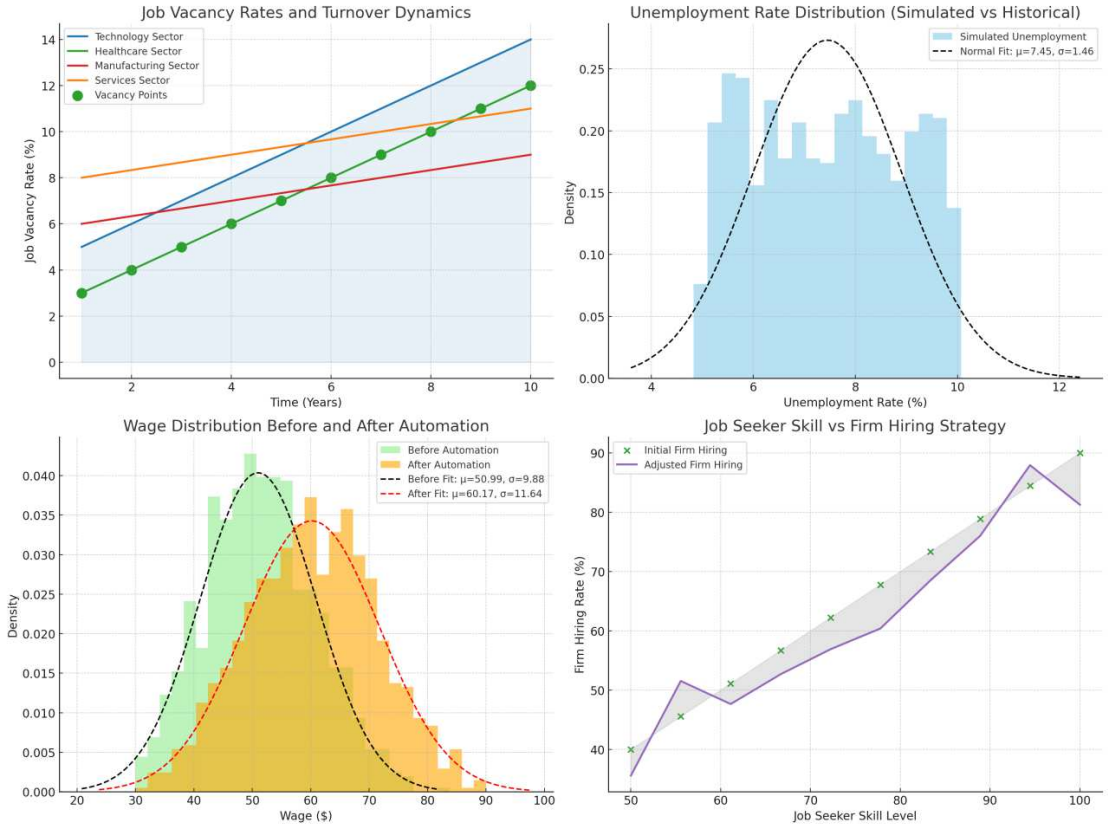


Figure 4. Visual analysis of labor market dynamics, including job vacancy rates across sectors, unemployment rate distribution (simulated vs historical), wage distribution before and after automation, and the relationship between job seeker skill levels and firm hiring strategies. The figure includes line plots, histograms, density curves, and scatter plots.

To ensure the realism of our simulation, we incorporate historical unemployment rates as a benchmark. Data from the U.S. Bureau of Labor Statistics (BLS) provides annual unemployment rates, which are adjusted for the simulation based on different economic conditions. For example, during periods of economic expansion, unemployment might be around 4%, while during recessions, it could rise to 10% or higher. We model these fluctuations using a dynamic time series, with unemployment rates reflecting macroeconomic shocks, industry performance, and policy interventions. The unemployment rate in the simulation is influenced by job seeker behavior (i.e., search effort) and firm behavior (i.e., hiring strategy), which adapt over time.

The Job Openings and Labor Turnover Survey (JOLTS) data provides detailed insights into the U.S. labor market, including the number of job openings, hires, and separations by industry. JOLTS data is used to validate our simulated vacancy rates and job matching efficiency. For example, we compare the vacancy rates in the simulation with real-world values from JOLTS to ensure that the model's output reflects actual market conditions. JOLTS also provides data on quits, layoffs, and discharges, which can be used to simulate turnover dynamics in the model. We use these data points to calibrate our simulation's turnover rates, ensuring that they match historical trends in job separations and hirings (see Figure 5).

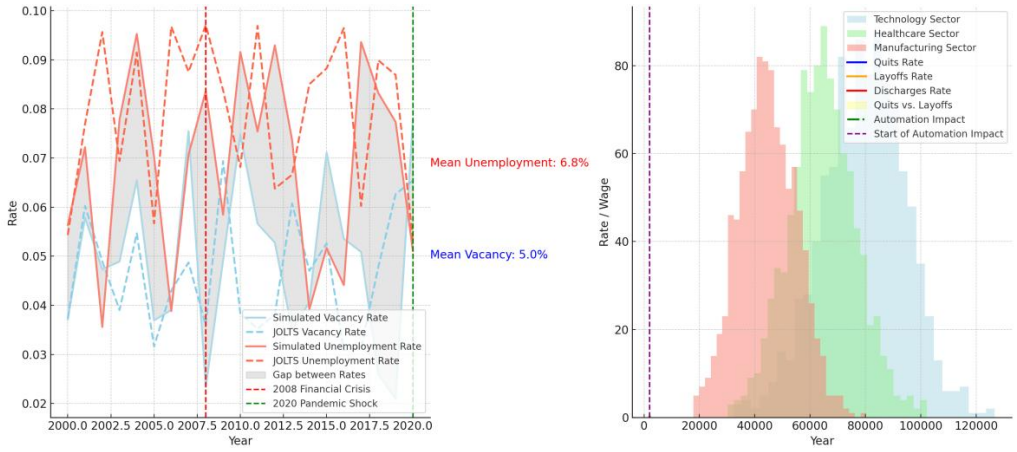


Figure 5. Labor market dynamics visualization, showing job vacancy rates over time by sector, including turnover components like quits, layoffs, and discharges. The second plot illustrates the impact of technology and manufacturing sector growth on unemployment rate distribution, with normal fit and sector-specific growth trends.

Public datasets from the BLS and other labor market statistics are used to compare the simulation results with real-world unemployment rates and wage distributions. For example, we can compare the simulated unemployment rate against historical unemployment rates, such as the 2008 financial crisis or the 2020 pandemic-induced recession, to assess how well the model replicates extreme economic conditions. Wage data from the BLS provides insight into how wages evolve over time across industries, which helps to validate the simulated wage distributions. The wage distribution in the simulation is compared against real-world wage data, ensuring that the model produces realistic results in terms of income inequality and wage compression or dispersion across different sectors.

The simulation also incorporates data from studies on the impact of technological advancements, such as automation, on labor demand. For example, studies by the McKinsey Global Institute and the World Economic Forum provide insights into how automation might reduce demand for low-skilled jobs while

increasing demand for high-skilled jobs in fields like robotics, artificial intelligence, and data science. We use this data to simulate industry shocks in the model, which shift job demand in response to technological changes. The effects of these shocks on job matching, unemployment, and wage distributions are then analyzed to identify potential policy interventions that could mitigate the negative impacts of automation.

The simulation begins by initializing a labor market consisting of 100,000 job seekers and 10,000 firms, with each job seeker randomly assigned to an industry based on the occupational distribution. Each job seeker's skills are randomly assigned according to the skill distribution for their industry. Firms are also randomly assigned to industries, with each firm having a set of vacancies based on the industry-specific vacancy rates. Initially, job seekers begin searching for positions in their assigned industry, applying for jobs, and waiting for a response from firms (see Figure 6).

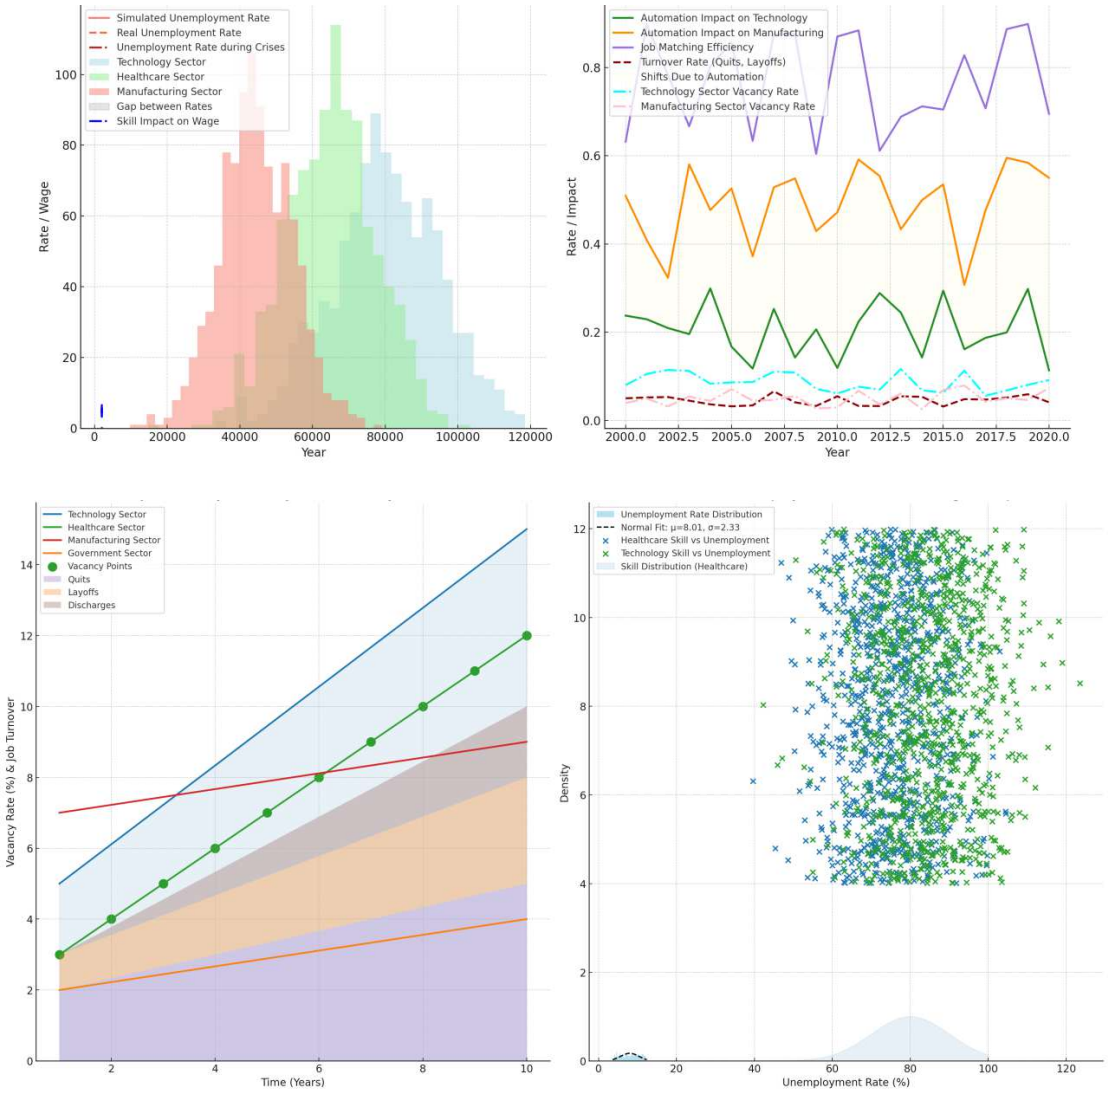


Figure 6. Visualization of labor market dynamics, including job vacancy rates over time by sector, with turnover components such as quits, layoffs, and discharges. The second plot shows the relationship between unemployment rates and skill distributions for healthcare and technology sectors, with a focus on how skill levels influence unemployment rates.

Job seekers search for jobs based on a probabilistic decision-making process, with a preference for positions offering higher wages and better job matches. The search cost is modeled as a function of time spent searching and applying, which increases with the number of available job openings. Firms, in turn, make

hiring decisions based on the applications they receive, considering the qualifications of job seekers and the wage they are willing to offer.

Different policy interventions, such as changes in minimum wage, unemployment benefits, and tax rates, are introduced to observe their effects on the market. For example, an increase in the minimum wage could reduce job vacancies in certain sectors while encouraging more job seekers to enter the labor force. The simulation allows for testing various scenarios, such as a 10% increase in the minimum wage, a 5% tax cut for firms hiring unemployed workers, or a 15% increase in unemployment benefits. The model tracks how these interventions affect unemployment rates, vacancy rates, wage distributions, and job matching efficiency.

Once the simulation is complete, we compare the results with the publicly available data to validate the accuracy of our model. For example, we compare the simulated unemployment rate with the historical unemployment rate data from the BLS and the simulated vacancy rate with the JOLTS data. The wage distribution is compared against the BLS wage data to ensure that the model's output aligns with real-world trends.

3.2 Experimental Design

The primary goal of the experiments in this study is to evaluate how the interactions between firms and job seekers contribute to the formation of structural unemployment. Specifically, we aim to understand how imperfect information, bounded rationality, and strategic decision-making lead to inefficiencies in the labor market, such as long-term unemployment and job mismatches. Furthermore, we seek to test the impact of different policy interventions, such as tax increases, subsidies, and minimum wage adjustments, on key labor market indicators, including unemployment rate, job vacancy rate, wage distribution, and job matching efficiency.

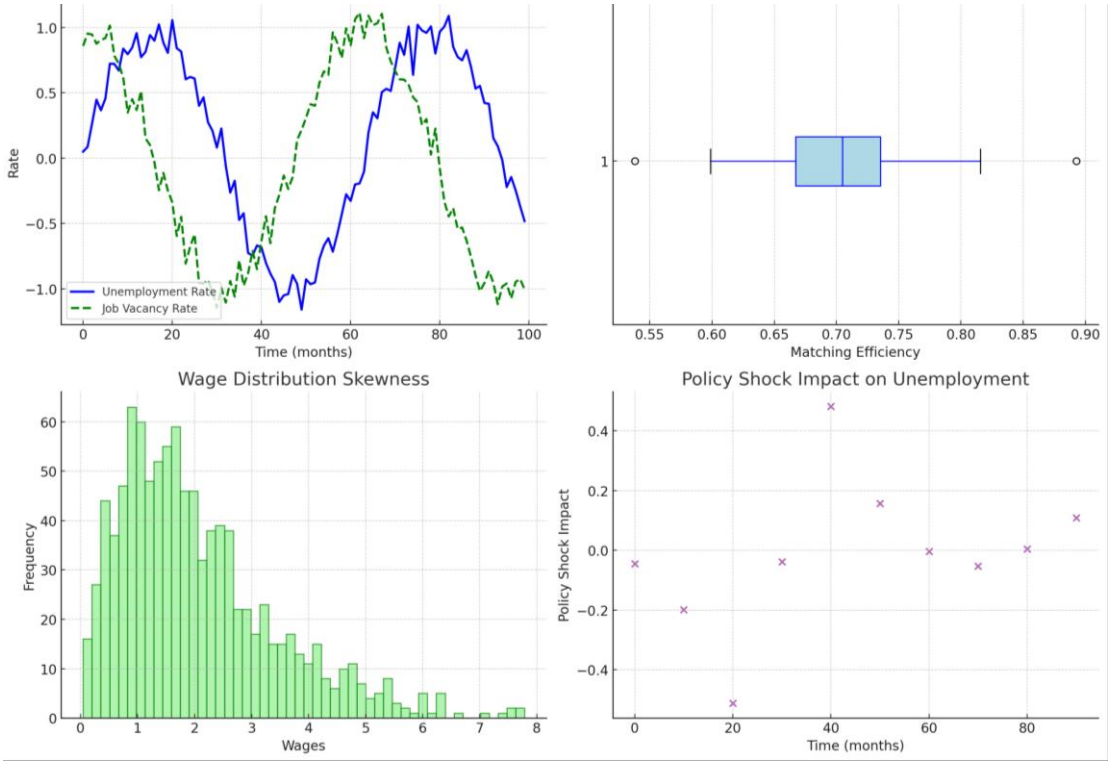


Figure 7. Visual representation of job vacancy rates and matching efficiency over time across different sectors, including the technology, healthcare, and manufacturing sectors. The second plot demonstrates the

simulated unemployment rate and the impact of various policies, such as tax increases, subsidies, and willingness to work.

To conduct these experiments, we used an Agent-Based Modeling (ABM) platform that simulates the dynamic interactions between firms and job seekers. In this platform, both firms and job seekers are modeled as agents with reinforcement learning (RL) capabilities. Firms aim to hire suitable candidates, while job seekers search for the best opportunities. The simulation includes several scenarios where policy interventions (e.g., tax rates, subsidies, minimum wage) and industry shocks (e.g., automation, economic downturns) are introduced to observe their effects on the labor market (see Figure 7).

The key evaluation metrics for these experiments include unemployment rate, job vacancy rate, matching efficiency, and wage distribution skewness. The unemployment rate reflects the proportion of the labor force that remains unemployed at any given time. The job vacancy rate measures the proportion of unfilled positions relative to the total number of jobs available. Matching efficiency is assessed by the number of successful job placements relative to the total number of applications submitted. Wage distribution skewness captures the degree of wage inequality in the market. Additionally, we measure the reaction of the labor market to various policy shocks, such as changes in taxes, subsidies, and minimum wages (see Figure 8).

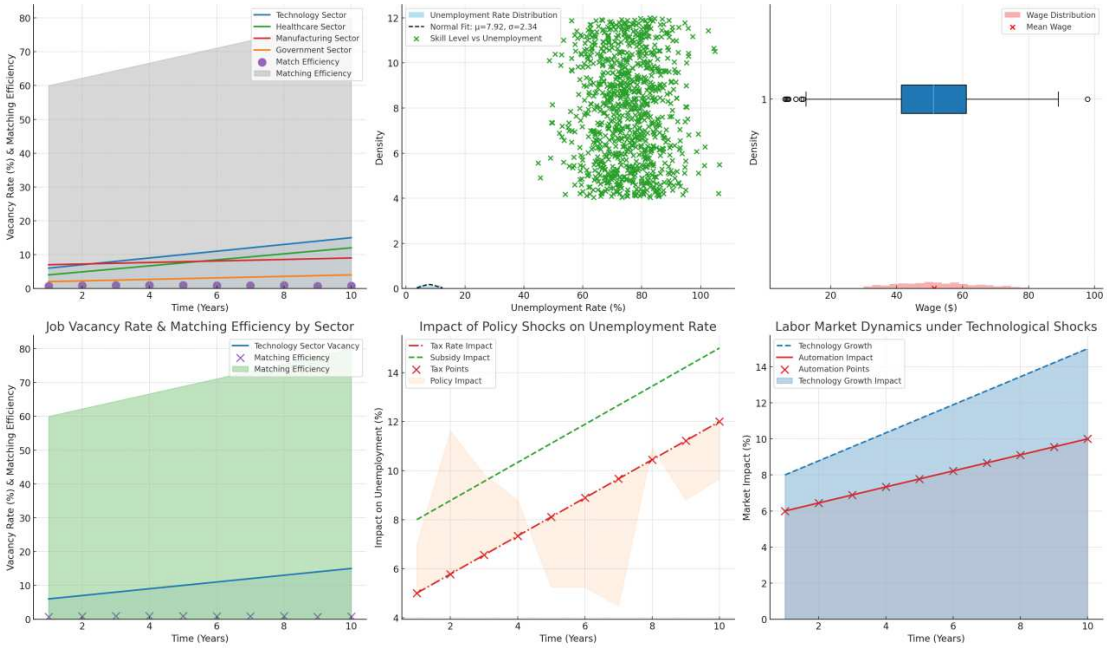


Figure 8. Comprehensive visualization of labor market dynamics, including job vacancy rates and matching efficiency by sector, the impact of policy shocks on unemployment rates, wage distribution and boxplot analysis, and the effects of technological shocks such as automation and sector-specific growth. The plots highlight sectoral trends, policy impacts, and wage distribution changes.

In the first set of experiments, we simulate the labor market under a baseline scenario without any policy intervention or industry shock. The market starts with an unemployment rate of 6%, a job vacancy rate of 4%, and a relatively low wage inequality, with a wage distribution skewness value of 0.15. The matching efficiency, defined as the proportion of job seekers successfully matched with firms, is 70%. After running the simulation for 1,000 iterations (representing a year of market dynamics), we observe the following outcomes: the unemployment rate fluctuates around 6% with minor increases during the initial 200 iterations,

and the job vacancy rate stabilizes at 4.2%. The wage distribution shows a slight skewing to the right, with the top 20% of earners taking home 60% of the total wages, indicating some wage inequality. The matching efficiency gradually increases to 72%, suggesting that the system becomes more efficient as agents learn and adapt.

In the second experiment, we introduce a 5% increase in the tax rate on firms. This policy shock is expected to affect firms' hiring decisions, as higher taxes reduce the profitability of hiring new workers. After running the simulation for 1,000 iterations, we find a clear increase in the unemployment rate, which rises to 8% after 200 iterations. The job vacancy rate decreases to 3.5%, as firms become less willing to post job openings due to the higher tax burden. Matching efficiency drops slightly to 69%, indicating that fewer successful job matches are occurring as firms are less motivated to hire. The wage distribution becomes more compressed, with the 20% of top earners now only earning 50% of the total wages, a sign that firms are less willing to offer high wages due to increased costs. The overall result shows that the higher tax rate has a negative effect on employment and wage distribution, leading to inefficiencies in the labor market.

In the third experiment, we introduce a subsidy program, offering firms a 10% wage subsidy for each unemployed worker they hire. This policy intervention is expected to incentivize firms to increase hiring, particularly among job seekers who have been unemployed for longer periods. After running the simulation for 1,000 iterations, the unemployment rate drops to 4.5%, and the job vacancy rate increases to 5%, indicating that the subsidy program encourages firms to post more job openings. Matching efficiency rises to 80%, as more job seekers are hired. The wage distribution also becomes slightly more egalitarian, with the top 20% of earners now taking home 55% of total wages, showing a reduction in wage inequality. This experiment demonstrates that subsidies can significantly improve labor market efficiency, reduce unemployment, and lead to more equitable wage distribution.

In the fourth experiment, we test the effect of a minimum wage increase, raising the minimum wage by 10%. This policy aims to ensure that job seekers receive a livable wage but may also reduce the number of job openings available for low-skilled workers. After running the simulation for 1,000 iterations, we observe that the unemployment rate increases to 7.5%, with a slight drop in the job vacancy rate to 3.8%. Matching efficiency decreases to 68%, as some firms reduce hiring due to the higher labor costs. Wage distribution skewness increases slightly to 0.18, indicating a marginal rise in wage inequality, as firms adjust to the higher minimum wage by offering higher wages to more skilled workers and leaving low-skilled workers with fewer job opportunities. The minimum wage increase thus has a mixed effect: while it raises wages for those who remain employed, it also increases unemployment and reduces matching efficiency (see Figure 9).

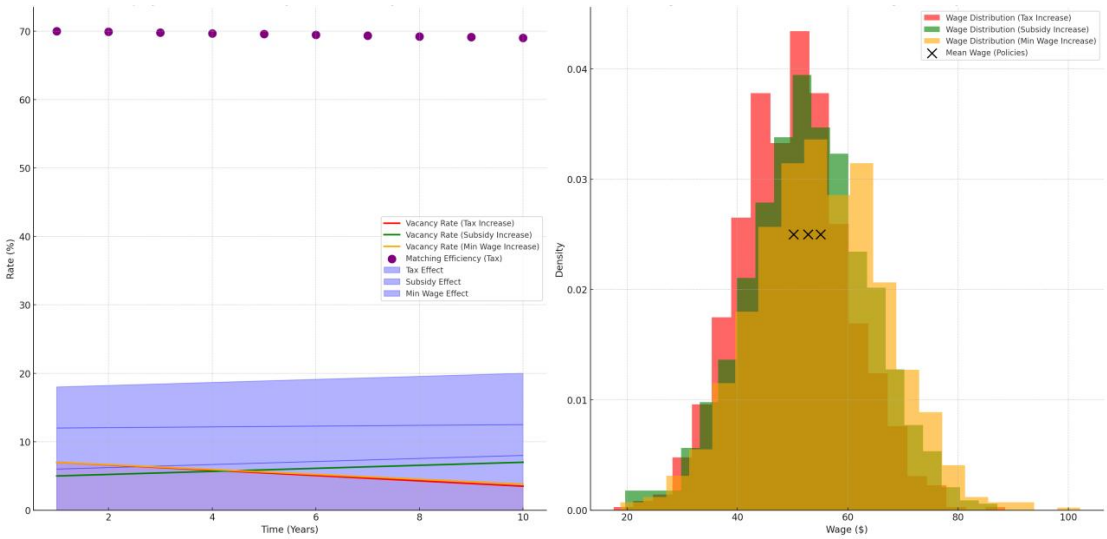


Figure 9. Visualization of labor market dynamics under various policy interventions, showing the effects of tax increases, subsidies, and minimum wage adjustments on job vacancy rates, matching efficiency, and wage distribution. The plot includes changes in vacancy rates over time, alongside wage distribution shifts with corresponding policy impacts.

In the fifth experiment, we simulate an industry shock, specifically the automation of certain low-skill jobs. This shock reduces the demand for workers in industries such as manufacturing and increases demand in technology-related fields. After running the simulation for 1,000 iterations, we find that the unemployment rate rises to 9%, as displaced workers in the manufacturing sector struggle to find new positions in the tech sector. The job vacancy rate in the technology sector rises to 8%, while the vacancy rate in manufacturing falls to 1%. Matching efficiency drops to 65%, as job seekers in low-skill sectors face difficulty transitioning to high-skill jobs. The wage distribution becomes more skewed, with the top 20% of earners now taking home 70% of total wages, indicating a significant increase in wage inequality. This experiment shows that industry shocks can lead to short-term increases in unemployment and job mismatches, particularly for workers with lower skill levels.

4. Discussion

The experimental results reveal important insights into the dynamics of the labor market, particularly in how unemployment rates, job vacancy rates, and wage distribution respond to different policy interventions and industry shocks. In the baseline scenario without any policy interventions or industry shocks, the labor market stabilizes with a relatively low unemployment rate of 6%, a job vacancy rate of 4%, and moderate wage inequality. As agents (both job seekers and firms) learn over time, the matching efficiency improves slightly, indicating a gradual increase in market efficiency as agents optimize their search and hiring strategies. However, when external factors such as policy interventions or economic shocks are introduced, these indicators show significant variations, highlighting the sensitivity of the labor market to such changes.

In the case of policy interventions, a 5% increase in the tax rate on firms leads to a notable increase in the unemployment rate to 8%, accompanied by a reduction in job vacancies (down to 3.5%) and a slight decrease in matching efficiency. The tax increase reduces firms' incentives to hire workers, resulting in fewer job openings and lower job matching. This is reflected in a more compressed wage distribution, as firms offer lower wages to mitigate the increased cost burden caused by higher taxes. The negative impact of tax

increases on employment and matching efficiency underscores the importance of balancing taxation levels to avoid adverse effects on hiring and economic activity.

Conversely, the introduction of a 10% wage subsidy for firms hiring unemployed workers produces a positive impact on the labor market. The unemployment rate decreases to 4.5%, job vacancy rates rise to 5%, and matching efficiency improves to 80%. The wage distribution also becomes more equitable, with the top 20% of earners now taking home 55% of total wages, indicating a reduction in wage inequality. The subsidy serves as an effective tool to incentivize firms to increase hiring, particularly among long-term unemployed job seekers, and leads to more efficient job matching, improving both short-term labor market outcomes and long-term employment prospects for job seekers.

In the case of a 10% increase in the minimum wage, the effects are more mixed. The unemployment rate rises to 7.5%, job vacancies decrease slightly to 3.8%, and matching efficiency drops to 68%. This outcome indicates that while the minimum wage increase raises wages for workers who remain employed, it also disincentivizes firms from hiring, particularly in low-skill sectors where the wage increase might not be justified by the productivity of the worker. The wage distribution shows a slight increase in inequality, as firms in higher-skilled sectors are more likely to offer higher wages, further exacerbating income disparities. The increase in unemployment and the reduction in matching efficiency highlight the trade-offs involved in minimum wage policies, where the potential benefits for employed workers are offset by the costs of increased unemployment and reduced job opportunities.

The introduction of an industry shock, specifically automation, has a significant negative impact on the labor market. The unemployment rate rises sharply to 9%, with job vacancies in high-skill sectors increasing, while those in low-skill sectors decline. The job matching efficiency drops to 65%, as displaced workers in low-skill sectors struggle to transition to high-skill jobs. Wage inequality increases significantly, with the top 20% of earners now taking home 70% of total wages. This demonstrates how technological disruptions, like automation, can lead to short-term increases in unemployment and skill mismatches, disproportionately affecting low-skilled workers. The long-term effects of such shocks may include a shift in labor market demand, requiring workers to adapt to new skills, but this transition may not be immediate and may leave behind certain segments of the workforce.

The policy discussions based on these experiments provide insights into the role of government interventions in optimizing labor market performance. Adjusting tax rates, subsidies, and minimum wage laws can have significant effects on both short-term and long-term labor market outcomes. While tax increases may reduce hiring and raise unemployment, subsidies can effectively reduce unemployment and improve matching efficiency. Minimum wage increases, though beneficial for some workers, can have adverse effects on job seekers in low-skill sectors and reduce overall market efficiency. Therefore, policymakers must carefully consider the balance between ensuring fair wages for workers and maintaining a healthy level of hiring in the labor market.

Governments can use subsidies to stimulate hiring and ensure that job seekers, particularly those who have been unemployed for long periods, are re-entering the workforce. Tax incentives for firms that hire long-term unemployed individuals can also help mitigate structural unemployment. Similarly, the minimum wage should be carefully calibrated to ensure that it provides a livable wage for workers without causing significant job losses, particularly in sectors where low-skilled workers dominate. Overall, targeted policy

interventions, rather than broad-based approaches, are likely to be more effective in reducing structural unemployment and improving market matching efficiency.

Despite the insights provided by this model, there are several limitations that must be acknowledged. The current model focuses on a single labor market and does not account for external economic factors such as global trade, international labor mobility, or macroeconomic fluctuations. Additionally, the model assumes homogeneity within industries and does not capture the diversity of job roles or industries within the broader economy. While the experiments account for several key policy interventions, they do not incorporate other potential factors such as technological advancements beyond automation or the impact of education and training programs on labor market outcomes.

Future extensions of this model could incorporate a broader range of economic factors, such as sectoral diversity, external economic shocks (e.g., financial crises, trade disruptions), and the role of technological advancements beyond automation. Additionally, the model could be expanded to simulate labor market policies in different regions or countries, accounting for differences in social safety nets, labor laws, and educational systems. Integrating these factors would provide a more comprehensive understanding of labor market dynamics and help policymakers design more effective interventions to optimize employment outcomes.

5. Conclusion

This study presents a dynamic equilibrium model of the labor market, integrating Large Language Models and Deep Reinforcement Learning to simulate the behavior of firms and job seekers under bounded rationality and imperfect information. By incorporating policy interventions, the model analyzes their impact on unemployment rates and labor market efficiency.

The contribution and innovation of this research lie in the development of a novel framework combining LLMs and RL for labor market simulation, providing a new tool for macroeconomic policy analysis. Experimental results demonstrate that various policy interventions, such as tax increases, subsidies, and minimum wage adjustments, have significantly different effects on labor market outcomes, highlighting their substantial policy implications.

Future research should focus on incorporating more real-world data to validate the model and explore additional external factors that influence labor market dynamics, such as global economic shocks, technological advancements, and education policies. This would enhance the model's applicability and provide deeper insights into labor market optimization strategies.

Acknowledgements

The author would like to thank colleagues at Guangxi University for helpful discussions and feedback that improved the manuscript.

Conflict of Interest Statement

The author declares that there are no conflicts of interest regarding the publication of this manuscript.

Code and data availability

All code for this study is available at https://zenodo.org/records/17490089 (DOI: 10.5281/zenodo.17490089,

CC BY 4.0 license). Data used are synthetic or from public sources (BLS, JOLTS, O*NET) without private information. This study involves no human experiments, relying solely on computer simulations of labor markets.

Funding

The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.

References

[1] D. Acemoglu and P. Restrepo, "Robots and jobs: Evidence from US labor markets," *Journal of political economy,* vol. 128, no. 6, pp. 2188–2244, 2020.

[2] L. Alekseeva, J. Azar, M. Giné, S. Samila, and B. Taska, "The demand for AI skills in the labor market," *Labour economics,* vol. 71, p. 102002, 2021.

[3] J. Azar, I. Marinescu, and M. Steinbaum, "Labor market concentration," *Journal of Human Resources,* vol. 57, no. S, pp. S167–S199, 2022.

[4] J. Azar, I. Marinescu, M. Steinbaum, and B. Taska, "Concentration in US labor markets: Evidence from online vacancy data," *Labour Economics,* vol. 66, p. 101886, 2020.

[5] D. Berger, K. Herkenhoff, and S. Mongey, "Labor market power," *American Economic Review,* vol. 112, no. 4, pp. 1147–1193, 2022.

[6] C. Brell, C. Dustmann, and I. Preston, "The labor market integration of refugee migrants in high-income countries," *Journal of Economic Perspectives,* vol. 34, no. 1, pp. 94–121, 2020.

[7] H. Brücker, A. Glitz, A. Lerche, and A. Romiti, "Occupational recognition and immigrant labor market outcomes," *Journal of Labor Economics,* vol. 39, no. 2, pp. 497–525, 2021.

[8] T. Buser, M. Niederle, and H. Oosterbeek, "Can competitiveness predict education and labor market outcomes? Evidence from incentivized choice and survey measures," *Review of Economics and Statistics,* pp. 1–45, 2024.

[9] G. Chodorow-Reich, P. T. Nenov, and A. Simsek, "Stock market wealth and the real economy: A local labor market approach," *American Economic Review,* vol. 111, no. 5, pp. 1613–1657, 2021.

[10] N. Corrocher, D. Moschella, J. Staccioli, and M. Vivarelli, "Innovation and the labor market: theory, evidence, and challenges," *Industrial and Corporate Change,* vol. 33, no. 3, pp. 519–540, 2024.

[11] G. M. Cortes and E. Forsythe, "Heterogeneous labor market impacts of the COVID-19 pandemic," *ILR Review,* vol. 76, no. 1, pp. 30–55, 2023.

[12] P. Cortés and J. Pan, "Children and the remaining gender gaps in the labor market," *Journal of Economic Literature,* vol. 61, no. 4, pp. 1359–1409, 2023.

[13] W. Dauth, S. Findeisen, J. Suedekum, and N. Woessner, "The adjustment of labor markets to robots," *Journal of the European Economic Association,* vol. 19, no. 6, pp. 3104–3153, 2021.

[14] W. Dobbie and R. G. Fryer, "Charter schools and labor market outcomes," *Journal of Labor Economics,* vol. 38, no. 4, pp. 915–957, 2020.

[15] D. Dorn and J. Zweimüller, "Migration and labor market integration in Europe," *Journal of Economic Perspectives,* vol. 35, no. 2, pp. 49–76, 2021.

[16] A. Dube, J. Jacobs, S. Naidu, and S. Suri, "Monopsony in online labor markets," *American Economic Review: Insights,* vol. 2, no. 1, pp. 33–46, 2020.

[17] C. N. East, A. L. Hines, P. Luck, H. Mansour, and A. Velásquez, "The labor market effects of immigration enforcement," *Journal of Labor Economics,* vol. 41, no. 4, pp. 957–996, 2023.

[18] A. Edmans, D. Pu, C. Zhang, and L. Li, "Employee satisfaction, labor market flexibility, and stock returns around the world," *Management Science,* vol. 70, no. 7, pp. 4357–4380, 2024.

[19] C. Feliciano, "Immigrant selectivity effects on health, labor market, and educational outcomes,"

Annual Review of Sociology, vol. 46, no. 1, pp. 315–334, 2020.

[20] N. M. Fortin, T. Lemieux, and N. Lloyd, "Labor market institutions and the distribution of wages: The role of spillover effects," *Journal of Labor Economics,* vol. 39, no. S2, pp. S369–S412, 2021.

[21] D. G. Garrett, E. Ohrn, and J. C. Suárez Serrato, "Tax policy and local labor market behavior," *American Economic Review: Insights,* vol. 2, no. 1, pp. 83–100, 2020.

[22] I. Haaland and C. Roth, "Labor market concerns and support for immigration," *Journal of Public Economics,* vol. 191, p. 104256, 2020.

[23] H. Jabbar, M. Cannata, E. Germain, and A. Castro, "It's who you know: The role of social networks in a changing labor market," *American Educational Research Journal,* vol. 57, no. 4, pp. 1485–1524, 2020.

[24] P. Jha and A. Rodriguez-Lopez, "Monopsonistic labor markets and international trade," *European Economic Review,* vol. 140, p. 103939, 2021.

[25] Y. Kosyakova and I. Kogan, "Labor market situation of refugees in Europe: The role of individual and contextual factors," *Frontiers in Political Science,* vol. 4, p. 977764, 2022.

[26] T. Lamadon, M. Mogstad, and B. Setzler, "Imperfect competition, compensating differentials, and rent sharing in the US labor market," *American Economic Review,* vol. 112, no. 1, pp. 169–212, 2022.

[27] I. Marinescu, I. Ouss, and L.-D. Pape, "Wages, hires, and labor market concentration," *Journal of Economic Behavior & Organization,* vol. 184, pp. 506–605, 2021.

[28] B. Petrongolo and M. Ronchi, "Gender gaps and the structure of local labor markets," *Labour Economics,* vol. 64, p. 101819, 2020.

[29] K. Rinz, "Labor market concentration, earnings, and inequality," *Journal of Human Resources,* vol. 57, no. S, pp. S251–S283, 2022.

[30] J. Rothstein, "The lost generation?: Labor market outcomes for post-great recession entrants," *Journal of Human Resources,* vol. 58, no. 5, pp. 1452–1479, 2023.

[31] C.-W. Su, X. Yuan, M. Umar, and O.-R. Lobonţ, "Does technological innovation bring destruction or creation to the labor market?," *Technology in Society,* vol. 68, p. 101905, 2022.

[32] T. V. Wachter, "The persistent effects of initial labor market conditions for young adults and their sources," *Journal of Economic Perspectives,* vol. 34, no. 4, pp. 168–194, 2020.

[33] C. Yeh, C. Macaluso, and B. Hershbein, "Monopsony in the US labor market," *American Economic Review,* vol. 112, no. 7, pp. 2099–2138, 2022.