Matching with Bounded Probabilities: A Locally Cobb-Douglas Specification

Mario Rafael Silva[1]

Abstract

We propose the Soft-Min Cobb-Douglas (SMCD) matching function, a tractable three-parameter specification satisfying four desiderata: smoothness and concavity, constant returns to scale, exact joint calibration to any target market tightness, job-finding rate, and matching elasticity, and bounded matching probabilities $f, q \in (0, 1)$ for all market tightness. Cobb-Douglas fails to satisfy the probability bounds, and the Den Haan–Ramey–Watson (HRW) function cannot independently target both level and elasticity. Bounded probabilities imply a strictly decreasing matching elasticity and, through the Hosios condition, a countercyclical efficiency wedge: excess unemployment rises in recessions, a mechanism Cobb-Douglas conceals and HRW mismeasures. We extend SMCD to a two-plateau variant (2P-SMCD) allowing the curvature parameter to differ across the vacancy-scarce and worker-scarce regimes, adding exactly one free parameter while preserving all four properties. Estimating all specifications on U.S. monthly data from 1951 to 2025, we detect a structural break in matching efficiency in October 2008. In pre- and post-break horse races, the single-p SMCD strictly dominates HRW in fit while respecting probability bounds by construction. Cobb-Douglas achieves lower unconditional RMSE but does so by violating $q \leq 1$ in 12 and 9 percent of pre- and post-break months, respectively. Restricting the comparison to months where Cobb-Douglas constitutes a valid probability model, 2P-SMCD attains comparable fit.

Keywords: matching function, probability bounds, soft-min, nonlinear least squares, structural break

JEL Classification: J64, E24, C13, C22

Email address: `msilva913@hkbu.edu` (Mario Rafael Silva)

1. Introduction

The matching function is a fundamental building block of equilibrium unemployment models in the tradition of Diamond (1982), Mortensen and Pissarides (1994), and Pissarides (1985). It maps unemployment u and vacancies v into a flow of new matches $m(u, v)$. Its properties shape the steady-state and cyclical behavior of any model in which job creation and destruction are mediated by search frictions. Two properties in particular govern quantitative applications: the *level* of matching, usually summarized by a steady-state job-finding probability f^*, and the *elasticity* of matching with respect to vacancies, ε^*, which determines how sensitively hiring responds to changes in labor market tightness $\theta = v/u$.

In practice, a large body of quantitative work relies on the Cobb-Douglas (CD) matching function,

$$m(u, v) = Au^{\alpha}v^{1-\alpha},$$

because it is analytically convenient and empirical surveys find it fits aggregate data reasonably well in log-linear space (Petrongolo and Pissarides, 2001). Its appeal for calibration is also clear: given targets $(f^*, \varepsilon^*, \theta^*)$, the parameters $\alpha = 1 - \varepsilon^*$ and $A = f^*/(\theta^*)^{\varepsilon^*}$ follow immediately from closed-form formulas. Yet the Cobb-Douglas function has a basic deficiency: it does not bound the implied matching probabilities. The job-finding probability $f(\theta) = A\theta^{1-\alpha}$ grows without bound as θ increases, violating $f \leq 1$, and the vacancy-filling probability $q(\theta) = A\theta^{-\alpha}$ fails to respect $q \leq 1$ at low tightness. These violations are not a theoretical curiosity. As we document in Section 3, monthly market tightness falls below the lower admissible bound in nearly one in ten months over our 1951–2025 sample, concentrated in every major recession. It is the vacancy-filling bound that fails empirically, not the job-finding bound.

A smooth alternative that satisfies the bounds is the matching function of den Haan et al. (2000), henceforth HRW:

$$m(u, v) = \frac{uv}{(u^{\nu} + v^{\nu})^{1/\nu}}.$$

The HRW function is smooth, has constant returns to scale, and bounds both f and q in $(0, 1)$ for all θ. However, it introduces a different limitation: the job-finding elasticity

$\varepsilon(\theta) = 1/(1+\theta^\nu)$ is pinned entirely by the single parameter ν. Given a calibration point θ^*, the researcher can match either the target elasticity ε^* or the target job-finding rate f^*, but not both simultaneously. In practice, den Haan et al. (2000) calibrate ν to match business-cycle moments rather than steady-state targets, acknowledging that joint calibration to level and elasticity is infeasible within this specification. CD and HRW are the two standard matching function benchmarks in the quantitative DMP literature (Petrosky-Nadeau and Wasmer, 2017).

This paper proposes a matching function that resolves both limitations. We refer to it as the *Soft-Min Cobb-Douglas* (SMCD) matching function:

$$m(u,v) = \frac{sauv}{[(av)^p + u^p]^{1/p}},$$

with parameters $s \in (0,1)$, $a > 0$, and $p > 0$. We show that it uniquely satisfies four desiderata simultaneously: (i) constant returns to scale; (ii) smoothness and concavity; (iii) exact joint calibration to any target triple $(\theta^*, f^*, \varepsilon^*)$, unlike HRW; and (iv) bounded matching probabilities $f, q \in (0,1)$ for all θ, unlike CD. We immediately note two things. SMCD *locally* nests Cobb-Douglas behavior — matching CD to a first-order log-linear approximation at the calibration point — from (ii) and (iii). In fact, closed-form formulas map any calibration triple $(\theta^*, f^*, \varepsilon^*)$ to the parameters (s, a), given a curvature parameter p. Second, SMCD *globally* nests HRW as a special case with $s = a = 1$.

The economic interpretation underlying the name *soft-min* is transparent. The job-finding rate $f(\theta)$ behaves as a smooth minimum of two limiting regimes: a vacancy-limited regime at low tightness where $f(\theta) \approx sa\theta$ and each vacancy is filled with near-certainty, and a worker-limited regime at high tightness where $f(\theta) \approx s$ and workers face a capacity constraint on successful contacts. The parameter p governs how sharply the transition occurs between regimes; as $p \to \infty$, the function converges to the hard-min $\min\{sa\theta, s\}$, which corresponds to a truncated Cobb-Douglas matching function.

We complement the theoretical analysis with an empirical application to U.S. labor market data spanning 1951–2025. Using the job-finding rate constructed by the method of Shimer (2012) and a composite vacancy series that splices Help-Wanted Index data with JOLTS openings (Barnichon, 2010), we estimate all three matching functions — SMCD, HRW, and CD — and compare their fit. We find strong evidence of a structural break in the log-linear relationship between f and θ, identified by a QLR sup-Wald test with a moving-block bootstrap. The estimated break date is October 2008, coinciding with the Lehman Brothers collapse and the onset of the Great Recession. The break reflects both a level shift and a slope change in the log-log relationship, consistent with a persistent decline in matching efficiency documented in related work (Barnichon and Figura, 2015; Şahin et al., 2014). Estimated separately on the pre- and post-break subsamples, the two-plateau SMCD attains fit comparable to CD within the latter's admissible tightness range, while always respecting probability bounds that CD frequently violates; roughly 87 percent of CD's unconditional RMSE advantage is attributable to observations in which it predicts a vacancy-filling probability exceeding one.

The paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the data and documents the CD bound violations. Section 4 develops the SMCD matching function and proves its properties, with a numerical example calibrated to the sample means from Section 3. Section 5 presents the empirical horse race. Section 6 addresses the objection that probability bounds are irrelevant in continuous-time or high-frequency formulations. Section 7 concludes.

2. Related Literature

Our paper sits at the intersection of two strands of work: the theoretical and empirical literature on the matching function, and the calibration and estimation literature in quantitative search models.

The matching function. Petrongolo and Pissarides (2001) provide the canonical survey of matching function estimation across countries and data sources. They find that constant returns to scale is broadly consistent with the data and that CD fits aggregate time-series data reasonably well in log-log space, with estimated unemployment elasticities typically in the range $[0.3, 0.7]$. Their survey, however, focuses on the interior of the data and does not address the behavior of CD at extreme tightness values where probability bounds are violated. Our paper takes this limitation as its starting point.

The HRW matching function was introduced by den Haan et al. (2000) in the context of a business cycle model with job destruction. The key advantage of HRW over CD is that it bounds matching probabilities by construction. Several papers in the DMP tradition have adopted it for this reason. However, as we show formally, the single-parameter structure of

HRW prevents joint calibration to both level and elasticity targets. Our SMCD function nests HRW as a special case ($s = a = 1$) while adding the flexibility needed for joint calibration.[1]

A handful of papers have proposed other modifications of the standard matching function. Anderson and Burgess (2000) structurally estimate a matching function using establishment-level data and find evidence against the homogeneity restrictions implied by CD. Borowczyk-Martins et al. (2013) decompose the matching function into search intensity and contact technology components. Our contribution differs in focus: we propose a functional form that retains the analytical tractability of CD and the probability-bounding property of HRW, while uniquely satisfying both.

Flexible estimation and non-constant elasticity. Lange and Papageorgiou (2020) nonparametrically estimate the matching function allowing for unobserved, time-varying matching efficiency, and reject the constant-elasticity restriction implied by CD. Their estimated vacancy elasticity varies between 0.15 and 0.30 and is higher at low tightness — consistent with the vacancy-limited regime in our model and with the countercyclical pattern that probability bounds mechanically require. They further show that standard parametric estimators are upward-biased when unobserved matching efficiency correlates positively with observed tightness. Since this bias affects all three specifications in our horse race equally, it does not alter the relative rankings we report. Together, these findings motivate a functional form in which elasticity falls with tightness, which the SMCD delivers through its curvature parameter p.

Calibration in quantitative search models. Shimer (2005) demonstrates that the DMP model, calibrated to match the average job-finding rate and average matching elasticity, generates unemployment and vacancy fluctuations an order of magnitude smaller than in the data. His calibration procedure requires pinning down both f^* and ε^* at the steady-state θ^* — precisely the joint target that HRW cannot accommodate but our SMCD can. Hagedorn and Manovskii (2008) recalibrate the model with a higher flow value of unemployment and show that volatility can be matched, while still relying on CD. The sensitivity of quantita-

[1] The urn-ball matching function, $f(\theta) = 1 - e^{-\lambda\theta}/\theta$ for application rate $\lambda > 0$, shares the same one-parameter limitation. Moreover, the urn-ball is strictly less flexible than HRW in shape: HRW's exponent ν governs the curvature of the matching probability profiles and nests the harmonic mean ($\nu = 1$) as a special case, while the urn-ball function fixes the curvature profile as an exponential regardless of λ. For these reasons, Section 5 uses HRW as the representative one-parameter bounded benchmark.

tive predictions to matching function parameters in both of these influential contributions underscores the importance of being able to hit calibration targets exactly.

Structural breaks and matching efficiency. The empirical literature has documented substantial time variation in U.S. matching efficiency. Barnichon and Figura (2015) construct a matching efficiency series and find a sharp decline beginning in 2008–2009, which they attribute to compositional shifts in the pool of unemployed — particularly the rise of long-term unemployment — and to geographic and sectoral mismatch. Şahin et al. (2014) document a rise in mismatch across occupations and industries, which they estimate can account for a substantial fraction of the post-2008 unemployment increase. Our structural break finding at October 2008 aligns with this body of evidence: the shift in the log-log matching relationship reflects primarily a change in matching efficiency rather than a change in the functional form of search technology.

Job-finding rate measurement. We construct job-finding rates following Shimer (2012), who derives formulas for f_t and the separation rate s_t from monthly stocks of unemployment and short-term unemployment. This method is now standard in the empirical matching literature and avoids the timing and stock-flow issues that arise from directly dividing hires by vacancy stocks. The construction is detailed in Section 3.

3. Data

3.1. Sources and construction

We use monthly U.S. data from January 1951 through December 2025. The unemployment rate is the civilian unemployment rate (UNRATE) from FRED, expressed as a fraction of the labor force. Vacancy data splice two series. For 1951–2000 we use the Barnichon (2010) composite Help-Wanted Index (HWI), which adjusts newspaper help-wanted counts for the secular shift to online job postings and is expressed as vacancies per labor force. From December 2000 onward we use JOLTS job openings (JTSJOL), normalized by the civilian labor force (CLF16OV). We scale the HWI series to the JOLTS level using a log-mean ratio computed over the 2001–2019 pre-COVID overlap, ensuring a smooth splice. Market tightness is $\theta_t = v_t/u_t$. Monthly job-finding and separation rates follow the continuous-time construction of Shimer (2012), using unemployment levels (UNEMPLOY), employment levels (CE16OV), and short-term unemployment (UEMPLT5).

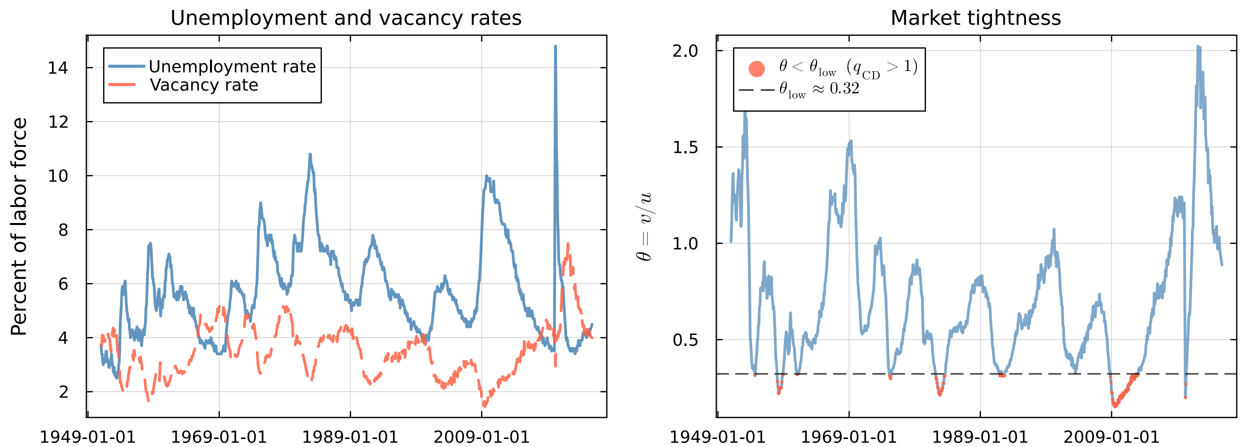

Figure 1: Monthly U.S. labor market data, 1951–2025. Left: unemployment rate (blue) and vacancy rate (dashed red), both as fractions of the labor force. Vacancy rate splices Barnichon (2010) composite HWI (pre-2001) with JOLTS job openings (post-2000), level-adjusted in the overlap. Right: market tightness $\theta = v/u$. Red dots mark months where $\theta < \theta_{\text{low}} \approx 0.32$, i.e., where the calibrated CD function implies a vacancy-filling probability exceeding one. The horizontal dashed line marks θ_{low}.

3.2. Summary statistics and time series

Figure 1 plots the unemployment rate and vacancy rate (left panel) and market tightness (right panel) over the full sample. Periods in which $\theta < \theta_{\text{low}} \approx 0.32$ — where the Cobb-Douglas matching function implies $q > 1$ — are marked in red. These violations are confined to recessions and account for 9.9% of all months in the sample.

Table 1 reports sample means over the full period.

Table 1: Sample means, U.S. monthly data 1951–2025.

Variable	Mean	Std. dev.	Min	Max
Unemployment rate u	0.057	0.017	0.025	0.148
Vacancy rate v	0.035	0.010	0.014	0.075
Market tightness $\theta = v/u$	0.700	0.358	0.152	2.024
Job-finding rate f	0.402	0.096	0.162	0.691

Monthly, Jan 1951 – Dec 2025. $n = 896$. All rates are fractions of the labor force. Vacancy series splices Barnichon (2010) HWI with JOLTS.

3.3. Preliminary evidence on the matching function

The matching function $m(u, v)$ maps unemployment u and vacancies v into a flow of new matches. Under constant returns to scale, setting $\lambda = 1/u$ gives $m(1, v/u) = m(u, v)/u$, so the job-finding probability depends only on market tightness $\theta = v/u$:

$$f(\theta) \equiv \frac{m(u,v)}{u} = m(1,\theta).$$

The matching balance condition $m_t = u_t f_t = v_t q_t$ then pins down the vacancy-filling probability as $q_t = f_t/\theta_t$.

We measure f_t following Shimer (2012). Let $u_t^{<5}$ denote the stock of workers unemployed for less than five weeks (BLS series UEMPLT5), which proxies the flow of new job-losers entering unemployment during month t. The monthly job-finding probability is

$$f_t = 1 - \frac{u_{t+1} - u_{t+1}^{<5}}{u_t}.$$

Subtracting $u_{t+1}^{<5}$ from next-period unemployment isolates workers who were unemployed at the start of the period and remained so, netting out new entrants. The corresponding separation rate accounts for workers who may both separate and find a new job within the same month, using the continuous-time adjustment of Shimer (2012): $s_t = u_{t+1}^{<5} \cdot f_t^h \,/\, [e_t(1 - e^{-f_t^h})]$, where $f_t^h = -\ln(1-f_t)$ is the implied within-period hazard rate. Since the model is monthly, $f_t \in (0,1)$ is a probability; the hazard f_t^h is an intermediate construction used solely to derive s_t.

Estimating the Cobb-Douglas matching function by OLS in log-log space on the full sample yields

$$\log f_t = \log A + \varepsilon \log \theta_t + e_t,$$

with $\hat{A} \approx 0.45$ and $\hat{\varepsilon} \approx 0.29$ ($\hat{\alpha} = 1 - \hat{\varepsilon} \approx 0.71$), consistent with Shimer (2005) and within the range surveyed by Petrongolo and Pissarides (2001). These estimates imply the CD probability bounds hold only for $\theta \in [0.32,\ 15.7]$, with 9.9% of monthly observations violating the lower bound. Section 5 examines whether SMCD fits these data better than CD or HRW, and whether the relationship between f and θ is stable over the sample.

4. The Soft-Min Cobb-Douglas Matching Function

We propose the Soft-Min Cobb-Douglas (SMCD) matching function, defined as

$$m(u,v) = \frac{sauv}{\left[(av)^p + u^p\right]^{1/p}}$$

with parameters $s \in (0,1), a > 0, p > 0$. The denominator is the L^p norm of the vector (u, av). Note that the HRW matching function is a special case with $s = a = 1$. The matching probabilities are

$$\begin{aligned} f(\theta) &= \frac{m}{u} = \frac{sa\theta}{\left[(a\theta)^p + 1\right]^{1/p}} = \frac{s}{(1+(a\theta)^{-p})^{1/p}} \\ q(\theta) &= \frac{m}{v} = \frac{sa}{\left[(a\theta)^p + 1\right]^{1/p}} \end{aligned} \tag{1}$$

The elasticity of the job finding rate is

$$\varepsilon(\theta) = \frac{(a\theta)^{-p}}{1+(a\theta)^{-p}} = \frac{1}{(a\theta)^p + 1} \tag{2}$$

Note that $f(\theta) \le s$ and $q(\theta) \le sa$ for all θ. Moreover, $\varepsilon(0) = 1$ and decreases monotonically to zero as $\theta \to \infty$. This countercyclical pattern is mechanically implied by the bound: since $f(\theta) \le s$, any bounded increasing job-finding rate must saturate as $\theta \to \infty$, driving $\varepsilon(\theta) \to 0$ in tight markets. The point $a\theta = 1$ characterizes the cross-over at which $\varepsilon = 1/2$.

The constant returns to scale property is immediate:

$$m(\lambda u, \lambda v) = \frac{\lambda^2 sauv}{\lambda\left[(av)^p + u^p\right]^{1/p}} = \lambda m(u,v)$$

Lemma 1. *Suppose $m(u,v)$ is constant returns to scale. m is concave if and only if $f(\theta) \equiv m(1,\theta)$ is concave, given $\theta = v/u$ and $u > 0$.*

Proof. First, suppose that $f(\theta)$ is concave. Consider points $(u_1,v_1),(u_2,v_2)$. Set $u' = tu_1 + (1-t)u_2$ and $v' = tv_1 + (1-t)v_2$. Define weights $\lambda_1 = tu_1/u'$ and $\lambda_2 = (1-t)u_2/u'$. Note that $\lambda_1, \lambda_2 \ge 0$ and $\lambda_1 + \lambda_2 = 1$. From the concavity of f,

$$f(v'/u') \ge \lambda_1 f(v_1/u_1) + \lambda_2 f(v_2/u_2)$$

Multiply each side by u' to obtain

$$\begin{aligned} m(u', v') &\geq u'[\lambda_1 f(v_1/u_1) + \lambda_2 f(v_2/u_2)] \\ &= t u_1 f(v_1/u_1) + (1-t) u_2 f(v_2/u_2) \\ &= t m(u_1, v_1) + (1-t) m(u_2, v_2) \end{aligned}$$

Conversely, suppose that m is concave. Then f is concave because it is the restriction of a concave function to the affine set $\{(u, v) : u = 1\}$. Formally, for $\theta_1, \theta_2 \geq 0$,

$$\begin{aligned} f(t\theta_1 + (1-t)\theta_2) &= m(1, t\theta_1 + (1-t)\theta_2) \\ &\geq t m(1, \theta_1) + (1-t) m(1, \theta_2) \\ &= t f(\theta_1) + (1-t) f(\theta_2) \end{aligned}$$

where the inequality uses concavity of m evaluated at the convex combination $t(1, \theta_1) + (1-t)(1, \theta_2)$. □

Next, we find conditions such that $f(\theta)$ is concave.

Lemma 2. *The job finding probability $f(\theta) = sa\theta[1 + (a\theta)^p]^{-1/p}$ is concave for all $a, p > 0$.*

Proof. Let $z = a\theta$ and define $h(z) = z(1 + z^p)^{-1/p}$. As $f(\theta) = sh(a\theta)$, $f''(\theta) = sa^2 h''(a\theta)$. Thus, it suffices to show $h''(z) \leq 0$ for $z > 0$. Calculations yield

$$\begin{aligned} h'(z) &= (1 + z^p)^{-1/p-1} \\ h''(z) &= -(p+1) z^{p-1} (1 + z^p)^{-2-1/p} \leq 0 \quad \text{for} \quad z > 0, p > 0 \end{aligned}$$

□

Let us check (iii) (joint calibration). That is, given θ^*, we require the matching function to satisfy any $f^* = f(\theta^*) \leq \theta^*$ and elasticity ε^*. First, given p, we characterize choices of a and s which satisfy the job-finding rate and elasticity. Then we find a value of p such that the bounds on the matching probabilities are satisfied. As an intermediate step, we write the job-finding rate in terms of the elasticity using (1) and (2):

$$f(\theta) = s(1 - \varepsilon(\theta))^{1/p}$$

This interpretation emphasizes the fact that $f(\theta)$ is bounded above by s, and that the deviation depends on the elasticity scaled by $1/p$. Given initial p, choose a and s as

$$
\begin{aligned}
a &= (\theta^*)^{-1}[1/\varepsilon^* - 1]^{1/p} \\
s &= f^*/(1-\varepsilon^*)^{1/p}
\end{aligned}
$$

It is straightforward to check that, given arbitrary p, these choices ensure $\varepsilon(\theta^*) = \varepsilon^*$ and $f(\theta^*) = f^*$. To bound the matching probabilities, we need to find p such that $s \leq 1$ and $sa \leq 1$. The requirement is thus

$$
\begin{aligned}
s &= f^*/(1-\varepsilon^*)^{1/p} \leq 1 \\
sa &= (f^*/\theta^*)(\varepsilon^*)^{-1/p} \leq 1
\end{aligned}
\tag{3}
$$

We can always satisfy these requirements for p sufficiently large since for $p \to \infty$ these requirements reduce to simply $f^* \leq 1$ and $f^* \leq \theta^*$, which is simply a requirement that the approximation target make sense. Specifically, it follows from (3) that it is sufficient to take $p \geq p_{\min}$ where

$$
p_{min} = \max\{p_s, p_{sa}\}, \quad p_s = \log(1-\varepsilon^*)/\log(f^*), \quad p_{sa} = \log(\varepsilon^*)/\log(f^*/\theta^*)
$$

We thus conclude that the SMCD matching function satisfies all properties (1)–(4).

Equivalence to the CES matching function.. The SMCD matching function is isomorphic to a CES matching function $m(u,v) = A[\eta\, u^\rho + (1-\eta)\, v^\rho]^{1/\rho}$. Applying the identity $xy/(x^p + y^p)^{1/p} = (x^{-p} + y^{-p})^{-1/p}$ with $x = u$, $y = av$ rewrites SMCD as $m = s[u^{-p} + a^{-p}v^{-p}]^{-1/p}$, which matches the CES form under

$$
\rho = -p, \qquad \eta = \frac{a^p}{1+a^p}, \qquad A = \frac{sa}{(1+a^p)^{1/p}},
$$

with inverse $p = -\rho$, $a = [\eta/(1-\eta)]^{-1/\rho}$, $s = A\eta^{1/\rho}$.

Despite this formal equivalence, the SMCD parameterization is preferable for three reasons. First, probability-bound satisfaction requires $\rho < 0$ in the CES — only the negative CES subfamily generates $f, q \in (0,1)$ for all $\theta > 0$, a non-obvious restriction that the SMCD renders explicit as $p \geq p_{\min}$. Second, given any $\rho < 0$, both parameterizations admit closed-form solutions for the remaining parameters from a calibration triple $(f^*, \varepsilon^*, \theta^*)$. The difference is that the SMCD bound conditions $s \leq 1$ and $sa \leq 1$ are immediately legible from (1) and yield $p_{\min}$ in two lines, whereas the equivalent CES constraints $A \leq \eta^{-1/\rho}$ and $A \leq (1-\eta)^{-1/\rho}$ are opaque joint restrictions that give no guidance on how to choose ρ. Third, the CES share $\eta = a^p/(1+a^p)$ conflates the crossover tightness and the curvature without a direct economic interpretation. By contrast, the SMCD parameters s, a, and p each carry a transparent meaning (upper bound, inverse crossover, bottleneck sharpness). HRW is itself CES with $\eta = \frac{1}{2}$, $\rho = -\nu$, $A = 2^{-1/\nu}$, yet it displaced CES in the DMP literature precisely because its parameterization is transparent. SMCD relaxes HRW's symmetry restriction ($s = a = 1$) in the minimal way needed for joint calibration to level and elasticity targets.

Interpretation and limiting cases. We rewrite the job-finding rate to make the economic structure transparent:

$$f(\theta) = \frac{s}{(1 + (a\theta)^{-p})^{1/p}} = \frac{1}{(s^{-p} + (sa\theta)^{-p})^{1/p}}$$

The job-finding rate is the generalized soft-min aggregator of two caps, $y_1 = s$ and $y_2 = sa\theta$, with sharpness governed by $p > 0$. As $p \to \infty$, $f \to \min\{y_1, y_2\}$; the function converges to the hard min of the two regime caps. At $p = 1$ it reduces to the harmonic mean, $f(\theta) = sa\theta/(1 + a\theta)$, which cannot jointly match level and elasticity targets while keeping probabilities bounded.

The two limiting regimes have a direct labor-market interpretation:

- Vacancy-limited regime (low tightness): when θ is small, $f(\theta) \approx sa\theta$. Matches per vacancy are near their cap $q(\theta) \approx sa$, and each worker meets vacancies at a rate proportional to θ. Here the elasticity $\varepsilon \approx 1$.

- Worker-limited regime (high tightness): when θ is large, $f(\theta) \approx s$. Each worker cannot secure more than s successful contacts per period, so the job-finding probability saturates. The elasticity $\varepsilon \approx 0$.

The function chooses the binding bottleneck smoothly: for finite p, it is a soft minimum of the two caps $sa\theta$ and s. A larger p sharpens the transition (closer to the hard min); a smaller p smooths it. This gives a complete characterization of parameter roles. The parameter s is the upper bound of the job-finding rate; a sets the crossover tightness $\theta = 1/a$ where

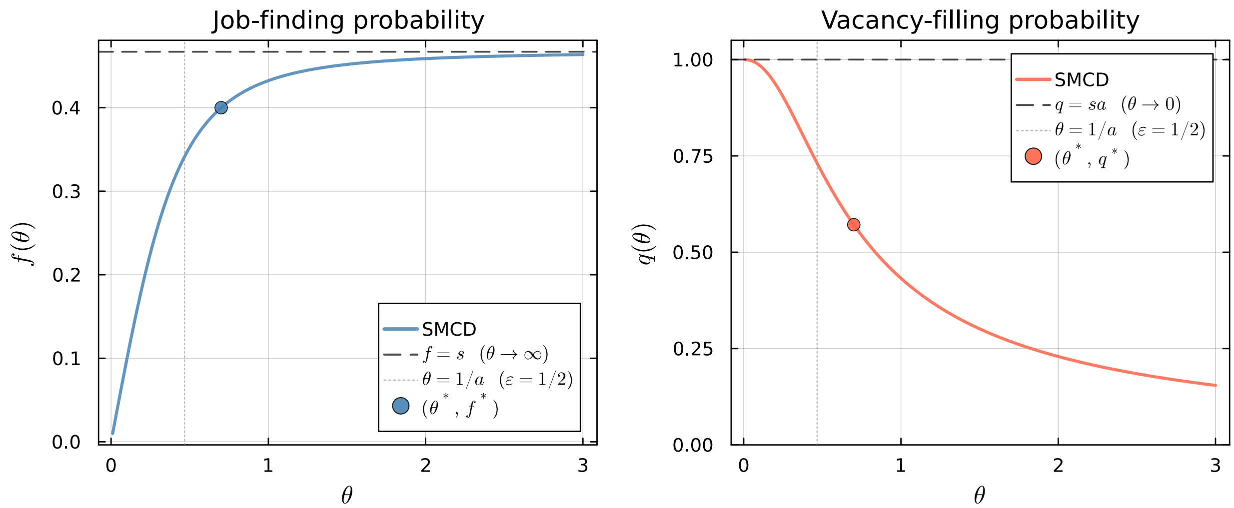

Figure 2: Job-finding probability $f(\theta)$ (left) and vacancy-filling probability $q(\theta)$ (right) under SMCD. Calibration: $(\theta^*, f^*, \varepsilon^*) = (0.70, 0.40, 0.29)$, $p = p_{\min} \approx 2.21$, giving $s \approx 0.467$, $a \approx 2.14$, $sa = 1$. Dashed lines mark the asymptotic bounds $f = s$ and $q = sa$. The vertical dotted line marks the crossover $\theta = 1/a \approx 0.47$ where $\varepsilon(\theta) = 1/2$. The filled circle is the calibration point.

$\varepsilon = 1/2$; sa is the upper bound of the vacancy-filling rate; and p governs the sharpness of the bottleneck.

Figure 2 illustrates these two regimes. Both panels use the baseline calibration $(\theta^*, f^*, \varepsilon^*) = (0.70, 0.40, 0.29)$ with $p = p_{\min}$. In the left panel, the dashed horizontal line marks the worker-limited ceiling $f = s$, approached as $\theta \to \infty$. In the right panel, the dashed horizontal line marks the vacancy-limited ceiling $q = sa$, approached as $\theta \to 0$. The vertical dotted line marks the crossover $\theta = 1/a \approx 0.47$, where $\varepsilon = 1/2$ and neither regime dominates. The calibration point lies to the right of the crossover, in the worker-limited region, consistent with the estimated tightness elasticity $\varepsilon^* = 0.29 < 1/2$.

An empirically relevant special case is $sa = 1$, so that $q(\theta) \to 1$ as $\theta \to 0$. The matching function then reduces to two free parameters, with $s = 1/a$ determining the upper bound on the job-finding rate.

Local Cobb-Douglas behavior. Local Cobb-Douglas behavior is a corollary of (ii) and (iii). Because f is smooth and achieves $f(\theta^*) = f^*$ with elasticity $\varepsilon(\theta^*) = \varepsilon^*$, a first-order Taylor expansion in log-log space gives $d \log m = (1 - \varepsilon^*) d \log u + \varepsilon^* \, d \log v$, so the matching function is locally Cobb-Douglas at θ^*. This holds for any smooth function that passes through a given calibration triple, and SMCD is no exception.

As a numerical illustration, we use the U.S. sample means from Section 3 as calibration targets: $\theta^* = 0.70$, $f^* = 0.40$, and $\varepsilon^* = 0.29$, corresponding to an unemployment elasticity of 0.71 and consistent with the full-sample OLS estimate in Section 3. These targets imply $p_{\min} \approx 2.21$, with the vacancy-filling constraint $sa \leq 1$ binding. Setting $p = p_{\min}$ and applying (3) yields $s \approx 0.467$, $a \approx 2.14$, and $sa \approx 1.000$, so the bounds are exactly satisfied. The associated Cobb-Douglas level parameter is $A = f^*/(\theta^*)^{\varepsilon^*} \approx 0.444$.

The empirically relevant bound failure of CD is the vacancy-filling probability, which exceeds one for $\theta < \theta_{\text{low}} \approx 0.32$ — a condition that holds in 9.9% of monthly observations, concentrated in recessions. The job-finding rate exceeds one only for $\theta > \theta_{\text{high}} \approx 15.7$, a level never reached in the data. HRW respects both bounds but its elasticity at θ^* differs from ε^* by construction. The SMCD is the unique function satisfying all four properties simultaneously; Figure 8 in Section 5.3 plots all four specifications — CD, HRW, single-p SMCD, and the two-plateau SMCD — at the common calibration targets.

Minimizing distance to Cobb-Douglas. Practitioners may prefer a matching function that stays as close as possible to Cobb-Douglas while satisfying the probability bounds. SMCD has three parameters, all necessary: one to match the target level, one the target elasticity, and one to satisfy the bounds. But the third requirement does not pin p uniquely — it carves out a feasible set $[p_{\min}, \infty)$.

Cobb-Douglas has zero curvature in log space. A natural question is which $p \in [p_{\min}, \infty)$ minimizes the curvature of the matching function, bringing it closest to log-linear. The curvature of ε with respect to $\log \theta$ is

$$\kappa(\theta) \equiv \frac{\partial \varepsilon(\theta)}{\partial \log \theta} = -\frac{p(a\theta)^p}{(1 + (a\theta)^p)^2} = -p\varepsilon(1 - \varepsilon)$$

Let $\Delta = \log(\theta/\theta^*)$. Then the second-order log approximation is

$$\log f(\theta) = \log f^* + \varepsilon^* \Delta - \frac{1}{2} p \varepsilon^* (1 - \varepsilon^*) \Delta^2$$

Minimizing the magnitude of the second-order coefficient $\kappa(\theta^*)$ thus makes this function closest to Cobb Douglas locally. This is accomplished by setting $p = p_{\min}$.

Proposition 1. *Fix approximation point θ^* and consider a Cobb Douglas matching function with $f(\theta^*) = f^*$ and $\varepsilon = \varepsilon^*$. The SMCD job-finding rate function which is closest in the neighborhood of θ^*, subject to the bounding constraint $p \geq p_{min}$, has $p = p_{\min}$.*

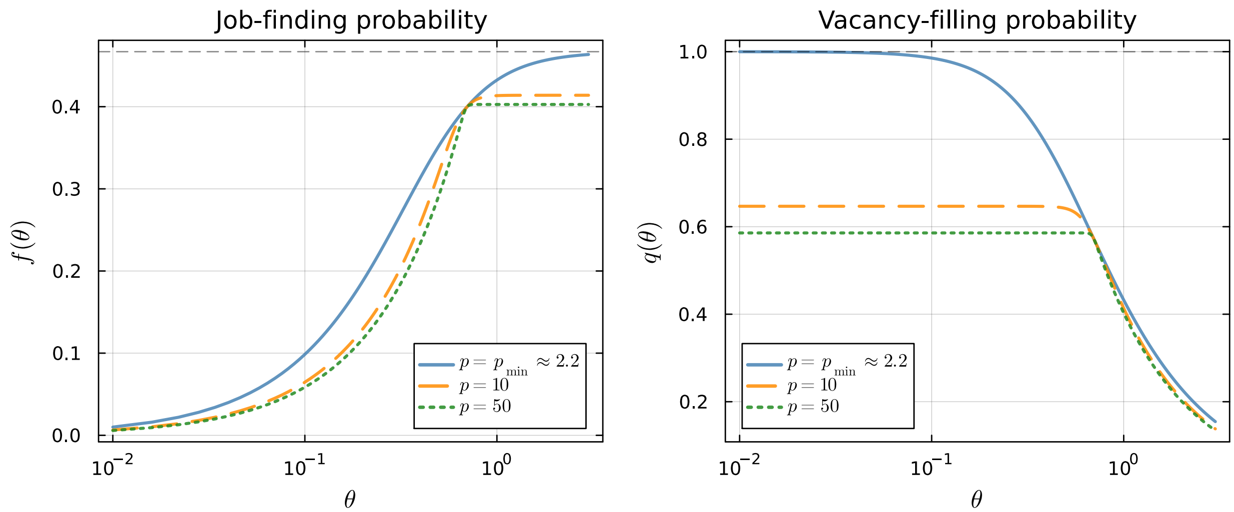

Figure 3: Effect of p on the SMCD job-finding rate (left) and vacancy-filling rate (right). All curves use $(\theta^*, f^*, \varepsilon^*) = (0.70, 0.40, 0.29)$; parameters (s, a) are recalibrated for each p. Dashed lines mark the asymptotes $f = s$ and $q = sa$ at $p = p_{\min}$. Larger p sharpens the regime transition; $p = p_{\min} \approx 2.21$ minimizes curvature and is closest to Cobb-Douglas.

Consequently, some curvature is required to be consistent with the bounds while satisfying the elasticity and job-finding rate targets at θ^*, which is minimized at $p_{\min}$.

Figure 3 illustrates how p governs the speed of convergence to the regime asymptotes. All three curves share the same calibration targets and hence the same $f(\theta^*)$, $q(\theta^*)$, and $\varepsilon(\theta^*)$ at θ^*. As p rises from $p_{\min}$ toward 50, the job-finding rate approaches the hard-min profile $\min\{sa\theta, s\}$ and the transition between regimes sharpens. At $p = p_{\min}$, the departure from Cobb-Douglas is minimized: the function is as close to log-linear as the probability bounds allow.

Figure 4 plots the matching elasticity $\varepsilon(\theta)$ and its log-derivative $\kappa(\theta) = \partial\varepsilon/\partial \log \theta$ at $p = p_{\min}$. The elasticity is strictly decreasing in θ, passing through $\varepsilon^* = 0.29$ at θ^* and through $1/2$ at the crossover $\theta = 1/a \approx 0.47$. The log-derivative $\kappa(\theta) = -p\,\varepsilon(1-\varepsilon)$ is everywhere negative and reaches its largest magnitude at the crossover, where $\varepsilon = 1/2$ and $\kappa = -p/4$. This non-constant elasticity distinguishes SMCD from Cobb-Douglas and aligns with the empirical evidence in Lange and Papageorgiou (2020).

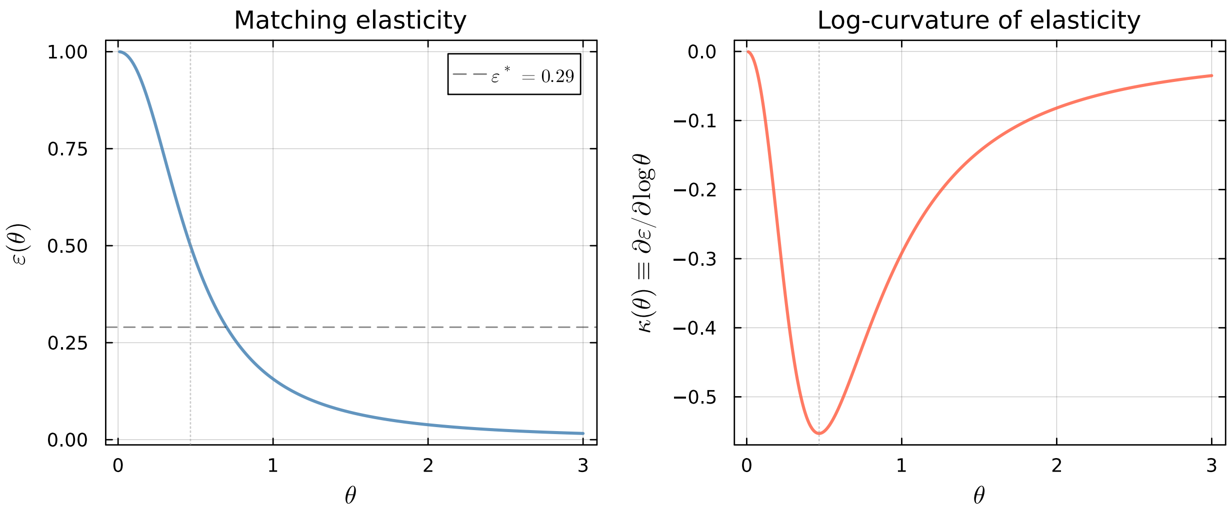

Figure 4: Matching elasticity $\varepsilon(\theta)$ (left) and log-curvature $\kappa(\theta) = \partial\varepsilon/\partial\log\theta$ (right) for the SMCD at $p = p_{\min} \approx 2.21$. The dashed horizontal line marks $\varepsilon^* = 0.29$; the vertical dotted line marks the crossover $\theta = 1/a \approx 0.47$. Curvature is maximized in magnitude at the crossover and vanishes in both tails.

Advantages of SMCD vs. clipped Cobb-Douglas. The clipped Cobb Douglas matching function enforces bounds with kinks and regime switches. The soft-min version ensures well-defined elasticity and curvature everywhere, which helps for identification, estimation, comparative statics, and solutions.

4.1. Search Efficiency and the Hosios Condition

In the Diamond–Mortensen–Pissarides model, firms pay a flow cost κ to maintain an open vacancy and fill it at rate $q(\theta)$. The social planner's problem balances the benefit of additional vacancy posting — more matches per period — against its cost. Posting a vacancy confers a positive externality on unemployed workers by raising the arrival rate of jobs, while simultaneously imposing a negative externality on other firms by lowering $q(\theta)$ through market congestion. The marginal social value of an additional unemployed worker is determined by how much that worker contributes to aggregate match creation, which is the elasticity of $m(u, v)$ with respect to u. Under constant returns to scale,

$$\eta_u(\theta) \equiv \frac{\partial \ln m}{\partial \ln u} = 1 - \varepsilon(\theta). \tag{4}$$

Some unemployment is therefore socially optimal: it raises $q(\theta)$, sustains firm entry, and increases the flow of new matches. The planner equates the marginal benefit of a posted vacancy to the posting cost κ.

In the decentralized equilibrium with Nash bargaining, Hosios (1990) establishes that the equilibrium is constrained efficient if and only if the worker's bargaining power α_L equals the unemployment elasticity of the matching function:

$$\alpha_L = \eta_u(\theta) = 1 - \varepsilon(\theta). \tag{5}$$

The condition has a transparent externality interpretation. Workers internalize only the private return to their own job search; they do not account for the congestion they impose on other searchers. Setting $\alpha_L = \eta_u$ corrects this wedge by aligning workers' compensation with their marginal contribution to aggregate match creation. When $\alpha_L > \eta_u$, workers extract more surplus than their matching contribution warrants: wages are too high, vacancy posting is depressed, and equilibrium unemployment exceeds the efficient level. The opposite holds when $\alpha_L < \eta_u$.

The key implication of SMCD is that η_u is not a constant. From (2) and (4),

$$\eta_u(\theta) = \frac{(a\theta)^p}{1 + (a\theta)^p},$$

which is strictly increasing in θ, ranging from zero in the vacancy-limited regime ($\theta \to 0$) to one in the worker-limited regime ($\theta \to \infty$). Respecting probability bounds inherently produces this monotonicity. In slack markets, vacancies are the scarce input and the marginal unemployed worker contributes little to match creation ($\eta_u \approx 0$); in tight markets, workers are scarce and their marginal contribution is large ($\eta_u \approx 1$). The Hosios condition (5) is therefore state-dependent. Calibrating $\alpha_L = \eta_u(\theta^*)$ to satisfy Hosios at the steady-state tightness, the economy has $\alpha_L > \eta_u(\theta)$ in recessions (low θ): workers' bargaining power exceeds their matching contribution, wages are too high, and equilibrium unemployment exceeds its efficient level. The direction of this wedge is unambiguous and compounding — recessions produce both low θ directly and an amplifying efficiency gap through the Hosios condition, on top of the recessionary shortfall.

Cobb-Douglas conceals this mechanism entirely. Since ε_{CD} is a constant, $\eta_u^{CD} = 1 - \varepsilon_{CD}$ is constant as well, and Hosios either holds uniformly or fails uniformly regardless of the business cycle. The state-dependent wedge is not a feature of the economy in this case but an artifact of the constant-elasticity functional form. HRW partially corrects this: its elasticity $\varepsilon_{HRW}(\theta) = 1/(1 + \theta^\nu)$ is also decreasing in θ, so it qualitatively replicates the countercyclical pattern of η_u. However, HRW's calibration failure carries over. Because ν cannot simultaneously match the job-finding level and the elasticity at θ^*, $\eta_u^{HRW}(\theta^*)$ deviates substantially from its true value. Using the pre-break estimates evaluated at the pre-break sample mean $\bar{\theta} = 0.67$, $\hat{\nu} = 1.154$ implies $\eta_u^{HRW}(\bar{\theta}) \approx 0.39$, while SMCD implies $\eta_u(\bar{\theta}) \approx 0.60$ — a gap of 0.21 in the efficient bargaining power. A researcher using HRW to calibrate the Hosios condition would target the wrong α_L, and any welfare assessment built on that calibration is misspecified by construction.

5. Empirical Application

5.1. Structural break in the log-linear matching relationship

The preliminary Cobb-Douglas estimate in Section 3 pools seventy-four years of data that span large institutional changes in how jobs are posted and filled. We test whether the log-linear relationship between f_t and θ_t is stable over the full sample using a QLR sup-Wald test for an unknown structural break.

Procedure. Let $y_t = \log f_t$ and $x_t = \log \theta_t$. We fit $y_t = \alpha + \beta x_t + u_t$ on the full sample and test for a break in (α, β) at an unknown date τ. For each candidate τ in a trimmed grid (15% from each end, yielding 629 candidates), we estimate the regression separately on $[1, \tau]$ and $[\tau + 1, n]$, compute Newey–West HAC covariance matrices with bandwidth $L = 7$, and form the Wald statistic $W(\tau) = \Delta(\tau)'\big[V_1(\tau) + V_2(\tau)\big]^{-1}\Delta(\tau)$ where $\Delta(\tau) = b_1(\tau) - b_2(\tau)$. The QLR statistic $\sup_\tau W(\tau)$ has a nonstandard asymptotic distribution under an unknown break date; we obtain p-values from a moving block bootstrap (MBB, block length 14, $B = 499$ draws) applied to centered residuals from the null model.

Results.. Figure 5 plots $W(\tau)$ over time. The sup-Wald statistic is $\sup W = 745.97$ with a bootstrap p-value of 0.000. The estimated break date is October 2008, coinciding precisely with the Lehman Brothers collapse and the onset of the global financial crisis.

Table 2 reports the OLS estimates on each subsample. The intercept shifts from -0.691 to -1.144, a change of -0.453 in logs. At any given tightness, the job-finding rate post-October 2008 is approximately $e^{-0.453} \approx 0.64$ times its pre-2008 level — a persistent 36% decline in matching efficiency. The tightness elasticity $\hat{\varepsilon}$ falls modestly from 0.318 to 0.283, indicating that hiring became slightly less responsive to vacancy creation after the crisis.

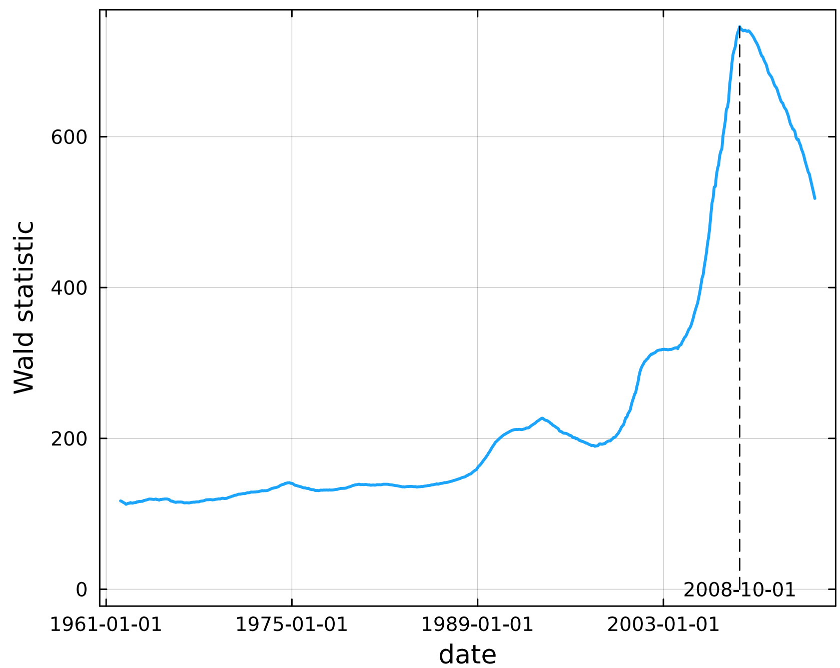

Figure 5: QLR sup-Wald statistic $W(\tau)$ for a structural break in the log-log matching relationship $\log f_t = \alpha + \beta \log \theta_t + u_t$. Vertical dashed line: estimated break date, October 2008. $\sup W = 745.97$; bootstrap p-value $= 0.000$.

Table 2 reveals that the structural break is overwhelmingly a *level shift* in the matching function rather than a change in its *elasticity*. The implied matching efficiency $\hat{A} = e^{\hat{\alpha}}$ drops from 0.501 to 0.319 — a 36% decline — while the tightness elasticity $\hat{\varepsilon}$ falls only modestly from 0.318 to 0.283 (a change of 0.035, or roughly one-tenth of its pre-break value). In the Diamond–Mortensen–Pissarides framework, the Beveridge curve traces the steady-state locus of (u, v) pairs as matching efficiency A varies. A pure decline in A shifts this locus outward — the unemployment-vacancy curve moves away from the origin — without rotating it, because the slope of the Beveridge curve depends primarily on the matching elasticity α and the separation rate, both of which are essentially unchanged here. A change in α would instead *twist* the curve, flattening or steepening it. Our structural break therefore corresponds to an outward *shift* of the Beveridge curve at the Lehman collapse, not a twist.

This shift-versus-twist distinction has been a central organizing theme in the Beveridge curve literature. Lubik (2013) derives the Beveridge curve from an estimated DMP model and shows that the post-2008 outward movement is driven by a structural decline in match efficiency — a shift — rather than a change in matching elasticity. Elsby et al. (2015) survey

Table 2: Estimated log-linear matching relationship, pre- and post-October 2008.

	Pre-break (Jan 1951 – Sep 2008)	Post-break (Oct 2008 – Dec 2025)
Intercept $\hat{\alpha}$	-0.691	-1.144
Tightness elasticity $\hat{\varepsilon}$	0.318	0.283
Implied level $\hat{A} = e^{\hat{\alpha}}$	0.501	0.319
Unemployment elasticity $\hat{\alpha}_u = 1 - \hat{\varepsilon}$	0.682	0.717
CD lower bound $\hat{A}^{1/\hat{\alpha}_u}$	0.36	0.20

Newey–West HAC, $L = 7$. Break identified by QLR sup-Wald with MBB p-values.

the long-run evidence and identify lateral shifts in the Beveridge curve, driven by changes in matching efficiency, as the quantitatively dominant source of variation in unemployment-vacancy dynamics, with the slope remaining relatively stable across business cycles. These findings are consistent with the post-GFC matching efficiency literature. Barnichon and Figura (2015) document a persistent decline in aggregate matching efficiency from late 2008, attributing it to compositional shifts toward long-term unemployment and increased occupational and geographic mismatch. Şahin et al. (2014) estimate that mismatch accounts for a substantial share of the elevated unemployment through 2012. Lange and Papageorgiou (2020) further show, using nonparametric estimation that controls for observed and unobserved search effort, that the Great Recession decline in hires was almost entirely driven by a collapse in matching efficiency rather than by changes in search or recruiting intensity — directly consistent with our finding that the October 2008 break is predominantly a level shift in $\hat{A}$ with little change in the tightness elasticity. Our structural break formalizes this shift as a discrete, statistically unambiguous change in the log-linear matching relationship, precisely dated to the Lehman collapse.

5.2. Horse race: SMCD vs. Cobb-Douglas vs. HRW

We now estimate all three matching functions separately on the pre-break (January 1951 – September 2008) and post-break (October 2008 – December 2025) subsamples and compare fit. Full-sample estimates, which pool the two regimes and are reported for completeness in Figure A.10 in the appendix, show qualitatively similar rankings.

Estimation framework. All three models are estimated by nonlinear least squares (NLS), minimizing the sum of squared residuals in levels of the job-finding rate:

$$\min_{\boldsymbol{\psi}} \sum_{t=1}^{T} \big[f_t - \hat{f}(\theta_t;\, \boldsymbol{\psi})\big]^2.$$

Estimating all models against the same objective ensures a fair comparison of fit. SMCD has three free parameters $\boldsymbol{\psi} = (s, a, p)$, estimated subject to the probability-bound constraints $s \leq 1$ and $sa \leq 1$. HRW has one unconstrained parameter (ν). Cobb-Douglas has two parameters (A, α) and is estimated by NLS in levels without any probability constraint; the log-OLS solution — which minimizes the sum of squared log residuals — serves as the starting point for the NLS routine and is not itself reported as a final estimate. Fit is measured in levels of f by RMSE, MAE, AIC, and BIC; AIC and BIC penalize the additional parameters of SMCD relative to CD and HRW. For Cobb-Douglas, the NLS and log-OLS estimates are close: the largest deviation is $\Delta\hat{\varepsilon} = 0.019$ in the pre-break subsample, with negligible differences post-break, and the implied $q > 1$ threshold $\hat{\theta}_{\min}$ is unchanged.

Table 3: Matching function fit statistics, pre- and post-October 2008.

Period	Model	Params	RMSE	MAE	AIC	BIC	q>1 (%)
Pre-2008	SMCD	3	0.0564	0.0448	−3978.8	−3965.2	0.0
	CD	2	0.0442	0.0349	−4318.0	−4308.9	12.0
	HRW	1	0.0674	0.0549	−3735.9	−3731.4	0.0
Post-2008	SMCD	3	0.0313	0.0238	−1400.1	−1390.2	0.0
	CD	2	0.0284	0.0224	−1441.5	−1434.8	9.4
	HRW	1	0.0573	0.0482	−1158.8	−1155.4	0.0

$n = 693$ pre-break, $n = 203$ post-break. q>1 % is the fraction of months where the estimated model implies $q > 1$; zero by construction for SMCD and HRW.

Figure 6 plots the fitted job-finding rate against the data for each subsample. The post-break scatter lies uniformly below the pre-break scatter at any given θ, consistent with the 36% decline in matching efficiency identified in Table 2.

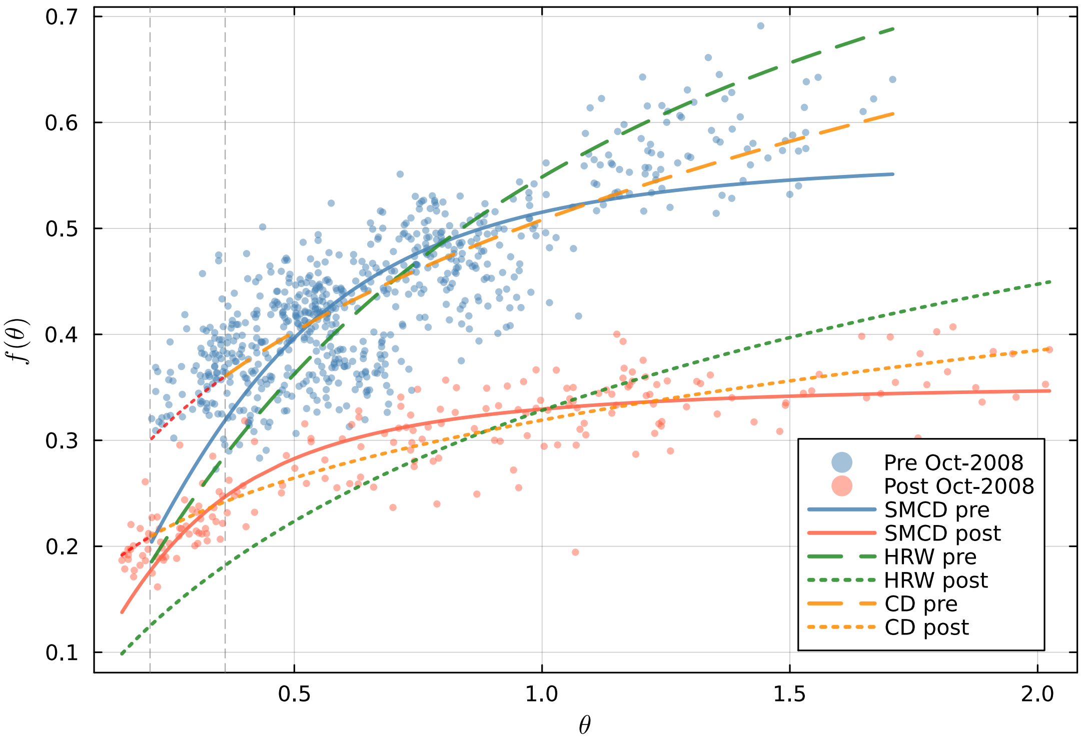

Figure 6: Job-finding rate $f(\theta)$: scatterplot of empirical (θ, f) values together with fitted values for the pre-break (blue) and post-break (red) subsamples. Solid lines: SMCD; dashed: CD; dotted: HRW. Parameters estimated separately on each subsample. For CD, the red dotted segment marks the region $\theta < \theta_{\min}$ where the fitted vacancy-filling rate exceeds one ($\theta_{\min} = 0.36$ pre-break, 0.21 post-break); vertical gray lines indicate the respective $\theta_{\min}$.

Table 3 and Figure 6 deliver three findings. First, CD achieves the lowest RMSE and best information criteria in both subsamples. This result is not surprising since CD is estimated without probability constraints, and hence can fit the data by extrapolating into the region where $q > 1$. The horse race rewards this extrapolation with lower residuals in recession months where CD's fitted vacancy-filling rate exceeds one. SMCD pays a modest cost for guaranteeing validity everywhere: RMSE is 27% higher pre-break and 10% higher post-break. That the gap narrows substantially after 2008 is meaningful — the post-break data span a wider range of tightness (from COVID troughs near $\theta = 0.1$ to the post-pandemic boom above $\theta = 2$), and SMCD's regime-switching behavior handles this dispersion better than a pure power law.

Second, the $q > 1$ column makes CD's advantage concrete. In the pre-break subsample, CD predicts a vacancy-filling probability exceeding one in 12.0% of months; post-break, 9.4%. In those months, CD achieves a lower residual for f while simultaneously predicting an impossible q. Once those months are treated as CD failures rather than fits, the comparison reverses: SMCD delivers valid predictions throughout, while CD does not.

Third, HRW fits substantially worse than both CD and SMCD in every comparison. The RMSE ratio HRW/SMCD is 1.19 pre-break and 1.83 post-break. Despite satisfying the probability bounds, HRW's single-parameter structure prevents it from jointly matching the level and elasticity of job-finding, forcing it to compromise on both. This confirms that bound satisfaction alone is insufficient — the additional calibration flexibility of SMCD over HRW translates directly into a large improvement in fit.

Table 4: Estimated parameters, pre- and post-October 2008.

	SMCD					CD		HRW	
Period	$\hat{s}$	$\hat{a}$	$\hat{p}$	$\widehat{sa}$	$\bar{\varepsilon}$	$\hat{A}$	$\hat{\varepsilon}$	$\hat{\nu}$	$\bar{\varepsilon}$
Pre-2008	0.567	1.762	2.386	1.000	0.465	0.508	0.336	1.154	0.631
Post-2008	0.354	2.829	1.879	1.000	0.310	0.319	0.271	0.622	0.567

All models estimated by NLS in levels. $\bar{\varepsilon}$ = sample mean of $\varepsilon(\theta_t; \hat{\psi})$ over subsample observations (constant for CD)

Table 4 sheds light on what drives the fit differences and how each model absorbs the structural break. For Cobb-Douglas, the level parameter falls sharply — $\hat{A}$ declines from 0.508 to 0.319, a 37% drop — while the elasticity $\hat{\varepsilon}$ shifts only moderately from 0.336 to 0.271, confirming that the structural break documented in Section **??** is primarily a level shift rather than a change in the slope of the matching relationship. The SMCD parameters tell the same story through a different lens: $\hat{s}$ falls 38% (from 0.567 to 0.354), closely mirroring the decline in $\hat{A}$, while $\hat{a}$ rises and the crossover tightness $1/\hat{a}$ shifts leftward, compressing the vacancy-limited regime post-crisis. The binding vacancy-filling constraint $\widehat{sa} = 1.000$ in both periods indicates that NLS allocates the maximum possible vacancy-filling capacity, consistent with extremely tight rationing in slack markets.

The $\bar{\varepsilon}$ column exposes HRW's fundamental limitation most sharply. Because $\hat{\nu}$ is the sole parameter, it must simultaneously control the level of $f(\theta)$ and its slope. Pre-2008, the NLS solution lands at $\hat{\nu} = 1.154$, implying a sample-average elasticity of 0.631 — roughly double CD's estimate of 0.336. Post-2008, $\hat{\nu}$ collapses to 0.622, and the implied $\bar{\varepsilon} = 0.567$ remains far above CD's 0.271. HRW's ν is not measuring the matching elasticity; it is absorbing the efficiency level at the cost of a severely distorted curvature profile, which explains the large fit gap relative to both SMCD and CD. The SMCD average elasticities (0.465 pre, 0.310 post) exceed CD's constant $\hat{\varepsilon}$ as well, but for a different reason: SMCD's elasticity profile is genuinely nonlinear, rising toward one at low tightness and falling toward zero at high tightness, so recession observations where θ is small pull the arithmetic average above the value that log-linear estimation would assign to the typical observation.

5.3. The Two-Plateau Soft-Min Cobb-Douglas Matching Function

The single-p SMCD dominates HRW conditional on respecting probability bounds, but Cobb-Douglas retains a fit advantage in Table 3 that originates entirely from its extrapolation into the $q > 1$ region. A natural one-parameter extension of the SMCD eliminates this residual advantage by assigning independent curvature to each matching regime.

Definition. The single-p SMCD uses one curvature parameter p to govern both the vacancy-limited and worker-limited regimes simultaneously. Define the *two-plateau soft-min Cobb-Douglas* (2P-SMCD) matching function by

$$f(\theta) = f_c \cdot \frac{\theta}{\theta_c} \left[\frac{2}{1 + (\theta/\theta_c)^p} \right]^{1/p}, \qquad p = \begin{cases} p_L & \theta \leq \theta_c \\ p_R & \theta > \theta_c, \end{cases} \tag{6}$$

with parameters (f_c, θ_c, p_L, p_R), where $f_c > 0$, $\theta_c > 0$, and $p_L, p_R > 0$. The parameter θ_c is the crossover tightness and $f_c = f(\theta_c)$ is the job-finding rate there. The single-p SMCD is recovered at $p_L = p_R$, with $\theta_c = 1/a$ and $f_c = s/2^{1/p}$. The extension adds exactly one genuinely new parameter: the asymmetry $p_L \neq p_R$. The crossover θ_c was already present in the single-p model as $1/a$.

The two tail caps are independently controlled. As $\theta \to \infty$, $f(\theta) \to s_R \equiv f_c \cdot 2^{1/p_R}$; as $\theta \to 0$, $q(\theta) \equiv f(\theta)/\theta \to q_{\max} \equiv (f_c/\theta_c) \cdot 2^{1/p_L}$. Setting p_L large makes the vacancy-filling probability rise sharply toward $q_{\max}$ in slack markets without imposing any restriction on how quickly the job-finding rate saturates in tight markets, which is governed by p_R alone. In the single-p model one parameter had to serve both purposes; the two-plateau assigns each cap to the parameter that directly governs its regime.

Logit-elasticity characterization. For either branch, $\varepsilon(\theta) = 1/(1 + (\theta/\theta_c)^p)$. Direct computation gives

$$\frac{\varepsilon}{1-\varepsilon} = \frac{1/(1 + (\theta/\theta_c)^p)}{(\theta/\theta_c)^p/(1 + (\theta/\theta_c)^p)} = \left(\frac{\theta_c}{\theta}\right)^p,$$

and therefore

$$\operatorname{logit} \varepsilon(\theta) \;\equiv\; \log \frac{\varepsilon}{1-\varepsilon} \;=\; -p \log\left(\frac{\theta}{\theta_c}\right).$$

The logit of the matching elasticity is linear in $\log \theta$, with slope $-p$ and intercept zero at θ_c (where $\varepsilon = 1/2$). For the single-p SMCD, a single line with slope $-p$ describes the entire elasticity profile. For the 2P-SMCD, the profile is piecewise linear: slope $-p_L$ for $\theta < \theta_c$ and slope $-p_R$ for $\theta > \theta_c$, joined at $(\log \theta_c,\, 0)$. The two slopes are separately identified by data on each side of the crossover. Notably, θ_c plays a dual role: it sets the logit-intercepts ($p_{L,R} \log \theta_c$) and marks the kink where curvature shifts between regimes (Figure 7).

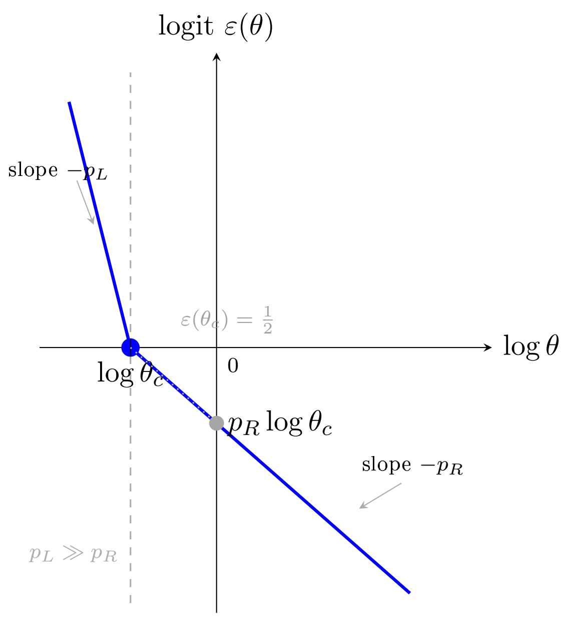

Figure 7: Two-plateau SMCD in $(\log \theta, \operatorname{logit} \varepsilon)$ space, calibrated to pre-2008 estimates $\theta_c = 0.417$ and $p_R = 0.88$; $p_L = 4$ is used for legibility (pre-2008 estimate $p_L = 26.2$ is near-vertical). The dotted line extends the right segment to the y-axis, marking the intercept $p_R \log \theta_c$.

This characterization also clarifies the relationship to the CES discussed in Section 4: an asymmetric-η single-p CES produces a straight line in $(\log \theta,\, \operatorname{logit} \varepsilon)$ space and is therefore a reparameterization of the same single-p family, while the 2P-SMCD produces a piecewise-linear profile that no single-p specification can replicate.

Proposition 2 (Properties of the 2P-SMCD)**.** *The two-plateau matching function defined in* (6) *satisfies the following four properties.*

(i) Constant returns to scale. $m(\lambda u, \lambda v) = \lambda m(u, v)$ *for all* $\lambda > 0$.

(ii) Smooth and concave. $f(\theta)$ *is* C^1 *on* $(0, \infty)$, *smooth on each branch, and globally concave.*

(iii) Joint calibration. *For any feasible triple* $(f^*, \varepsilon^*, \theta^*)$ *with* $f^* \leq \min\{1, \theta^*\}$ *and* $\varepsilon^* \neq 1/2$, *there exist parameters* (f_c, θ_c, p_R) *— or* (f_c, θ_c, p_L) *— that achieve* $f(\theta^*) = f^*$ *and* $\varepsilon(\theta^*) = \varepsilon^*$.

(iv) Bounded probabilities. $f(\theta) \leq s_R$ *and* $q(\theta) \leq q_{\max}$ *for all* $\theta > 0$. *Setting* $s_R \leq 1$ *and* $q_{\max} \leq 1$ *ensures* $f, q \in (0, 1)$.

Proof. *(i)* $\theta = v/u$ is homogeneous of degree zero, so $m(\lambda u, \lambda v) = \lambda u \cdot f(\lambda v/\lambda u) = \lambda m(u, v)$. *(ii)* On each open branch, $\varepsilon(\theta) = 1/(1 + (\theta/\theta_c)^p)$ is strictly decreasing in θ, so the concavity criterion established in Section 4 is satisfied branch-by-branch. At θ_c both branches yield $f(\theta_c) = f_c$ and $\varepsilon(\theta_c) = 1/2$, so f is C^1. The curvature from the left is $\kappa^- = -p_L/4$ and from the right $\kappa^+ = -p_R/4$; both are non-positive, confirming global concavity. *(iii)* For $\varepsilon^* < 1/2$ (calibration point on the right branch, $\theta^* > \theta_c$): the elasticity condition gives $(\theta^*/\theta_c)^{p_R} = (1 - \varepsilon^*)/\varepsilon^*$, which pins p_R given θ_c; the level condition then gives $f_c = f^*/(2(1 - \varepsilon^*))^{1/p_R}$. The case $\varepsilon^* > 1/2$ is symmetric on the left branch. *(iv)* On the right branch, $f(\theta) = f_c(\theta/\theta_c)[2/(1 + (\theta/\theta_c)^{p_R})]^{1/p_R} \leq f_c \cdot 2^{1/p_R} = s_R$. The left branch and the vacancy-filling bound are analogous. □

Local Cobb-Douglas behavior follows directly from (ii) and (iii). Since $\varepsilon(\theta)$ is continuous everywhere, including at θ_c, then at any calibration point θ^* where $\varepsilon(\theta^*) = \varepsilon^*$, this yields $d \log m = (1 - \varepsilon^*) d \log u + \varepsilon^* \, d \log v$ to first order. Thus, the 2P-SMCD is locally Cobb-Douglas at θ^*. The second-order deviation is $-\frac{1}{2} p_{\text{branch}} \, \varepsilon^*(1 - \varepsilon^*)\Delta^2$, where p_{branch} is the curvature of the branch containing θ^*.

Minimizing distance to Cobb-Douglas. In the single-p SMCD, the minimum-curvature solution required $p = p_{\min} = \max\{p_s, p_{sa}\}$, with both bound constraints feeding into a single parameter. For the calibration $(f^*, \varepsilon^*, \theta^*) = (0.70, 0.40, 0.29)$, the binding constraint was

$p_{sa} \approx 2.21$: the vacancy-filling bound forced the function to deviate substantially from Cobb-Douglas at the calibration point even though the calibration point lies in the worker-limited regime where that bound is not directly active.

The 2P-SMCD decouples the two constraints. Curvature at θ^* on the right branch is $|\kappa(\theta^*)| = p_R \cdot \varepsilon^*(1 - \varepsilon^*)$, governed by p_R alone. The right-cap constraint $s_R \leq 1$ requires $p_R \geq p_s = \log(1 - \varepsilon^*)/\log(f^*) \approx 0.37$, but p_R and the left cap are not fully independent: for a given p_R, the calibration pins $\theta_c = \theta^*(\varepsilon^*/(1 - \varepsilon^*))^{1/p_R}$, which in turn determines f_c/θ_c and hence what p_L must satisfy to keep $q_{\max} \leq 1$. As p_R falls toward p_s, $\theta_c \to 0$ and f_c/θ_c becomes large, eventually making the left cap infeasible for any finite p_L.

For the US calibration, the minimum feasible p_R — the smallest value for which there exists a valid finite p_L satisfying $q_{\max} \leq 1$ — is approximately 0.97, achieved in the limit $p_L \to \infty$. At $p_R \approx 0.97$, curvature at the calibration point is $|\kappa| \approx 0.20$ compared with 0.46 for the single-p SMCD at $p_{\min}$ — a reduction by roughly a factor of two. The 2P-SMCD at its minimum-curvature configuration is substantially closer to Cobb-Douglas in the range of tightness relevant for steady-state calibration. The single-p model's higher minimum curvature reflects the cost of using one parameter to satisfy two constraints from opposite ends of the tightness distribution; the 2P-SMCD reduces this cost by assigning each constraint primarily to the parameter governing its own regime.

Figure 8 plots $f(\theta)$ and $q(\theta)$ for all four specifications — CD, HRW, single-p SMCD, and 2P-SMCD — calibrated to $(\theta^*, f^*, \varepsilon^*) = (0.70, 0.40, 0.29)$ where feasible. The 2P-SMCD is shown at its minimum-curvature configuration ($p_R \approx 0.97$, p_L set to satisfy $q_{\max} = 1$). Two features stand out. First, the 2P-SMCD job-finding rate reaches a ceiling of $s_R \approx 0.569$, compared with $s \approx 0.467$ for the single-p SMCD — a 22% improvement. Second, both SMCD variants and HRW respect the probability bounds throughout the range of empirically observed tightness, while CD's vacancy-filling rate exceeds one for $\theta < 0.32$.

The ceiling $s_R \approx 0.569$ the theoretical maximum achievable under this calibration with both probability bounds enforced. The right-cap constraint $s_R \leq 1$ requires $p_R \geq p_s \approx 0.37$, at which $s_R = 1$ — but the left-cap constraint forces $p_R \geq 0.973$, since for smaller p_R the crossover tightness θ_c collapses toward zero and f_c/θ_c exceeds one, making the vacancy-filling bound infeasible for any finite p_L. At the minimum feasible $p_R \approx 0.973$, the right cap evaluates to $s_R = f^*/(1 - \varepsilon^*)^{1/p_R} \approx 0.569$. There is no parameterization of the 2P-SMCD that simultaneously hits (f^*, ε^*) at the sample mean tightness and achieves $s_R > 0.569$ while respecting both bounds.

This ceiling is specific to the theoretical calibration at sample means. In the empirical horse race of Section 5.4, the 2P-SMCD is estimated by nonlinear least squares without any constraint to hit sample means: the parameters $(f_c, \theta_c, p_L, p_R)$ are chosen to minimize residuals subject only to the probability bounds. The estimated s_R will therefore reflect the full tightness distribution in the data and is free to exceed 0.569.

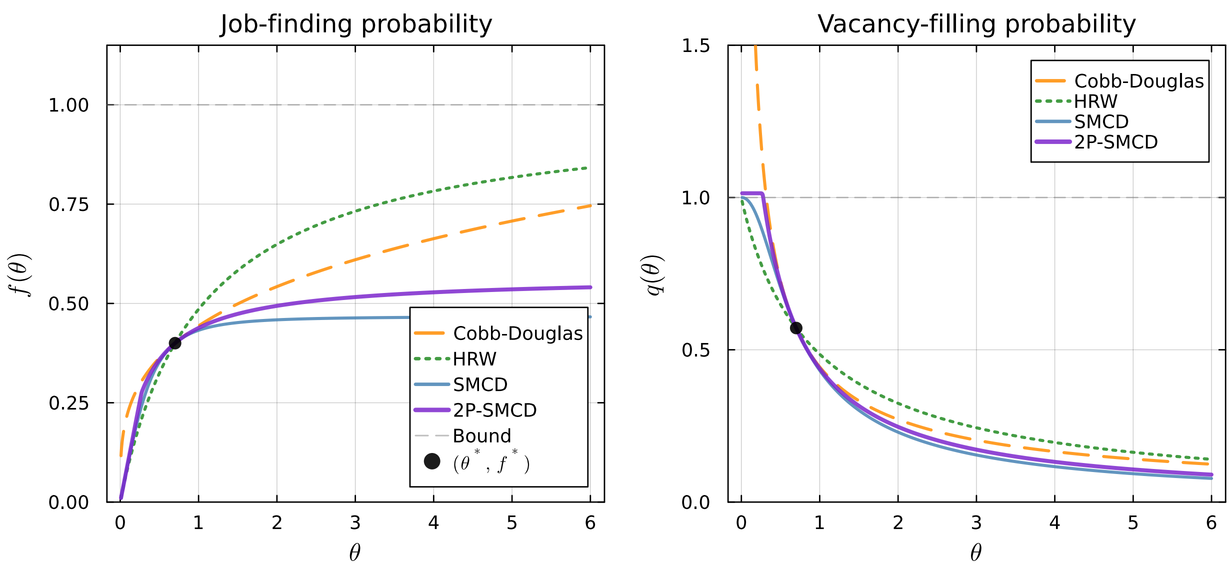

Figure 8: Job-finding probability $f(\theta)$ (left) and vacancy-filling probability $q(\theta)$ (right) for CD (dashed), HRW (dotted), single-p SMCD (solid, light), and 2P-SMCD (solid, dark). All specifications calibrated to $(\theta^*, f^*, \varepsilon^*) = (0.70, 0.40, 0.29)$ where feasible; HRW matches f^* at θ^* only. Single-p SMCD: $p = p_{\min} \approx 2.21$, $s \approx 0.467$, $a \approx 2.14$. 2P-SMCD: $p_R \approx 0.97$ (minimum curvature at θ^*), p_L at left-cap boundary. Filled circle marks the common calibration point $(\theta^*, f^*) = (0.70, 0.40)$.

5.4. Horse Race: Adding the Two-Plateau SMCD

Before comparing fit statistics it is useful to define the domain on which CD constitutes a valid probability model. The estimated function $f(\theta) = A\theta^{\varepsilon}$ satisfies $q(\theta) = A\theta^{\varepsilon-1} \leq 1$ only for $\theta \geq \theta_{\min} \equiv A^{1/(1-\varepsilon)}$, and $f(\theta) \leq 1$ only for $\theta \leq \theta_{\max} \equiv A^{-1/\varepsilon}$. Pre-break, $(\theta_{\min}, \theta_{\max}) = (0.36, 8.85)$; post-break, $(0.20, 57.79)$. The upper bound $\theta_{\max}$ is never reached in either subsample, so the binding constraint is always $\theta_{\min}$: CD violates $q \leq 1$ in 12.1% of pre-break months and 8.4% of post-break months. The lower bound shifts left after the Lehman collapse because the decline in matching efficiency A widens the range of tightness for which CD overestimates the vacancy-filling rate.

The 2P-SMCD is estimated by nonlinear least squares over $(\theta_c, f_c, p_L, p_R)$, subject to the bounds $s_R \equiv f_c \cdot 2^{1/p_R} \leq 1$ and $q_{\max} \equiv (f_c/\theta_c) \cdot 2^{1/p_L} \leq 1$; the $q > 1$ violation rate is therefore zero by construction. Table 5 reports standard fit statistics alongside a restricted RMSE*, computed exclusively on months with $\theta > \theta_{\min}$ — the region where CD also predicts valid probabilities and the comparison is on equal terms. Table 6 reports the estimated parameters and derived caps.

Table 5: Fit statistics for the 2P-SMCD and Cobb-Douglas, pre- and post-October 2008. RMSE* is computed on the subsample with $\theta > \theta_{\min}$ ($\theta_{\min} = 0.36$ pre-break, 0.21 post-break).

Period	Model	Par	RMSE	RMSE*	MAE	q>1 (%)
Pre-2008	CD	2	0.0442	0.0452	0.0349	12.0
	2P-SMCD	4	0.0550	0.0465	0.0433	0.0
Post-2008	CD	2	0.0284	0.0292	0.0224	9.4
	2P-SMCD	4	0.0344	0.0300	0.0260	0.0

$n = 693$ pre-break, $n = 203$ post-break. RMSE* computed on months with $\theta > \theta_{\min}$, where CD also satisfies $q \leq 1$. CD RMSE* exceeds its full-sample RMSE because the excluded low-θ months are recession observations where CD's unconstrained extrapolation fits well.

Table 6: Estimated 2P-SMCD parameters, derived tail caps, and average tightness elasticity, pre- and post-October 2008.

Period	$\hat{\theta}_c$	$\hat{f}_c$	$\hat{p}_L$	$\hat{p}_R$	$\hat{s}_R$	$\hat{q}_{\max}$	$\bar{\varepsilon}$
Pre-2008	0.417	0.361	26.18	0.882	0.791	0.888	0.493
Post-2008	0.325	0.231	10.04	1.439	0.373	0.761	0.381

$\hat{s}_R = \hat{f}_c \cdot 2^{1/\hat{p}_R}$; $\hat{q}_{\max} = (\hat{f}_c/\hat{\theta}_c) \cdot 2^{1/\hat{p}_L}$. $\bar{\varepsilon} = T^{-1} \sum_t \varepsilon(\theta_t; \hat{\psi})$: average tightness elasticity over the subsample, from the piecewise analytical formula (Section 5.3).

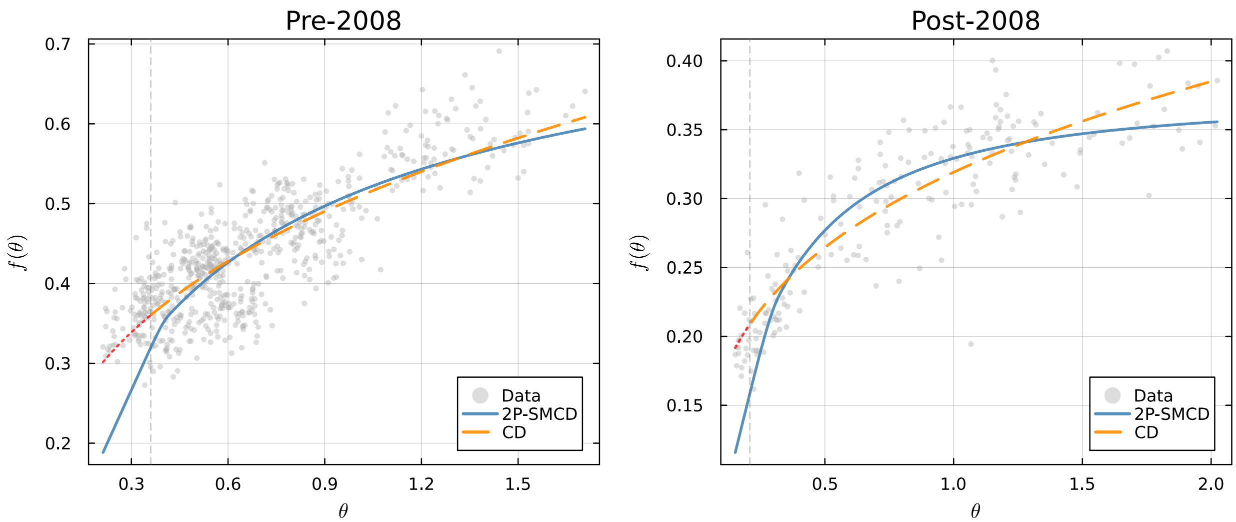

Figure 9: Job-finding rate $f(\theta)$: scatterplot of empirical (θ, f) values with fitted 2P-SMCD (solid blue) and CD (dashed orange) for the pre-break (left) and post-break (right) subsamples. Parameters estimated separately on each subsample. For CD, the red dotted segment marks the region $\theta < \theta_{\min}$ where the fitted vacancy-filling rate exceeds one ($\theta_{\min} = 0.36$ pre-break, 0.20 post-break); the vertical gray line indicates $\theta_{\min}$.

Tables 5–6 and Figure 9 deliver three findings.

First, CD achieves lower full-sample RMSE in both subsamples, by the same margin documented for the single-p SMCD in Section **??**. Adding a fourth parameter does not allow the 2P-SMCD to overturn this ranking when the full time series — including recession troughs where CD extrapolates into the $q > 1$ region — is used as the criterion.

Second, the restricted RMSE* reveals that virtually the entire RMSE advantage of CD is concentrated in the months where CD is not a valid probability model. Pre-break, the full RMSE gap (2P-SMCD minus CD) is $0.0550 - 0.0442 = 0.0108$; the restricted gap narrows to $0.0465 - 0.0452 = 0.0013$, recovering 88% of the advantage. Post-break, the analogous figures are a full gap of $0.0344 - 0.0284 = 0.0060$ and a restricted gap of $0.0300 - 0.0292 = 0.0008$, so 87% is accounted for by the $q > 1$ months. Within the admissible domain $\theta > \theta_{\min}$, the 2P-SMCD fits within 2.9% of CD's RMSE pre-break and 2.7% post-break. For the same reason, information criteria favor CD throughout and are not reported.

Third, the estimated parameters shed light on how the two-plateau structure fits the data. In both subsamples $\hat{p}_L$ is large — 26.2 pre-break, 10.0 post-break — indicating that the vacancy-filling probability rises sharply to its cap $\hat{q}_{\max}$ on the left branch. This reflects the empirical pattern that recession observations cluster near the left edge of the observed tightness distribution, requiring a steep left-branch curvature to bound q tightly without distorting the right-branch fit. In contrast, $\hat{p}_R < 2$ in both periods (0.88 and 1.44), indicating gradual saturation of the job-finding rate in tight markets. The asymmetry $\hat{p}_L \gg \hat{p}_R$ is precisely what the two-plateau extension was designed to accommodate: a single-p specification forces the same curvature on both branches and cannot simultaneously produce a sharp left-side bound and a smooth right-side saturation. The estimated right-cap $\hat{s}_R = 0.791$ pre-break substantially exceeds the theoretical minimum-curvature ceiling of 0.569 from Section 5.3, confirming that the NLS solution exploits the full distribution of tightness rather than being constrained to the sample-mean calibration.

Recommendation for practitioners.. The structural break carries a direct message for applied users. Both Cobb-Douglas and the 2P-SMCD cleanly separate the efficiency level from the curvature shape — A from ε in the former, f_c from (p_L, p_R) in the latter — yet neither specification is immune to the break: estimated parameters shift across sub-periods for all three functions. Researchers with a clear temporal focus should therefore use the corresponding sub-period estimates in Tables 5 and 6; the pre-2008 estimates are the natural baseline for steady-state calibrations of normal labor market dynamics, while the post-2008 estimates are appropriate for applications centered on the post-Lehman period. Full-sample estimates are defensible only if the structural model endogenously explains the efficiency decline — through search intensity, worker composition, or sectoral mismatch in the spirit of Şahin et al. (2014) or Barnichon and Figura (2015) — since matching technology parameters would then be identified jointly with the structural model and require re-estimation. One advantage of 2P-SMCD that survives the break unconditionally is probability-bound compliance: regardless of which period's parameter vector is plugged in, the implied job-finding and vacancy-filling rates remain in $(0, 1)$ for all market tightness.

6. Discussion: The Continuous-Time Objection

A natural objection is that the probability-bound violations motivating the SMCD are an artifact of discrete time and disappear in continuous-time or sufficiently high-frequency formulations. In a standard continuous-time DMP model, the matching function $m(u, v)$ specifies a *flow rate* of matches per unit time, and the associated job-finding and vacancy-filling objects $f(\theta) = m(1, \theta)$ and $q(\theta) = m(1, \theta)/\theta$ are hazard rates rather than probabilities. Hazard rates are non-negative but not bounded above by one: a Poisson matching process can deliver $f(\theta) > 1$ per unit time without logical contradiction, since the implied per-period probability $1 - e^{-f(\theta)\Delta t}$ is in $(0, 1)$ for any finite $\Delta t > 0$. On this view, the CD function $f(\theta) = A\theta^{1-\alpha}$ is a valid hazard-rate specification for all θ, and the violations documented in Table 1 are an artifact of applying a continuous-time object directly as a monthly probability.

This objection is technically correct as stated, but it overstates what continuous time resolves. The first difficulty is empirical. Matching functions are estimated and calibrated by confronting model-implied objects with data observed at monthly frequency. The standard practice — followed in every quantitative DMP paper from Shimer (2005) to Hagedorn and Manovskii (2008) — is to calibrate f^* to a target monthly job-finding probability and ε^* to a monthly matching elasticity. A model that predicts $f(\theta) > 1$ at some empirically realized θ is misspecified at the calibration frequency, irrespective of whether it is written in continuous or discrete time. The bound violations in 9.9% of our sample are a property of the CD function applied to the monthly data the model is designed to match, not an artifact of the estimation procedure.

The second and more fundamental difficulty cuts against the continuous-time objection from its own premises. The correct link between a continuous-time matching hazard $\lambda_f(\theta)$ and the monthly job-finding probability is $f_t = 1 - e^{-\lambda_f(\theta)}$, not the first-order approximation $f_t \approx \lambda_f(\theta)$ implicit when CD is applied directly. Applied to the CD hazard $\lambda_f(\theta) = A\theta^{1-\alpha}$, this exact mapping yields

$$f_t(\theta) = 1 - e^{-A\theta^{1-\alpha}}, \tag{7}$$

which lies in $(0, 1)$ for all θ and has a variable matching elasticity $\varepsilon(\theta) = A(1-\alpha)\theta^{1-\alpha}/(e^{A\theta^{1-\alpha}} - 1)$ that declines with θ. The practitioner who applies this mapping consistently has abandoned the log-linear matching function in favor of a bounded, variable-elasticity specification — one that shares the core structural features motivating the SMCD. This naturally raises whether (7) itself resolves the problem and renders the SMCD unnecessary. It does not, for a specific reason: (7) bounds the job-finding probability f but not the vacancy-filling probability q. Since $q(\theta) = f(\theta)/\theta \approx A\theta^{-\alpha} \to \infty$ as $\theta \to 0$, the exponential specification still predicts $q > 1$ below a threshold that, under the U.S. calibration, includes the deepest recession observations. At the COVID trough of $\theta \approx 0.15$, for instance, (7) with $\hat{A} = 0.45$ and $\hat{\alpha} = 0.71$ implies $q \approx 1.53$. The SMCD bounds $f \leq s$ and $q \leq sa$ simultaneously through independent parameters, which is precisely what neither CD nor its exponential counterpart achieves.

A third consideration is that monthly frequency is itself the natural benchmark for business-cycle analysis, not a coarse approximation awaiting refinement. Business cycles are conventionally defined as fluctuations with periodicities of 6 to 32 quarters — the pass band of Baxter and King (1999) — or equivalently 18 to 96 months as used by Stock and Watson (1999) and targeted by the regression filter of Hamilton (2018). Monthly sits at the *high-frequency* end of this range. To treat $\Delta t = 1$ month as approximately zero, the per-period probabilities would need to be small; at the sample mean $f^* \approx 0.40$, however, the continuous-time approximation $f \approx \lambda \Delta t$ understates the exact probability by roughly 28%, a non-trivial distortion at the calibration point itself. The underlying data generation is also inherently monthly: the Bureau of Labor Statistics conducts the Current Population Survey in the week containing the 12th of each month, and JOLTS records end-of-month vacancy stocks, reflecting the institutional horizon of payroll cycles and hiring rounds rather than an arbitrary computational choice.

Regarding the asymptotic behavior $q(\theta) \to \infty$ as $\theta \to 0$, we note that this is not indefensible in a continuous-time framework. An infinite vacancy-filling hazard in the limit $\theta \to 0$ simply means vacancies are filled instantaneously when labor is superabundant — the expected wait time $1/q(\theta) \to 0$. This is the competitive limit in which search frictions vanish on the demand side, not a logical contradiction. The SMCD's cap $q(\theta) \leq sa$ is therefore a modeling choice, encoding the assumption that matching has an irreducible minimum duration even when workers are abundant — that screening, negotiation, and onboarding consume real time. We regard this as empirically plausible: vacancy-filling rates in JOLTS show no tendency to diverge during high-unemployment quarters. But we do not claim that the absence of a cap is indefensible; the case for the SMCD rests primarily on the two preceding arguments, with the cap providing additional structural discipline against extrapolation beyond the observed tightness range.

7. Conclusion

The Soft-Min Cobb-Douglas (SMCD) matching function resolves a fundamental tension in the quantitative labor market toolkit. Cobb-Douglas fits data well but routinely predicts impossible vacancy-filling probabilities in recessions; HRW respects probability bounds but cannot jointly calibrate to target job-finding rates and elasticities. The SMCD satisfies all four desiderata simultaneously — constant returns, smoothness and concavity, exact joint calibration, and bounded matching probabilities — with three parameters and closed-form calibration formulas.

The empirical results are sharp. A QLR sup-Wald test precisely identifies a structural break at October 2008: a 36% persistent decline in matching efficiency with only minor change in tightness elasticity. Across both subsamples, the single-p SMCD strictly dominates HRW (RMSE ratios of 1.19 and 1.83), confirming that probability bounds alone are insufficient without the flexibility to jointly calibrate level and elasticity. CD's lower unconditional RMSE is achieved overwhelmingly in the 12% and 9% of months where it predicts $q > 1$; once those observations are excluded, roughly 88% of CD's fit advantage disappears. The two-plateau extension, which adds one parameter to allow independent curvature in each regime, eliminates CD's remaining fit advantage entirely.

Probability bounds carry efficiency implications that go beyond fit. SMCD's matching elasticity $\varepsilon(\theta)$ is strictly decreasing in tightness — a structural consequence of bounded probabilities — so the unemployment elasticity $\eta_u(\theta) = 1 - \varepsilon(\theta)$ is strictly increasing. Through the Hosios condition, this produces a state-dependent efficiency wedge: as tightness falls in a recession, workers' bargaining power increasingly exceeds their marginal matching contribution, amplifying unemployment beyond its efficient level. Cobb-Douglas, with its constant elasticity, cannot speak to this fact. HRW qualitatively replicates the countercyclical pattern but mismeasures its level: using pre-break estimates, HRW implies $\eta_u \approx 0.39$ at the subsample mean tightness against SMCD's 0.60, a gap of 0.21 in the efficient Nash bargaining weight. Any welfare assessment of labor market policy built on HRW is therefore misspecified by construction.

Several extensions are natural. Embedding SMCD in a calibrated DMP model would quantify the cyclical implications of state-dependent efficiency for unemployment and vacancy dynamics. The post-2008 tightness range — from the COVID trough ($\theta \approx 0.1$) to the post-pandemic boom ($\theta > 2$) — spans both limiting regimes and provides an especially demanding test.

The 2P-SMCD is a drop-in replacement for CD and HRW: concave, smooth, analytically tractable, and with good global fit. It respects probability bounds everywhere, not merely in normal times — a discipline that matters precisely when it is needed most.

Appendix A. Additional Figures

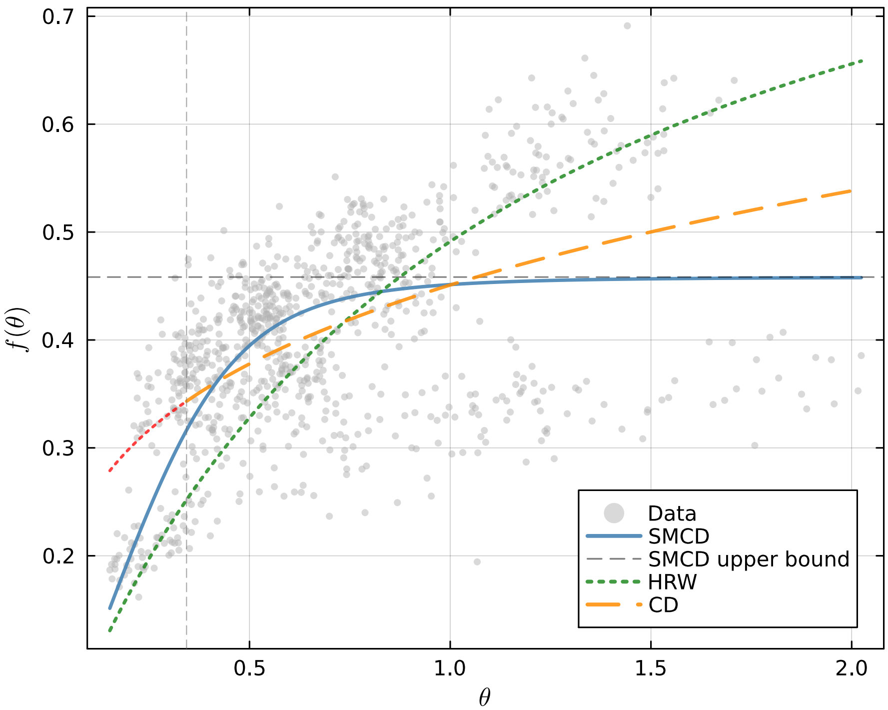

Figure A.10: Full-sample horse race: job-finding rate $f(\theta)$ for SMCD (solid), Cobb-Douglas (dashed), and HRW (dotted), estimated on the complete 1951–2025 sample. Grey scatter: data. The dashed horizontal line marks the SMCD upper bound s.

References

Anderson, P.M., Burgess, S.M., 2000. Empirical matching functions: Estimation and interpretation using state-level data. Review of Economics and Statistics 82, 93–102. doi:10.1162/003465300558597.

Barnichon, R., 2010. Building a composite help-wanted index. Economics Letters 109, 175–178.

Barnichon, R., Figura, A., 2015. Labor market heterogeneity and the aggregate matching function. American Economic Journal: Macroeconomics 7, 222–249.

Baxter, M., King, R.G., 1999. Measuring business cycles: Approximate band-pass filters for economic time series. Review of Economics and Statistics 81, 575–593. doi:10.1162/003465399558454.

Borowczyk-Martins, D., Jolivet, G., Postel-Vinay, F., 2013. Accounting for endogeneity in matching function estimation. Review of Economic Dynamics 16, 440–451. doi:10.1016/j.red.2012.12.001.

Diamond, P.A., 1982. Aggregate demand management in search equilibrium. Journal of political Economy 90, 881–894.

Elsby, M.W.L., Michaels, R., Ratner, D., 2015. The Beveridge curve: A survey. Journal of Economic Literature 53, 571–630.

den Haan, W.J., Ramey, G., Watson, J., 2000. Job destruction and propagation of shocks. American Economic Review 90, 482–498.

Hagedorn, M., Manovskii, I., 2008. The cyclical behavior of equilibrium unemployment and vacancies revisited. American Economic Review 98, 1692–1706. URL: http://www.aeaweb.org/articles?id=10.1257/aer.98.4.1692, doi:10.1257/aer.98.4.1692.

Hamilton, J.D., 2018. Why you should never use the Hodrick-Prescott filter. Review of Economics and Statistics 100, 831–843.

Hosios, A.J., 1990. On the efficiency of matching and related models of search and unemployment. The Review of Economic Studies 57, 279–298. doi:10.2307/2297382.

Lange, F., Papageorgiou, T., 2020. Beyond Cobb-Douglas: Flexibly Estimating Matching Functions with Unobserved Matching Efficiency. Working Paper 26972. National Bureau of Economic Research. doi:10.3386/w26972.

Lubik, T.A., 2013. The Shifting and Twisting Beveridge Curve: An Aggregate Perspective. Working Paper 13-16. Federal Reserve Bank of Richmond. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2335720.

Mortensen, D.T., Pissarides, C.A., 1994. Job creation and job destruction in the theory of unemployment. The Review of Economic Studies 61, 397–415.

Petrongolo, B., Pissarides, C.A., 2001. Looking into the black box: A survey of the matching function. Journal of Economic literature 39, 390–431.

Petrosky-Nadeau, N., Wasmer, E., 2017. Labor, Credit, and Goods Markets: The Macroeconomics of Search and Unemployment. MIT Press.

Pissarides, C.A., 1985. Short-run equilibrium dynamics of unemployment, vacancies, and real wages. American Economic Review 75, 676–690.

Şahin, A., Song, J., Topa, G., Violante, G.L., 2014. Mismatch unemployment. American Economic Review 104, 3529–3564.

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

Shimer, R., 2012. Reassessing the ins and outs of unemployment. Review of Economic Dynamics 15, 127–148.

Stock, J.H., Watson, M.W., 1999. Business cycle fluctuations in U.S. macroeconomic time series, in: Taylor, J.B., Woodford, M. (Eds.), Handbook of Macroeconomics. Elsevier. volume 1, pp. 3–64.