AI, Labor Adjustment, and Spatial Equilibrium*

Ying Fan† Yiran Niu‡

June 5, 2026

Abstract

AI penetration jointly reshapes job opportunities, work intensity, and labor sorting between and within cities through extensive and intensive labor adjustments. We develop a dynamic spatial equilibrium model in which AI enters goods firms through a task-based production structure. Goods firms jointly choose output scale, employment headcounts, and average effective work intensity while allocating tasks between human labor and rented AI services; AI changes the frontier of AI-feasible tasks, AI task productivity, and human productivity. Workers respond by optimizing over labor supply, job/status choices, and residential locations, with actual hours affected by retention risk and fallback values. The model therefore transmits AI-induced shocks into urban rents and labor sorting through retention competition, utility, budget, time, and equilibrium labor-market constraints. Quantitatively, the framework identifies efficiency losses, unemployment/search margins, welfare costs, and inequality associated with AI-induced labor adjustment. Counterfactual analyses evaluate how technology improvements, labor-market and housing frictions shape these distortions.

Keywords: Artificial intelligence; labor adjustment; spatial equilibrium; housing; commuting; work intensity.
JEL Classification: J22; J23; R13; R23; O33.

*The work described in this paper was supported by the National Natural Science Foundation of China (No. 72203191), the Public Policy Research Funding Scheme of The Government of the HKSAR (Project Number: 2026.A6.312.26A), and the Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. PolyU 25226423).

†The Hong Kong Polytechnic University. Email: yingfan@polyu.edu.hk
‡The Hong Kong Polytechnic University. Email: yiran.niu@connect.polyu.hk

1 Introduction

Artificial intelligence changes not only which tasks firms perform with machines, but also how workers allocate time and where they choose to live. Task-based models show that new technologies can substitute for some forms of labor while complementing others (Autor, Levy and Murnane, 2003; Autor, 2015; Acemoglu and Restrepo, 2019*b*). AI-specific evidence measures exposure across occupations, industries, and skills, while recent studies document productivity gains that vary sharply across tasks (Webb, 2020; Felten, Raj and Seamans, 2021, 2023; Acemoglu et al., 2022; Alekseeva et al., 2021; Noy and Zhang, 2023; Brynjolfsson, Li and Raymond, 2025; Dell'Acqua et al., 2026). Yet this evidence leaves open a spatial-equilibrium question: when AI changes firms' demand for human task time and employment positions, how do workers adjust both their labor allocation and their residential choices?

The answer cannot be inferred from employment and wages alone. General-equilibrium studies of AI have begun to discipline the gap between local labor-market effects and aggregate effects (Huang, 2025), and empirical work documents AI-related hiring, skill demand, workforce reorganization, and task-level reallocation (Alekseeva et al., 2021; Acemoglu et al., 2022; Babina et al., 2023; Hampole et al., 2025; Bonney et al., 2026). The missing margin is how firm-side reorganization is passed to households: when AI changes firms' desired mix of employment positions and required human task time, workers face altered retention risks and fallback values, and they may adjust through hours, fallback jobs, search, and location choices. This is consistent with evidence that adjustment costs and firm responses can make individual labor-supply responses differ from aggregate labor-market elasticities (Chetty et al., 2011). These margins matter for rent burdens, commuting valuation, welfare, and within-city sorting even when aggregate employment or average wages move little.

The mechanism begins on the firm side. Goods firms choose output scale, employment positions, and average effective work intensity while allocating tasks between human task capacity and rented AI services. Following task-based models of automation and AI (Acemoglu and Restrepo, 2018, 2019*a*,*b*), AI changes production by shifting task feasibility, AI task productivity, and human productivity. Firms internalize these task-level changes when deciding how many workers to retain and how intensively those workers are asked to work under fixed worker costs, overtime premia, and hiring and firing frictions. The firm decision determines both the number of retained positions and the performance cutoff workers must clear.

Workers respond to these firm decisions through extensive and intensive margins. Retained workers may adjust hours defensively when the cutoff becomes harder to clear. Displaced workers do not disappear from the economy; they can move to other jobs, enter

city unemployment/search, take residual or developer work, or leave for the outside option. These labor-market responses interact with housing and commuting choices because workers also choose whether to remain in the city, where to live within the city, and whether to rent or own.

Hours connect AI to space through three channels. The income channel arises because wages multiplied by hours determine housing affordability and the ability to remain in high-cost locations. The work-intensity disutility channel captures the direct cost of high hours and work pressure beyond the mechanical loss of leisure. The commuting channel arises because work and commuting jointly use scarce time, so a change in work time changes the value of living closer to job access. As a result, AI can reduce hours for some workers through income effects while increasing hours for others through performance pressure, weaker fallback opportunities, or high housing costs.

We embed these forces in a dynamic spatial general equilibrium model with a two-zone city and an outside option. Goods firms have a task-based AI production technology; developers produce private housing; households choose job/status, hours, savings, private housing, migration, and residential zones. Equilibrium clears employment slots, firm work-intensity demand, city unemployment/search, and housing markets. The baseline economy is calibrated to the observed AI-use, labor-market, housing, and spatial environment, and counterfactuals focus on three large experiments: stronger AI technology, converting defensive excess time into skill investment, and employer-provided housing support.

We contribute to the literature in three ways. First, we show that firm-level labor demand operates on both an extensive margin (headcount) and an intensive margin (human task time). While existing task-based studies document how new technologies substitute or complement specific tasks (Autor, Levy and Murnane, 2003; Autor, 2015; Acemoglu and Restrepo, 2019*a*,*b*, 2022; Webb, 2020; Felten, Raj and Seamans, 2021; Hampole et al., 2025; Bonney et al., 2026; Huang, 2025), they largely overlook how firms reorganize the remaining human work post-adoption. We bridge this gap by embedding task-level AI adoption into a unified framework where firms jointly optimize output scale, headcount, and work intensity. This allows us to capture how AI-induced firm changes in task allocation and productivity translate into workers' changes in the labor adjustment and subsequent spatial sorting behavior.

Second, we separate firm-required work intensity from household defensive hours, addressing a key limitation in current studies that view technology-driven hours adjustments as a firm-side phenomenon. While recent evidence shows that AI exposure extends the workday through firm-driven channels like productivity and monitoring (Jiang et al., 2025), and the labor-supply literature establishes that optimization frictions separate individual hours

choices from aggregate market elasticities (Chetty et al., 2011), this research overlooks how workers strategically alter their labor supply to mitigate the displacement risks induced by AI. By ignoring the worker's defensive response, the literature risks misattributing unpaid, costly worker effort to productive firm demand. We bridge this gap by modeling the interaction between firm-required hours and the worker's optimal hours under retention risk. This distinction is vital for welfare analysis: a longer workday driven by defensive effort yields no extra output or wages, but generates substantial welfare losses that depend heavily on workers' job security, outside options, and local labor-market attachment.

Finally, we contribute to urban and spatial economics by demonstrating that neglecting intensive-margin labor adjustments leads to an incomplete understanding of housing markets and residential sorting. Current spatial frameworks typically channel localized shocks through extensive margins like employment positions (Moretti, 2011; Redding and Rossi-Hansberg, 2017; Monte, Redding and Rossi-Hansberg, 2018; Hsieh and Moretti, 2019) or focus on telecommuting patterns (Davis, Ghent and Gregory, 2024; Delventhal, Kwon and Parkhomenko, 2022; Brueckner, Kahn and Lin, 2023). However, no existing study examines how shifts in required working hours translate into spatial and housing outcomes. We address this omission by analyzing how technology-induced work intensity propagates through three distinct household channels: overwork, time disutility, and commuting demand. We reveal that a longer workday alters workers' spatial valuation of rent burdens and commuting times, driving a distinct pattern of urban sorting.

The remainder of the paper is organized as follows. Section 2 provides background and motivating facts. Section 3 introduces the model. Section 4 describes the quantitative implementation. Section 5 presents baseline mechanisms. Section 6 presents the counterfactual experiments, and Section 7 concludes.

2 Background and Motivating Facts

2.1 AI Use, Labor Demand, and Work-Time Adjustment

AI use changes firms' production tradeoff without mechanically requiring either layoffs or longer workdays. In some tasks, AI raises the amount or quality of output a worker can produce in an hour; in other tasks, it substitutes for routine task capacity or changes the standards against which a worker's output is evaluated (Autor, Levy and Murnane, 2003; Acemoglu and Restrepo, 2019*b*; Hampole et al., 2025). Evidence on generative AI adoption and experimental use further shows that these productivity effects can be large but task-dependent (Bick, Blandin and Deming, 2024; Noy and Zhang, 2023; Dell'Acqua et al.,

2026). Firm-level evidence also shows that AI diffusion is multi-layered: overall adoption, deployment across business functions, and worker-task use do not always coincide (Bonney et al., 2026). The relevant shock is therefore not simply a change in employment status. It is a firm-side change in task costs and in the joint return to output scale, headcount, and work intensity.

This distinction matters because the direction of hours adjustment is not pinned down ex ante. Workers in jobs where firms choose higher work intensity may work more when the return to time rises, but they may also work less if higher income relaxes the need to supply hours. Workers in jobs whose positions contract may work more defensively before layoffs, switch jobs after displacement, enter city unemployment/search, or leave the city; workers facing high housing costs may instead supply more hours in fallback jobs to preserve consumption and location access. Evidence on AI and workday adjustment highlights firm-side productivity and contracting channels (Jiang et al., 2025); the model below allows household-side responses to emerge from task-level AI use, employment slots, firm work intensity, performance cutoffs, fallback states, and equilibrium prices.

2.2 Work Intensity in Urban Labor Markets

Work intensity is especially relevant in urban labor markets because time, income, and housing costs are jointly determined. Spatial equilibrium research shows that commuting, local prices, and housing constraints shape how workers sort across locations (Moretti, 2011; Monte, Redding and Rossi-Hansberg, 2018; Hsieh and Moretti, 2019). A change in hours affects labor income and thus the ability to pay rent or remain in the city. It also affects utility directly through the disutility of high work effort, beyond the mechanical reduction in leisure. Finally, it changes the value of commuting-time savings because commuting and work both use scarce time.

These forces imply that work intensity is not just a residual labor-supply outcome after wages and employment are determined. It is one margin through which workers adjust to AI and one channel through which AI can affect urban sorting. The same AI shock can therefore generate productivity gains, time pressure, and redistribution across locations at the same time.

2.3 Spatial Structure, Commuting, and Proximity

The spatial setting distinguishes the city from an outside option and divides the city into a core and a periphery. This structure captures broad migration and residential sorting while keeping the model close to quantitative spatial and dynamic urban frameworks (Redding

and Rossi-Hansberg, 2017; Diamond, 2016; Greaney, Parkhomenko and Van Nieuwerburgh, 2025). The core and periphery differ in rents, commuting time, and amenities, but the model does not impose a direct AI premium for the core.

This structure is enough to discipline the main spatial mechanism. A household displaced from an AI-affected job may leave the city, move to the lower-rent zone, enter city unemployment/search, or work more hours in a fallback job to remain in place. A retained household may work more defensively, bid for the core, reduce commuting time, or consume more leisure. These adjustments are equilibrium responses to firm-side task allocation, employment positions, performance cutoffs, housing costs, and commuting time.

3 Model

3.1 Environment

Time is discrete. The economy has a fixed population $\bar{N}$. The population can live in a city, denoted by $K = C$, or in an outside option, denoted by $K = O$. The city contains two residential zones, indexed by $k \in \{1, 2\}$, where Zone 1 is the core and Zone 2 is the periphery. The outside option has no internal zone choice and is represented by reduced-form outside values, income, and housing costs. City labor, search, housing, and development-input markets clear endogenously; the outside option disciplines migration by providing an alternative value of living and working outside the city.

Production has three blocks. Goods firms produce final output in job groups $g \in \mathcal{G}$. Local service firms S produce nontradable city services with little direct AI exposure. Housing developers D produce private housing in each city zone. AI affects the economy through goods firms' task allocation, employment positions, and work intensity. Local services and developers do not face a direct AI task shock, although they respond to wages, population, income, rents, and housing demand in equilibrium. Let $\mathcal{J}_C \equiv \{\mathcal{G}, S, D\}$ denote city wage-paying jobs and sectors, and let $\mathcal{J} \equiv \mathcal{J}_C \cup \{U_C, O\}$ denote the full job/status set, where U_C is city unemployment/search and O is the outside option.

Household i enters period t with state

$$s_{it} = (age_i, a_{it}, z_{it}, \{\chi_{ij,t}\}_{j \in \mathcal{J}_C}, j_{i,t-1}, K_{i,t-1}, k_{i,t-1}, r_{i,t-1}) . \tag{1}$$

Here a_{it} is assets, z_{it} is individual productivity, $\chi_{ij,t}$ is worker-job fit, $j_{i,t-1}$ is the previous-period job or status, and $(K_{i,t-1}, k_{i,t-1}, r_{i,t-1})$ record previous-period region, zone, and tenure states. City working jobs have job-specific effective-hour wages W_{jCt} for $j \in \mathcal{J}_C$. These wages are calibrated to observed baseline pay differences and adjust in equilibrium with job-level

effective labor demand and supply. City unemployment/search receives search income or transfers rather than wage income. These objects allow employment, unemployment/search, and outside residence to adjust separately.

Throughout the model, n_{it} denotes household actual hours. Employed city jobs require at least the common normal full-time benchmark $\bar{h}^F$, so $n_{it} \geq \bar{h}^F$ for $j \in \mathcal{J}_C$, while city unemployment/search has $n_{it} = 0$. For a retained goods worker, firm-paid required work intensity is h^F_{gt}, with $\bar{h}^F \leq h^F_{gt} \leq n_{it}$. The difference $n_{it} - h^F_{gt}$ is unpaid defensive overwork. In local service and developer work, actual hours are paid hours; the formal paid-hours definition appears in the household problem below.

3.2 A Motivating Illustration

Goods firms choose employment headcount, work intensity, output scale, and task allocation jointly. The first step is task-side: AI changes feasible machine tasks, AI task costs, and human task productivity before firms decide how to organize residual human task demand into positions and hours. Figure 1 gives the roadmap. Panel (a) separates AI task feasibility, rented AI services, and human task productivity. Panel (b) then places this task-side shock in the paper's mechanism: firms translate human task demand into output scale, employment headcount, and work intensity, and households respond through hours, fallback choices, housing demand, commuting, and spatial sorting.

Consider a representative task bundle $\mathcal{T}$. Before AI adoption, the firm produces baseline output Y_0 with baseline headcount N_0, each retained worker supplying the normal full-time benchmark $\bar{h}^F$, so baseline human labor input is $N_0\bar{h}^F$. The task bundle describes the composition of production activities, while output Y is the firm's scale or value choice. After the task-side shock changes the need for human task time, Figure 2 illustrates the joint choice space for firm headcount N_{gt}, average effective work intensity h^F_{gt}, and output scale Y_{gt} under AI-adjusted technology and labor costs. This follows the task-based view that AI changes how production tasks are performed, while evidence on AI diffusion shows that adoption, business-function deployment, and worker-task use need not translate mechanically into layoffs (Acemoglu and Restrepo, 2019*b*; Demirer et al., 2026; Bonney et al., 2026).

The illustration separates task composition from output scale and labor adjustment. The same representative task bundle can be paired with different output contours because the firm jointly chooses task allocation, employment positions, and work intensity. Workers then respond to the employment positions and performance cutoff created by the firm decision.

3.3 Task-Based AI Goods Production

The goods-production block builds on the task-based approach to automation and AI in Acemoglu and Restrepo (2018, 2019*a*,*b*). In that approach, technology affects labor demand by changing which tasks are performed by machines and by changing the productivity of human task inputs. This paper uses the task block to map AI service prices, task feasibility, and task productivity into human task demand. The resulting human task demand is then organized into employment positions and required productive work intensity, which shape household actual hours, fallback choices, housing demand, and spatial sorting.

Let A_t denote the aggregate AI environment and let r_t^A denote the rental price of AI services. A goods firm in group $g \in \mathcal{G}$ operates a fixed unit-measure continuum of tasks

$$x \in \mathcal{X}_g \equiv [0,1].$$

The aggregate AI environment and the AI service rental price are taken as given by the firm. The firm does not choose a separate AI adoption stock. Instead, it chooses AI use endogenously through task-level allocation between rented AI services and human task capacity.

The task side of technology is summarized by an AI-feasible frontier and two productivity schedules:

$$I_{gt} = I_g(A_t), \qquad 0 \le I_{gt} \le 1, \tag{2}$$

$$\gamma_{gt}^A(x) \equiv \gamma_g^A(x; A_t), \tag{3}$$

$$\gamma_{gt}^H(x) \equiv \gamma_g^H(x; A_t). \tag{4}$$

The frontier I_{gt} is a feasibility frontier: tasks $x \le I_{gt}$ can be performed by AI, while tasks $x > I_{gt}$ must be performed by humans. Feasibility is not the same as realized AI use. The schedule $\gamma_{gt}^A(x)$ is AI productivity on AI-feasible tasks, and $\gamma_{gt}^H(x)$ is effective human productivity. Realized task allocation is determined by relative task costs.

Task allocation. For any candidate output scale Y_{gt} and firm-organized human task capacity H_{gt}^{cap}, the task block asks how the firm allocates production tasks between human capacity and rented AI services. The key object is not the feasibility frontier I_{gt} itself, but the realized AI allocation chosen from the feasible tasks. Formally, define the conditional task-allocation cost

$$\mathcal{C}_g^T(Y_{gt}, H_{gt}^{cap}; A_t, r_t^A) = \min_{\{y_{gt}(x), \ell_{gt}(x), a_{gt}(x)\}_{x\in[0,1]}} \quad r_t^A \int_0^1 a_{gt}(x)dx, \tag{5}$$

subject to the task production technology

$$y_{gt}(x) = \gamma_g^H(x; A_t)\ell_{gt}(x) + \mathbf{1}\{x \le I_{gt}\}\gamma_g^A(x; A_t)a_{gt}(x), \tag{6}$$

the output aggregator

$$Y_{gt} = \left[\int_0^1 y_{gt}(x)^{\frac{\sigma_g - 1}{\sigma_g}} dx\right]^{\frac{\sigma_g}{\sigma_g - 1}}, \qquad \sigma_g > 0, \quad \sigma_g \neq 1, \tag{7}$$

AI feasibility, $a_{gt}(x) = 0$ for $x > I_{gt}$, and the human-capacity constraint

$$\int_0^1 \ell_{gt}(x)dx \le H_{gt}^{cap}. \tag{8}$$

The conditional problem treats H_{gt}^{cap} as given. Hence it does not charge the wage bill internally; wages, fixed worker costs, and overtime costs enter the outer firm problem below. Let λ_{gt}^H denote the multiplier on the human-capacity constraint. By the envelope theorem,

$$\lambda_{gt}^H \equiv -\frac{\partial \mathcal{C}_g^T(Y_{gt}, H_{gt}^{cap}; A_t, r_t^A)}{\partial H_{gt}^{cap}}. \tag{9}$$

Thus λ_{gt}^H is the firm's shadow value of human task capacity, not the market wage. The KKT conditions imply that a task can be produced by human capacity at unit cost $c_{gt}^H(x)$ and, if AI-feasible, by rented AI services at unit cost $c_{gt}^A(x)$:

$$c_{gt}^A(x) = \frac{r_t^A}{\gamma_g^A(x; A_t)}, \qquad c_{gt}^H(x) = \frac{\lambda_{gt}^H}{\gamma_g^H(x; A_t)}. \tag{10}$$

The realized AI task set is therefore the subset of feasible tasks for which AI is weakly cheaper than human task capacity:

$$\mathcal{A}_{gt}^* = \{x \le I_{gt} : c_{gt}^A(x) \le c_{gt}^H(x)\}, \qquad \mathcal{H}_{gt}^* = [0, 1] \setminus \mathcal{A}_{gt}^*. \tag{11}$$

This distinction is central: I_{gt} is only an upper bound on tasks that AI can perform, while realized AI use is endogenous and depends on relative task costs. When relative costs satisfy a monotone ordering over tasks, the same allocation can be summarized by a realized cutoff I_{gt}^* satisfying

$$c_{gt}^A(I_{gt}^*) = c_{gt}^H(I_{gt}^*), \qquad I_{gt}^* \le I_{gt}, \tag{12}$$

with the set notation in eq. (11) covering the more general non-monotone case. Given the realized task allocation, the unit task price is

$$p_{gt}(x) = \begin{cases} c_{gt}^{A}(x), & x \in \mathcal{A}_{gt}^{*}, \\ c_{gt}^{H}(x), & x \in \mathcal{H}_{gt}^{*}. \end{cases} \tag{13}$$

The corresponding unit-cost index is

$$C_{gt}^{Y} = \left[\int_{0}^{1} p_{gt}(x)^{1-\sigma_g} dx \right]^{\frac{1}{1-\sigma_g}}, \tag{14}$$

and task demand for output scale Y_{gt} is

$$y_{gt}^{*}(x) = Y_{gt} \left(\frac{p_{gt}(x)}{C_{gt}^{Y}} \right)^{-\sigma_g}. \tag{15}$$

The resulting human task demand and AI service demand are

$$H_{gt}^{D} = \int_{\mathcal{H}_{gt}^{*}} \frac{y_{gt}^{*}(x)}{\gamma_{g}^{H}(x; A_t)} dx, \tag{16}$$

$$A_{gt}^{D} = \int_{\mathcal{A}_{gt}^{*}} \frac{y_{gt}^{*}(x)}{\gamma_{g}^{A}(x; A_t)} dx. \tag{17}$$

At the optimum of the conditional task-allocation problem, $\mathcal{C}_{g}^{T}(Y_{gt}, H_{gt}^{cap}; A_t, r_t^A) = r_t^A A_{gt}^{D}$. The task block therefore maps $(Y_{gt}, H_{gt}^{cap}, A_t, r_t^A)$ into the shadow value of human task capacity, realized AI use, AI service demand, and human task demand. Here H_{gt}^{cap} is the capacity input organized by the firm, while H_{gt}^{D} is the resulting human task demand implied by task allocation. When the human-capacity constraint binds, $H_{gt}^{D} = H_{gt}^{cap}$.

3.4 Firms

Goods Firms Given A_t and r_t^A, a goods firm jointly chooses output scale, employment positions, and work intensity:

$$\mathcal{P}_{gt} = \{Y_{gt}, N_{gt}, h_{gt}^{F}\}.$$

The task block determines human task demand H_{gt}^{D}. The firm organizes this demand into headcount and work intensity by choosing human task capacity according to

$$H_{gt}^{cap} = N_{gt} h_{gt}^{F}. \tag{18}$$

The firm earns revenue from output through a concave demand or revenue schedule,

$$\mathcal{R}_g(Y_{gt}) = P_{gt}\left(Y_{gt}\right)^{1-\frac{1}{\eta_g}}, \qquad \eta_g > 1. \tag{19}$$

This reduced-form curvature captures downward-sloping demand for the job-group output and gives the firm an interior scale choice. Firm overtime effective hours are

$$O^F_{gt} = N_{gt} \max\{h^F_{gt} - \bar{h}^F, 0\}. \tag{20}$$

The parameter $\rho^O_g \geq 0$ is the overtime wage premium.[1] The firm also pays a fixed cost per actual worker, $F_g N_{gt}$.[2] Employment adjustment costs depend on actual headcount:

$$C^N_g(N_{gt}, N_{g,t-1}) = \frac{\phi^+_g}{2}\left[\max\{N_{gt} - N_{g,t-1}, 0\}\right]^2 + \frac{\phi^-_g}{2}\left[\max\{N_{g,t-1} - N_{gt}, 0\}\right]^2. \tag{21}$$

The parameter ϕ^+_g governs hiring costs and ϕ^-_g governs firing costs. The fixed worker cost, hiring and firing costs, and overtime premium make headcount and work intensity imperfect substitutes, consistent with labor-demand and search models in which firms adjust employment and hours jointly under shocks, adjustment costs, and hours constraints (Caballero, Engel and Haltiwanger, 1997; Cooper, Haltiwanger and Willis, 2015; Trapeznikova, 2017).

The goods firm solves

$$\begin{aligned} \max_{\mathcal{P}_{gt}} \quad & \mathcal{R}_g(Y_{gt}) - \mathcal{C}^T_g(Y_{gt}, N_{gt}h^F_{gt}; A_t, r^A_t) \\ & - W_{gCt}N_{gt}h^F_{gt} - \rho^O_g W_{gCt}O^F_{gt} - F_g N_{gt} - C^N_g(N_{gt}, N_{g,t-1}) \end{aligned} \tag{22}$$

subject to the conditional task-allocation problem in eq. (5), firm overtime accounting in eq. (20), and the organization of human capacity in eq. (18). This problem makes the firm labor organization margin explicit. The task block determines A^D_{gt}, H^D_{gt}, and λ^H_{gt} from output scale, organized human capacity, AI service prices, task feasibility, AI productivity, and human productivity. The firm then decides how much output to produce and how to organize human task capacity into headcount N_{gt} and work intensity h^F_{gt}.

For an interior solution away from hiring, firing, and overtime kinks, the following first-order conditions jointly characterize simultaneous choices. Output scale satisfies

$$\mathcal{R}'_g(Y_{gt}) = \frac{\partial \mathcal{C}^T_g(Y_{gt}, N_{gt}h^F_{gt}; A_t, r^A_t)}{\partial Y_{gt}}. \tag{23}$$

[1] For example, if overtime is paid at 1.5 times the base wage, then $\rho^O_g = 0.5$.

[2] The fixed worker cost captures headcount-related costs such as benefits, workplace capacity, and baseline hiring or administration.

The output condition equates marginal revenue to the marginal task-allocation cost of expanding output, evaluated at the optimal allocation of human task capacity and rented AI services. AI use is disciplined by the task-level cost condition in eq. (11): feasible tasks are performed by AI only when rented AI services are weakly cheaper than human task capacity for that task.

For headcount and work-intensity choices, λ_{gt}^H connects the task-allocation shadow value to actual labor organization costs:

$$\lambda_{gt}^H h_{gt}^F = W_{gCt} h_{gt}^F + \rho_g^O W_{gCt} \max\{h_{gt}^F - \bar{h}^F, 0\} + F_g + \frac{\partial C_g^N}{\partial N_{gt}}, \tag{24}$$

$$\lambda_{gt}^H = W_{gCt} + \rho_g^O W_{gCt} \mathbf{1}\{h_{gt}^F > \bar{h}^F\}. \tag{25}$$

These conditions describe the firm's interior optimality system: output scale is disciplined by task input costs; task allocation is disciplined by relative AI and human unit costs; fixed worker costs and hiring/firing frictions discipline N_{gt}; and wage and overtime costs discipline h_{gt}^F.[3] The AI margin operates through task feasibility, AI task productivity, human task productivity, and the rental price of AI services, which jointly change H_{gt}^D and the organization of human work.

3.5 Layoff Competition and Performance

When AI changes task feasibility, AI task costs, human-task productivity, and output incentives, the firm chooses employment positions N_{gt}. If the firm reduces positions,

$$N_{gt} < N_{g,t-1}.$$

incumbent workers compete for fewer remaining positions. Here n_i denotes planned actual hours as defined above. Let individual performance in job group g be

$$q_{i,g,t} = z_i \chi_{ig,t} m_g(n_i) + \iota_i, \tag{26}$$

where $m_g'(n) > 0$ and $m_g''(n) \leq 0$. Individual productivity z_i, job fit $\chi_{ig,t}$, and planned actual hours n_i raise performance, while ι_i is an arbitrarily small deterministic tie-breaker. The firm retains the N_{gt} workers with the highest performance among the relevant incumbent and applicant pool, following the rank-order competition logic of tournament models (Lazear and Rosen, 1981). Let q_{gt}^* denote the endogenous cutoff for remaining in job group g. The

[3] When overtime costs, hiring/firing adjustment costs, nonnegativity constraints, or task-allocation ties bind, the problem is characterized by the corresponding KKT conditions and solved numerically.

retention rule (i remains in g) is $q_{i,g,t} \geq q^*_{gt}$. And the cutoff satisfies:

$$\int \mathbf{1}\{q_{i,g,t} \geq q^*_{gt}\} d\mu_{gt}(i) = N_{gt}, \tag{27}$$

where μ_{gt} is the pool of workers competing for job group g. The cutoff is not an arbitrary retention probability. It is implied by the firm-chosen number of positions and the distribution of worker performance. When households choose hours, the computational equilibrium uses the current cutoff belief, denoted q^b_{gt}, because the ex post ranking has not yet been realized. The perceived probability of clearing the cutoff is represented by a smooth increasing function of the distance between expected performance and the cutoff belief:

$$p^{ret}_{igt}(n_i) = \mathcal{S}\left(\frac{z_i \chi_{ig,t} m_g(n_i) - q^b_{gt}}{\sigma^q_g}\right). \tag{28}$$

Here $\mathcal{S}(\cdot)$ is a smooth approximation used in the household problem, and σ^q_g governs how sharply expected performance maps into retention risk. When AI reduces positions, the realized cutoff q^*_{gt} and the corresponding belief q^b_{gt} rise. Workers close to the cutoff may raise n_i to increase $p^{ret}_{igt}(n_i)$. This probability is not chosen by the firm; it is the worker's perceived probability of clearing a cutoff generated by firm positions and performance ranking. The simulation then applies the realized deterministic ranking cutoff in eq. (27). This is the model's pre-displacement defensive work-time channel, complementing recent evidence and theory on AI and workday adjustment (Jiang et al., 2025).

Local Service Firms Local service firms provide a non-AI reallocation margin inside the city. The sector is not assumed to absorb displaced workers mechanically. Its capacity to hire workers is determined by local service demand, service wages, and workers' outside options. Following the spatial-equilibrium logic in which routine-task displacement can reallocate labor toward local service work (Autor and Dorn, 2013), the service sector is modeled as low-AI-exposed and nontradable.

To avoid introducing a separate service-consumption clearing condition, the baseline uses a reduced-form service labor demand schedule:

$$L^{D,S}_t = A_S Z^S_t (W^S_{Ct})^{-\eta_S}, \qquad \eta_S > 0. \tag{29}$$

Here $L^{D,S}_t$ is local service demand for paid effective human hours, W^S_{Ct} is the service wage,

and η_S is the service labor-demand elasticity. The local demand shifter is

$$Z_t^S = (N_{Ct})^{\omega_N}(Y_{Ct}^{inc})^{\omega_Y}, \qquad \omega_N, \omega_Y \geq 0, \tag{30}$$

where N_{Ct} is city population and Y_{Ct}^{inc} is aggregate city disposable income. Local service demand expands when the city retains population or income and contracts when AI-induced labor reallocation reduces them. Thus the service sector absorbs only the workers who choose service employment and can be supported by local service demand in equilibrium.

Developers Developers produce private housing in city zone k using construction human time and non-labor development inputs:

$$Q_{kt}^H = Z_k^D (L_{Dkt}^D)^{\alpha_D} (N_{kt}^D)^{1-\alpha_D}. \tag{31}$$

Here Q_{kt}^H is new private housing construction, Z_k^D is developer productivity, L_{Dkt}^D is developer demand for paid construction human time, N_{kt}^D is a local non-labor development input, and α_D is the labor share in housing production. The local supply of non-labor development input is fixed at $\bar{N}_k^D$.

Developer profits are

$$\Pi_{kt}^D = P_{kt}^H Q_{kt}^H - W_{Ct}^D L_{Dkt}^D - r_{kt}^N N_{kt}^D, \tag{32}$$

where P_{kt}^H is the housing asset price and r_{kt}^N is the clearing price of non-labor development inputs. The labor first-order condition is

$$W_{Ct}^D = P_{kt}^H \frac{\partial Q_{kt}^H}{\partial L_{Dkt}^D}. \tag{33}$$

Developers have no direct AI task shock. They are affected indirectly through construction wages, housing demand, rents, and population sorting.

3.6 Households

States and Choices Households choose non-housing consumption C_{it}, housing, next-period bond holdings $b_{i,t+1}$, job/status $j \in \mathcal{J}$, actual hours n_{it} when working, region $K \in \{C, O\}$, zone $k \in \mathcal{K}(K)$, and private tenure $r \in \{R, O\}$, where R denotes private renting and O denotes owner occupancy. All employed city jobs require at least the normal full-time benchmark, $n_{it} \geq \bar{h}^F$, while city unemployment/search has $n_{it} = 0$. For goods applicants,

actual hours are planned before the retention outcome is realized; realized hours are the planned hours if retained and the fallback hours if not retained.

The model distinguishes actual hours from paid hours. For a retained goods worker, the firm pays for the required work intensity h^F_{gt}, while any additional hours $n_{it} - h^F_{gt}$ are unpaid defensive overwork. For local service and developer work, all actual hours are paid:

$$n^p_{ij,t} = \begin{cases} h^F_{gt}, & j = g \in \mathcal{G}, \text{ retained in } g, \\ n_{it}, & j \in \{S, D\}. \end{cases} \tag{34}$$

For retained goods workers, $\bar{h}^F \leq h^F_{gt} \leq n_{it}$. The unpaid defensive component is

$$n^{def}_{ig,t} = \max\{n_{it} - h^F_{gt}, 0\}. \tag{35}$$

It affects retention performance, leisure, and work-intensity disutility, but not wage income, firm labor cost, or paid effective-hours clearing. Paid hours generate paid effective hours

$$e^p_{ij,t} = z_i \chi_{ij,t} n^p_{ij,t}, \tag{36}$$

which are the units priced by the job-specific wage. To avoid confusion with work hours, housing services are denoted by s. Renters choose rental housing services s^r_{it}. Owners choose owner-occupied housing s^o_{it} and investment housing $\hat{s}^o_{it}$ supplied to the rental market. The owner-tenure notation is distinct from the outside option because it appears under the tenure index r rather than the region index K.

Preferences, Time, and Budget For a worker applying to a risky goods job $g \in \mathcal{G}$, the value before the retention outcome is

$$V^G_{igt}(s_{it}) = \max_{\substack{C, s^r, s^o, \hat{s}^o, b_{i,t+1}, n, K, k, r \\ n \geq h^F_{gt}}} u(C, s, \ell; k) - \Psi_g(n) - \kappa^W_{j_{i,t-1},g} + p^{ret}_{igt}(n) V^E_{igt} + \left[1 - p^{ret}_{igt}(n)\right] V^F_{it}, \tag{37}$$

where V^E_{igt} is the continuation value if retained in job group g and V^F_{it} is the fallback value if not retained. The switching cost $\kappa^W_{j_{i,t-1},g}$ is normalized to zero when the household remains in the same job/status. This specification makes defensive hours endogenous: increasing n above the firm-paid requirement h^F_{gt} raises the probability of retention, but the extra defensive hours are unpaid. For an interior defensive-hours choice with $n_{it} > h^F_{gt}$, the condition balances

retention gains against lost leisure and direct work-intensity disutility:

$$0 = -u_\ell(C, s, \ell; k) - \Psi'_g(n) + \frac{\partial p^{ret}_{igt}(n)}{\partial n}\left(V^E_{igt} - V^F_{it}\right), \tag{38}$$

where $\partial\ell/\partial n = -1$. Equivalently, the first term is $-u_\ell(C, s, \ell; k)$. The final term is the marginal value of defensive work time through job retention. For local service and developer jobs, all actual hours are paid. Away from the overtime kink, their paid-hours condition is

$$0 = u_C(C, s, \ell; k)W_{jCt}z_i\chi_{ij,t}\left[1 + \rho^O_j \mathbf{1}\{n > \bar{h}^F\}\right] - u_\ell(C, s, \ell; k) - \Psi'_j(n), \qquad j \in \{S, D\}, \tag{39}$$

up to any taxes or wedges included in the household budget constraint.[4]

The fallback value after non-retention is

$$V^F_{it} = \max\left\{\max_{g' \neq g} V^G_{ig't}, V^S_{it}, V^D_{it}, V^U_{it}, V^O_{it}\right\}, \tag{40}$$

net of the relevant switching and moving costs. Thus a displaced worker can switch to another goods job, move into local service work, move into developer work, remain in the city while unemployed or searching, or leave for the outside option. Entering local service employment is a household choice, and the number of such jobs is disciplined by service labor-market clearing.

If the fallback option is another goods job $g' \neq g$, the worker is not mechanically absorbed into that job group. The worker enters the residual slot competition in g' and must clear that group's remaining capacity and cutoff. If the worker does not clear the secondary goods cutoff, the household moves to its next non-goods fallback option, including local service, developer work, city search, or the outside option.

The values of local service work, developer work, city unemployment/search, and the outside option use the same consumption, housing, and location structure but do not involve the goods-job retention cutoff. In compact form,

$$V_t(s_{it}) = \max\left\{\max_{g \in \mathcal{G}} V^G_{igt}, V^S_{it}, V^D_{it}, V^U_{it}, V^O_{it}\right\}. \tag{41}$$

For all city choices, $s = s^r$ for renters and $s = s^o$ for owners. The parameter β is the discount factor, $\Psi_j(n)$ is work-intensity disutility, and $b_{i,t+1} \in \mathcal{B}$ is the next-period bond/assets choice. Actual hours satisfy $n \in [\bar{h}^F, 1 - \tau^{time}_k]$ for employed city workers and

[4] At the overtime threshold or when the hours choice is on a grid boundary, the corresponding discrete/KKT condition replaces the differentiable interior FOC.

$n = 0$ in city unemployment/search. Leisure is

$$\ell = 1 - n - \tau_k^{time}. \tag{42}$$

Commuting time τ_k^{time} is location-specific and is distinct from direct work-intensity disutility. Normal work time reduces leisure through the time constraint. Direct work-intensity disutility applies to actual hours above the normal full-time benchmark:

$$\Psi_j(n) = \frac{\varphi_j}{1+\nu_n} \max\{n - \bar{h}^F, 0\}^{1+\nu_n}, \qquad \bar{h}^F > 0, \quad \varphi_j > 0, \quad \nu_n > 0. \tag{43}$$

For a city worker in job $j \in \mathcal{J}_C$, individual labor income includes the base effective-hour wage paid on $e^p_{ij,t}$ and the overtime premium paid on paid hours above the normal full-time benchmark:

$$y^{lab}_{ijCt} = W_{jCt} e^p_{ij,t} + \rho^O_j W_{jCt} z_i \chi_{ij,t} \max\{n^p_{ij,t} - \bar{h}^F, 0\}, \qquad W_{jCt} = \begin{cases} W_{gCt}, & j = g \in \mathcal{G}, \\ W^S_{Ct}, & j = S, \\ W^D_{Ct}, & j = D. \end{cases} \tag{44}$$

For goods jobs, $\rho^O_j = \rho^O_g$ when $j = g \in \mathcal{G}$; for service and developer work, ρ^O_S and ρ^O_D are the corresponding overtime premia. For retained goods workers, $n_{it} - h^F_{gt}$ is unpaid defensive overwork and is excluded from eq. (44). Local service work does not involve the AI-induced retention cutoff. For households in city unemployment/search, set $n = 0$ and $y^{lab}_{ijCt} = 0$. These households receive a city unemployment/search transfer or search income and do not receive wage income. The model parameterizes this support as a replacement-rate object tied to service-sector normal-hours earnings:

$$y^U_{Ct} = \rho^U W^S_{Ct} \bar{h}^F, \qquad \rho^U \in [0, 1). \tag{45}$$

The outside option uses reduced-form outside income and values. Private renters pay market rent and do not own housing capital:

$$C + R_{kt} s^r + Q_t b_{i,t+1} = a_{i,t} + (1 - \tau^w_t) y^{lab}_{ijCt} + y^{other}_{it} + \mathbf{1}\{j = U_C\} y^U_{Ct} - \tau^{comm}_k - \tau^{move}_{K,k}. \tag{46}$$

Private owners purchase owner-occupied housing and investment housing. Investment housing is supplied to the rental market and earns rental income:

$$C + P^H_{kt} s^o + P^H_{kt} \hat{s}^o + Q_t b_{i,t+1} = a_{i,t} + (1 - \tau^w_t) y^{lab}_{ijCt} + y^{other}_{it} + \mathbf{1}\{j = U_C\} y^U_{Ct} + R_{kt} \hat{s}^o - \tau^{comm}_k - \tau^{move}_{K,k}. \tag{47}$$

Here R_{kt} is rent, P^H_{kt} is the private housing asset price, $Q_t = 1/(1 + r^B_t)$ is an exogenous risk-free bond price, τ^w_t is the labor income tax rate, y^{other}_{it} includes transfers, pensions, profit rebates, and land dividends but excludes owner rental income, y^U_{Ct} is the city unemployment/search transfer or search-income object, τ^{comm}_k is financial commuting cost, and $\tau^{move}_{K,k}$ is moving cost. The current-period choice $b_{i,t+1}$ enters next period's asset state. End-of-period assets are

$$a_{i,t+1} = \begin{cases} b_{i,t+1}, & r = R, \\ b_{i,t+1} + (1 - \delta^o_k - \tau^H_k) P^H_{kt} s^o + (1 - \delta^I_k - \tau^H_k) P^H_{kt} \hat{s}^o, & r = O, \end{cases} \tag{48}$$

where δ^o_k and δ^I_k are depreciation rates for owner-occupied and investment housing, and τ^H_k is the property tax rate. Owner borrowing is subject to the collateral constraint

$$-Q_t b_{i,t+1} \leq \theta P^H_{kt} \left(s^o + \hat{s}^o\right), \tag{49}$$

where θ is the maximum loan-to-value ratio.

State Transitions The continuation value is evaluated over next-period states after current choices, layoff outcomes, and shocks. If the household applies to goods job g and is retained, the realized current-period job/status carried into period $t + 1$ is g. If the household is not retained, the realized current-period job/status carried into period $t + 1$ is the fallback status chosen in eq. (40). In compact notation,

$$j_{it} = \mathcal{T}_j(j, \mathbf{1}\{q_{i,j,t} \geq q^*_{jt}\}, j^F_i), \tag{50}$$

where j^F_i is the fallback status selected if not retained. Productivity and job fit evolve as predetermined states:

$$z_{i,t+1} \sim \Pi_z(z_{it}, age_i), \qquad \chi_{ij,t+1} \sim \Pi_\chi(\chi_{ij,t}, age_i). \tag{51}$$

The kernel Π_χ applies componentwise to the fit vector $\{\chi_{ij}\}_{j \in \mathcal{J}_C}$.

3.7 Equilibrium Overview

The equilibrium consists of job-specific wages, goods-job performance cutoffs, rents, housing prices, development-input prices, firm policies, household policies, and a household distribution. Goods-job cutoffs clear the headcount margin, while job-specific wages clear effective work-intensity markets. Service, developer, city search, outside, and private housing mar-

kets clear jointly with migration and residential sorting. Formal equilibrium conditions are collected in Appendix A.1.

AI affects households through goods firms' task allocation, employment positions, firm work intensity, performance cutoffs, fallback states, job-group wages, rents, and commuting costs. Spatial sorting is therefore an equilibrium outcome rather than an imposed AI premium for any city zone.

4 Calibration

This section describes how the baseline economy is disciplined quantitatively. Parameters with direct empirical counterparts are assigned using external evidence, policy rules, observed institutional features, or standard normalizations. Parameters that govern the AI, retention, and defensive-hours mechanisms are calibrated to AI-use, task-exposure, hours, displacement, and job-transition moments. Table 1 reports the compact set of structural and counterfactual-relevant parameters used in the main text. Numerical grids, solver tolerances, damping rules, runtime flags, and diagnostic-output controls are implementation objects and are not included in the main calibration table.

4.1 Parametrization

Spatial and Housing Environment The city is calibrated to a two-zone Beijing spatial structure with an outside option. Zone capacities discipline the relative tightness of residential land supply, while zone productivity captures the spatial wage and productivity gradient. Rents and housing prices are set to match observed residential moments. The commuting-time calibration uses a Beijing travel-time matrix from residential subdistrict centroids to employment-weighted workplace centroids constructed from firm-location employment density. The model therefore keeps the main rent-commuting tradeoff at the core-periphery level.

Households and Time Use Household preferences use standard values for risk aversion, discounting, housing expenditure, and time-use weights. Hours are normalized by the period time endowment. The normal full-time benchmark $\bar{h}^F$ is the threshold used for overtime premia and direct work-intensity disutility; for employed city workers, it is also the minimum normal-hours requirement. The work-intensity disutility parameters are disciplined by moments on hours, overwork incidence, and defensive excess hours.

Labor Retention and Fallback The retention block is calibrated around two objects: the performance effect of actual hours and the smooth approximation to cutoff risk used in the household problem. The baseline sets a positive performance-hours slope so that households near the cutoff can choose defensive hours. Moments on layoffs, retention, job-to-job transitions, and hours responses around displacement risk discipline this block. Switching costs and the city search replacement rate discipline fallback choices across other goods jobs, service work, developer work, city search, and the outside option.

Goods Firms and AI Task Technology Goods firms are disciplined by demand curvature, task substitution, fixed worker costs, hiring and firing costs, overtime premia, and the group-specific task technology. The AI-feasible task frontier, AI task productivity, human task productivity, AI diffusion path, and AI service rental price path are calibrated by group-specific baseline values. The benchmark uses $A_t = 0.20 + 0.002t$ and $r_t^A = \max\{0.05, 0.35/(1 + 0.60A_t)\}$. These parameters are targeted to evidence on AI use, task exposure, output, employment positions, firm required hours, and human task demand.

Non-Goods Sectors and Fiscal Objects The local service block is disciplined by service employment, wage, population, and income moments. Developer production is tied to housing supply elasticity and local housing demand. Fiscal and support parameters include payroll taxes and the city search replacement rate. These objects matter for counterfactuals because they determine how displaced workers are absorbed by fallback sectors and how much income support buffers AI-induced labor reallocation.

Benchmark Moments The benchmark economy is disciplined to reproduce the joint environment in which the defensive-hours mechanism operates. The targeted moments cover employment positions by job group, service and developer employment shares, city search and outside shares, average hours, overwork incidence, wage dispersion, rent burdens, residential sorting, migration, and fallback transitions. These moments anchor the labor, time-use, housing, and spatial blocks without requiring the main text to repeat the calibration targets in a separate validation table.

5 Baseline Results

This section reports the baseline transmission mechanism before moving to counterfactual policy analysis. The baseline starts from the calibrated AI, labor-market, housing, and

spatial environment and follows three connected margins. First, goods firms organize AI-adjusted human task demand into output scale, employment positions, and required productive work intensity. Second, workers facing firm-generated retention cutoffs choose actual hours, including defensive excess hours, when additional time raises the probability of keeping a job. Third, housing costs, commuting time, fallback/search values, and migration options shape the retained-fallback value gap and therefore amplify or dampen work-time pressure.

5.1 Firm-Side Labor Organization

The firm-side baseline reports how AI-adjusted task requirements are organized into production and labor inputs. Across goods job groups, the same AI environment can be associated with different combinations of output scale, employment positions, and required productive work intensity. A reduction in per-unit human task requirements can lower labor demand, expand output, or shift the firm toward a different mix of headcount and required hours, with the response governed by fixed worker costs, hiring and firing costs, overtime premia, and revenue curvature.

Figure 3 summarizes this margin across goods job groups sorted by baseline AI exposure. The panels compare AI exposure, employment positions, required productive work intensity, and human task demand. The figure shows whether baseline AI exposure is absorbed primarily through output expansion, fewer positions, higher required work intensity, or a combination of these margins.

5.2 Retention Risk and Defensive Hours

Firm labor organization becomes a household hours problem through retention competition. A reduction in positions raises the performance cutoff q^*_{gt}. A higher firm requirement h^F_{gt} raises the productive time a retained worker must supply. Households then choose actual hours n_{it} before the retention outcome is realized. The central mechanism is the marginal retention value of additional hours from the household first-order condition, which is strongest when additional time substantially improves retention prospects and when fallback values are weak relative to the retained job value.

Figure 4 summarizes this channel across worker/job cells sorted by baseline retention risk. The panels report cutoff pressure, the retained-fallback value gap, actual hours, and defensive excess hours. The baseline identifies where defensive hours are concentrated and whether the concentration lines up with high cutoff pressure and weak fallback values.

5.3 Spatial Amplification

The spatial block determines how costly firm-side labor adjustment is for households. Housing costs, commuting time, city search income, local service absorption, and migration costs all enter continuation values. They therefore affect the value gap $V_{igt}^E - V_{it}^F$ that governs defensive hours. A worker facing high rent burdens or costly migration may supply more defensive time to preserve a retained goods job, while a worker with strong fallback options or lower commuting burden may absorb the same firm-side shock with less work-time pressure.

Figure 5 summarizes this spatial incidence. The panels connect work-time outcomes to rent burden, work-related time, residential sorting, and welfare across zone categories, worker-location cells, location choices, and worker groups sorted by the relevant baseline outcome. This baseline evidence motivates the counterfactuals: policies can operate by changing the headcount margin, the intensity margin, or the spatial environment that amplifies the retained-fallback value gap.

The baseline therefore delivers the objects needed for policy analysis. Firm labor organization determines the productive demand for positions and required work intensity. Retention competition maps those firm decisions into household defensive hours. Spatial frictions determine how costly that adjustment is and which households bear it. The counterfactuals below ask which interventions relax these margins without simply shifting adjustment into unemployment, city exit, unpaid defensive time, or housing-cost pressure.

6 Counterfactuals

The baseline results establish that AI-induced labor adjustment is not summarized by employment or output alone. Firms translate the task-side shock into positions N_{gt}, required productive work intensity h_{gt}^F, and AI task use. Workers then respond to the retained-fallback value gap by choosing actual hours n_{it}, including defensive excess hours above the firm requirement. The counterfactuals therefore focus on three experiments. The first changes the technology frontier itself. The second asks whether time otherwise spent in defensive competition can be redirected into individual skill investment. The third studies employer-provided housing as a spatial insurance policy that lowers the cost of staying attached to the city labor market.

6.1 AI Technology Counterfactual

The first experiment studies a stronger AI technology. This experiment answers how the economy changes when AI becomes more capable, cheaper to rent, and more productive

on AI-feasible tasks. On the firm side, the experiment reports changes in AI task use, human task demand, output scale, employment positions N_{gt}, and required productive hours h^F_{gt}. On the household side, it reports whether the stronger technology increases retention competition, raises actual hours n_{it} among retained workers, expands defensive excess hours, or shifts displaced workers into other goods jobs, local services, developer work, city search, and the outside option.

This counterfactual is the technological benchmark for the paper. It separates aggregate productivity gains from their incidence across positions, hours, and locations. A stronger AI technology can raise output while reducing the number of retained positions, or it can expand output enough to preserve employment while still increasing cutoff pressure within exposed job groups. The key empirical object is therefore not AI adoption alone, but the joint movement of A^D_{gt}, H^D_{gt}, N_{gt}, h^F_{gt}, and household defensive excess time.

6.2 Skill-Investment Counterfactual

The second experiment asks whether the time absorbed by defensive competition can instead become individual human-capital investment. In the baseline, defensive excess hours above h^F_{gt} are costly to workers, enter retention performance, and do not enter production or wages. The counterfactual redirects a fraction $\lambda_E \in [0, 1]$ of retained goods workers' excess hours above the firm requirement into education or skill accumulation. The remaining excess component is still defensive work time. Equivalently, for a retained goods worker with baseline excess time, the policy reduces the purely defensive component by the share λ_E while preserving the firm-required paid assignment h^F_{gt}. Skill investment raises next-period productivity relative to the baseline transition $z^0_{i,t+1}$ with a concave return governed by $\psi_E > 0$ and $\chi_E \in (0, 1]$.

This experiment is not a simple reduction in hours. It changes the use of time at the margin that the baseline identifies as unpaid and nonproductive for the firm. The economic question is whether an economy facing AI-induced time pressure can turn part of that pressure into long-run resilience. The relevant channels are higher future productivity, stronger fallback values, improved job mobility, and lower welfare losses for AI-exposed workers. The results report changes in skill growth, future wages or productivity, fallback values, actual hours, defensive excess hours, retention probabilities, displacement losses, and welfare.

6.3 Employer-Provided Housing Counterfactual

The third experiment introduces employer-provided housing support for workers exposed to AI labor adjustment. Housing support is modeled as an employment-linked rent discount for

eligible workers in AI-exposed goods firms. Assigned renter households pay $\widetilde{R}_{kt} = (1-\eta_k^E)R_{kt}$, with $\eta_k^E \in [0, 1)$, and assignment and resource accounting are described in eqs. (75) and (76). The policy is targeted to workers whose labor-market attachment is directly affected by the AI shock, rather than to the full city population.

The mechanism is spatial but not only a housing mechanism. By lowering the cost of remaining in the city, employer housing raises the value of continued labor-market attachment and can reduce the fear that job loss forces a costly residential move or exit to the outside option. If housing pressure is an important component of the retained-fallback value gap, then housing support can reduce defensive excess hours, lower city exit, and improve welfare without directly changing AI technology. The experiment reports rent burdens, location choices, commuting time, defensive excess hours, retention competition, city search, outside exit, and welfare for exposed and non-exposed workers.

6.4 Summary of Counterfactual Results

Together, the counterfactuals ask whether the costs of AI-induced labor adjustment are best understood as a technology problem, a time-allocation and skill-accumulation problem, or a spatial insurance problem. The first experiment identifies how stronger AI technology changes the firm's demand for human task time and the household retention environment. The second experiment asks whether defensive time can be converted into future productivity and stronger fallback values. The third experiment evaluates whether lowering housing pressure can reduce the welfare cost of labor-market attachment under AI exposure.

The counterfactuals are evaluated using a common set of outcomes: output, AI task use, employment positions, required productive hours, actual hours, defensive excess hours, skill growth, fallback values, rent burdens, city exit, and welfare. This common set of outcomes separates productivity gains from their incidence across employment, time use, housing pressure, and spatial sorting.

7 Conclusion

Artificial intelligence is changing not only which tasks are performed by machines, but also how human work is organized across jobs, hours, and space. A focus on employment counts alone misses an important part of this adjustment. Firms can absorb AI-induced changes in human task demand through output scale, headcount, and work intensity, while workers respond through actual hours, fallback choices, migration, and residential location. This paper addresses this gap by developing a spatial general equilibrium model in which task-

based AI production is linked to firm labor organization and household spatial decisions.

The model highlights three connected mechanisms. First, the AI task block maps AI feasibility, AI service prices, AI task productivity, and human productivity into human task demand. Second, goods firms organize this human task demand into employment positions and required productive work intensity under fixed worker costs, overtime premia, and hiring and firing frictions. Third, households facing retention cutoffs choose hours and fallback options while also choosing where to live and whether to remain attached to the city labor market. This structure turns AI exposure into a joint extensive-margin and intensive-margin labor adjustment problem rather than a pure displacement shock.

The framework also identifies key trade-offs in AI-driven productivity gains. Stronger AI can expand output and reduce the human time required for some tasks, but it can also shift pressure onto retained workers when employment positions, performance cutoffs, and fallback values move together. Defensive hours are most costly when job loss is tied to weak fallback options, high rent burdens, long commutes, or costly migration. Spatial frictions therefore shape the incidence of AI adjustment: two workers exposed to the same firm-side technology shock can face different hours, welfare, and location outcomes because their housing and fallback environments differ.

The counterfactual exercises separate three policy margins. The AI technology counterfactual distinguishes aggregate productivity gains from their incidence across employment positions, required hours, and household time pressure. The skill-investment counterfactual asks whether excess defensive time can be redirected toward human capital formation, strengthening future productivity and fallback values rather than only raising current retention effort. The employer-provided housing counterfactual studies whether reducing housing pressure for exposed workers can lower the welfare cost of remaining attached to the city labor market. Together, these exercises show how technology, time allocation, and spatial insurance operate through different margins of the same equilibrium adjustment.

Overall, the paper argues that AI policy should not be evaluated only through adoption, productivity, wages, or employment headcounts. The relevant adjustment also includes hours, retention risk, search and fallback insurance, commuting costs, housing affordability, and the spatial attachment of workers to productive cities. Policies that ignore these margins may preserve aggregate efficiency while shifting costs into unpaid time pressure, displacement risk, or residential sorting. A more complete approach to AI and labor-market policy should therefore treat human task demand, headcount, work intensity, and spatial frictions as jointly determined objects.

Table 1: Calibration Parameters of the Baseline Model

Description	Parameter	Value	Target/Source
Panel A: Geography, Housing, and Spatial Environment			
Spatial supply and prices	$\overline{H}_1, \overline{H}_2, \overline{H}_O; R_{kt}, P^H_{kt}$	0.3; 1.5; 7	Beijing zoning, residential rents, and house prices
Zone productivity	Z_1, Z_2	1.09487; 1.00	Spatial wage/productivity gradient
Commuting costs	$\tau^{time}_k, \tau^{comm}_k$	matrix; normalized	Beijing travel-time matrix and travel-cost data
Panel B: Households: Preferences and Lifecycle			
Risk aversion and discounting	γ, β	5; 0.995	Standard macro calibration and risk-free rate
Utility weights	α_H, α_N	0.41; 0.63	Housing expenditure and time-use moments
Elasticities and time bound	$\epsilon_{CH}, \epsilon_{CN}, n_{\min}, Q$	-0.5; 0.0 0.15; 0.9833	Housing-demand evidence, normalization, and Treasury rate
Panel C: Labor, Retention, and Defensive Hours			
Hours benchmark and overtime	$\bar{h}^F, \rho^O_j$	0.40; 0.50	Normal full-time benchmark and overtime premium
Work-intensity disutility	φ_j, ν_n	0.20; 2.00	Hours and overwork moments
Retention technology	$\theta^m_g, \kappa^m_g, \sigma^q_g$	0.50; 4.00; 0.25	retention, layoff, hours response, and cutoff-risk sensitivity
Job/status switching cost	κ^W	0.04	Job mobility and fallback transitions
Panel D: Goods Firms, Task Technology, and AI			
Goods demand and task substitution	η_g, σ	3.0; 2.0	Demand elasticity and task-substitution evidence
Task frontier and productivity	I_{0g}, θ^I_g $\gamma^A_{0g}, \theta^A_g$ $\gamma^H_{0g}, \theta^H_g$	group-specific	Task exposure, AI productivity, and human complementarity
AI paths	A_t, r^A_t	$0.20 + 0.002t$ declining in A_t	AI diffusion and service cost paths
Worker and adjustment costs	F_g, ϕ^+_g, ϕ^-_g	0.04; 0.02; 0.04	Employment adjustment and separation moments
Panel E: Non-Goods Sectors, Search, and Fiscal Objects			
Service demand	$A_S, \eta_S, \omega_N, \omega_Y$	0.10; 1.25 0.60; 0.40	Service employment, wage, population, and income shares
Production returns to scale	$RTSH, RTSC$	0.66; 0.66	Housing supply elasticity, labor share, and production evidence
Search support and payroll tax	ρ^U, τ^{ss}	0.35; 0.10	City search/unemployment support and tax data
City productivity	$T_0, \alpha, \gamma_c, \delta$	0.2; 0.3 1.2; 2.6×10^{-5}	Urban productivity and agglomeration-congestion evidence

Note: This table summarizes the calibration parameters used in the baseline model. Values are sourced from various references as indicated in the "Target" column.

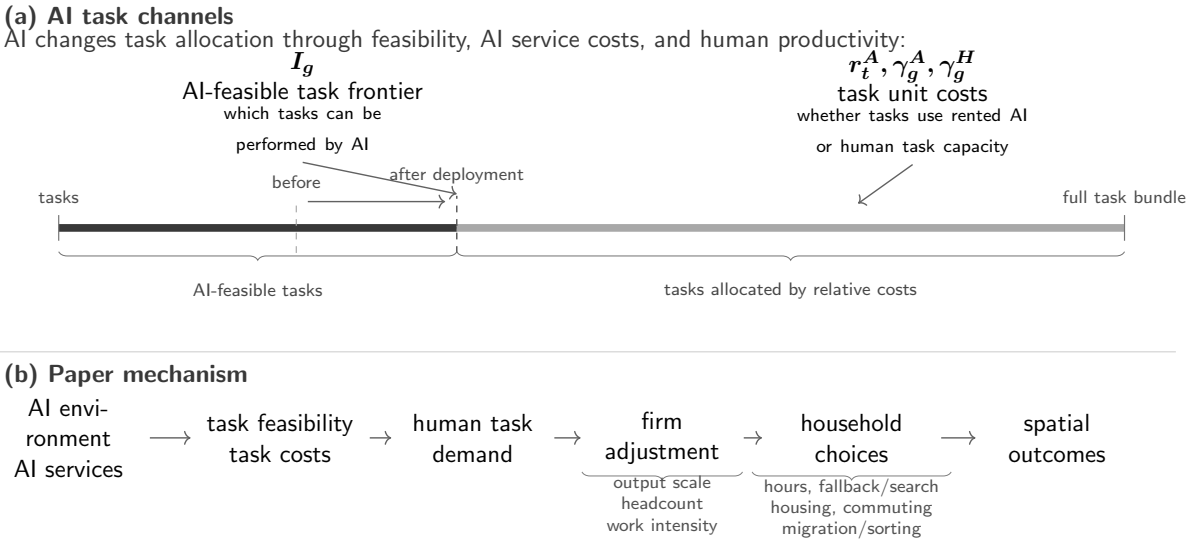


Figure 1: AI Task Channels and the Paper Mechanism

Note: Panel (a) gives an intuitive task-space representation of AI feasibility and task-level cost allocation between rented AI services and human task capacity. Panel (b) summarizes the paper's transmission mechanism from task allocation to firm labor organization and household spatial choices.

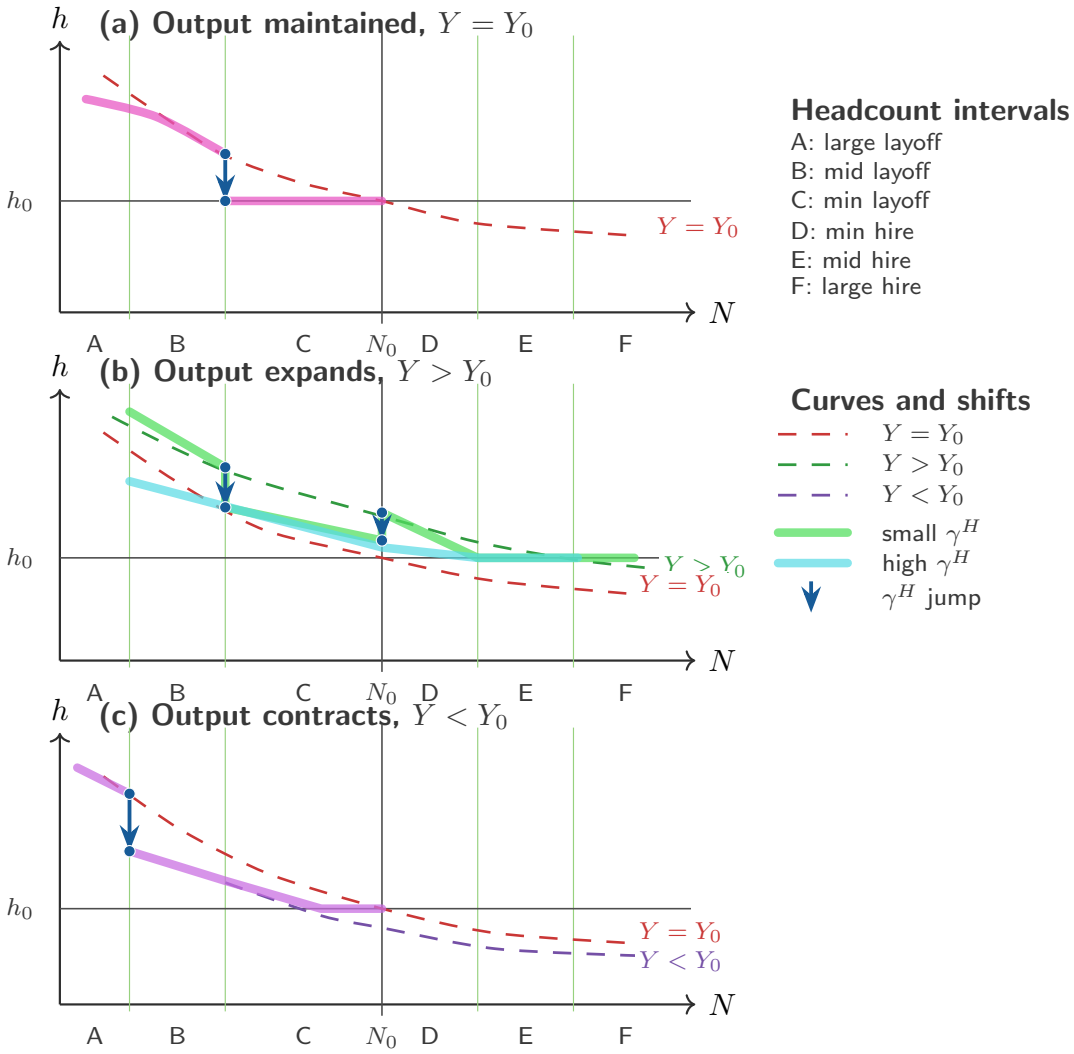
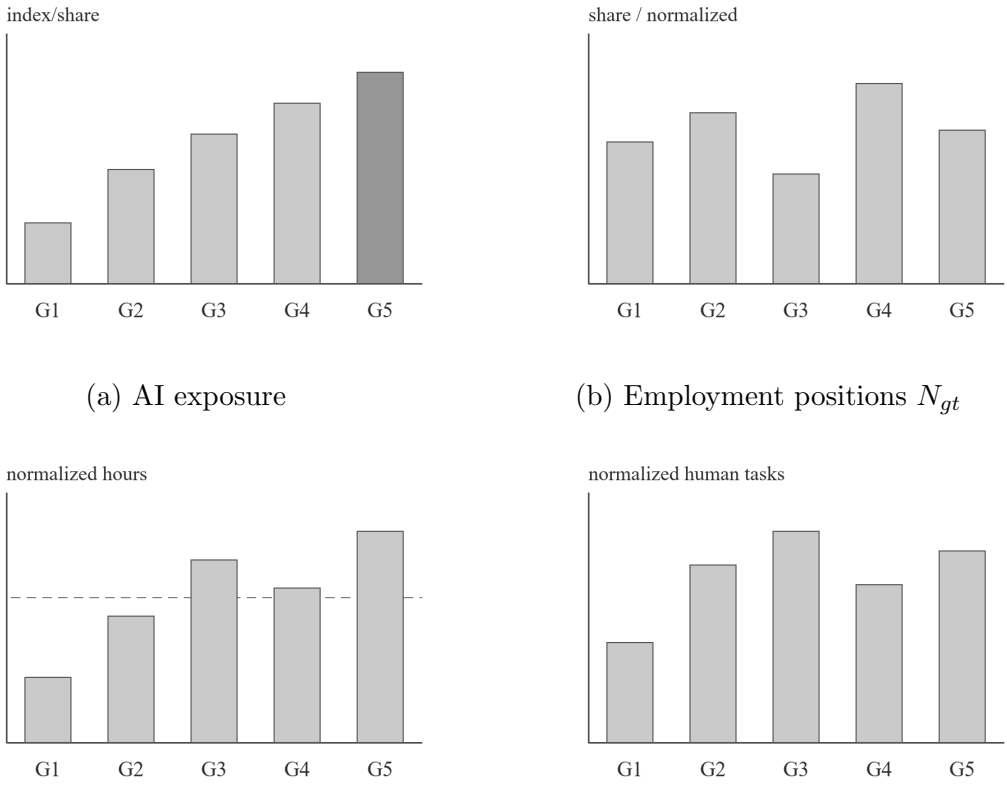


Figure 2: Firm Joint Choice of Headcount, Work Intensity, and Output Scale

Note: The figure is a conceptual illustration, not a calibrated object. A–F label headcount intervals from large layoffs to large hires. Dashed curves are output-scale contours, with $Y = Y_0$ as the baseline reference and $Y > Y_0$ or $Y < Y_0$ showing expansion or contraction. The horizontal and vertical reference lines mark $h_0 \equiv \bar{h}^F$ and N_0. Changes in human-task productivity γ^H move firms across the headcount–intensity space; household-level hours are determined later through household choices, retention/fallback, and equilibrium clearing.



(a) AI exposure

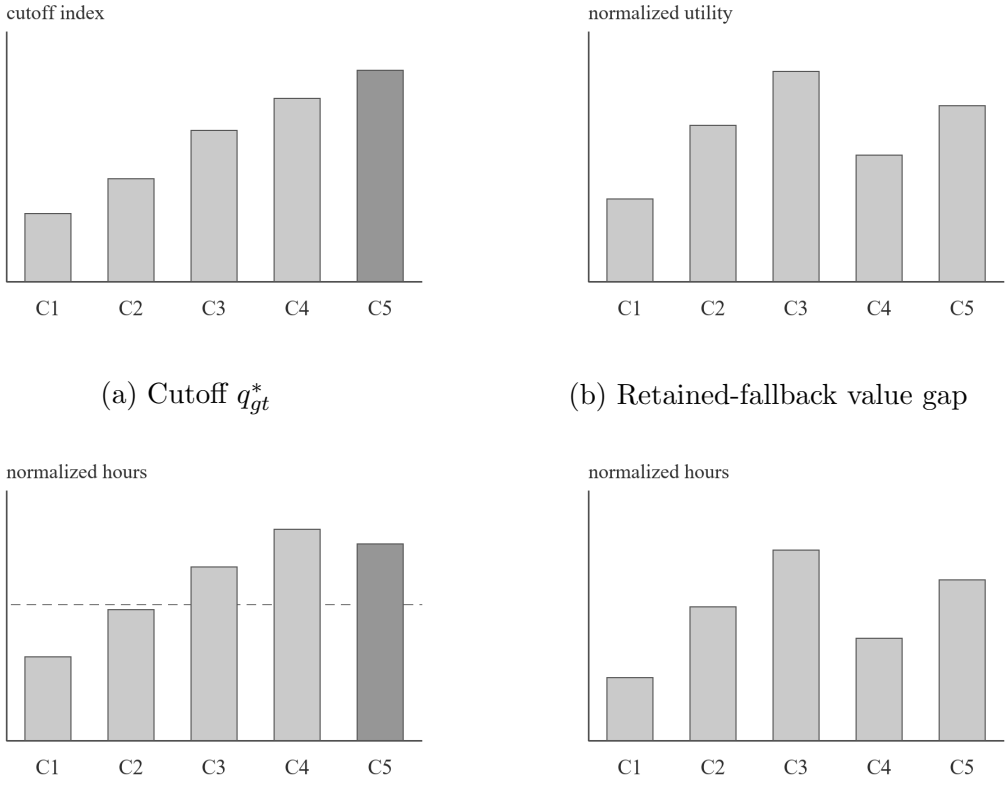
(b) Employment positions N_{gt}

(c) Required work intensity h_{gt}^F

(d) Human task demand H_{gt}^D

Figure 3: Baseline Firm Labor Organization by Job Group

Note: Each bar is a goods job-group cell. Labels $G1$–$G5$ denote anonymized job groups sorted by baseline AI exposure. The horizontal axis is categorical, not a time-series or continuous ordering. Vertical units are shown inside each panel.



(a) Cutoff q^*_{gt}

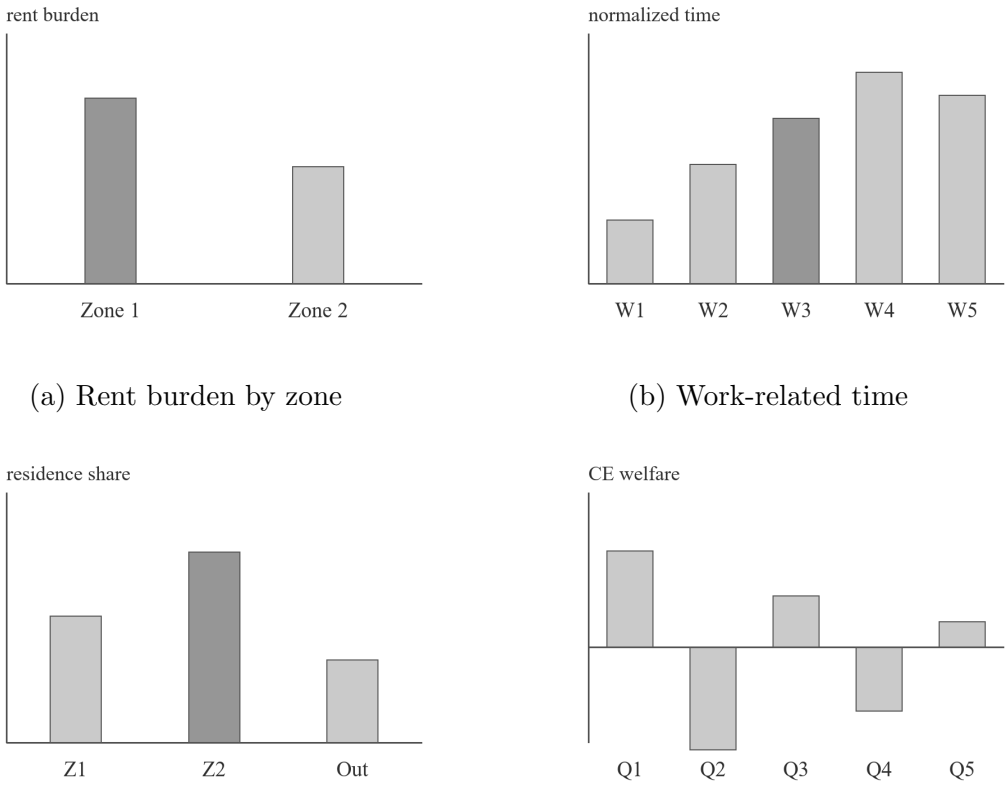
(b) Retained-fallback value gap

(c) Actual hours n_{it}

(d) Defensive excess hours

Figure 4: Retention Risk and Defensive Hours

Note: Each bar is a worker/job cell, not a time-series observation. Labels $C1$–$C5$ denote anonymized cells sorted by baseline retention risk. Vertical units are shown inside each panel.



(a) Rent burden by zone

(b) Work-related time

(c) Residential sorting

(d) Welfare incidence

Figure 5: Spatial Incidence of Work-Time Pressure

Note: Each bar is a categorical baseline model output, not a time-series observation. Panels (a) and (c) report zone or location categories; panel (b) uses worker/location cells $W1$–$W5$ sorted by work-related time; panel (d) uses worker groups $Q1$–$Q5$ sorted by welfare incidence. Vertical units are shown inside each panel.

References

Acemoglu, Daron, and Pascual Restrepo. 2018. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review*, 108(6): 1488–1542.

Acemoglu, Daron, and Pascual Restrepo. 2019*a*. "Artificial Intelligence, Automation, and Work." In *The Economics of Artificial Intelligence: An Agenda*. , ed. Ajay Agrawal, Joshua Gans and Avi Goldfarb, 197–236. Chicago:University of Chicago Press.

Acemoglu, Daron, and Pascual Restrepo. 2019*b*. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives*, 33(2): 3–30.

Acemoglu, Daron, and Pascual Restrepo. 2022. "Tasks, Automation, and the Rise in U.S. Wage Inequality." *Econometrica*, 90(5): 1973–2016.

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2022. "AI and Jobs: Evidence from Online Vacancies." National Bureau of Economic Research Working Paper 28257. Revised February 2022.

Alekseeva, Liudmila, Jose Azar, Mireia Gine, Sampsa Samila, and Bledi Taska. 2021. "The Demand for AI Skills in the Labor Market." *Labour Economics*, 71: 102002.

Autor, David H. 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives*, 29(3): 3–30.

Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review*, 103(5): 1553–1597.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics*, 118(4): 1279–1333.

Babina, Tania, Anastassia Fedyk, Alex X. He, and James Hodson. 2023. "Firm Investments in Artificial Intelligence Technologies and Changes in Workforce Composition." National Bureau of Economic Research Working Paper 31325.

Bick, Alexander, Adam Blandin, and David J. Deming. 2024. "The Rapid Adoption of Generative AI." National Bureau of Economic Research Working Paper 32966.

Bonney, Kathryn, Cory L. Breaux, Emin Dinlersoz, Lucia S. Foster, John C. Haltiwanger, and Aditya A. Pande. 2026. "The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks." National Bureau of Economic Research Working Paper 35141. Also circulated as U.S. Census Bureau CES Working Paper CES-26-25.

Brueckner, Jan K., Matthew E. Kahn, and Gary C. Lin. 2023. "A New Spatial Hedonic Equilibrium in the Emerging Work-from-Home Economy?" *American Economic Journal: Applied Economics*, 15(2): 285–319.

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. 2025. "Generative AI at Work." *The Quarterly Journal of Economics*, 140(2): 889–942.

Caballero, Ricardo J., Eduardo M. R. A. Engel, and John C. Haltiwanger. 1997. "Aggregate Employment Dynamics: Building from Microeconomic Evidence." *American Economic Review*, 87(1): 115–137.

Chetty, Raj, John N. Friedman, Tore Olsen, and Luigi Pistaferri. 2011. "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records." *Quarterly Journal of Economics*, 126(2): 749–804.

Cooper, Russell, John C. Haltiwanger, and Jonathan L. Willis. 2015. "Dynamics of Labor Demand: Evidence from Plant-Level Observations and Aggregate Implications." *Research in Economics*, 69(1): 37–50.

Davis, Morris A., Andra C. Ghent, and Jesse M. Gregory. 2024. "The Work-From-Home Technology Boon and Its Consequences." *The Review of Economic Studies*, 91(6): 3362–3401.

Dell'Acqua, Fabrizio, III McFowland, Edward, Ethan Mollick, Hila Lifshitz, Katherine C. Kellogg, Saran Rajendran, Lisa Krayer, Francois Candelon, and Karim R. Lakhani. 2026. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality." *Organization Science*, 37(2): 403–423.

Delventhal, Matthew J., Eunjee Kwon, and Andrii Parkhomenko. 2022. "How Do Cities Change When We Work from Home?" *Journal of Urban Economics*, 127: 103331.

Demirer, Mert, John J. Horton, Nicole Immorlica, Brendan Lucier, and Peyman Shahidi. 2026. "Chaining Tasks, Redefining Work: A Theory of AI Automation." National Bureau of Economic Research Working Paper 34859.

Diamond, Rebecca. 2016. "The Determinants and Welfare Implications of U.S. Workers' Diverging Location Choices by Skill: 1980–2000." *American Economic Review*, 106(3): 479–524.

Felten, Edward W., Manav Raj, and Robert Seamans. 2021. "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses." *Strategic Management Journal*, 42(12): 2195–2217.

Felten, Edward W., Manav Raj, and Robert Seamans. 2023. "Occupational Heterogeneity in Exposure to Generative AI." *SSRN Electronic Journal*.

Greaney, Brian, Andrii Parkhomenko, and Stijn Van Nieuwerburgh. 2025. "Dynamic Urban Economics." National Bureau of Economic Research Working Paper 33512.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller. 2025. "Artificial Intelligence and the Labor Market." National Bureau of Economic Research Working Paper 33509.

Hsieh, Chang-Tai, and Enrico Moretti. 2019. "Housing Constraints and Spatial Misallocation." *American Economic Journal: Macroeconomics*, 11(2): 1–39.

Huang, Yueling. 2025. "The Labor Market Impact of Artificial Intelligence: Local vs. Aggregate Effect." *SSRN Electronic Journal*.

Jiang, Wei, Junyoung Park, Rachel Jiqiu Xiao, and Shen Zhang. 2025. "AI and the Extended Workday: Productivity, Contracting Efficiency, and Distribution of Rents." National Bureau of Economic Research Working Paper 33536.

Lazear, Edward P., and Sherwin Rosen. 1981. "Rank-Order Tournaments as Optimum Labor Contracts." *Journal of Political Economy*, 89(5): 841–864.

Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg. 2018. "Commuting, Migration, and Local Employment Elasticities." *American Economic Review*, 108(12): 3855–3890.

Moretti, Enrico. 2011. "Local Labor Markets." In *Handbook of Labor Economics*. Vol. 4B, , ed. Orley Ashenfelter and David Card, Chapter 14, 1237–1313. Amsterdam:Elsevier.

Noy, Shakked, and Whitney Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science*, 381(6654): 187–192.

Redding, Stephen J., and Esteban Rossi-Hansberg. 2017. "Quantitative Spatial Economics." *Annual Review of Economics*, 9: 21–58.

Trapeznikova, Ija. 2017. "Employment Adjustment and Labor Utilization." *International Economic Review*, 58(3): 889–922.

Webb, Michael. 2020. "The Impact of Artificial Intelligence on the Labor Market." *SSRN Electronic Journal*.

Online Appendix
(Not for Publication)

A Additional Model Details

This appendix collects formal equilibrium conditions, fallback-state accounting, and additional counterfactual details.

A.1 Formal Equilibrium Conditions

Given exogenous sequences for the AI environment and AI service rental price, an equilibrium is a sequence of prices

$$\{\{W_{gCt}\}_{g\in\mathcal{G}}, W^S_{Ct}, W^D_{Ct}, R_{kt}, P^H_{kt}, r^N_{kt}, Q_t\},$$

firm policies, household policies, performance cutoffs, support-income objects, and household distributions μ_t such that: the following five blocks are satisfied.

Optimization and policy consistency. Goods firms solve the joint task-allocation, output, headcount, and work-intensity problem in eq. (22); local service labor demand is given by eq. (29); developers solve eq. (32); and households solve the Bellman problem in eq. (37). These optimality conditions determine firm policies, household application and fallback policies, and household choices over city status, tenure, location, and hours.

Goods jobs: task allocation, headcount, work intensity, and overtime. Goods-sector equilibrium has two margins on the labor side: the cutoff clears employment positions, and job-specific wages clear effective human task input. The goods firm obtains human task demand and the associated multiplier λ^H_{gt} from the conditional task-allocation problem in eq. (5). Employment slots clear through the endogenous cutoff:

$$\sum_i \mu_{it}\mathbf{1}\{K_i = C\}\mathbf{1}\{j_i = g\}\mathbf{1}\{q_{i,g,t} \geq q^*_{gt}\} = N_{g,t}, \qquad g \in \mathcal{G}. \tag{52}$$

The cutoff q^*_{gt} is the equilibrium rationing object for firm positions and satisfies eq. (27).

Household paid effective human work supplied to each retained goods job group is

$$E^S_{gCt} = \sum_i \mu_{it} \mathbf{1}\{K_i = C\} \mathbf{1}\{j_i = g\} \mathbf{1}\{q_{i,g,t} \geq q^*_{gt}\} z_i \chi_{ig,t} h^F_{gt}, \qquad g \in \mathcal{G}, \tag{53}$$

and goods work-intensity markets clear by job group when retained workers' paid effective hours equal the human task demand required by production and the capacity provided by firm headcount and work intensity:

$$E^S_{gCt} = H^D_{gt} = H^{cap}_{gt} = N_{gt} h^F_{gt}, \qquad g \in \mathcal{G}. \tag{54}$$

Job-specific wages W_{gCt} are the equilibrium prices for these effective-hour markets. Equivalently, for each goods job group,

$$Z_{gCt}(W_{gCt}) \equiv E^S_{gCt}(W_{gCt}, \mu_t, \pi_t) - H^D_{gt}\left(Y^*_{gt}, H^{cap*}_{gt}, A_t, r^A_t\right) = 0, \qquad g \in \mathcal{G}, \tag{55}$$

where $H^{cap*}_{gt} = N^*_{gt} h^{F*}_{gt}$, and π_t collects household policy functions. Here H^D_{gt} is task-side human demand and H^{cap}_{gt} is firm-organized capacity; equilibrium makes them coincide. The model does not require a closed-form wage equation because household hours, fallback choices, migration, residential sorting, and housing demand jointly determine effective work supply. The wage instead solves this fixed point together with firm policies and the household distribution. Overtime is reported as an accounting diagnostic rather than as a separate market-clearing condition. Paid goods overtime effective hours are

$$O^S_{gt} = \sum_i \mu_{it} \mathbf{1}\{K_i = C, j_i = g\} \mathbf{1}\{q_{i,g,t} \geq q^*_{gt}\} z_i \chi_{ig,t} \max\{h^F_{gt} - \bar{h}^F, 0\}, \tag{56}$$

while firm-side paid overtime accounting is

$$O^F_{gt} = N_{gt} \max\{h^F_{gt} - \bar{h}^F, 0\}, \qquad g \in \mathcal{G}. \tag{57}$$

Unpaid defensive work in goods is monitored separately:

$$U^{def}_{gt} = \sum_i \mu_{it} \mathbf{1}\{K_i = C, j_i = g\} \mathbf{1}\{q_{i,g,t} \geq q^*_{gt}\} z_i \chi_{ig,t} \max\{n_i - h^F_{gt}, 0\}. \tag{58}$$

The model monitors paid overtime and unpaid defensive work as diagnostics. The binding labor-market condition remains effective-hours clearing in eq. (54).

Non-goods city labor and status accounting. Local service and developer markets clear in paid effective hours, while city search and outside status enter the population accounting.

Since these sectors do not have a goods-style retention tournament, actual hours are paid hours. Effective human hours supplied to local service work are

$$L_t^{S,S} = \sum_i \mu_{it} \mathbf{1}\{K_i = C, j_i = S\} z_i \chi_{iS,t} n_i, \tag{59}$$

and the local service labor market clears according to

$$L_t^{S,S} = L_t^{D,S}. \tag{60}$$

The number of households in local service work is

$$N_t^S = \sum_i \mu_{it} \mathbf{1}\{K_i = C, j_i = S\}. \tag{61}$$

Effective human hours supplied to construction work are

$$H_{DCt}^S = \sum_i \mu_{it} \mathbf{1}\{K_i = C\} \mathbf{1}\{j_i = D\} z_i \chi_{iD,t} n_i, \tag{62}$$

with developer employment

$$N_t^{D,work} = \sum_i \mu_{it} \mathbf{1}\{K_i = C\} \mathbf{1}\{j_i = D\}, \tag{63}$$

and construction human-time clearing

$$H_{DCt}^S = \sum_{k \in \{1,2\}} L_{Dkt}^D. \tag{64}$$

City unemployment/search and outside accounting are

$$N_{Ct}^U = \sum_i \mu_{it} \mathbf{1}\{K_i = C, j_i = U_C\}, \tag{65}$$

$$N_t^O = \sum_i \mu_{it} \mathbf{1}\{K_i = O\}. \tag{66}$$

Housing and development markets. Housing markets remain private-only in the paper1b baseline. City housing stocks evolve according to developer production:

$$H_{k,t+1}^S = (1 - \delta_k^H) H_{kt}^S + Q_{kt}^H, \qquad k \in \{1, 2\}. \tag{67}$$

Private rental demand, private rental supply, and privately owned housing stock are

$$H_{kt}^{R,D} = \sum_i \mu_{it} \mathbf{1}\{K_i = C, k_i = k, r_i = R\} s_i^r, \tag{68}$$

$$H_{kt}^{R,S} = \sum_i \mu_{it} \mathbf{1}\{K_i = C, k_i = k, r_i = O\} \hat{s}_i^o, \tag{69}$$

$$H_{kt}^{O} = \sum_i \mu_{it} \mathbf{1}\{K_i = C, k_i = k, r_i = O\} \left(s_i^o + \hat{s}_i^o\right). \tag{70}$$

City rental and ownership housing markets clear by zone:

$$H_{kt}^{R,D} = H_{kt}^{R,S}, \tag{71}$$

$$H_{kt}^{S} = H_{kt}^{O}, \qquad k \in \{1, 2\}. \tag{72}$$

Non-labor development input markets clear according to

$$N_{kt}^{D} = \bar{N}_k^D, \qquad k \in \{1, 2\}. \tag{73}$$

Population and distribution dynamics. The fixed population is allocated across city employment, city unemployment/search, and the outside option:

$$\bar{N} = \sum_{g \in \mathcal{G}} N_{gt} + N_t^S + N_t^{D,work} + N_{Ct}^U + N_t^O, \tag{74}$$

where $N_t^{D,work}$ is employment in developer work and is distinct from the non-labor development input N_{kt}^D. The distribution of households evolves according to policy functions, demographic transitions, productivity and fit transitions, migration, job choices, hours choices, and layoff outcomes.

A.2 Fallback-State Accounting

The fixed population is allocated across employed goods workers, local service workers, developer workers, city unemployment/search, and the outside option. The accounting identity in eq. (74) ensures that workers displaced from AI-affected goods jobs remain in the economy rather than disappearing from the model.

The city unemployment/search state is not a residual error term. It is a household status with its own income y_{Ct}^U, housing choice, and continuation value. Local service work is also not a residual assignment: workers enter it only when its value exceeds other options and the local service labor market clears. These states allow the model to distinguish workers

who remain in the city while searching from those who leave for the outside option or move into fallback employment.

A.3 Employer-Provided Housing Counterfactual

A policy counterfactual introduces employment-linked housing units as a firm- or policy-provided support that lowers effective housing costs for assigned workers. Let E_{kt} denote the number of employment-linked housing units in zone k. Assignment probability for eligible workers is

$$\pi_{kt}^{E} = \min\left\{1, \frac{E_{kt}}{\sum_i \mu_{it} I_{it}^{E}}\right\}, \tag{75}$$

where I_{it}^{E} is an eligibility indicator based on employment status or job group. Assigned renter households face an effective rent

$$\widetilde{R}_{kt} = (1 - \eta_k^{E}) R_{kt}, \qquad \eta_k^{E} \in [0, 1). \tag{76}$$

The parameter η_k^{E} is the rent discount rate for assigned units. The counterfactual records the fiscal or employer resource cost of these units separately from the baseline allocation. This experiment asks whether lowering housing costs for workers exposed to AI-induced labor adjustment reduces work-intensity pressure and improves welfare.