An Interactive Introduction to Search and Matching Models

A Large-Firm Approach for Undergraduate Teaching

Jose Silva

May 10, 2026

Abstract

This paper develops a simple search-and-matching model designed for undergraduate teaching. The framework combines labor-market frictions, vacancy creation, wage bargaining, and equilibrium unemployment in a transparent large-firm environment with constant worker productivity. Equilibrium is characterized geometrically as the intersection of an upward-sloping wage curve and a downward-sloping job creation curve. The paper is accompanied by an interactive computational implementation in Python and Google Colab, which lets students visualize equilibrium, the Beveridge curve, and comparative statics through live simulations. The exposition in the main text emphasizes economic intuition and the graphical analysis. The technical derivations of the job creation condition and the Nash-bargaining wage are collected in two appendices, together with a third appendix describing the companion notebook.

1 Introduction

In real labor markets, unemployed workers do not instantly find jobs and firms do not instantly hire workers. Workers search for employment opportunities while firms post vacancies and recruit employees, and because matching takes time and resources, unemployment and vacancies coexist in equilibrium. The empirical regularity that emerges from this coexistence is the **Beveridge curve**: a downward-sloping relationship between the unemployment rate and the vacancy rate, traced out over the business cycle. The U.S. Bureau of Labor Statistics publishes an interactive version of this curve, updated monthly: *The Beveridge Curve (job openings rate vs. unemployment rate), seasonally adjusted*.[1] The model presented in this paper is, in essence, a theory of why this curve exists and what shifts it.

Search-and-matching models provide a unified framework for studying equilibrium unemployment, job creation, labor-market tightness, wage determination, the Beveridge curve, and labor-market policy. The version presented here uses a *large-firm* formulation in which a representative

[1] https://www.bls.gov/charts/job-openings-and-labor-turnover/job-openings-unemployment-beveridge-curve.htm

multi-worker firm chooses how many vacancies to post, taking the wage and the vacancy-filling rate as given.

After reading this paper and running the companion notebook, a student should be able to:

(i) explain why unemployment and vacancies coexist in equilibrium;

(ii) locate equilibrium graphically as the intersection of the wage curve (WC) and the job creation curve (JC);

(iii) compute equilibrium tightness, wages, and unemployment numerically;

(iv) interpret the Beveridge curve;

(v) perform comparative statics with respect to productivity, separations, matching efficiency, vacancy costs, bargaining power, and unemployment benefits;

(vi) run policy experiments using a live simulator.

A note on what this paper does and doesn't do. The model presented here is continuous-time and dynamic at the aggregate-flow level: the stocks of employment, unemployment, and vacancies evolve according to differential equations, and the Beveridge curve and the steady-state unemployment rate are derived from these flow equations. The firm's hiring problem, however, is solved as a *steady-state reduced-form optimization* rather than as a fully intertemporal dynamic program with asset values for filled jobs and vacancies. This is a deliberate pedagogical choice: the reduced form delivers the same job creation condition as the dynamic-programming formulation in steady state under our assumptions, but it is accessible without the Bellman-equation apparatus. The same applies to the Nash-bargaining wage equation: we work directly with the reduced-form match surplus rather than with worker and firm asset values. Students who want the fully dynamic version are referred to Pissarides (2000), which is the canonical treatment.

The paper is organized as follows. Section 2 sets up the environment. Sections 3–6 introduce the matching function, firms, the job creation condition, and the Nash-bargaining wage. Section 7 characterizes the equilibrium as the intersection of the WC and JC curves. Section 8 derives the Beveridge curve. Section 9 works through two comparative-statics examples. Section 10 describes the interactive notebook and proposes a sequence of computer labs. Section 11 sketches policy applications. Section 12 connects the model to key references in the literature. Section 13 catalogues the main features the model abstracts from and points to active research areas. Section 14 concludes. Appendix A derives the job creation condition from the firm's optimization problem. Appendix B derives the wage equation from Nash bargaining. Appendix C explains how the benchmark parameter values are calibrated to U.S. data. Appendix D describes the structure of the companion notebook.

2 Environment

Time is continuous. The economy is populated by a unit mass of identical workers. The labor market features search frictions: workers and firms cannot meet instantaneously.

At each instant, a worker is either *employed* or *unemployed*. Let e denote the aggregate employment rate and u the aggregate unemployment rate. Because the labor force is normalized to one,

$$e + u = 1.$$

We use lowercase letters throughout the body for aggregate quantities (e for employment, u for unemployment, v for vacancies). The matching function below uses these aggregates. Appendix A introduces uppercase E and V for the firm-level optimization problem and shows how the per-firm and aggregate quantities are related.

Intuition

Why a representative firm? The classical Pissarides exposition uses many small one-job firms with free entry. Here we use a representative large firm that employs many workers and decides how many vacancies to post. The two formulations deliver the same equilibrium under our assumptions, but the large-firm version is closer to how real employers behave—they post many vacancies at once—and is the natural starting point for richer extensions.

3 Matching Technology

Matches between unemployed workers and vacancies are created according to the matching function

$$M(u, v) = \phi\, u^{1-\alpha}\, v^{\alpha}, \qquad \phi > 0, \;\; \alpha \in (0, 1), \tag{1}$$

where v denotes vacancies, ϕ measures matching efficiency, and α is the elasticity of new matches with respect to vacancies.

Define labor-market tightness

$$\theta \equiv \frac{v}{u}.$$

The probability per unit of time that an unemployed worker finds a job (the *job-finding rate*) is

$$p(\theta) \equiv \frac{M}{u} = \phi\, \theta^{\alpha}, \tag{2}$$

and the probability per unit of time that a vacancy is filled (the *vacancy-filling rate*) is

$$q(\theta) \equiv \frac{M}{v} = \phi\, \theta^{-(1-\alpha)}. \tag{3}$$

Key Result

Tightness simultaneously affects both sides of the market:

- tighter markets help *workers* find jobs (p rises with θ),
- tighter markets make hiring *harder for firms* (q falls with θ).

The accounting identity $p(\theta) = \theta\, q(\theta)$ holds at every θ.

4 Firms

Each employed worker produces a constant flow of output equal to A, so aggregate output is $Y = A\, e$. Firms collectively pay a wage w to each employed worker and post v vacancies in total at flow cost c per vacancy. Filled jobs separate exogenously at rate s.

Aggregate employment evolves according to

$$\dot{e} \;=\; q(\theta)\, v \;-\; s\, e. \tag{4}$$

The first term captures successful hiring (the vacancy-filling rate times the stock of vacancies) and the second job destruction (the separation rate times the stock of employed workers). Aggregate flow profits are $\Pi = A\, e - w\, e - c\, v$.[2]

Intuition

Aggregates and firms. Throughout the body we use *lowercase* letters for aggregate quantities: e for employment, u for unemployment, v for vacancies, $\theta = v/u$ for tightness. The matching function $M(u, v)$ takes these aggregates as inputs. Equation (4) describes the aggregate flow balance; because the labor force is normalized to one, e doubles as the employment rate and v as the vacancy rate.

The optimization problem of an individual firm is set up at the firm level, with *uppercase* letters E and V for its own employment and vacancies. With N_f identical firms in equilibrium, $e = N_f \cdot E$ and $v = N_f \cdot V$. Each firm takes the aggregate tightness θ (and hence $q(\theta)$) as given. Appendix A works out the firm's problem explicitly and shows that the per-firm and aggregate flow balances are equivalent.

[2] The firm's optimization problem in the next section is solved as a steady-state reduced-form rather than as a fully intertemporal dynamic program with asset values for filled jobs and vacancies. This delivers the same job creation condition as the dynamic-programming formulation under our assumptions but avoids Bellman-equation machinery. See Pissarides (2000) for the dynamic version.

5 The Job Creation Condition

The firm posts vacancies optimally, taking the wage and the vacancy-filling rate as given. Posting a vacancy costs c per unit of time but is filled only at rate $q(\theta)$, so the expected cost of recruiting one worker is $c/q(\theta)$. Once filled, a job survives until it is destroyed at rate s, giving an expected duration of $1/s$. At the optimum the flow surplus from a worker, capitalized over the expected duration of the match, must equal the expected hiring cost.

The optimization problem leads to the **job creation condition**:

$$\boxed{A - w \;=\; \frac{s\,c}{q(\theta)}.} \tag{5}$$

A formal derivation is given in Appendix A.

Solving (5) for w gives the *job creation curve* in (θ, w) space:

$$w^{JC}(\theta) \;=\; A \;-\; \frac{s\,c}{q(\theta)}. \tag{6}$$

The JC curve slopes *downward* in θ: a tighter market makes vacancies harder to fill, raising hiring costs and forcing wages down.

Intuition

Read (5) as: *flow surplus per worker = separation rate × expected hiring cost*. The presence of s on the right-hand side has a clean interpretation: because filled jobs separate at rate s, the firm must repeatedly replace workers over time, so the relevant cost of holding one worker is the per-period flow cost of *re-hiring* that worker, $s \cdot c/q(\theta)$. If hiring becomes cheaper (low c or high $q(\theta)$), the firm can afford a higher wage. If matches break up quickly (high s), the firm needs a larger surplus to make vacancy posting worthwhile.

6 Wage Determination

When a worker and the firm meet, the match generates a surplus that did not exist before they met. Wages are determined by Nash bargaining: the worker and the firm split the surplus according to the worker's bargaining power $\beta \in (0,1)$. Workers receive flow value b when unemployed (unemployment benefits plus the value of leisure or home production).

The bargaining outcome is the **wage curve**:

$$\boxed{w^{WC}(\theta) \;=\; (1-\beta)\,b \;+\; \beta\,(A + c\,\theta).} \tag{7}$$

A formal derivation is given in Appendix B.

The wage is a convex combination of the worker's outside option b and the firm's full productive value $A + c\theta$. The term $c\theta$ captures the fact that, in a tight market, the firm has been saving on hiring costs (it could find a worker quickly) and the worker can credibly demand a share of those savings. The WC curve is therefore *upward*-sloping in (θ, w) space.

Exercise

Show that $\partial w^{WC}/\partial b > 0$, $\partial w^{WC}/\partial A > 0$, $\partial w^{WC}/\partial c > 0$, and $\partial w^{WC}/\partial \theta > 0$. Discuss in one sentence each why the sign is intuitive.

7 Equilibrium: Where WC Meets JC

The equilibrium of the model is the unique point at which the wage curve and the job creation curve cross. Setting $w^{WC}(\theta) = w^{JC}(\theta)$:

$$(1-\beta)\, b + \beta\,(A + c\,\theta) \;=\; A - \frac{s\,c}{q(\theta)}. \tag{8}$$

Equation (8) pins down equilibrium tightness θ^*. Since the WC curve is upward-sloping and the JC curve is downward-sloping, the solution is unique whenever $A > b$. Once θ^* is known, the equilibrium wage w^* is read off either curve.

Figure 1 shows what this looks like at the benchmark calibration. The blue WC curve rises with θ because workers can demand a larger share of the firm's productive contribution when the market is tight. The red JC curve falls with θ because vacancies take longer to fill in tight markets, eating into firms' surplus. The two curves cross exactly once, at $(\theta^*, w^*) \approx (0.63, 0.996)$.

Key Result

The equilibrium is the point where two opposing forces balance. Firms find it harder to fill vacancies as θ rises, which limits how high wages can go (the JC curve falls). Workers can demand higher wages as θ rises, because the firm has been saving on hiring costs (the WC curve rises). Equilibrium θ^* is the unique tightness consistent with both forces.

8 Steady-State Unemployment and the Beveridge Curve

Unemployment evolves according to

$$\dot{u} \;=\; s\,e \;-\; p(\theta)\,u, \qquad \text{where } e = 1 - u. \tag{9}$$

The first term captures inflows into unemployment (the separation rate times the stock of employed workers); the second captures outflows (the job-finding rate times the stock of unemployed workers).

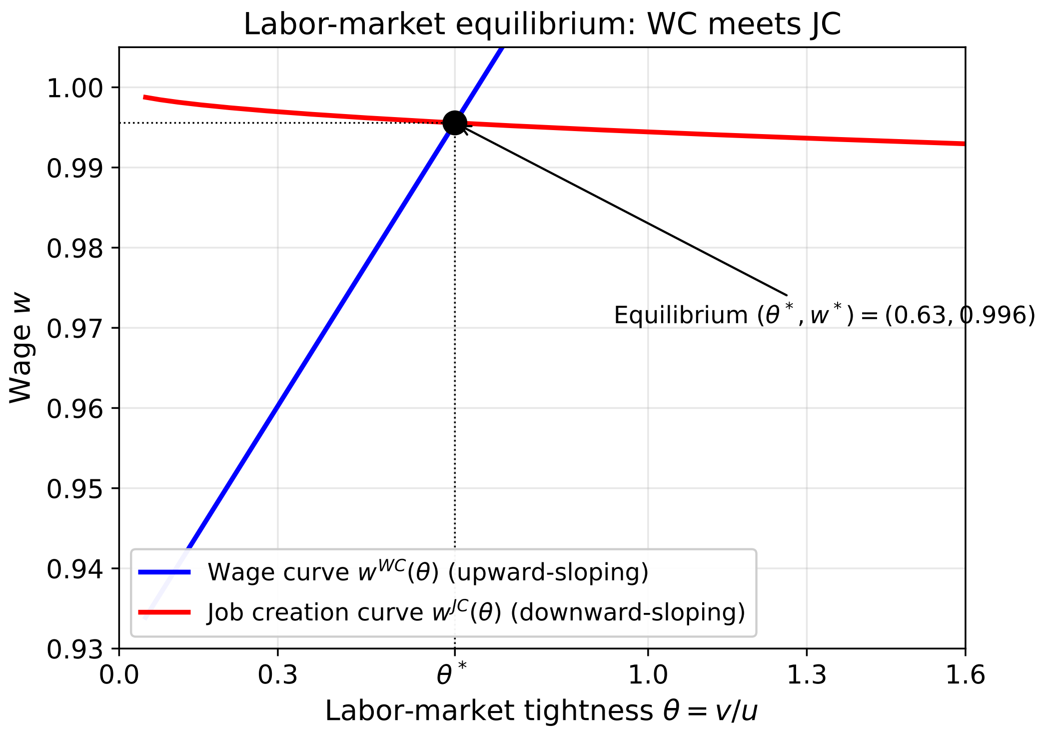

Figure 1: Equilibrium as the intersection of the wage curve (blue, upward-sloping) and the job creation curve (red, downward-sloping). Generated by the companion notebook at the benchmark calibration $A = 1.00$, $b = 0.41$, $\beta = 0.879$, $c = 0.12$, $s = 0.035$, $\phi = 0.754$, $\alpha = 0.5$.

Setting $\dot{u} = 0$ gives the steady-state unemployment rate

$$\boxed{u^* = \frac{s}{s + p(\theta^*)}.} \tag{10}$$

Vacancies are then $v^* = \theta^* u^*$.

Equation (10) traces out the **Beveridge curve**: the locus of (u, v) pairs consistent with steady state for different values of tightness. It is downward-sloping and convex. Higher separation rates or lower matching efficiency shift it outward (more unemployment for any given vacancy level).

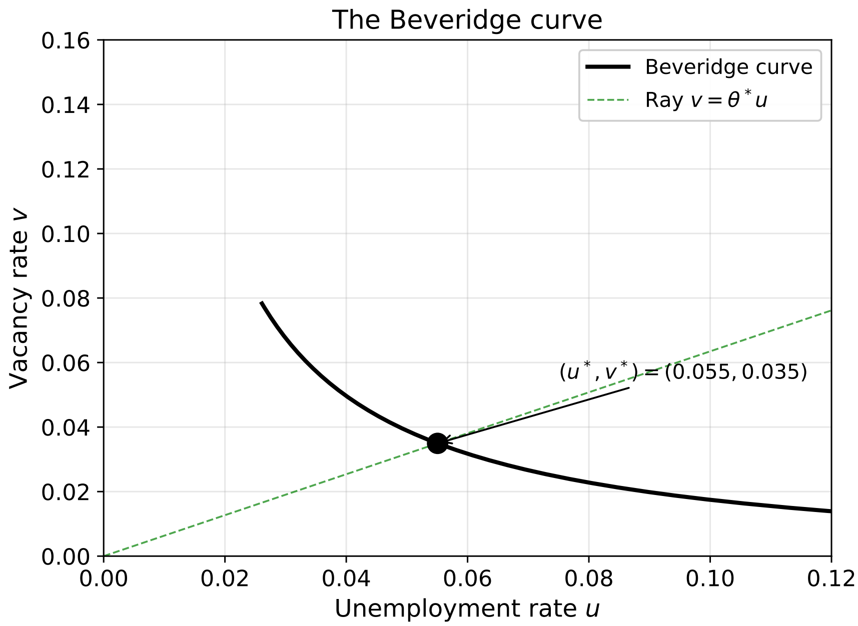

Figure 2: The Beveridge curve. Each point on the black curve is a steady-state (u, v) pair for a given tightness θ. The equilibrium of Figure 1 is the dot. The dashed green line is the ray $v = \theta^* u$: every (u, v) pair on that ray has the same tightness θ^*.

Intuition

The Beveridge curve in real data. The model's predicted downward-sloping convex Beveridge curve is one of the most robust empirical regularities in modern macro. The U.S. Bureau of Labor Statistics maintains a live, monthly-updated chart of the U.S. Beveridge curve at bls.gov. Hovering over the chart shows the recession periods (the curve shifts out and the economy moves up the curve as the labor market tightens during recoveries) and the post-2009 outward shift that motivated a large empirical literature on changes in matching efficiency ϕ. Students are encouraged to compare the qualitative shape of Figure 2 to what the BLS shows in real time.

Benchmark numbers

Solving (8) numerically at the benchmark calibration gives

$$\theta^* \approx 0.635, \quad w^* \approx 0.996, \quad p(\theta^*) \approx 0.601, \quad q(\theta^*) \approx 0.946,$$

$$u^* \approx 0.055, \quad v^* \approx 0.035.$$

That is, the benchmark economy has roughly 5.5% unemployment and a vacancy rate of 3.5%.

9 Comparative Statics: Two Worked Examples

The clearest way to understand the model is to see how the equilibrium moves when a parameter changes. We work through two leading examples.

9.1 A productivity shock

Suppose productivity rises from $A = 1.00$ to $A = 1.15$. The wage curve (7) shifts upward (workers can demand a higher share of a larger pie), and the job creation curve (6) also shifts upward (firms can offer higher wages and still break even). Both curves move up. Equilibrium tightness rises and the equilibrium wage rises; unemployment falls because $p(\theta^*)$ is larger.

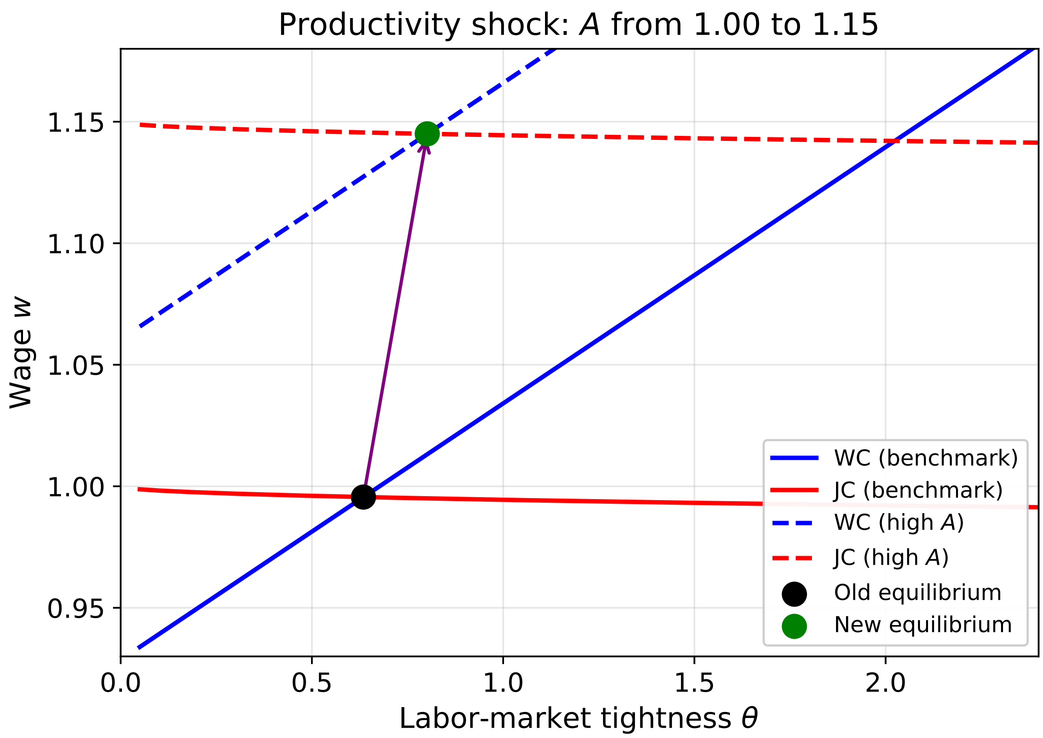

Figure 3: Effect of a productivity shock. Both the WC and JC curves shift upward. The new equilibrium (green) features higher tightness, a higher wage, lower unemployment, and more vacancies.

9.2 Higher unemployment benefits

Now suppose unemployment benefits rise from $b = 0.41$ to $b = 0.65$. This shifts the wage curve upward (workers' outside option is better), but *leaves the job creation curve unchanged* (firms' surplus and hiring costs are unaffected). The new intersection lies up and to the *left*: tightness falls, the wage rises, and unemployment goes up.

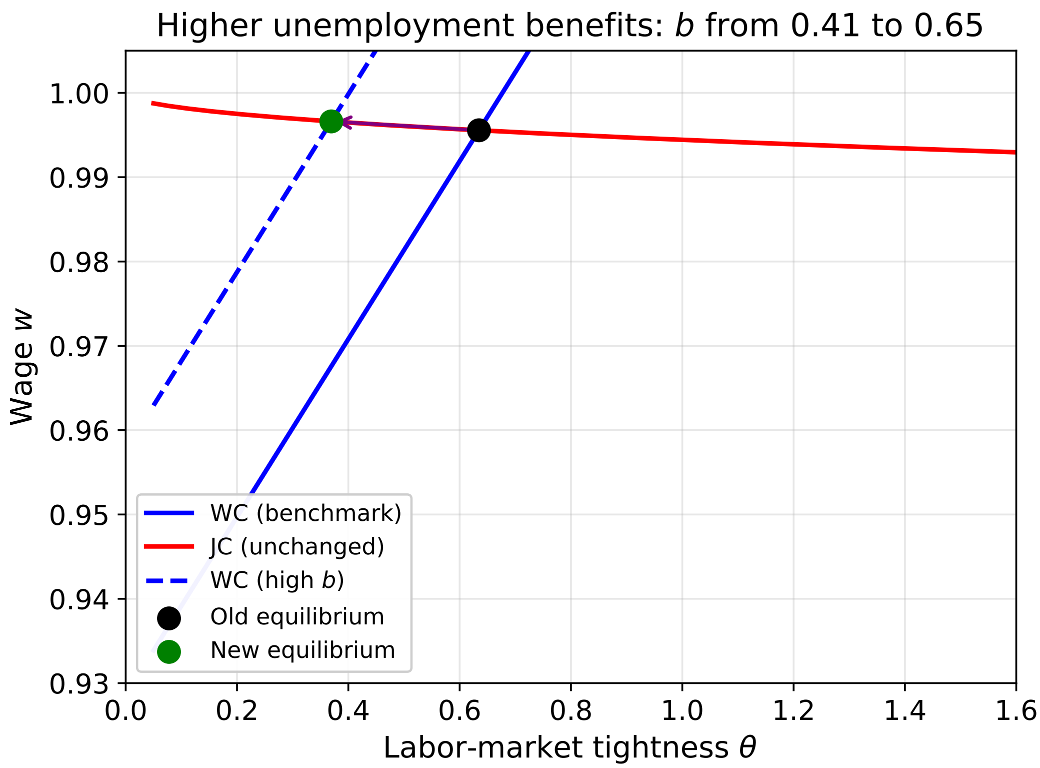

Figure 4: Effect of higher unemployment benefits. Only the WC curve shifts up; JC is unchanged. The new equilibrium (green) features lower tightness, a higher wage, and higher unemployment.

9.3 The full picture

Table 1 summarizes how each of the seven parameters affects the equilibrium.

Table 1: Comparative statics of the model.

Shock	Effect on θ^*	Effect on w^*	Effect on u^*
Higher productivity A	$\uparrow$	$\uparrow$	$\downarrow$
Higher vacancy cost c	$\downarrow$	ambiguous	$\uparrow$
Higher separation rate s	$\downarrow$	$\downarrow$	$\uparrow$
Higher unemployment benefit b	$\downarrow$	$\uparrow$	$\uparrow$
Higher matching efficiency ϕ	$\uparrow$	$\uparrow$	$\downarrow$
Higher bargaining power β	$\downarrow$	$\uparrow$	$\uparrow$
Higher elasticity α	ambiguous	ambiguous	ambiguous

> **Intuition**
>
> The results in Table 1 fall into two families. Anything that makes hiring more profitable for firms (higher A, lower c, lower s, better matching ϕ) shifts JC upward, raises tightness, and lowers unemployment. Anything that strengthens workers' bargaining position (higher b, higher β) shifts WC upward, raises wages, but lowers tightness and increases unemployment.

10 The Interactive Notebook

The model is implemented as an interactive Python notebook on Google Colab. When you run the notebook, you see the interface shown in Figure 5: seven sliders for the parameters, a "Reset to benchmark" button, the WC–JC equilibrium plot on the left, the Beveridge curve on the right, and a panel below reporting the equilibrium values. Solid lines correspond to the benchmark calibration (frozen); dashed lines correspond to the counterfactual generated by the slider values. The black dot is the benchmark equilibrium and the green dot is the counterfactual equilibrium. The technical structure of the notebook is described in Appendix D.

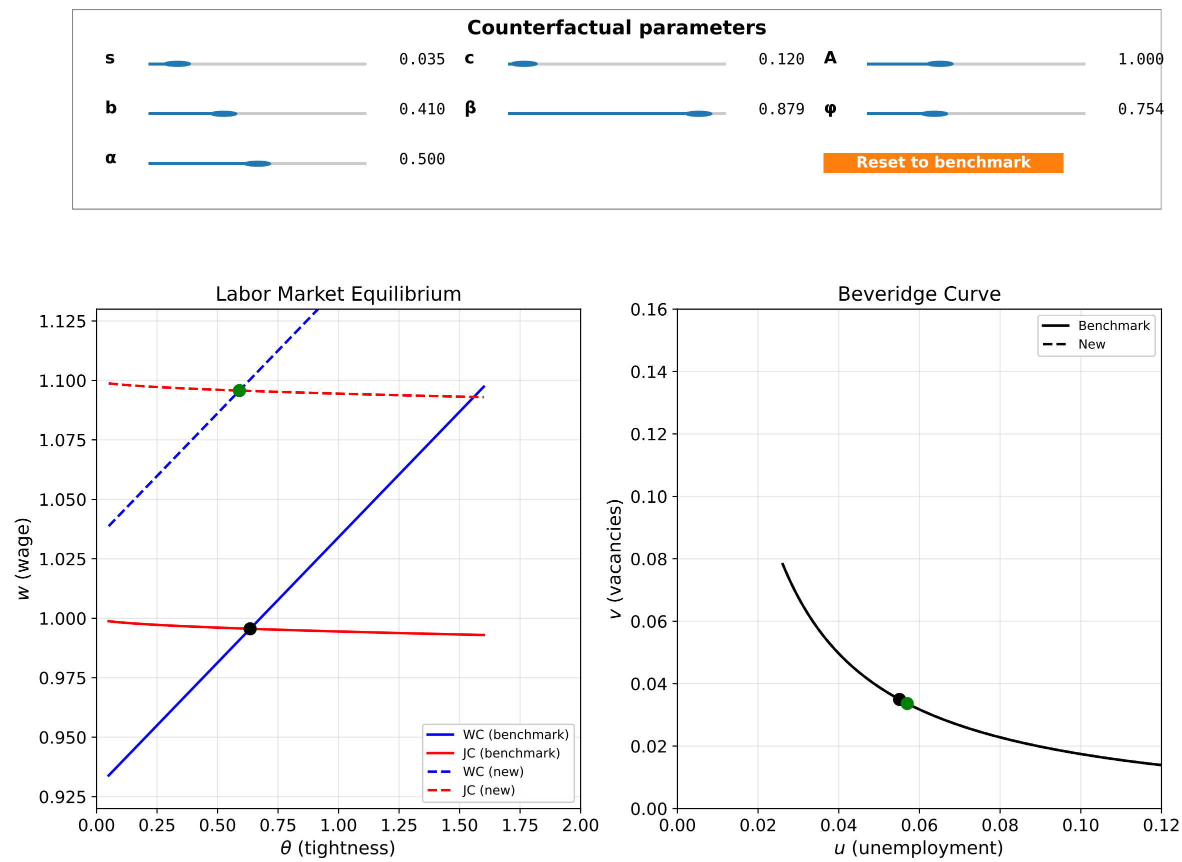

Figure 5: The interactive interface. Sliders set the seven parameters of the counterfactual economy. Solid lines correspond to the benchmark calibration; dashed lines correspond to the counterfactual. Here both A and b have been raised relative to the benchmark.

Suggested labs

The following four labs walk students through the main mechanisms of the model.

Computer Lab

Lab 1 – Productivity shocks. Starting from the benchmark, raise A from 1.00 to 1.20 in steps of 0.05. Verify that you reproduce the qualitative pattern of Figure 3. Track θ^*, w^*, and u^* in the results panel. Where does the equilibrium travel along the Beveridge curve?

Computer Lab

Lab 2 – Unemployment insurance. Raise b from 0.41 to 0.70 in steps of 0.05. Compare your screen with Figure 4. Document the effect on w^*, θ^*, and u^*. Are your findings consistent with the "moral hazard" view of unemployment insurance?

Computer Lab

Lab 3 – Matching efficiency. Reset to benchmark and lower ϕ from 0.754 to 0.50. The Beveridge curve shifts *outward*. Why? Can you reproduce a stylised "Beveridge curve outward shift" similar to the one observed in the U.S. after 2009?

Computer Lab

Lab 4 – Bargaining power. Vary β between 0.30 and 0.95. At which β does the wage approach A? Why is this limit economically sensible?

11 Policy Applications

The framework can be used to analyze several policy questions:

- **Unemployment insurance**: higher b shifts the wage curve upward (Figure 4), lowering tightness, raising wages, and increasing unemployment.
- **Hiring subsidies**: lower c shifts the JC curve upward, raising tightness and reducing unemployment.
- **Firing costs**: in this simplified framework, firing costs can be represented as raising the effective cost of maintaining employment relationships, with broadly negative effects on net job creation. In richer models with endogenous separations, however, firing costs also reduce the destruction margin and alter how match surplus is split, so the net labor-market effect is more nuanced than a higher effective s alone would suggest.
- **Matching technology improvements**: higher ϕ (better employment services, online platforms, retraining) raises both $p(\theta)$ and $q(\theta)$ and reduces unemployment.

Exercise

Suppose a government wants to reduce equilibrium unemployment by 1 percentage point. Compare two policies in the notebook: (i) raising ϕ from 0.754 to a higher value, and (ii) reducing b. Which policy delivers the largest reduction in u^* for a given change in wages? What does each policy "cost" the worker?

12 Connections to the Literature

The teaching model in this paper is a deliberately simplified version of a much richer research literature. The box below sketches how each piece of the model connects to a key reference, so that interested students know where to go next.

Connections to the literature

- **The framework.** Mortensen and Pissarides (1994) introduce the modern search-and-matching model with endogenous job destruction. They share the 2010 Nobel Prize with Peter Diamond for this body of work. Pissarides (2000) is the canonical textbook treatment of the model and is the standard reference for the large-firm formulation used here.
- **Calibration to steady-state moments.** Shimer (2005) provides the calibration strategy followed in our Appendix C: target steady-state unemployment and vacancy rates rather than business-cycle volatilities.
- **Vacancy posting costs.** Silva and Toledo (2009) estimates that hiring costs represent about 4% of quarterly wages in the U.S., the basis for the benchmark value $c = 0.12$.
- **Wage rigidity and the volatility puzzle.** Shimer (2005) also shows that, under Nash bargaining, the model generates too little volatility in unemployment and vacancies relative to the data. Hall (2005) responds by replacing flexible Nash bargaining with sticky wages, which substantially improves the model's cyclical performance. These papers motivate one of the main extensions of the simple framework presented here.

13 What This Model Leaves Out

The framework presented here is deliberately simplified for teaching. Students should be aware of the main features it abstracts from, both to avoid drawing over-strong conclusions from the simulator and to see where the modern research frontier lies. The most important omissions are:

- **Endogenous separations.** The job-destruction rate s is exogenous. In Mortensen and Pissarides (1994), s is determined by an idiosyncratic productivity shock and a reservation-productivity threshold, so policies that affect the separation margin (such as employment protection) have first-order effects that this model cannot capture.

- **Worker heterogeneity.** Workers are identical here. Real labor markets feature heterogeneous skills, demographics, and reservation wages, which in turn generate composition effects in unemployment dynamics.

- **On-the-job search.** Employed workers cannot search here; in reality much of the labor market's reallocation happens through job-to-job transitions, which moderate the wage and tightness responses to shocks.

- **Wage rigidity and the volatility puzzle.** Nash bargaining delivers wages that are too flexible to match the cyclical volatility of unemployment and vacancies in the data (Shimer, 2005). Hall (2005) addresses this by replacing Nash bargaining with sticky wages.

- **Business cycles and aggregate dynamics.** The model is purely steady-state. Transitional dynamics, cyclical comovement, and the interaction with monetary or fiscal policy are outside its scope.

- **Firm dynamics.** Firms are identical and infinitely lived; there is no entry, exit, or firm size distribution. Models of labor-market dynamics with firm heterogeneity (e.g., Elsby and Michaels, 2013) generate richer flows.

- **Directed search.** Workers and vacancies meet randomly through the matching function $M(u, v)$. In directed-search models, workers and firms condition their search on observable wages or types, which changes the welfare properties of the equilibrium.

- **Monopsony and labor-market concentration.** Firms here take the wage as the outcome of bilateral bargaining. A growing literature (e.g., Berger, Herkenhoff, and Mongey, 2022) emphasizes that firms also have wage-setting power, which depresses wages and employment relative to the bargaining benchmark.

- **Imperfect product-market competition.** Firms here are price-takers. Allowing firms to charge a markup over marginal cost (Dixit and Stiglitz, 1977) modifies both the JC and WC conditions and links the labor-market block to the modern markup literature (De Loecker, Eeckhout, and Unger, 2020).

These extensions are active research areas. A natural next step for students who finish this paper is to choose one of them and explore how it modifies the simple WC–JC equilibrium developed here.

14 Conclusion

This paper presented a simple search-and-matching model designed for undergraduate teaching. The framework combines labor-market frictions, vacancy creation, wage bargaining, and equilibrium unemployment in a transparent large-firm environment. Equilibrium is characterized geometrically as the intersection of an upward-sloping wage curve and a downward-sloping job creation curve. The interactive interface lets students see the equilibrium move in real time as parameters change, turning comparative statics from an algebraic exercise into a visual one.

A Deriving the Job Creation Condition

This appendix derives equation (5) from the firm's profit-maximization problem in steady state. We make the firm/aggregate distinction explicit: *uppercase* letters E and V denote the firm's own employment and vacancies, while *lowercase* letters e, u, v denote aggregates (as in the body). With N_f identical firms in equilibrium,

$$e \;=\; N_f \cdot E, \qquad v \;=\; N_f \cdot V, \qquad u \;=\; 1 - e, \qquad \theta \;=\; v/u.$$

Each firm takes the aggregate tightness θ (and hence the matching rates $p(\theta)$, $q(\theta)$) as given.

The firm

A representative firm produces a homogeneous good according to

$$Y \;=\; A\,E, \tag{11}$$

where A is productivity per worker and E is the firm's employment. A fraction s of employed workers separate at each instant, so the firm must post vacancies V to replace them. Posting a vacancy entails a flow cost c, and a posted vacancy is filled at the rate $q(\theta)$ defined in equation (3).

The firm's flow profits are

$$\Pi \;=\; A\,E \;-\; w\,E \;-\; c\,V. \tag{12}$$

Steady-state flows

At the firm level, in steady state the inflow into employment must equal the outflow. Inflows come from successful matches at rate $q(\theta)V$; outflows come from separations at rate sE. Balancing them gives the firm-level flow condition

$$q(\theta)\,V \;=\; s\,E. \tag{13}$$

Aggregating across the N_f identical firms (multiplying both sides by N_f) recovers the aggregate flow balance $q(\theta)\,v = s\,e$, equivalent to $p(\theta)\,u = s\,e$ by the definition of tightness, since $q(\theta)\,v = q(\theta)\,u\,\theta = p(\theta)\,u$. This is exactly equation (4) in the body, set to $\dot{e} = 0$.

Solving (13) for V pins down the number of vacancies the firm must post in order to maintain employment E:

$$V \;=\; \frac{s\,E}{q(\theta)}. \tag{14}$$

Profit maximization

Substituting (14) into (12), the firm's problem reduces to a one-variable maximization in E:

$$\max_{E} \; \Pi \;=\; A\,E \;-\; w\,E \;-\; c\,\frac{s\,E}{q(\theta)}, \tag{15}$$

taking w and θ as given. The first-order condition with respect to E delivers the **job creation condition**:

$$\boxed{A \;=\; w \;+\; \frac{s\,c}{q(\theta)}.} \tag{16}$$

Rearranging gives equation (5) of the main text, $A - w = s\,c/q(\theta)$. The economic content is exactly the same: the flow surplus $A - w$ that the firm earns from a worker must equal the expected hiring cost incurred to keep that position filled in the face of separations.

B Deriving the Wage Equation from Nash Bargaining

This appendix derives equation (7) from the Nash-bargaining problem.

The bargaining set-up

When a match is formed, the job produces output A. The worker receives wage w, while the firm retains the surplus $A - w$. If negotiations break down the worker receives the outside option b (unemployment benefits plus the value of leisure or home production), and the firm must reopen the position. In a tighter labor market—a higher tightness θ—it becomes more difficult for firms to find workers, raising the expected vacancy cost $c\,\theta$ of hiring again. This reduces the firm's fallback value and strengthens the worker's bargaining position.

Intuition

A note on the surplus formulation. The expression $A - w + c\,\theta$ should be read as a *reduced-form* representation of the firm's surplus, designed for undergraduate exposition. The term $c\,\theta$ captures the recruiting-cost savings the firm enjoys by retaining the current match; it is not literally the firm's continuation value in a fully dynamic Bellman formulation. The fully dynamic asset-value derivation (see, e.g., Pissarides 2000) yields the same wage equation (7) in steady state under our assumptions, but is more notationally demanding. We adopt the reduced-form version here because it preserves the economic content while remaining accessible.

Let $\beta \in (0,1)$ denote the worker's bargaining power. The Nash solution chooses w to maximize the weighted product of the two parties' surpluses:

$$\max_{w} \;\; (w-b)^{\beta}\left[A - w + c\,\theta\right]^{1-\beta}. \tag{17}$$

Taking logs and differentiating

Because $\ln(\cdot)$ is strictly increasing, we can equivalently maximize the log of the Nash product:

$$\max_{w} \;\; \beta \ln(w-b) \;+\; (1-\beta) \ln\big(A - w + c\,\theta\big). \tag{18}$$

Differentiating with respect to w:

$$\frac{\partial}{\partial w}\Big[\beta \ln(w-b) \;+\; (1-\beta)\ln(A-w+c\,\theta)\Big] \;=\; \frac{\beta}{w-b} \;-\; \frac{1-\beta}{A-w+c\,\theta}.$$

Setting the derivative to zero gives the first-order condition

$$\frac{\beta}{w-b} \;=\; \frac{1-\beta}{A-w+c\,\theta} \quad\iff\quad \beta\,\big(A-w+c\,\theta\big) \;=\; (1-\beta)\,(w-b). \tag{19}$$

Solving for the wage

Solving (19) for w delivers the **wage curve**:

$$\beta\,(A-w+c\,\theta) \;=\; (1-\beta)\,(w-b),$$
$$\boxed{w \;=\; (1-\beta)\,b \;+\; \beta\,\big[A + c\,\theta\big].} \tag{20}$$

This is equation (7) of the main text.

Reading the wage equation

The wage is a convex combination of the worker's outside option b and the firm's full productive value $A + c\,\theta$. Setting $\beta = 0$ gives $w = b$: with no bargaining power, the worker receives only the outside option. Setting $\beta = 1$ gives $w = A + c\,\theta$: the worker captures the entire productive value, including the recruiting cost the firm has saved by filling the vacancy. For intermediate β, the wage lies between these two extremes.

C Calibration of the Benchmark Economy

This appendix explains where the benchmark numbers used throughout the paper come from. The seven structural parameters of the model are A, b, β, c, s, ϕ, and α (Sections 3–6). The model is calibrated to U.S. data averaged over 2002–2025 and to the steady-state conditions of the model. The calibration follows the spirit of Shimer (2005), targeting empirically plausible steady-state labor-market outcomes rather than business-cycle dynamics. Throughout this appendix we set the elasticity of the matching function to $\alpha = 0.5$, the standard value in the empirical literature.

Externally set parameters and data targets

Six quantities are taken as given, either directly from the data or from external evidence. Four are structural parameters of the model:

- separation rate $s = 0.035$;
- vacancy posting cost $c = 0.12$, following Silva and Toledo (2009), who shows that hiring costs represent about 4% of quarterly wages ($c = 0.04 \times 3$);
- productivity $A = 1$ (a normalization);
- flow value of unemployment $b = 0.41$, set to match the proportion of earnings lost to taxes and benefit reduction when a jobseeker starts a new job in the U.S. (OECD, 2023).

Two are observed labor-market moments used as calibration targets:

- unemployment rate $u_{\text{data}} = 0.05$;
- vacancy rate $v_{\text{data}} = 0.035$,

which together imply a target tightness $\theta_{\text{data}} = v_{\text{data}}/u_{\text{data}} = 0.636$.

Internally calibrated parameters

Two structural parameters remain to be pinned down: matching efficiency ϕ and worker bargaining power β. Together with the equilibrium objects w^*, $p(\theta^*)$, and $q(\theta^*)$, they form the unknowns of the calibration problem

$$\{\, \phi,\ \beta,\ p(\theta^*),\ q(\theta^*),\ w^* \,\}.$$

These quantities must satisfy the five steady-state equilibrium conditions of the model: the steady-state unemployment relation (10), the job creation condition (5), the matching rates (2)–(3), and the wage curve (7). With $\alpha = 0.5$ this gives five equations in five unknowns.

The system is nonlinear but small, and can be solved with Excel Solver (or any equivalent root finder) by minimizing the sum of squared residuals to zero. The solution is

$$\phi = 0.754, \quad \beta = 0.879, \quad p(\theta^*) = 0.601, \quad q(\theta^*) = 0.946, \quad w^* = 0.996.$$

These are the benchmark values used in the body of the paper (Section 8) and in the companion notebook (Appendix D).

Intuition

On the magnitude of β. The calibrated value $\beta \approx 0.88$ should not be interpreted literally as "workers capture 88% of the surplus" in the empirical sense. It is the value required for the simple model to match the steady-state targets simultaneously, and it partly compensates for features the benchmark abstracts from (no on-the-job search, no endogenous separations, no wage rigidity). Shimer (2005) and the subsequent literature have emphasized that calibrating β to match steady-state moments tends to produce values much higher than direct empirical estimates of bargaining power. For counterfactual exercises in the simulator, students should treat β as a model parameter rather than an empirical quantity.

Validation

The model implications at the calibrated parameters are very close to the data targets. The benchmark equilibrium delivers $\theta^* \approx 0.635$, $u^* \approx 0.055$, and $v^* \approx 0.035$ (see Section 8), against the data targets $\theta_{\text{data}} = 0.636$, $u_{\text{data}} = 0.05$, and $v_{\text{data}} = 0.035$. The calibrated job-finding rate $p(\theta^*) = 0.601$ is also close to the average value $p_{\text{data}} \approx 0.657$ observed in the U.S. economy between 2002 and 2025. The small residual discrepancies reflect the simplicity of the model—a single sector, constant productivity, exogenous separations—and are acceptable for the teaching purposes of this paper.

D The Companion Notebook

The model is implemented in a single Google Colab cell that takes only a few seconds to run. This appendix documents the structure of the notebook and reproduces its core code.

Slider mapping

Each slider in the notebook corresponds to one parameter of the model, as shown in Table 2.

Table 2: Notebook sliders and corresponding model parameters.

Slider	Meaning	Benchmark	Range
`s`	Job separation rate	0.035	[0.01, 0.20]
`c`	Vacancy cost	0.12	[0.05, 1.00]
`A`	Productivity per worker	1.00	[0.50, 2.00]
`b`	Flow value of unemployment	0.41	[0.10, 1.00]
`beta`	Worker bargaining power	0.879	[0.10, 0.99]
`phi`	Matching efficiency	0.754	[0.20, 2.00]
`alpha`	Vacancy elasticity in matching	0.50	[0.10, 0.90]

Core functions

The economic content of the simulator is contained in five short functions. The matching rates (2) and (3) are

```
def p(theta, phi, alpha):
    return phi * theta**alpha

def q(theta, phi, alpha):
    return phi * theta**(-(1 - alpha))
```

The wage curve (7) and the job creation curve (6) are

```
def wage_curve(theta, A, b, beta, c):
    return beta*A + beta*c*theta + (1-beta)*b

def job_creation(theta, A, c, s, phi, alpha):
    return A - c * (s / q(theta, phi, alpha))
```

Equilibrium is computed by solving (8) numerically:

```
from scipy.optimize import brentq

def solve_equilibrium(A, b, beta, c, s, phi, alpha):
    def diff(theta):
        return (wage_curve(theta, A, b, beta, c)
                - job_creation(theta, A, c, s, phi, alpha))
    theta_star = brentq(diff, 1e-4, 100)
    w_star = wage_curve(theta_star, A, b, beta, c)
    return theta_star, w_star
```

The Beveridge curve (10) together with $v = \theta u$ is

```
def bev(theta, s, phi, alpha):
    p_val = p(theta, phi, alpha)
    u = s / (s + p_val)
    v = theta * u
    return u, v
```

Interface and plotting

The remaining code in the notebook is purely cosmetic: it builds the `ipywidgets` sliders and the "Reset" button, draws the WC–JC equilibrium plot and the Beveridge curve in a side-by-side layout, and renders the results panel reporting θ^*, w^*, p, q, u^*, and v^*. None of it performs any economic operation beyond what is described in the main text.

References

Berger, D., Herkenhoff, K. and Mongey, S. (2022). Labor market power. *American Economic Review*, 112(4), 1147–1193.

De Loecker, J., Eeckhout, J. and Unger, G. (2020). The rise of market power and the macroeconomic implications. *Quarterly Journal of Economics*, 135(2), 561–644.

Dixit, A. and Stiglitz, J. (1977). Monopolistic competition and optimum product diversity. *American Economic Review*, 67(3), 297–308.

Elsby, M. W. L. and Michaels, R. (2013). Marginal jobs, heterogeneous firms, and unemployment flows. *American Economic Journal: Macroeconomics*, 5(1), 1–48.

Hall, R. (2005). Employment fluctuations with equilibrium wage stickiness. *American Economic Review*, 95(1), 50–65.

Mortensen, D. and Pissarides, C. (1994). Job creation and job destruction in the theory of unemployment. *Review of Economic Studies*, 61(3), 397–415.

Pissarides, C. (2000). *Equilibrium Unemployment Theory*. MIT Press.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1), 25–49.

Silva, J. and Toledo, M. (2009). Labor turnover costs and the cyclical behavior of vacancies and unemployment. *Macroeconomic Dynamics*, 13(S1), 76–96.