Bank of England

When vacancies become cheap: vacancy posting costs and equilibrium labour market tightness

Staff Working Paper No. 1,215
November 2026

Michal Stelmach, James Kensett and Philip Schnattinger

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or any of its committees, or to state Bank of England policy.



Bank of England

Staff Working Paper No. 1,215

When vacancies become cheap: vacancy posting costs and equilibrium labour market tightness

Michal Stelmach,[1] James Kensett[2] and Philip Schnattinger[3]

Abstract

Labour market tightness is a key determinant of wage pressure and an important input into monetary policy decisions. A commonly used measure of labour market tightness is the vacancy-to-unemployment ratio (V/U), but this ratio has trended upwards over recent decades. This raises an important measurement challenge: if the equilibrium level of vacancies has changed over time, movements in the raw V/U ratio may provide a misleading signal about labour market slack. This paper argues that falling vacancy posting costs are an important part of the explanation for the secular rise in vacancies observed in the UK labour market. We construct a new long-run measure of vacancy posting costs using data on recruitment advertising expenditure spanning the period since 1982. The resulting series documents a substantial decline in real vacancy posting costs following the transition from print-based recruitment to online hiring platforms. Incorporating this measure into estimates of equilibrium vacancy creation suggests that a substantial share of the increase in vacancies reflects structural rather than cyclical developments, materially affecting estimates of equilibrium labour market tightness.

Key words: Vacancies, labour market tightness, vacancy posting costs, recruitment advertising, search and matching, Beveridge curve.

JEL classification: E24, J23, J63, J64.

---

(1) Bank of England. Email: michal.stelmach@bankofengland.co.uk
(2) Bank of England. Email: james.kensett@bankofengland.co.uk
(3) Bank of England. Email: philip.schnattinger@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. We are indebted to Suzy Young and Chelsea Nobre at WARC for sharing the data on recruitment advertising expenditure, and to James Cockett at the CIPD for sharing the data on recruitment costs. We also thank Mike Elsby for helpful comments and suggestions. Any remaining errors are our own.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email: enquiries@bankofengland.co.uk

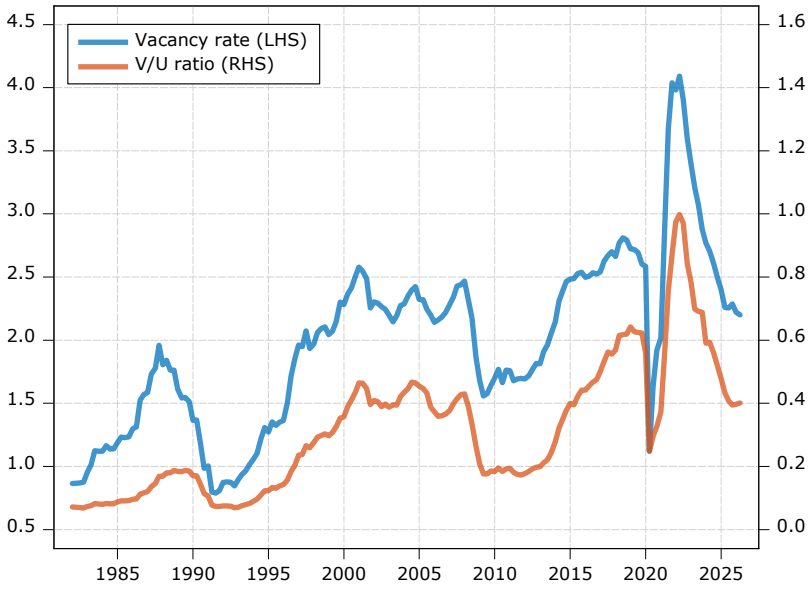
©2026 Bank of England
ISSN 1749-9135 (on-line)

1 Introduction

Labour market tightness plays a central role in the transmission of monetary policy. Tight labour markets are typically associated with greater competition for workers, stronger wage growth and, ultimately, domestic inflationary pressure. As a result, policymakers devote considerable attention to assessing the degree of slack in the labour market and its implications for the outlook for pay and inflation. A key indicator of labour market tightness is the ratio of vacancies to unemployment (V/U), which summarises the balance between firms' demand for labour and the availability of workers.

A challenge, however, is that both the vacancy rate and the vacancy-to-unemployment ratio have exhibited a pronounced upward trend over the past four decades (Figure 1). Whereas vacancies once represented a relatively scarce signal of labour demand, vacancy rates today are substantially higher than those observed in earlier decades, even during periods generally considered to be characterised by balanced labour market conditions. This raises an important measurement question. If the equilibrium level of vacancies has itself changed over time, comparisons against historical averages or previous cycles may provide a misleading signal about the true degree of labour market tightness.

Figure 1: The vacancy rate and the V/U ratio



Sources: Office for National Statistics; Bank of England (A millennium of macroeconomic data).

The existing literature offers several explanations for long-run changes in vacancy creation. Search-and-matching models emphasise the role of productivity, matching efficiency and recruitment costs in determining firms' willingness to create vacancies. At the same time, technological changes have transformed recruitment practices. Over the past two decades, hiring has increasingly shifted away from traditional print advertising and towards low-cost online platforms. While a growing literature has examined the effects of online recruitment technologies and recruiting intensity, relatively little attention has been paid to the implications of falling vacancy posting costs for the measurement of

equilibrium labour market tightness.

This paper argues that changes in vacancy posting costs are an important part of the explanation for the secular rise in vacancies observed in the UK labour market. We construct a new long-run measure of vacancy posting costs using data on recruitment advertising expenditure, spanning the period from 1982 onwards. The resulting series documents a substantial decline in real vacancy posting costs, particularly following the transition from print-based recruitment advertising to online job boards and digital recruitment platforms. We then use this information to estimate the equilibrium level of vacancies.

The paper makes three contributions. First, it provides, to our knowledge, the first long-run time series measure of vacancy posting costs for the UK. Second, it documents the extent to which recruitment technology has reduced the cost of vacancy creation over time. Third, it demonstrates that accounting for these structural changes materially alters the estimated equilibrium level of vacancies and, consequently, the interpretation of labour market tightness. Our results suggest that a substantial share of the long-run increase in vacancies reflects structural rather than cyclical forces, implying that the raw vacancy-to-unemployment ratio can provide a misleading signal of labour market conditions when viewed over a long horizon.

The remainder of this paper is organised as follows. Section 2 reviews the relevant literature. Section 3 outlines the theoretical framework linking vacancy posting costs to equilibrium labour market tightness. Section 4 introduces the vacancy cost dataset and describes the construction of the vacancy cost index. Section 5 presents the empirical methodology and estimates of equilibrium vacancies. Section 6 examines the implications of falling vacancy posting costs in a calibrated Diamond-Mortensen-Pissarides model. Section 7 concludes.

2 Related literature

2.1 Equilibrium labour market tightness and structural shifts in vacancies

The search-and-matching literature views labour market tightness as an equilibrium outcome determined by firms' incentives to create vacancies and workers' incentives to search for jobs. A central implication of this framework is that equilibrium labour market tightness need not be constant over time. In the Diamond-Mortensen-Pissarides (DMP) model, changes in labour productivity, matching efficiency and vacancy creation costs can all alter firms' incentives to post vacancies and therefore affect the equilibrium relationship between vacancies and unemployment (Mortensen and Pissarides, 1994; Pissarides, 2000). Hagedorn and Manovskii (2008) further show that seemingly modest changes in hiring incentives can generate sizeable movements in equilibrium vacancy creation.

While much of the literature on shifts in labour market tightness has focused on changes in matching efficiency and labour market flows, comparatively little attention has been paid to the role of recruitment costs. This omission is potentially important given the substantial technological changes that have transformed recruitment practices over recent decades.

Recent work has also questioned whether raw vacancy measures provide a stable indicator of labour market tightness over time. For example, Mongey and Horwich (2023) argue that the upward trend in US vacancies can distort assessments of labour market conditions and propose adjusting vacancy measures for long-run trends. While they do not take a definitive view on the source of this drift, they note that digital technologies may have reduced the cost of job posting, recruitment and candidate evaluation, contributing to a gradual increase in measured vacancies. This interpretation is closely related to the mechanism explored in this paper.

2.2 Digitalisation and the changing cost of recruitment

The emergence of internet-based recruitment fundamentally altered the process through which firms advertise vacancies and search for workers. Autor (2001) argued that online job search substantially reduced information frictions relative to traditional newspaper advertising, lowering search costs for both employers and job seekers. Subsequent work by Kuhn and Skuterud (2004) and Kuhn and Mansour (2014) documented how online search technologies increasingly displaced traditional job-search methods and became an important component of labour market matching.

Technological progress has also transformed the economics of vacancy advertising. Online job boards allow firms to reach a much larger pool of candidates at a fraction of the cost of print advertising, while platform-based recruitment has further reduced the marginal cost of posting vacancies. Evidence from the United States suggests that these changes have had economically meaningful effects on labour market outcomes. For example, studies exploiting variation in online advertising prices have found that changes in the cost of posting job adverts can affect vacancy creation, while Djourelova et al. (2025) show that the emergence of Craigslist fundamentally disrupted traditional newspaper markets that had historically relied on classified advertising revenues.

At the same time, the increasing use of online vacancy data has generated a growing literature on vacancy measurement. Cajner and Ratner (2016), for example, highlight both the opportunities and challenges associated with using online vacancy data as measures of labour demand. Despite this growing body of work, there has been relatively little attempt to measure how recruitment costs themselves have evolved over time and what implications this has for equilibrium labour market tightness.

2.3 Recruitment costs, recruiting intensity and vacancy duration

A related literature emphasises that firms adjust recruitment effort as well as vacancy creation. Recruitment costs encompass not only the direct cost of advertising vacancies but also expenditure on screening, interviewing and matching workers to jobs. Consequently, firms can respond to changing labour market conditions along several recruitment margins simultaneously.

Faberman (2020) argues that recruiting intensity is an important dimension of labour demand that cannot be inferred from vacancy data alone. Similarly, Davis et al. (2013) show that hiring outcomes depend both on the number of vacancies posted and the effort devoted to filling them. Gavazza et al. (2018) further demonstrate that firms can achieve desired employment targets through adjustments in recruiting effort rather than vacancy creation itself.

These findings are particularly relevant for understanding the implications of falling vacancy costs. Lower advertising costs may increase the observed stock of vacancies even if firms' desired employment changes only modestly, for example by reducing the urgency with which firms seek to fill positions or by encouraging vacancies to remain advertised for longer periods. As a result, observed vacancies may evolve differently from underlying labour demand.

2.4 Changes in recruitment behaviour

Alongside technological developments, there is increasing evidence that recruitment practices themselves have evolved. Firms increasingly use vacancies for purposes beyond the immediate filling of positions, including pipeline recruitment, talent pooling and gathering information about labour market conditions. Cappelli (2012) argues that changes in employer recruitment practices have contributed to the coexistence of large vacancy stocks alongside reports of hiring difficulties.

More recently, attention has focused on so-called 'ghost jobs' – vacancies that remain advertised despite the absence of an immediate hiring intention.[1] Using interview-review data from Glassdoor, Ng (2024) estimates that around 21% of job postings may exhibit characteristics consistent with ghost hiring, although the precise prevalence remains uncertain and depends on the classification methodology employed. The author further argues that the rise of ghost jobs may help explain the recent disconnect in the Beveridge curve and highlights falling online recruitment costs as a potential driver. More recently, Singer and Oktay (2025) argue that ghost postings can distort labour market signals and vacancy-based indicators by weakening the link between advertised vacancies and actual hiring activity.

Taken together, the literature suggests several channels through which technological change may have altered the relationship between vacancies and labour market tightness. Falling recruitment costs may directly increase vacancy creation, indirectly affect recruiting intensity and vacancy duration, and facilitate broader changes in recruitment behaviour. However, direct evidence on the evolution of vacancy posting costs remains scarce, largely reflecting the absence of long-run measures of such costs. This paper seeks to bridge that gap.

3 Theoretical framework

3.1 Vacancy creation in a search-and-matching framework

A central feature of search-and-matching models is that hiring is not instantaneous. Firms must first create vacancies and incur recruitment costs before a worker can be hired. Vacancy creation can therefore be viewed as an investment decision: firms compare the expected benefits of employing an additional worker with the costs of attracting and recruiting that worker.

In the Diamond-Mortensen-Pissarides (DMP) framework, the value of a filled job can be expressed

[1] 'Ghost jobs' encompass a range of behaviours with no widely accepted academic definition. Recent contributions distinguish between evergreen postings maintained to build candidate pipelines, market research postings used to gauge labour supply or salary expectations, vacancies posted for branding or compliance purposes, and postings that remain active despite having already been filled or withdrawn (Santhosh and Jennifer, 2025).

as:

$$J = \frac{x - w}{r + s} \tag{1}$$

where x denotes the marginal product of labour, w the wage, r the discount rate and s the separation rate. The value of a job therefore rises with labour productivity and falls with labour costs.

A firm that wishes to hire an additional worker must first post a vacancy. The value of an unfilled vacancy is:

$$V = q(\theta)J - \kappa \tag{2}$$

where $q(\theta)$ is the probability that the vacancy is successfully filled, θ denotes labour market tightness, and κ denotes the cost of posting a vacancy. The first term represents the expected benefit of finding a match, while the second term captures the cost of creating a vacancy.

Under free entry, firms continue to create vacancies until the expected value of posting a vacancy is driven to zero ($V = 0$). This implies the following job creation condition:

$$\kappa = q(\theta)J \tag{3}$$

The equilibrium level of labour market tightness therefore depends on both the expected value of employment and the cost of recruitment. Substituting the expression for J into the job creation condition gives:

$$\kappa = q(\theta)\frac{x - w}{r + s} \tag{4}$$

This expression illustrates the key trade-off facing firms. For a given wage, an increase in labour productivity raises the expected value of a filled job and therefore increases vacancy creation. Conversely, higher recruitment costs reduce the attractiveness of posting vacancies.

Since the probability of filling a vacancy decreases as labour market tightness rises ($q'(\theta) < 0$), the job creation condition implies a positive relationship between labour productivity and equilibrium tightness, and a negative relationship between vacancy posting costs and equilibrium tightness. In other words, for a given level of unemployment, a fall in vacancy posting costs leads firms to create more vacancies.

Wages respond to the same forces. With Nash bargaining, a worker bargaining weight η and a flow value of non-employment b, the wage is $w = \eta(x + \kappa\theta) + (1 - \eta)b$ (Appendix C.1), and the job creation condition becomes

$$\frac{\kappa}{q(\theta)} = \frac{(1 - \eta)(x - b) - \eta\kappa\theta}{r + s}. \tag{5}$$

Productivity therefore raises equilibrium tightness only to the extent that it outpaces b; if b and κ grow in line with productivity, balanced growth leaves tightness unchanged. A fall in posting costs relative to productivity raises tightness in either case. The empirical analysis accordingly conditions on productivity alongside posting costs (Section 5), and the model of Section 6 works in units of output per worker.

3.2 Structural determinants of equilibrium vacancies

The framework above suggests that the equilibrium level of vacancies may change over time in response to structural changes affecting either the value of employment or the cost of recruitment. These include:

- **Labour productivity growth.** Higher productivity raises the asset value of the job, boosting job creation. This holds to the extent that productivity outpaces the flow value of non-employment; balanced growth that lifts b in step leaves equilibrium tightness unchanged.

- **Changes in matching efficiency.** Improvements in the efficiency with which workers and firms are matched increase the likelihood that vacancies are successfully filled, raising the expected return from recruitment and supporting higher levels of vacancy creation.

- **Vacancy posting costs.** A reduction in the cost of advertising and maintaining vacancies directly lowers the cost of recruitment and therefore increases firms' willingness to create vacancies. In the presence of large and persistent declines in vacancy posting costs, the equilibrium level of vacancies may therefore rise even if underlying labour demand conditions remain unchanged.

While the first two drivers primarily reflect underlying labour market fundamentals, lower vacancy posting costs may also influence how vacancies are created and used by firms. As discussed in the following section, changes in recruitment behaviour may affect the relationship between observed vacancies and desired employment, with implications for the interpretation of vacancy-based measures of labour market tightness.

Taken together, these mechanisms suggest that changes in the vacancy-to-unemployment ratio need not solely reflect cyclical labour market conditions. Instead, structural changes in recruitment costs may shift the equilibrium level of vacancies over time. Motivated by this framework, the empirical analysis that follows focuses on labour productivity and vacancy posting costs as two key determinants of equilibrium vacancies.

3.3 Are vacancies comparable over time?

Given the discussion above, an important question is whether a measured vacancy today is directly comparable to one observed forty years ago. If the cost of creating and maintaining vacancies has changed materially over time, movements in vacancy stocks may reflect not only firms' demand for labour, but also changes in the way vacancies are used and advertised.

When vacancy advertising is costly, firms may be expected to post vacancies selectively and remove them promptly once recruitment efforts cease. By contrast, when advertising costs are very low, firms may be willing to advertise positions with a lower probability of being filled or keep vacancies active for longer periods. For example, low advertising costs may encourage firms to search more selectively ('fish') for suitable candidates, accepting that not every advertised position will be filled. Firms may also maintain active vacancies to signal growth or build a pool of potential applicants.

Such behaviour can increase measured vacancy stocks without implying a corresponding increase in desired employment.

An important implication is that falling vacancy posting costs may influence labour market tightness through two distinct channels. The first is the conventional channel captured by standard search-and-matching models: lower posting costs increase firms' incentives to create vacancies. The second is a behavioural channel, whereby lower posting costs alter the relationship between observed vacancies and desired employment. In this sense, a vacancy observed in a low-cost online recruitment environment may not be directly comparable to a vacancy observed in an earlier period when advertising was substantially more expensive.

This distinction may crucially have implications for measured matching efficiency. If firms increasingly post vacancies for exploratory recruitment, candidate pipelining or highly selective search, the number of hires generated by a given stock of vacancies may decline over time. Measured matching efficiency could therefore weaken even in the absence of any deterioration in the underlying matching technology. The model of Section 6 separates the two channels in its accounting. The conventional channel works through free entry, while the behavioural channel registers as lower matching efficiency per measured vacancy, which moves unemployment far more than tightness (Appendix C.4).

These considerations do not imply that vacancy measures have become uninformative. Rather, they help explain why the *equilibrium level* of vacancies may itself evolve over time and, therefore, why changes in vacancy stocks should not automatically be interpreted as proportional changes in underlying labour demand. The ONS Vacancy Survey counts positions for which employers are actively seeking recruits, not advertisements, and US evidence suggests that the link between postings and openings held up as recruitment moved online: Dalton et al. (2026) find openings per posting broadly stable as more establishments posted online.

4 Measuring vacancy posting costs

4.1 Data

To examine the role of recruitment costs in shaping labour market tightness, we combine data on aggregate recruitment advertising expenditure with estimates of the stock of vacancies. The Advertising Association/WARC Expenditure Report provides quarterly estimates of aggregate UK expenditure on recruitment advertising, disaggregated by medium, from 1982 onwards. To our knowledge, this is the longest available time series of recruitment advertising expenditure in the UK. The series combines information collected from media owners with estimates of expenditure for the wider market and records advertising revenue after negotiated discounts. It therefore measures aggregate expenditure on recruitment advertising rather than the price of an individual vacancy advertisement. We construct a measure of expenditure per vacancy in Section 4.3.

Our vacancy series also combines two sources. From 2001 onwards, we use the ONS Vacancy Survey, a monthly business survey measuring positions for which employers are actively seeking recruits from outside their organisation. Earlier observations are taken from *A Millennium of*

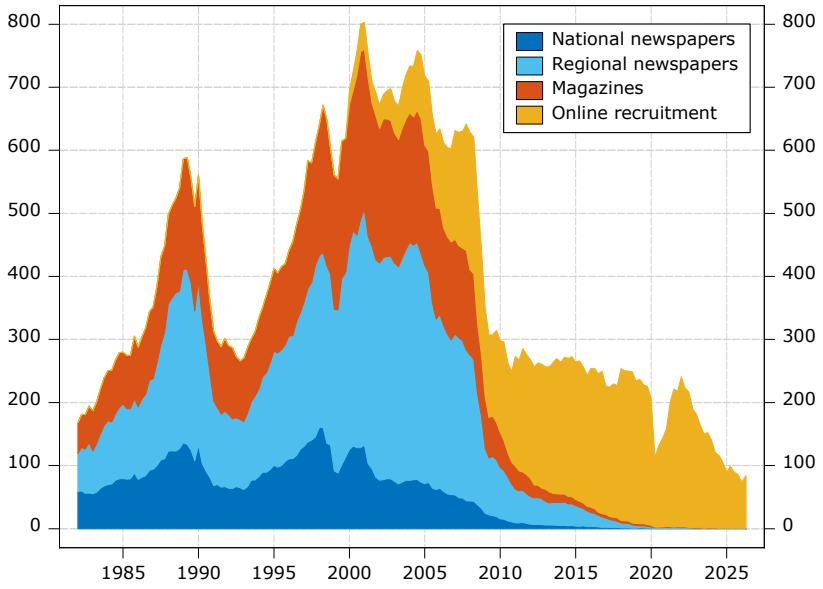
Macroeconomic Data (Thomas and Dimsdale, 2017) and measure vacancies notified to Jobcentres. The historical series is spliced to the ONS measure using a constant multiplier estimated over the overlap period. Appendix A.1 provides further details on the construction, coverage and limitations of both datasets and examines the sensitivity of the calibration to alternative splice assumptions.

4.2 Recruitment advertising and the shift to online hiring

Historically, recruitment advertising relied heavily on classified advertisements in national and regional newspapers. These advertisements were typically sold on a *pay-per-word* basis, making vacancy advertising relatively expensive and naturally limiting both the number and duration of vacancies. The emergence of online job boards during the late 1990s fundamentally altered this model. Instead of purchasing newspaper space, employers increasingly advertised vacancies through dedicated recruitment websites using *pay-per-post* pricing structures. More recently, social media and platform-based recruitment have introduced *pay-per-click* and *pay-per-applicant* pricing models, while some providers offer basic vacancy listings free of charge.

Figure 2 shows the scale of this transition. Recruitment expenditure shifted away from traditional print advertising and towards online channels, with national and regional newspapers accounting for an ever-smaller share of recruitment spending after the mid-2000s. This coincided with a marked decline in aggregate recruitment advertising expenditure in real terms, driven largely by the collapse in newspaper advertising following the Global Financial Crisis.

Figure 2: Composition of aggregate spending on recruitment advertising
(GBP millions, 2025 prices)



Notes: Recruitment advertising expenditure is expressed in 2025 prices using the GDP deflator.
Sources: Advertising Association and WARC Expenditure Report; Office for National Statistics; authors' calculations.

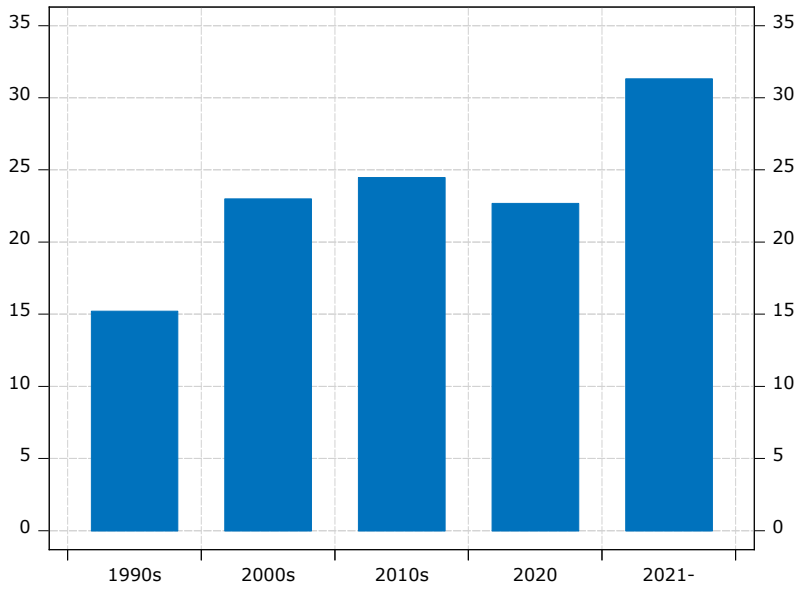
This pattern is consistent with developments observed in advertising markets more broadly. For example, Mandel (2019) documents a substantial shift from print to digital advertising across several

advanced economies and argues that digital advertising has become both cheaper and more effective than traditional print media. The result has been a sustained decline in advertising expenditure relative to economic activity, reflecting lower advertising costs rather than weaker demand for advertising itself.

Evidence from the early stages of online recruitment suggests that the associated reduction in advertising costs was substantial. Autor (2001) reports that in 2000 a 30-day vacancy advertisement on Monster.com cost $137, compared with $4,500 for a comparable listing in the Sunday *New York Times*. Platforms such as Indeed and LinkedIn continue to charge employers for premium placements and enhanced visibility, but many now allow basic vacancy listings at little or no monetary cost.

A decline in advertising costs is likely to affect not only the number of vacancies posted but also firms' recruitment behaviour. Consistent with this hypothesis, the average duration of job vacancies has increased substantially over time, rising from around three weeks in the 1990s to more than six weeks in recent years (Figure 3).[2] While vacancy duration is influenced by many factors, the rise is consistent with lower costs of maintaining active vacancies and a reduced urgency to fill positions quickly.[3]

Figure 3: Average vacancy duration
(working days)



Notes: Vacancy duration is approximated as the ratio of the stock of vacancies to quarterly hiring flows; see Appendix A.2 for details.
Sources: Authors' calculations based on ONS data.

[2]Vacancy duration is approximated as the ratio of the stock of vacancies to quarterly hiring flows (see Appendix A.2 for further detail). Although historical vacancy duration data are not available for the UK over the full sample, our estimates are broadly consistent with comparable measures for the US and with recent survey evidence for the UK, which reported an average time-to-fill of around 35 days (Adzuna, 2026).

[3]We discuss evergreen vacancies and ghost jobs in Section 2.4.

Taken together, these developments suggest that technological change has altered both the cost of vacancy creation and the behaviour of firms once vacancies are posted.

4.3 Constructing a vacancy posting cost index

To quantify these changes, we construct a measure of real vacancy posting costs by relating aggregate recruitment advertising expenditure to the stock of vacancies in the economy.

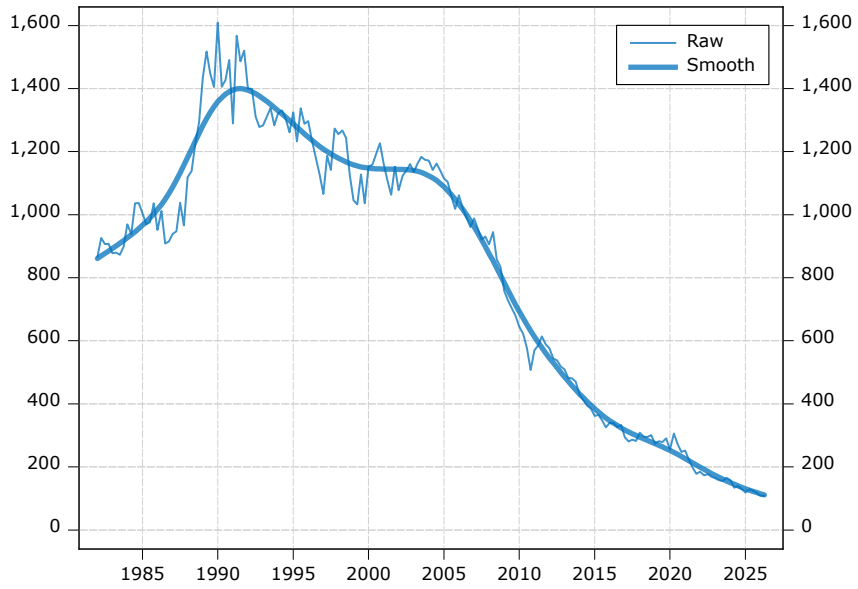
Specifically, we define the vacancy cost index as:

$$\kappa_t = \frac{A_t / P_t}{V_t} \tag{6}$$

where κ_t denotes the real vacancy posting cost, A_t is aggregate recruitment advertising expenditure obtained from the WARC dataset, P_t is the GDP deflator, and V_t is the stock of vacancies. The resulting measure therefore captures real advertising expenditure per vacancy and serves as a proxy for the cost of creating and maintaining vacancies.

Figure 4 presents the resulting vacancy cost index. The decline in vacancy posting costs over the past two decades is striking. Real advertising expenditure per vacancy has fallen by more than 80% since 2000, implying that firms can advertise substantially larger numbers of vacancies for a given recruitment budget than was previously possible.

Figure 4: Unit vacancy posting cost
(GBP, 2025 prices)



Notes: Vacancy posting costs are constructed as in equation (6).
Sources: Authors' calculations based on data from the Advertising Association and WARC Expenditure Report and the Office for National Statistics.

The measure should not be interpreted as a comprehensive estimate of all recruitment costs. Firms also incur expenditure associated with recruitment agencies, candidate sourcing, screening, interviewing, onboarding and other hiring activities. Evidence from the United States suggests

that vacancy advertising accounts for only one component of overall recruitment expenditure, alongside agency fees, professional networking platforms, company websites, employee referrals and recruitment events (Krider et al., 2015). Nevertheless, advertising is one of the few components of recruitment expenditure for which a long historical series is available. In the absence of comparable long-run measures of broader recruitment costs, it provides a useful proxy for changes in the economics of vacancy creation over the past four decades. The evidence presented in the next section suggests that broader recruitment costs have also declined materially over time.

Under the *per-post* pricing that now dominates, the index approximates the relevant object. Spend per vacancy in the stock then equals the fee per posting times the rate at which vacancies are filled, which is the flow-equivalent cost of a fee paid at posting (Appendix C.2). Because that rate fell as vacancies stayed open longer, part of the measured decline reflects duration: at a constant fee, longer durations account for about a seventh of the fall between 2001–09 and the subsequent period, implying that the estimated fee per posting spell nevertheless fell by a factor of around 2.5.

Viewed through the lens of the theoretical framework developed in Section 3, the observed decline in advertising expenditure per vacancy therefore corresponds to a substantial decline in the vacancy posting cost parameter (κ). The magnitude of the decline suggests that technological change has materially altered the incentives governing vacancy creation and may therefore have affected the equilibrium level of vacancies observed in the labour market.

4.4 Recruitment costs beyond advertising

An important question is whether the decline in vacancy posting costs documented above reflects a broader reduction in recruitment costs, or simply a reallocation of expenditure away from advertising towards other hiring activities.

Since the mid-2000s, the Chartered Institute of Personnel and Development (CIPD) has collected information on hiring costs through its *Recruitment, Retention and Turnover Survey* and its successor, the *Resourcing and Talent Planning Report* (Hogarth and McCartney, 2024). These surveys report median recruitment costs per hire, covering advertising costs and agency or search fees and, from 2011, in-house recruitment time. Unlike the WARC vacancy cost series, which focuses specifically on recruitment advertising expenditure, the CIPD measure captures a substantially broader range of recruitment activities.

Table 1 compares real vacancy posting costs and recruitment costs per hire over time. While the two series measure different subsets of hiring activity, both point to a substantial decline in the cost of recruitment.

Vacancy posting costs fell by around 87% between 2005 and 2024. Over the same period, the median recruitment cost per hire for managerial roles fell by around 88%. For non-managerial roles, for which comparable data are available from 2011 onwards, recruitment costs fell by around 57% between 2011 and 2024. The similarity of these declines is notable because the CIPD measure incorporates a much broader range of recruitment activities than advertising alone.

Table 1: Vacancy posting costs and recruitment costs per hire
(GBP, 2025 prices)

Year	Vacancy posting cost	Median recruitment cost per hire	
		Senior managers / directors	Other employees
2005	1,070	16,800	–
2011	590	11,000	3,700
2024	140	2,100	1,600
2005–2024 change (%)	-87	-88	–
2011–2024 change (%)	-76	-81	-57
Source	WARC	CIPD	CIPD

Notes: All figures are expressed in 2025 prices using the GDP deflator. WARC figures denote real recruitment advertising expenditure per vacancy. CIPD figures denote median recruitment cost per hire. For senior managers and directors, the 2005–10 series includes advertising costs and agency or search fees only; from 2011 onwards it also includes resourcing technology, in-house resourcing time, psychometric and selection tests, vetting, and onboarding costs. Data for other employees are available from 2011 onwards only. WARC figures are rounded to the nearest £10 and CIPD figures to the nearest £100.

This evidence suggests that the decline in vacancy posting costs documented in the WARC data reflects more than a simple shift in expenditure away from advertising. Instead, it is consistent with a broader reduction in the overall cost of recruitment. This interpretation accords with wider technological changes affecting candidate sourcing, screening and matching. Online hiring platforms have reduced the cost of advertising vacancies, while professional networking sites, applicant-tracking systems and digital recruitment tools have also lowered the resources required to identify and process candidates.

Recruitment costs are themselves only part of the cost of a hire. US survey evidence compiled by Manning (2011, Table 2) from Barron et al. (1997) values the hours spent training a new hire at several times those spent recruiting one, and Manning concludes that training accounts for the bulk of hiring costs. The CIPD's 2007 survey provides figures for both in the UK: in 2006 the median recruitment cost per hire was £4,333, against a median labour-turnover cost of £7,750 that adds vacancy cover, redundancy, training and induction (Chartered Institute of Personnel and Development, 2007). Training is paid for once a match has formed, so the model of Section 6 treats it as a one-off hiring cost, separate from the flow cost of keeping a vacancy open that the advertising index proxies (Appendix C.6).

Four stylised facts emerge from the evidence presented in this section: recruitment has shifted from print to online channels; vacancies remain open for longer; advertising expenditure per vacancy has fallen substantially; and broader recruitment costs reported by employers have declined in parallel. While each piece of evidence is imperfect in isolation, collectively they suggest that technological change has materially altered the economics of vacancy creation and reduced the cost of creating and maintaining vacancies.

Although the vacancy cost index captures only one component of recruitment expenditure, its

evolution appears broadly consistent with independent survey-based measures of employers' overall recruitment costs. The evidence therefore points to falling recruitment costs as an economically meaningful driver of vacancy creation and suggests that at least part of the secular rise in vacancies may reflect structural rather than cyclical developments. The next section examines this possibility by estimating the equilibrium level of vacancies and labour market tightness.

5 Estimating equilibrium labour market tightness

5.1 Equilibrium vacancies (V_t^*)

The theoretical framework in Section 3 suggests that the equilibrium level of vacancies depends on both labour productivity and vacancy posting costs. Both variables exhibit persistent long-run movements, while actual vacancies are influenced by short-term cyclical fluctuations. A natural empirical framework is therefore an error-correction model (ECM), which allows vacancies to deviate temporarily from equilibrium while maintaining a stable long-run relationship with their structural determinants.

Our baseline specification takes the following functional form:

$$\Delta V_t = \alpha_1 \Delta \kappa_t + \alpha_2 \Delta y_t + \gamma \left(V_{t-1} - V_{t-1}^* \right) + \varepsilon_t \tag{7}$$

where κ_t denotes the vacancy cost index developed in Section 4, y_t denotes hourly labour productivity, and V_t^* represents the equilibrium level of vacancies. The corresponding long-run relationship is:

$$V_t^* = \beta_0 + \beta_1 \kappa_t + \beta_2 y_t \tag{8}$$

The model is estimated using quarterly data over the period 1982–2025. Labour productivity and vacancy posting costs are included because they correspond directly to the key determinants of job creation identified in the theoretical framework. While matching efficiency is also an important determinant, long-run measures are only available over a shorter sample. The baseline long-run specification is deliberately parsimonious, focusing on slow-moving structural determinants of vacancy creation. The main results are robust to alternative specifications, including the addition of a deterministic trend, controls for matching efficiency, and estimation over an earlier sample period.[4]

5.2 Equilibrium labour market tightness (θ_t^*)

Vacancies alone provide only a partial measure of labour market conditions. To assess labour market tightness, we combine the estimated equilibrium level of vacancies with an estimate of equilibrium unemployment (U_t^*).

[4]Regression estimates, cointegration tests and additional robustness checks are reported in Appendix B.

Observed labour market tightness is defined as:

$$\theta_t = \frac{V_t}{U_t} \tag{9}$$

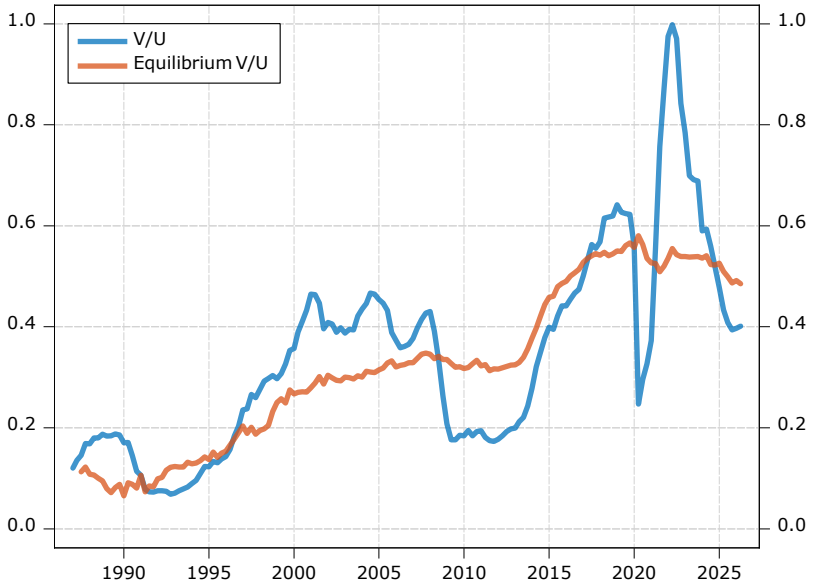
while equilibrium labour market tightness is given by:

$$\theta_t^* = \frac{V_t^*}{U_t^*} \tag{10}$$

For U_t^*, we use the Bank of England's existing estimate of equilibrium unemployment (the NAIRU) rather than attempting to estimate a separate measure for the purposes of this paper. This allows the analysis to isolate the contribution of equilibrium vacancies while remaining consistent with the broader framework used in policy assessments.

Figure 5 compares the observed V/U ratio with its estimated equilibrium level. The results suggest that equilibrium labour market tightness has increased materially over the past four decades, reflecting the combined effects of rising productivity and falling vacancy posting costs. As a result, an increase in the observed V/U ratio does not necessarily imply a progressively tighter labour market. Instead, part of the rise reflects a higher equilibrium level of vacancies.

Figure 5: The observed V/U ratio (θ_t) and its estimated equilibrium (θ_t^*)



Notes: Observed and equilibrium V/U ratios are defined as in equations (9) and (10), respectively. Equilibrium vacancies are estimated using the ECM described in Section 5.1.

The gap between observed and equilibrium labour market tightness (θ_t^{gap}) provides a measure of labour market slack relative to its estimated equilibrium level. A decomposition of this gap into vacancy and unemployment contributions is reported in Appendix A.3. The results suggest that a large proportion of the variation in labour market tightness is explained by the vacancy component, with deviations in unemployment generally playing a smaller role. This finding is consistent with the

central mechanism emphasised throughout the paper: structural changes in recruitment technology primarily affect labour market tightness through vacancy creation.

The framework also provides a useful lens through which to interpret recent labour market developments. The sharp increase in labour market tightness following the pandemic appears to have reflected a genuine cyclical rise in vacancies above their equilibrium level, rather than a further increase in equilibrium tightness itself. As labour demand subsequently eased, vacancies fell back and the gap between observed and equilibrium labour market tightness narrowed materially. More recently, the vacancy-to-unemployment ratio has moved below its estimated equilibrium level, suggesting some degree of spare capacity in the labour market. This interpretation is consistent with other indicators of labour market slack, including measures of underemployment (Bank of England, 2026).[5]

6 A search-and-matching application

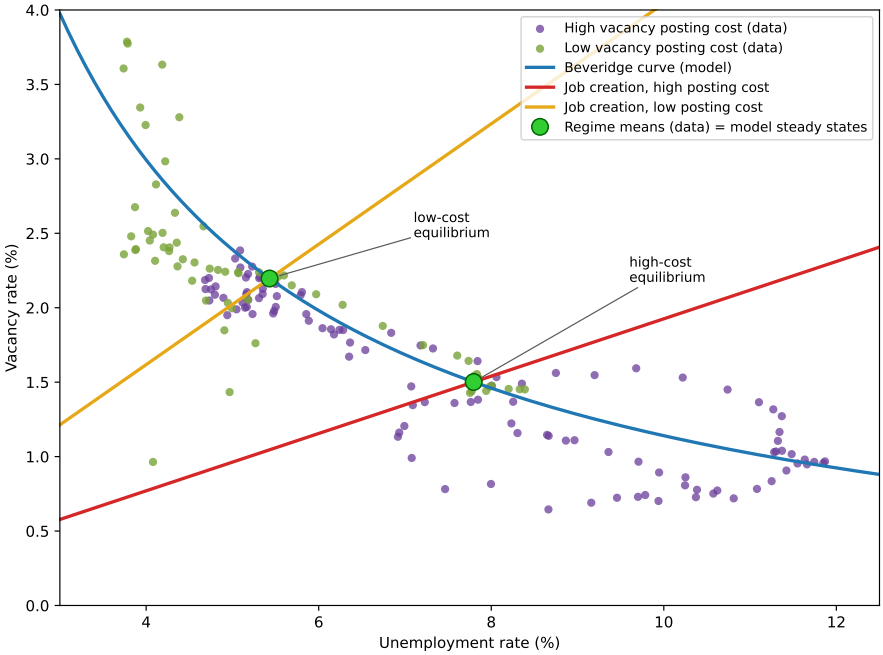
The empirical framework developed in Section 5 suggests that falling vacancy posting costs have contributed to a structural increase in equilibrium labour market tightness. As a complementary robustness exercise, we examine whether a standard search-and-matching model produces a similar response to observed changes in vacancy posting costs. Using a calibrated Diamond-Mortensen-Pissarides (DMP) model, we assess whether the magnitude of the increase in equilibrium tightness implied by the theory is broadly consistent with the empirical estimates presented above.

We split the sample into two vacancy-cost regimes using the unit vacancy cost series developed in Section 4: a high-cost, print-dominated regime (1982 Q2 to 2009 Q4) and a low-cost, online regime (2010 Q1 onwards). The model is calibrated at monthly frequency so that its two steady states reproduce the average unemployment and vacancy rates in each regime exactly, with no rescaling of either axis. The calibration is exactly identified, and two of its outputs are informative in their own right. First, requiring both regime means to lie on a single Beveridge curve returns a matching elasticity of about 0.49, close to the value of 0.5 commonly assumed in the literature. This is a check rather than an estimate, since it depends on the pre-2001 vacancy splice (Appendix A.1) and assumes that matching efficiency and separations did not change (Appendix C.4). Second, the free-entry condition implies that total vacancy posting costs must roughly halve (a factor of about 2) to move the labour market between the two regimes, compared with a measured fall in advertising expenditure per vacancy by a factor of around 3.2 – consistent with advertising being one component, and a proxy, of broader recruitment costs. Appendix C sets out the model, the data and the full calibration.

The impact of a fall in vacancy posting costs is illustrated in Figure 6. A fall in the unit vacancy cost raises the net benefit from posting a vacancy, holding productivity and other variables unchanged. Graphically, the job creation curve rotates anti-clockwise, so that vacancies are higher for a given

[5] As with all equilibrium concepts, substantial uncertainty remains around the precise level of slack. In particular, developments affecting labour supply or the equilibrium rate of unemployment could alter the benchmark against which labour market tightness is assessed.

Figure 6: The Beveridge curve and the job creation curves in a high- and a low-cost regime



Notes: Quarterly observations, 1982 Q2–2024 Q3, split into the high-cost regime (1982 Q2–2009 Q4) and the low-cost regime (2010 Q1 onwards) according to real recruitment advertising expenditure per vacancy. Solid lines are the calibrated model's Beveridge curve and the job creation curves under the high- and low-cost regime posting costs ($\kappa_H = 0.216$ and $\kappa_L = 0.110$ months of output per vacancy per month); the two intersections reproduce the regime means of the data exactly. Both axes are in natural units. See Appendix C for the model and calibration. Sources: ONS Vacancy Survey; Bank of England (A millennium of macroeconomic data); ONS Labour Force Survey; Advertising Association and WARC Expenditure Report; and authors' calculations.

level of unemployment, and the new job creation curve forms a new equilibrium with the Beveridge curve at a higher V/U ratio. An attractive feature of the model is that the vacancy posting cost does not enter the Beveridge curve at all: the entire equilibrium adjustment is a movement *along* a stable Beveridge curve (Appendix C.3).[6] The sensitivity of equilibrium tightness to posting costs is substantial: at the calibrated steady state the elasticity of tightness with respect to the posting cost is about −1.1, so halving posting costs slightly more than doubles equilibrium tightness. Treating training as a one-off hiring cost of the size UK survey evidence suggests lowers the elasticity to between −0.9 and −0.7 (Appendix C.6). As Figure 6 shows, the two model steady states coincide with the regime means of the data (by construction), and the informative check is that the within-regime quarterly observations scatter around the same downward-sloping curve. The implied regime shift in equilibrium tightness – from 0.19 to 0.40 – is of the same order as the rise in the ECM-based equilibrium tightness estimates between the two regime periods (Figure 5), which also roughly double. Letting matching efficiency and separations differ across regimes, with the matching elasticity set from outside, leaves posting costs accounting for essentially all of the regime change in tightness (Appendix C.4). Feeding the full time path of observed vacancy posting costs through the model's transition dynamics reproduces the low-frequency drift in tightness over

[6] Steady-state tightness roughly doubles – unemployment falls by over two percentage points and the vacancy rate rises by half – entirely along the fixed curve; see Appendix C.3.

the past four decades (Appendix C.7).

7 Conclusions

This paper examines whether changes in vacancy posting costs can help explain the secular rise in vacancies observed in the UK labour market over recent decades. We make three main contributions. First, we construct a new long-run measure of vacancy posting costs using recruitment advertising expenditure data spanning the period since 1982. Second, we document a substantial decline in vacancy posting costs following the transition from print-based recruitment to online hiring platforms. Third, we show that accounting for these changes materially affects estimates of equilibrium vacancies and equilibrium labour market tightness.

The findings help address an important puzzle in the interpretation of labour market conditions. While the vacancy-to-unemployment ratio has risen substantially over time, our results suggest that a significant share of this increase reflects structural rather than cyclical developments. In particular, lower vacancy posting costs appear to have increased the equilibrium level of vacancies, implying that comparisons of labour market tightness across time must account for changes in recruitment technology and hiring behaviour.

These issues are particularly relevant for policymakers. Vacancy-based indicators play a prominent role in assessments of labour market slack, wage pressures and inflationary risks. By distinguishing between cyclical fluctuations and longer-run shifts in equilibrium labour market tightness, the framework developed in this paper provides a complementary perspective on the interpretation of labour market conditions and the signals contained in vacancy data.

Several limitations remain. The vacancy cost index captures recruitment advertising expenditure rather than the full set of recruitment costs faced by firms, and the relative importance of alternative recruitment channels may have changed over time. Future research could extend the analysis by incorporating broader measures of recruitment activity, exploring industry-level differences in vacancy creation, or extending the steady-state accounting of posting costs, matching efficiency and separations in Appendix C.4 to recruiting intensity and productivity. Nevertheless, the evidence presented in this paper suggests that vacancy posting costs are an important determinant of equilibrium labour market tightness and should be considered when interpreting vacancy-based indicators of labour market conditions.

References

Adzuna. 2026. "UK Job Market Report." May, https://www.adzuna.co.uk/job-market-report/, 30 May 2026.

Autor, David H. 2001. "Wiring the Labor Market." *Journal of Economic Perspectives* 15 (1): 25–40.

Bank of England. 2026. "Monetary Policy Report – July 2026."

Barnichon, Regis, and Andrew Figura. 2015. "Labor Market Heterogeneities and the Aggregate Matching Function." *American Economic Journal: Macroeconomics* 7 (4): 222–249.

Barron, John M., Mark C. Berger, and Dan A. Black. 1997. *On-the-Job Training*. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research, . 10.17848/9780585262369.

Cajner, Tomaz, and David Ratner. 2016. "A Cautionary Note on the Help Wanted Online Data." FEDS Notes, Board of Governors of the Federal Reserve System, Washington, DC.

Cappelli, Peter. 2012. *Why Good People Can't Get Jobs: The Skills Gap and What Companies Can Do About It*. Wharton School Press.

Chartered Institute of Personnel and Development. 2007. *Recruitment, Retention and Turnover 2007: Annual Survey Report*. London: CIPD.

Dalton, Michael, Lisa B. Kahn, and Andreas I. Mueller. 2026. "Do Online Job Postings Capture Job Vacancies?." NBER Working Paper 35697, National Bureau of Economic Research.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2013. "The Establishment-Level Behavior of Vacancies and Hiring." *The Quarterly Journal of Economics* 128 (2): 581–622.

Djourelova, Milena, Ruben Durante, and Gregory J. Martin. 2025. "The Impact of Online Competition on Local Newspapers: Evidence from the Introduction of Craigslist." *The Review of Economic Studies* 92 (3): 1738–1772.

Elsby, Michael W. L., Jennifer C. Smith, and Jonathan Wadsworth. 2011. "The Role of Worker Flows in the Dynamics and Distribution of UK Unemployment." *Oxford Review of Economic Policy* 27 (2): 338–363.

Faberman, R. Jason. 2020. "Recruiting Intensity." *IZA World of Labor* 21 1–21.

Gavazza, Alessandro, Simon Mongey, and Giovanni L. Violante. 2018. "Aggregate Recruiting Intensity." *American Economic Review* 108 (8): 2088–2127.

Hagedorn, Marcus, and Iourii Manovskii. 2008. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited." *American Economic Review* 98 (4): 1692–1706.

Hall, Robert E. 2017. "High Discounts and High Unemployment." *American Economic Review* 107 (2): 305–330.

Hogarth, Annette, and Claire McCartney. 2024. *Resourcing and Talent Planning Report 2024*. London: Chartered Institute of Personnel and Development (CIPD).

Krider, Jennifer, Karen O'Leonard, and Robin Erickson. 2015. *Talent Acquisition Factbook 2015: Benchmarks and Trends in Spending, Staffing, and Key Recruiting Metrics*. Bersin by Deloitte.

Kuhn, Peter, and Hani Mansour. 2014. "Is Internet Job Search Still Ineffective?." *Economic Journal* 124 (581): 1213–1233.

Kuhn, Peter, and Mikal Skuterud. 2004. "Internet Job Search and Unemployment Durations." *American Economic Review* 94 (1): 218–232.

Linckh, Carolin, Samuel Muehlemann, and Harald Pfeifer. 2026. "Beggars Cannot Be Choosers: Labor Market Tightness, Hiring Standards, Wages, and Hiring Costs." *Labour Economics* 102 102941.

Luke, Anthony. 2005. "Advertising Placement Corporate Services Price Index." *Economic Trends* 625 50–53, Office for National Statistics. https://escoe-website.s3.amazonaws.com/wp-content/uploads/2020/09/05175305/ET-625-Advertising-placement-CSPI-Tony-Luke-Dec-2005.pdf.

Mandel, Michael. 2019. "The Declining Cost of Advertising: Policy Implications."Technical report, Progressive Policy Institute.

Manning, Alan. 2011. "Imperfect Competition in the Labor Market." In *Handbook of Labor Economics*, edited by Ashenfelter, Orley, and David Card Volume 4B. 973–1041, Elsevier, . 10.1016/S0169-7218(11)02409-9.

Martellini, Paolo, and Guido Menzio. 2020. "Declining Search Frictions, Unemployment, and Growth." *Journal of Political Economy* 128 (12): 4387–4437.

Mongey, Simon, and Jeff Horwich. 2023. "Are Job Vacancies Still as Plentiful as They Appear? Implications for the "Soft Landing"." Federal Reserve Bank of Minneapolis.

Mortensen, Dale T., and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies* 61 (3): 397–415.

Ng, Hunter. 2024. "Why Is It So Hard to Find a Job Now? Enter Ghost Jobs." Working paper.

Office for National Statistics. 2026. "Labour Force Survey Two-Quarter Longitudinal Datasets, 1993–2026." UK Data Service, Accessed under the End User Licence and pooled by the authors.

Petrongolo, Barbara, and Christopher A. Pissarides. 2001. "Looking into the Black Box: A Survey of the Matching Function." *Journal of Economic Literature* 39 (2): 390–431.

Pissarides, Christopher A. 2000. *Equilibrium Unemployment Theory*. Cambridge, MA: MIT Press.

Pissarides, Christopher A. 2009. "The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?" *Econometrica* 77 (5): 1339–1369.

Pizzinelli, Carlo, and Bradley Speigner. 2017. "Matching Efficiency and Labour Market Heterogeneity in the United Kingdom." Staff Working Paper 667, Bank of England.

Santhosh, Stephen, and Gabriela Jennifer. 2025. "Ghost Jobs in the Digital Labor Market: An Empirical Analysis of Hiring Intent Signals Using a Novel Eight-Category Classification Framework." working paper, SSRN.

Sedláček, Petr. 2014. "Match Efficiency and Firms' Hiring Standards." *Journal of Monetary Economics* 62 123–133.

Singer, Steven, and Derya Oktay. 2025. "Ghost Jobs, Real Costs: Labor Market Distortions from Ghost Job Postings." *Business Economics* 60 216–229.

Smith, Jennifer C. 2011. "The Ins and Outs of UK Unemployment." *The Economic Journal* 121 (552): 402–444.

Stelmach, Michal, James Kensett, and Philip Schnattinger. 2025. "What Can 40 Years of Data on Vacancy Advertising Costs Tell Us about Labour Market Equilibrium?." Bank Underground.

Thomas, Ryland, and Nicholas Dimsdale. 2017. "A Millennium of UK Data." Bank of England OBRA dataset, version 3.1, https://www.bankofengland.co.uk.

Technical appendix

A Data and measurement

A.1 Data sources

Table A.1 summarises the variables and data sources used throughout the paper.

Table A.1: Data sources and description

Variable	Notation	Source	Description
Observed variables			
Vacancies (2001–)	V_t	ONS Vacancy Survey	ONS series AP2Y
Historical vacancies (pre-2001)	V_t	Bank of England (A millennium of macroeconomic data)	Spliced to the ONS Vacancy Survey series
Unemployment	U_t	ONS Labour Force Survey	ONS series MGSC
Recruitment advertising expenditure	A_t	Advertising Association and WARC Expenditure Report	Aggregate expenditure on recruitment advertising
GDP deflator	P_t	ONS	ONS series YBGB
Hourly labour productivity	y_t	ONS	ONS series ABMM / YBUS
Derived variables			
Labour market tightness	θ_t	Authors' calculations	$\theta_t = V_t/U_t$
Hiring flows	H_t	ONS Labour Force Survey	$H_t = UE + NE + J2J$ flows
Vacancy duration	D_t	Authors' calculations	$D_t = V_t/H_t$, working days
Vacancy posting cost	κ_t	Authors' calculations	$\kappa_t = (A_t/P_t)/V_t$
Equilibrium variables			
Equilibrium unemployment	U_t^*	Bank of England	Bank of England NAIRU estimate
Equilibrium vacancies	V_t^*	Authors' estimates	Long-run ECM estimate
Equilibrium tightness	θ_t^*	Authors' estimates	$\theta_t^* = V_t^*/U_t^*$

Notes: All series are quarterly and seasonally adjusted. An asterisk denotes an equilibrium concept.

Recruitment advertising expenditure. The Advertising Association/WARC Expenditure Report provides quarterly estimates of aggregate UK recruitment advertising expenditure by medium. Revenue is recorded before the deduction of agency commission but after all negotiated discounts and generally excludes production costs. National newsbrands are covered in full. For regional newsbrands, quarterly returns are obtained from a panel of large and small publishers accounting for an estimated 60% of total regional newsbrand advertising revenue. These returns are grossed up by applying the annual change reported by the panel to the market estimate for the corresponding quarter of the previous year.

The expenditure series is not a direct price index. It records aggregate recruitment advertising

expenditure, which may change because of movements in the price, number or type of advertisements purchased. In Section 4.3, we divide real aggregate expenditure by the vacancy stock to construct expenditure per vacancy. The resulting measure is therefore an estimate of the unit cost of advertising and maintaining vacancies, rather than the observed price of a standardised advertisement.

The US series in Appendix D.4 is constructed using list prices rather than realised expenditure and is therefore not directly comparable with the UK series in levels. We consequently use it only to compare ratios and growth rates.

Vacancies from 2001 onwards. From 2001 onwards, vacancies are measured using the ONS Vacancy Survey (series AP2Y). This is a statutory, monthly, sample-based survey in which businesses report the number of vacancies on a specified date for which they are actively seeking recruits from outside their organisation. The sampling frame is the Inter-Departmental Business Register. The monthly sample comprises approximately 6,100 businesses, including around 1,400 large businesses surveyed each month and approximately 4,700 smaller enterprises sampled quarterly. Headline estimates are published as seasonally adjusted three-month moving averages.

The survey covers businesses across Great Britain, excluding agriculture, forestry and fishing; activities of households as employers; and employment agencies, the latter to avoid double-counting vacancies advertised on behalf of other businesses. UK estimates are obtained by weighting the Great Britain results to account for Northern Ireland. Although the survey is designed to provide a comprehensive measure of vacancies across the economy, its estimates are subject to sampling and non-sampling error. The ONS reports that the three-month average of total vacancies typically has a coefficient of variation of around 1.3%.

Vacancies before 2001. Earlier quarters are from *A Millennium of Macroeconomic Data* (Thomas and Dimsdale, 2017), which counts vacancies notified to Jobcentres, an administrative subset of all vacancies, and splices them to the Vacancy Survey with a constant multiplier of 1.7576 estimated over the overlap. The multiplier exceeds one because Jobcentres saw only part of the market, and holding it constant assumes a constant notification rate. It matters for one output of the calibration in Appendix C.2: across multipliers from 1.3 to 2.3, the matching elasticity identified from the two regime means ranges from 0.58 to 0.35, while the regime ratio of tightness moves only from 2.46 to 1.80.

A.2 Vacancy duration measure

Vacancy duration can be approximated using the relationship between vacancy stocks and vacancy outflows. In steady state, the stock of vacancies equals the flow of vacancies leaving the stock multiplied by the average duration for which vacancies remain open. This implies:

$$D_t = \frac{V_t}{O_t} \tag{A.1}$$

where D_t denotes average vacancy duration, V_t the stock of vacancies, and O_t vacancy outflows.

Direct measures of vacancy outflows are not available for the UK over a sufficiently long sample. In Figure 3 we therefore proxy vacancy outflows using aggregate hires from the LFS flows dataset:

$$D_t \approx \frac{V_t}{H_t} \tag{A.2}$$

where H_t is the sum of unemployment-to-employment (UE), inactivity-to-employment (NE) and job-to-job (J2J) transitions. The vacancy stock is the spliced series of Appendix A.1, Jobcentre-based before 2001. Hires come from LFS two-quarter longitudinal data throughout: from 2001 Q4 they follow the published ONS flows (X02), and before that they are derived from the microdata (Appendix A.4).

This approximation is likely to overstate the level of vacancy duration to some extent, because not all vacancies leave the stock through successful hires. Some vacancies are withdrawn or cancelled before being filled, implying that total vacancy outflows are likely to exceed observed hiring flows. However, the focus of this paper is on changes in vacancy duration over time rather than its precise level. To the extent that the relationship between hiring flows and total vacancy outflows has remained broadly stable, this measurement issue should have only a limited effect on the estimated trend.

A.3 Tightness gap decomposition

The tightness gap is defined as:

$$\theta_t^{gap} = \theta_t - \theta_t^* \tag{A.3}$$

To assess the relative contribution of vacancies and unemployment to fluctuations in labour market tightness, the tightness gap is decomposed into vacancy and unemployment components using Shapley values. This approach accounts for the non-linear relationship between vacancies, unemployment and labour market tightness.

The vacancy and unemployment contributions are given, respectively, by:

$$\theta_{V,t}^{gap} = \frac{1}{2}\left[\frac{V_t - V_t^*}{U_t} + \frac{V_t - V_t^*}{U_t^*}\right] \tag{A.4}$$

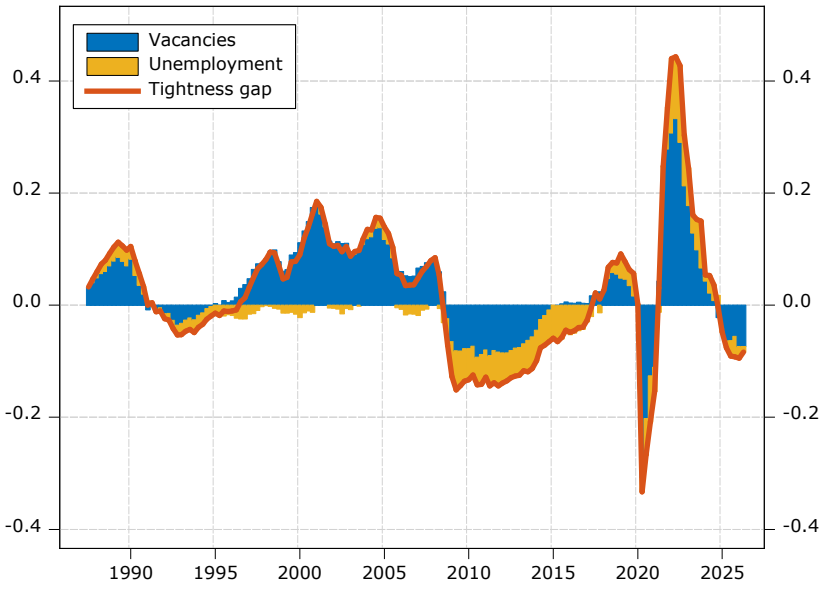
$$\theta_{U,t}^{gap} = \frac{1}{2}\left[V_t\left(\frac{1}{U_t} - \frac{1}{U_t^*}\right) + V_t^*\left(\frac{1}{U_t} - \frac{1}{U_t^*}\right)\right] \tag{A.5}$$

such that:

$$\theta_t^{gap} = \theta_{V,t}^{gap} + \theta_{U,t}^{gap} \tag{A.6}$$

As shown in Figure A.1, movements in the tightness gap are driven predominantly by the vacancy component, particularly during periods of elevated labour demand such as the post-pandemic expansion.

Figure A.1: Decomposition of the tightness gap into vacancy and unemployment contributions



Notes: The tightness gap is decomposed into vacancy and unemployment contributions as defined in equations (A.3)–(A.6).

A.4 The LFS two-quarter longitudinal microdata

The hires flow of Appendix A.2 before 2001 and the transition rates of Appendices C.4 and C.5 come from the ONS Labour Force Survey two-quarter longitudinal datasets (Office for National Statistics, 2026), pooled over 1993 Q1–2026 Q2: 5.2 million person-quarter pairs after restricting to ages 16–64, valid economic activity at both waves and a positive longitudinal weight. Each record observes one person in two consecutive quarters, so transitions are measured on the individual rather than inferred from stocks. Elapsed unemployment duration at the first wave is collapsed to under six months, six to under twelve, and twelve months and over; 0.4% of unemployed first-wave records lack a band. All rates are weighted by the ONS longitudinal weight `LGWT`; the weighted first-wave stocks reproduce published employment and unemployment levels, and transitions rebuilt from the activity variables match the ONS `FLOW` classification to within 0.02%. Cells with fewer than 50 unweighted observations are flagged and excluded from window means. Two caveats carry through: the comparison windows exclude the 2020–21 response collapse, and from 2023 the transformed LFS leaves roughly 250–400 unemployed respondents per quarter, so recent quarters are volatile. Longitudinal weights calibrate to the first-wave cross-section, which understates flows to inactivity. The microdata are available from the UK Data Service under the standard End User Licence.

B Econometric results and diagnostics

B.1 Cointegration tests

Table B.1 reports several residual-based cointegration tests based on the preferred long-run vacancy equation. The long-run residuals are stationary according to an Augmented Dickey-Fuller test, while both the Engle-Granger and Phillips-Ouliaris tests reject the null hypothesis of no cointegration at the 5% significance level. Taken together, these results support the use of an error-correction specification.

Table B.1: Cointegration tests

Test	Statistic	p-value
Augmented Dickey-Fuller on long-run residuals	-3.605	0.007
Engle-Granger (tau-statistic)	-3.666	0.021
Phillips-Ouliaris (tau-statistic)	-3.336	0.049

Notes: Tests are based on the preferred long-run vacancy equation estimated over 1982Q1–2025Q4, with vacancies as the dependent variable.

B.2 Regression results

Table B.2 reports estimates from the baseline specification. Vacancy posting costs enter negatively and labour productivity positively in both the long-run equation and the error-correction model. The estimated error-correction coefficient is negative and statistically significant, indicating that deviations from the long-run relationship are gradually corrected over time.

Table B.2: Baseline specification

Variable	Coefficient	Std. Error	p-value
	Long-run equation		
Vacancy posting cost	-0.133	0.035	0.000
Labour productivity	1.783	0.175	0.000
Constant	-209.0	109.2	0.057
	Error-correction model		
Δ Vacancy posting cost	-0.197	0.065	0.003
Δ Labour productivity	2.831	0.701	0.000
Error-correction term	-0.060	0.028	0.032
Sample	1982Q1–2025Q4		
Observations	176		
Long-run equation R^2	0.687		

Notes: The ECM is estimated in a single step, with the change in vacancies as the dependent variable. Long-run coefficients are additionally reported from the corresponding static long-run equation.

B.3 Robustness checks

Table B.3 assesses the sensitivity of the baseline long-run relationship to alternative specifications. Specifically, we consider an earlier estimation sample ending before the sharpest decline in vacancy posting costs, the inclusion of a deterministic linear trend, and the addition of a measure of matching efficiency available from 1994 onwards. The vacancy posting cost coefficient retains its expected negative sign across all specifications and remains close to its baseline estimate when matching efficiency is included. Cointegration is supported in the baseline, trend and matching efficiency specifications, although the evidence is weaker over the truncated pre-2008 sample.

Table B.3: Alternative specifications of the long-run equation

	(1)	(2)	(3)	(4)
	Long-run equation			
Vacancy posting cost	-0.133	-0.188	-0.079	-0.126
	(0.035)	(0.043)	(0.052)	(0.056)
Labour productivity	1.783	2.071	0.960	3.863
	(0.175)	(0.111)	(0.612)	(0.517)
Deterministic trend			1.767	
			(1.260)	
Matching efficiency				2.492
				(0.358)
Sample	1982–2025	1982–2007	1982–2025	1994–2025
Observations	176	104	176	128
Cointegration (tau-statistic)	-3.666	-2.747	-3.777	-4.419
p-value	0.021	0.176	0.049	0.009
R^2	0.687	0.785	0.691	0.633

Notes: Vacancies are the dependent variable throughout. Column (1) reports the baseline specification. Column (2) restricts the sample to 1982–2007. Column (3) adds a deterministic linear trend. Column (4) adds matching efficiency. Standard errors are reported in parentheses. Cointegration statistics are based on the Engle–Granger test with vacancies as the dependent variable.

C The search-and-matching model

This appendix sets out the search-and-matching model behind the exercise in Section 6: the environment and equilibrium conditions (Appendix C.1), the data and calibration (Appendix C.2), the comparative statics of a change in the vacancy posting cost (Appendix C.3), and the transition dynamics implied by the observed path of posting costs (Appendix C.7).

C.1 Model setup

A standard 'search and matching' model specifies the job creation curve as a function of the expected returns from employing a worker and the average cost of creating a vacancy. One can therefore think of posting a vacancy as an investment decision: hiring an additional worker depends on the expected profit from filling the vacancy. This appendix states the discrete-time, monthly version of the Diamond-Mortensen-Pissarides (DMP) model that we take to the data; it is the discrete-time counterpart of the exposition in Section 3.

Time is monthly. Unemployed workers and vacancies are matched according to a Cobb-Douglas matching function $M = m\,V^{1-\xi}U^{\xi}$, where m is matching efficiency and ξ the matching elasticity. With labour market tightness $\theta = V/U$, a vacancy is filled within the month with probability $q(\theta) = m\,\theta^{-\xi}$ and an unemployed worker finds a job with probability $p(\theta) = \theta q(\theta) = m\,\theta^{1-\xi}$. Matches produce output x (normalised to one), pay a wage w, and are destroyed with exogenous probability λ each month – the discrete-time counterpart of the separation rate s in Section 3. Firms discount the future with the factor β, corresponding to the discount rate $r \approx 1/\beta - 1$. Unemployed workers receive the flow value b, and wages are set by Nash bargaining with worker bargaining weight η.

Unemployment evolves according to

$$u_{t+1} = u_t + \lambda(1 - u_t) - (1 - \lambda)\,p(\theta_t)\,u_t, \tag{C.1}$$

so that in steady state the flows into and out of unemployment balance, giving the Beveridge curve:

$$u = \frac{\lambda}{\lambda + (1 - \lambda)\,p(\theta)}. \tag{C.2}$$

Values and free entry. Let J_t denote the value to a firm of a job filled at the start of period t, and V_t the value of an open vacancy. A filled job yields $x - w_t$ this period and survives into the next with probability $1 - \lambda$:

$$J_t \;=\; x - w_t + \beta\,(1 - \lambda)\,J_{t+1}. \tag{C.3}$$

A vacancy costs κ_t to maintain for the period and is filled with probability $q(\theta_t)$, in which case the firm holds a job that survives to produce with probability $1 - \lambda$:

$$V_t \;=\; -\kappa_t + q(\theta_t)\,(1 - \lambda)\,\beta\,J_{t+1} + \left[1 - q(\theta_t)\right]\beta\,V_{t+1}. \tag{C.4}$$

Free entry drives the value of an open vacancy to zero in every period, $V_t = 0$ for all t, so (C.4)

reduces to the requirement that the expected cost of filling a vacancy equals the discounted value of the resulting match,

$$\frac{\kappa_t}{q(\theta_t)} \;=\; (1-\lambda)\,\beta\,J_{t+1}. \tag{C.5}$$

Wages. Nash bargaining with worker weight η, combined with free entry, delivers the standard wage equation

$$w_t \;=\; \eta\,(x+\kappa_t\theta_t) + (1-\eta)\,b, \tag{C.6}$$

in which the vacancy posting cost enters through the worker's threat point: a tighter market with costlier vacancies improves the worker's outside option.

Job creation. Leading (C.3) one period, substituting the wage (C.6) to obtain $x - w_{t+1} = (1-\eta)(x-b) - \eta\,\kappa_{t+1}\theta_{t+1}$, and using (C.5) at $t+1$ to replace $\beta(1-\lambda)J_{t+2}$ with $\kappa_{t+1}/q(\theta_{t+1})$, gives the *dynamic* job creation condition

$$\frac{\kappa_t}{q(\theta_t)} \;=\; (1-\lambda)\,\beta\left[(1-\eta)(x-b) - \eta\,\kappa_{t+1}\theta_{t+1} + \frac{\kappa_{t+1}}{q(\theta_{t+1})}\right]. \tag{C.7}$$

Equations (C.1) and (C.7) are the full dynamic system: unemployment is a state variable that evolves according to (C.1), while tightness is a forward-looking jump variable satisfying (C.7). Appendices C.7 and C.8 solve this system directly.

Steady state. Imposing $\kappa_t = \kappa$ and $\theta_t = \theta$ for all t in (C.7) and collecting the $\kappa/q(\theta)$ terms gives the stationary job creation condition,

$$\frac{\kappa}{q(\theta)}\Big[1-(1-\lambda)\beta\Big] \;=\; (1-\lambda)\,\beta\Big[(1-\eta)(x-b) - \eta\,\kappa\,\theta\Big], \tag{C.8}$$

which is equivalent to $J = (x-w)/[1-\beta(1-\lambda)]$ together with $\kappa/q(\theta) = (1-\lambda)\beta J$. The calibration in Appendix C.2 and the comparative statics in Appendix C.3 use this stationary form; the transition exercises use (C.7).

Two properties of the *steady-state* system organise everything that follows. First, the job creation condition (C.8) pins down equilibrium tightness $\theta^*(\kappa)$ on its own: in (u,v) space the job creation curve is the ray $v = \theta^*(\kappa)\,u$. Second, the Beveridge curve (C.2) depends on (λ, m, ξ) but *not* on κ. The steady-state equilibrium is the intersection of the ray with the curve. Neither property survives outside the steady state: along a transition the economy is generally off the Beveridge curve, and (C.7) makes current tightness depend on the whole anticipated path of posting costs rather than on κ_t alone – a distinction that turns out to matter in Appendix C.8.

C.2 Data and calibration

Data. The model is confronted with the quarterly dataset described in Appendix A.1, covering 1982 Q2 to 2024 Q3 (170 quarters). The unemployment rate is LFS unemployment divided by the sum of employment and unemployment (ONS series MGSC, MGRZ); the vacancy rate is the vacancy stock (ONS Vacancy Survey, spliced to *A Millennium of UK Data* before 2001) divided by the same denominator; and the vacancy posting cost series is real recruitment advertising

expenditure per vacancy, $\kappa_t = (A_t/P_t)/V_t$, from the Advertising Association / WARC Expenditure Report as constructed in Section 4. We split the sample into two cost regimes according to whether real advertising expenditure per vacancy is above or below £660 per vacancy per quarter (constant prices). The series crosses this threshold exactly once, so the split is also a date split: a high-cost regime from 1982 Q2 to 2009 Q4 (111 quarters, mean £1,066 per vacancy-quarter) and a low-cost regime from 2010 Q1 onwards (59 quarters, mean £335). The regime means of the data are $(u_H, v_H) = (7.80\%,\ 1.49\%)$ and $(u_L, v_L) = (5.46\%,\ 2.20\%)$, implying tightness at the regime means of $\theta_H = 0.191$ and $\theta_L = 0.403$. To be precise about units and vintage: the analysis dataset is the November 2024 research extract of the Advertising Association / WARC recruitment-spend data, with nominal amounts deflated by the GDP deflator normalised to 2024 Q3 = 1, so monetary amounts in this appendix are in constant 2024 prices; the advertising-cost series is quoted per vacancy per *quarter* (the threshold of £660 per vacancy-quarter corresponds to roughly £220 per vacancy per month, and the regime means to roughly £355 and £112 per vacancy per month). The series shown in Figure 4 is constructed from the updated 2025-price data vintage and differs in level; the regime split and the calibration depend only on relative movements in the series, and all pound conversions in this appendix are internally consistent within the analysis dataset.

Calibration. The calibration is exactly identified: each parameter is pinned by one moment or one external piece of evidence, and the model's two steady states then reproduce the two regime means exactly.

1. *Matching elasticity ξ.* Requiring both regime means to lie on a single Beveridge curve (C.2) identifies ξ independently of m and λ:

$$\xi = 1 - \frac{\ln(p_H/p_L)}{\ln(\theta_H/\theta_L)}, \qquad p_i = \frac{\lambda(1-u_i)}{(1-\lambda)\,u_i}, \tag{C.9}$$

   where p_i is the job-finding rate required by (C.2) at each regime mean (the λ terms cancel in the ratio). The data deliver $\xi = 0.487$, close to the value of 0.5 commonly assumed in the literature. We impose the Hosios condition $\eta = \xi$. This identification maintains that matching efficiency and the separation rate are unchanged across regimes, and it depends on the pre-2001 vacancy splice (Appendix A.1), so it is a check rather than an estimate. Appendix C.4 reports the variant that takes ξ from outside and lets matching efficiency and separations differ across regimes: matching efficiency fell rather than rose, and posting costs still account for essentially all of the regime change in tightness.

2. *Separation rate λ.* We set $\lambda = 0.006$ per month, the order of magnitude of UK employment-to-unemployment transition rates; the LFS microdata put the monthly rate at about 0.004 (Appendix C.5), and Appendix C.4 uses the measured rates. This is considerably smaller than the values used in US calibrations, and it is what allows the model to match the low level of measured UK tightness.

3. *Matching efficiency m.* Given ξ and λ, $m = p_H/\theta_H^{1-\xi} = 0.167$ matches the level of the Beveridge curve.

4. *Posting costs κ_H, κ_L.* The job creation condition (C.8), inverted at each regime's tightness, delivers $\kappa_H = 0.216$ and $\kappa_L = 0.110$ (at the baseline b set in step 5) as flow costs per vacancy per month in units of the monthly output of one worker – a required fall by a factor of 1.96.

5. *Flow value of non-employment b.* The moments used so far pin down tightness, the flow rates and the matching function, but they cannot separate the level of posting costs from the level of match rents: at the observed tightness, the free-entry condition (C.8) determines only the *combination* of κ and the surplus $x - b$, with κ proportional to $x - b$. Intuitively, the same observed tightness is consistent with a high-rent, high-cost labour market and a low-rent, low-cost one: if filled jobs are very profitable, only high posting costs can hold entry down to the observed level of vacancies. The parameter b is the worker's outside value – the total flow value of non-employment, comprising unemployment income, home production and the value of leisure and search, not the benefit replacement rate alone – and we calibrate it directly on those grounds, setting $b = 0.95$ following Hagedorn and Manovskii (2008). This choice sets the level of posting costs, and the implied cost per hire offers an order-of-magnitude check against evidence not used in the calibration. The flow cost per hire is $\kappa_H/q_H = 0.58$ months of output, the expected number of vacancy-months to fill a post times the flow cost. The CIPD's 2007 survey puts the median recruitment cost per hire in 2006, covering advertising and agency or search fees, at £4,333, or 0.87 months of that year's output per worker (Chartered Institute of Personnel and Development, 2007). The two overlap only in advertising: the survey adds one-off fees that the flow cost excludes and omits in-house recruiting time, which is a flow cost. The comparison therefore does not pin b tightly ($b = 0.90$ gives 1.2 months), but it rules out low outside values, which overshoot it several-fold (3.4 months at $b = 0.71$) or by an order of magnitude (7.0 at $b = 0.40$; Appendix C.6). Only the level of κ – in output units, and hence in pounds – depends on b; ξ, m, the cost ratio κ_H/κ_L, the tightness elasticity and every object plotted in Figure 6 are invariant to it.

Table C.1 summarises the calibration.

External checks. Three implications of the calibration can be checked against evidence not used in its construction. First, the required fall in total posting costs (factor 1.96) is smaller than the measured fall in advertising expenditure per vacancy (factor 3.18); converting each regime's advertising spend and posting costs into pounds with that regime's average output per full-time-equivalent worker, the advertising decline accounts for roughly half of the required decline in total posting costs, consistent with advertising being one component – and a proxy – of broader recruitment costs such as screening and applicant processing, which fell alongside it. Under the same regime-specific conversion, advertising expenditure accounts for around a quarter (26%) of total posting costs in the high-cost regime and around an eighth (13%) in the low-cost regime. Second, the model's vacancy-filling rates imply vacancy durations that *rise* as posting costs fall, from 2.7 to 3.9 months against the unemployed pool; scaled by the share of hires that come from unemployment (roughly one third of total hires H_t in the LFS flows data, Appendix A.2), this corresponds to all-pools durations of around 4 to 6 weeks – in line with the level and, importantly, the upward drift of

Table C.1: Calibration of the DMP model (monthly frequency)

Parameter	Value	Description	Source / target
β	$0.96^{1/12}$	Discount factor	4% annual discount rate
ξ	0.487	Matching elasticity	Both regime means on one Beveridge curve
η	0.487	Worker bargaining weight	Hosios condition ($\eta = \xi$)
λ	0.006	Separation rate	UK E→U flows (LFS)
m	0.167	Matching efficiency	Level of the Beveridge curve
x	1	Labour productivity	Normalisation
b	0.95	Flow value of non-employment	Hagedorn–Manovskii (2008)
κ_H	0.216	Posting cost, high-cost regime	Free entry at $\theta_H = 0.191$
κ_L	0.110	Posting cost, low-cost regime	Free entry at $\theta_L = 0.403$

Notes: Posting costs are flow costs per vacancy per month, expressed in units of the monthly output of one worker. The calibration is exactly identified: the model's two steady states reproduce the regime means $(u_H, v_H) = (7.80\%, 1.49\%)$ and $(u_L, v_L) = (5.46\%, 2.20\%)$ exactly. The outside value b sets the level of posting costs; calibration step 5 checks the implied cost per hire against UK survey evidence to an order of magnitude.

measured vacancy duration in Figure 3. Third, the implied job-finding rates (0.071 and 0.105 per month in the high- and low-cost regimes) correspond to average completed unemployment durations of around 10–14 months, of the same order as LFS-based estimates over the sample.

Fee timing. Under a fee F paid when a vacancy is posted and a filling hazard q, steady-state advertising expenditure is F times the flow of new postings, FqV, so measured spend per vacancy in the stock is Fq. Free entry requires $F = qJ/(r+q)$ against the fee and $\kappa = qJ/(1+r)$ against a flow cost, so a posted fee is equivalent to a flow cost $\kappa = F(r+q)/(1+r)$, and Fq equals that flow equivalent to within one percent at the calibrated hazards. Measured spend per vacancy is therefore the flow equivalent of the fee rather than an accounting artefact. Its endogenous component is q: at a constant fee, the fall in q as durations lengthened accounts for about a seventh of the decline in spend per vacancy between 2001–09 and the years since, and about a quarter from 1993, when the duration data begin.

C.3 The effect of a fall in the vacancy posting cost

The comparative statics of a fall in κ can be read directly off the two organising properties of Appendix C.1. Because κ enters the job creation condition (C.8) but not the Beveridge curve (C.2), a fall in the posting cost rotates the job creation ray anti-clockwise while leaving the Beveridge *locus* unchanged: the entire equilibrium adjustment is a movement along a fixed Beveridge curve. The equilibrium change itself is far from trivial – the steady state moves, and materially so – but the model confines all of it to the curve.

Differentiating (C.8), the elasticity of equilibrium tightness with respect to the posting cost is

$$\frac{d\ln\theta^*}{d\ln\kappa} \;=\; -\frac{A+C}{\xi A + C} \;\in\; \Big(-\frac{1}{\xi},\,-1\Big), \qquad A = \big[1-(1-\lambda)\beta\big]\frac{\kappa\,\theta^{\xi}}{m}, \quad C = (1-\lambda)\,\beta\,\eta\,\kappa\,\theta, \tag{C.10}$$

where A is the flow-cost component of the marginal entry condition and C the wage-channel

component (posting costs raise the bargained wage through the threat point, which dampens the response and pushes the elasticity towards -1; as $\eta \to 0$ the elasticity approaches the $-1/\xi$ bound). Movements along the Beveridge curve then split the adjustment between the two margins according to $d \ln u^* = -(1-\xi)(1-u)\, d \ln \theta^*$ and $d \ln v^* = [\xi + (1-\xi)u]\, d \ln \theta^*$.

A single feature of (C.10) signs the entire model. Writing $\varepsilon \equiv d \ln \theta^* / d \ln \kappa$, note that $(A + C) - (\xi A + C) = A(1-\xi) > 0$ whenever $\xi < 1$, so $|\varepsilon| > 1$ *strictly*: a fall in the posting cost raises tightness more than one-for-one in logs. Everything else follows.

Proposition 1 (Direction of change)**.** *Consider a permanent fall in the vacancy posting cost κ, with $0 < \xi < 1$ and $0 < \eta < 1$. Then equilibrium tightness θ^*, the job-finding rate $p(\theta)$, the vacancy rate v^*, expected vacancy duration $1/q(\theta)$ and* the wage *w all rise, while the unemployment rate u^*, the vacancy-filling rate $q(\theta)$, the value of a filled job J and the cost per hire $\kappa/q(\theta)$ all fall.*

Proof. $\varepsilon < 0$ gives θ^*; $d \ln p = (1-\xi)\, d \ln \theta$ and $d \ln q = -\xi\, d \ln \theta$ give p, q and duration; the Beveridge elasticities above give u^* and v^*. For the cost per hire, $d \ln(\kappa/q) = (1 + \xi\varepsilon)\, d \ln \kappa$, and $1 + \xi\varepsilon > 0$ because $|\varepsilon| < 1/\xi$, so it moves with κ. For the wage, (C.6) makes w a function of the product $\kappa\theta$, and $d \ln(\kappa\theta)/d \ln \kappa = 1 + \varepsilon < 0$ because $|\varepsilon| > 1$; hence $\kappa\theta$ moves *opposite* to κ and the wage rises. Finally $J = (x - w)/[1 - \beta(1-\lambda)]$ falls because w rises. □

The wage result deserves emphasis, because it is the one case in which two forces offset. A fall in κ lowers the posting cost that enters the worker's threat point directly, but it raises tightness, which enters the same term. Proposition 1 shows that the tightness effect *always* dominates: cheaper vacancy posting raises bargained wages. The result is unconditional in this class of models – it requires only $\xi, \eta \in (0,1)$ – though its *magnitude* is governed by η, since $dw = \eta\, d(\kappa\theta)$, so under the Hosios condition a modest matching elasticity mutes it. At the calibration of Table C.1 the regime change raises the wage by about 0.15%.

Two qualifications matter for interpretation. The result concerns the *bargained* wage in a model with flexible Nash bargaining; wage rigidity or on-the-job search could overturn it. And it is a comparison of *steady states*: as Appendix C.8 shows, along a gradual transition the wage first dips below its initial level before rising above it, because tightness itself dips while firms delay hiring. For a paper concerned with wage pressure, the distinction between the long-run and transitional signs is worth keeping explicit.

At the calibrated high-cost steady state the elasticity is -1.12: halving the posting cost multiplies equilibrium tightness by 2.15. This is the main result of the exercise. Appendix C.6 shows that it is smaller, between -0.9 and -0.7, if training is treated as a one-off hiring cost of the size UK survey evidence suggests. The calibrated fall in posting costs between the two regimes ($\kappa_H \to \kappa_L$, a factor of 1.96) carries the steady state from $(u, v) = (7.80\%,\ 1.49\%)$ to $(5.46\%,\ 2.20\%)$ – tightness roughly doubles, from 0.19 to 0.40 – travelling down the fixed Beveridge curve exactly as in Figure 6. The model thus attributes the secular rise in the V/U ratio to the rotation of the job creation curve, complementing the ECM-based estimates of Section 5. Its stable-Beveridge-curve restriction can also be confronted with the data. The regime means lie on the calibrated curve by

construction; the informative check is that the within-regime quarterly scatter in Figure 6 traces the same downward-sloping locus, with no visible level shift between the two regime clouds. This is a visual check rather than a formal test, and it abstracts from shorter-lived Beveridge-curve movements within regimes (the outward drift of the 1980s, the subsequent inward shift, and the post-pandemic outward move); but it is consistent with the model's restriction that changes in posting costs move the economy along, rather than off, the curve. Appendix C.4 shows that this stability is the net of two offsetting shifts: lower separations move the curve in and lower matching efficiency moves it out, by similar amounts over the full regimes.

A natural question is whether this invariance is an artefact of treating separations as exogenous. It is not innocuous: if job destruction is endogenous – matches below a reservation productivity are dissolved – then cheaper vacancy posting weakens the incentive to hoard marginal workers (replacement hiring is cheaper), the destruction margin rises, and the Beveridge curve shifts *outward* when posting costs fall. A quantitatively dominant shift of this kind, however, would imply *rising* separation rates and rising unemployment alongside falling posting costs – the opposite of the UK experience, in which unemployment fell across regimes and employment-to-unemployment separation rates fell, by about a sixth between the 2000s and 2010s expansions in the LFS microdata (Appendix C.5). The data therefore support the configuration in which the job-creation rotation dominates, and the stability of the empirical Beveridge curve across regimes provides corroboration.

C.4 Posting costs, matching efficiency and separations

The calibration of Appendix C.2 attributes the regime change in tightness to posting costs by holding matching efficiency and the separation rate fixed. This subsection relaxes both. We take the matching elasticity from outside, $\xi = \eta = 0.5$ (Petrongolo and Pissarides, 2001), and set the separation rate in each regime to the monthly employment-to-unemployment hazard in the LFS two-quarter longitudinal microdata (Appendix A.4): 0.0043 over 2001–09, applied to the whole high-cost regime, and 0.0037 since. The Beveridge curve then implies matching efficiency in each regime, free entry implies the posting cost, and the observed change in $\ln \theta$ splits into the parts each driver delivers alone from the high-cost steady state (Table C.2).

Across the full regimes, posting costs fall by about half and account for the whole rise in tightness, while matching efficiency and separations fall by 13 and 14 percent and each contribute about five percent or less, in opposite directions. Between 2001–09 and the years since, posting costs fall by about a tenth and matching efficiency by a sixth, which pulls tightness back by four log points. The asymmetry is a property of free entry at these parameters: the elasticity of θ^* is -1.13 with respect to the posting cost but only 0.26 with respect to matching efficiency, because the wage term dominates the vacancy-cost term in the entry condition. Shifts of the Beveridge curve therefore move unemployment at first order and tightness at second order. Over the full regimes the two shifts nearly cancel: at the low-cost tightness, lower separations alone would lower unemployment by 13% and lower matching efficiency alone would raise it by 14%, a net shift of under one percent, which is why the empirical Beveridge curve looks stable across regimes (Appendix C.3). Since 2001 the matching-efficiency shift dominates, and the curve moves out by about 4%.

This accounting separates the two channels of Section 3. The conventional channel works through free entry and the posting cost. The behavioural channel, postings that are cheap to hold and seldom filled, registers as lower matching efficiency per measured vacancy, and it moves unemployment rather than tightness. The posting cost is inverted from observed tightness here, so the decomposition shows that the fall free entry requires does not depend on the measured changes in matching efficiency and separations; the independent evidence that posting costs fell is the advertising series and the CIPD comparison of Section 4.

Table C.2: The regime change in tightness by driver

Comparison; separation rate	κ	m	$\Delta \ln \theta$	Posting cost	Matching eff.	Separations
Full regimes; measured	−49%	−13%	+0.75	+0.76	−0.04	+0.02
Full regimes; common 0.006	−49%	+1%	+0.75	+0.74	+0.00	+0.00
2001–09 vs since; measured	−9%	−17%	+0.08	+0.10	−0.04	+0.01
2001–09 vs since; common	−8%	−4%	+0.08	+0.09	−0.01	+0.00

Notes: $\xi = \eta = 0.5$, $b = 0.95$. Full regimes: high-cost against low-cost regime; 2001–09 vs since: 2001 Q2 to the end of the high-cost regime against the low-cost regime, both on the ratio of mean rates. Measured separation rates are monthly employment-to-unemployment hazards from the LFS two-quarter longitudinal microdata. The last three columns give the log change in θ^* when each driver alone moves from its high-cost to its low-cost value; interactions, not shown, are at most 0.01.
Sources: ONS; Advertising Association and WARC Expenditure Report; authors' calculations.

C.5 Evidence from the LFS microdata

Person-level transitions in the LFS two-quarter longitudinal data (Appendix A.4) show where the two Beveridge-curve shifts of Appendix C.4 come from. We compare the expansion windows 2001 Q2–2008 Q2 and 2014–19, which exclude recessions and the pandemic.

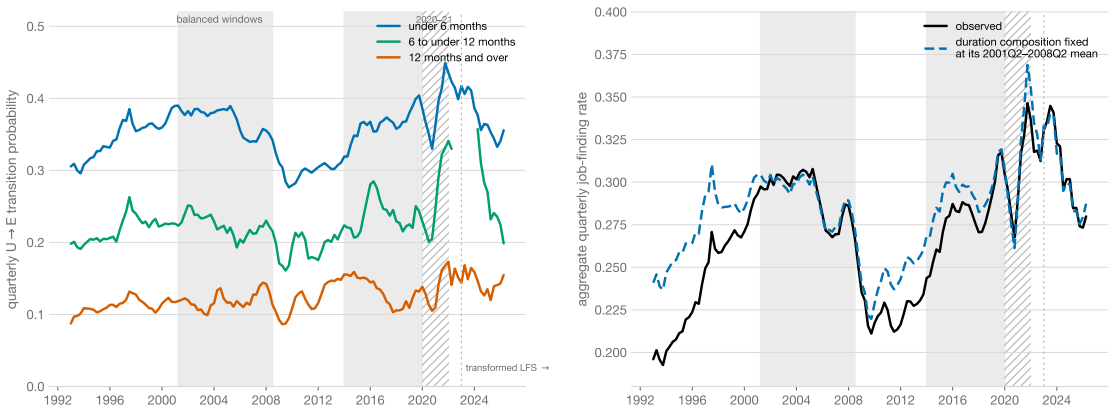
Job-finding. Between these windows the quarterly job-finding probability moves from 0.369 to 0.366 for those unemployed under six months, from 0.225 to 0.239 at six to twelve months, and from 0.120 to 0.133 at twelve months and over, while the twelve-plus share of the unemployed rises from 22 to 29% (Figure C.1). Holding the duration composition at its 2001–08 mean, the aggregate rate would have risen about 1% instead of falling 4%: composition accounts for the entire fall, the mechanism by which Barnichon and Figura (2015) show a deteriorating composition of the unemployed depresses measured aggregate matching efficiency, and the duration gradient is the one Elsby et al. (2011) document to 2010. The fall in measured matching efficiency is therefore not a fall in how quickly a given unemployed worker finds a job. Part of it reflects the changing composition of the unemployed, and the rest more vacancies standing behind each hire, the behavioural channel of Section 3. That remainder is consistent with a genuine decline in matching (Pizzinelli and Speigner, 2017) and with selectivity absorbing cheaper search: with near-costless posting and application, firms and workers meet more and accept less (Martellini and Menzio, 2020; Sedláček, 2014; Linckh et al., 2026).

Separations and flows. The employment-to-unemployment probability falls 17% between the windows, and 14% when the age composition of employment is held at its 2001–08 mean: separation

hazards fall by nearly a fifth for those under 25 and a seventh for those aged 25–49, while the hazard for those aged 50–64 is unchanged, and tenure composition accounts for the same fifth as age. The decline is therefore not demographic. Evaluating $u = s/(s + f)$ at the measured quarterly rates, the inflow margin accounts for 124% of the implied unemployment decline between the windows and the outflow margin for −29%; purged of duration composition, the outflow margin accounts for only 7%. This is the margin that Smith (2011) documents as central to UK unemployment's variation.

Robustness. The shift-share attributes 130% of the change in the aggregate job-finding rate to composition, −22% to within-group rates, which rose, and −8% to the interaction; reweighting by duration, age and sex rather than duration alone moves the fixed-composition change by less than 0.2 log points. Destinations are identified: the unemployment-to-inactivity probability is flat (0.188 against 0.192), and the job-to-job share of hires moves only from 44 to 43%. A synthetic-cohort construction on published duration stocks (`MGSC`, `YBWG`, `YBWH`) gives a relative total-exit rate of 0.35–0.37 over the same windows, matching the microdata.

Figure C.1: Job-finding rates by unemployment duration and the composition counterfactual



Notes: Panel (a): quarterly U→E transition probabilities by elapsed unemployment duration, LGWT-weighted, four-quarter centred moving averages; cells below the 50-observation floor are masked. Panel (b): the aggregate rate, observed and with the duration composition of unemployment held at its 2001 Q2–2008 Q2 mean. Shaded spans: the comparison windows and the 2020–21 response collapse; the dashed rule marks the transformed LFS from 2023, after which unemployed samples are small and the quarterly rates volatile.
Sources: ONS Labour Force Survey two-quarter longitudinal microdata, derived aggregates (Appendix A.4); authors' calculations.

C.6 One-off hiring costs and training

The flow cost κ covers the costs of keeping a vacancy open. Other hiring costs are paid once a match forms, chief among them training. Following Pissarides (2009), let a one-off cost H be paid when the match forms, on top of the flow cost. Because H is sunk once the match forms, it raises the value of a filled job and therefore the worker's threat point, so it enters both free entry and the wage, $w = \eta[x + \kappa\theta + p(\theta)H] + (1 - \eta)b$, and the steady-state entry condition becomes

$$\left[\frac{\kappa}{q(\theta)} + H\right]\left[1 - (1 - \lambda)\beta\right] = (1 - \lambda)\beta\left[(1 - \eta)(x - b) - \eta\,\kappa\,\theta - \eta\,p(\theta)\,H\right]. \tag{C.11}$$

This makes the role of b explicit. With the flow cost held at its calibrated value, reaching the same steady states at $b = 0.40$ would require $H \approx 6.4$ months of output per hire, four times the

CIPD's median labour-turnover cost, which includes training and induction; $b = 0.95$ requires none. Omitting the $\eta\, p(\theta)H$ term would overstate the required H by a factor of about 4.7.

Survey evidence sizes training. US survey evidence compiled by Manning (2011, Table 2) from Barron et al. (1997) values the hours spent recruiting a new hire at 5–14% of a month's pay and those spent training one at several times that. The CIPD's 2007 survey (Chartered Institute of Personnel and Development, 2007) prices both for the UK in 2006: a median recruitment cost per hire of £4,333, covering advertising and agency or search fees (337 responses), and a median labour-turnover cost of £7,750, covering vacancy cover, redundancy, recruitment and selection, training and induction (24 responses). The gap, 0.68 months of 2006 output per worker, overstates training, since it also contains vacancy cover and redundancy.

Treating half or all of that gap as a one-off cost, with the flow cost held at its calibrated value, requires $b = 0.92$ or 0.89 and lowers the elasticity of tightness with respect to the posting cost from -1.12 to -0.89 or -0.73: halving the posting cost then raises tightness 1.8-fold or 1.7-fold rather than 2.2-fold. The fall in the flow posting cost needed to move between the regimes rises correspondingly, from 1.96-fold to 2.5-fold or 3.4-fold, the upper figure close to the measured 3.18-fold fall in advertising expenditure per vacancy. One-off costs also raise the response of the bargained wage to tightness, since $p(\theta)H$ enters the wage, which matters for how tightness passes into marginal cost.

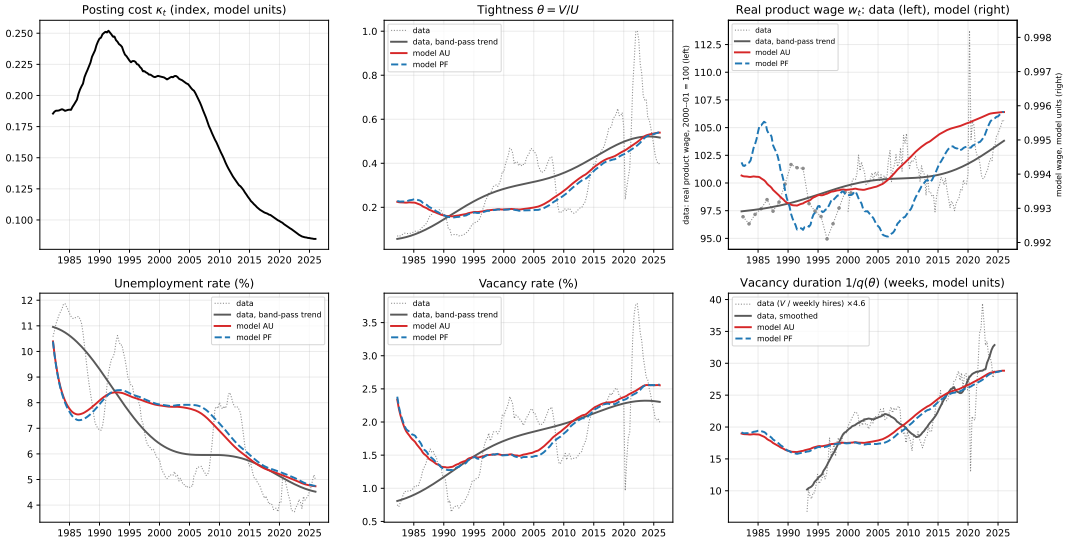
C.7 Transition dynamics

The two-regime comparison treats each regime as a steady state. As a complementary exercise, we feed the full observed time path of vacancy posting costs through the model's transition dynamics. The (smoothed) real advertising cost per vacancy c_t is mapped into the model's posting cost as an index, $\kappa_t = a + s\, c_t$, with (a, s) pinned by the two regime means of κ from Appendix C.2. The mapping is a proxy relationship rather than a level decomposition: the slope s exceeds the direct money-to-output conversion of advertising spend, reflecting that advertising expenditure proxies broader recruitment costs that co-move with it.

We solve the dynamic system of Appendix C.1 directly: the forward-looking job creation condition (C.7) is solved backwards from the terminal steady state under perfect foresight (PF), while unemployment evolves as the state variable (C.1) from its 1982 starting value. We also compute an anticipated-utility (AU) variant in which firms treat the current posting cost as permanent, so that $\theta_t = \theta^*(\kappa_t)$ period by period. The two paths nearly coincide, because the smoothed cost path is close to permanent relative to the firm's effective discounting horizon; the economy tracks the Beveridge curve closely because unemployment converges at rate $\lambda + (1 - \lambda)p \approx 0.09$ per month (a half-life of around eight months).

Driven by this single input, the model reproduces the low-frequency drift of UK labour market tightness: the correlation of model and data tightness is 0.71 quarter by quarter and around 0.8 for the five-year trend, with unemployment drifting down to just below 5% and the vacancy rate rising towards 2.6% (Figure C.2). Two caveats apply. The two regime means are matched by construction,

Figure C.2: Model transition driven by the observed path of vacancy posting costs



Notes: The posting-cost index is a five-year centred moving average of real recruitment advertising expenditure per vacancy (truncated at the sample endpoints), mapped into model units as described in the text. 'Model AU' (anticipated utility) treats the current posting cost as permanent period by period; 'model PF' solves the perfect-foresight path backwards from the terminal steady state. Monthly model output is averaged to quarterly frequency.
Sources: ONS; Advertising Association and WARC Expenditure Report; and authors' calculations.

so a smooth monotone path between them would achieve much of the trend fit mechanically; the model's incremental content lies in the timing and shape of the transition implied by the observed cost path, and in the split of the tightness drift between unemployment and vacancies. And the model is driven by posting costs alone: it is not designed to fit business-cycle fluctuations, which would require cyclical drivers such as productivity or separation-rate shocks (Appendix C.8).

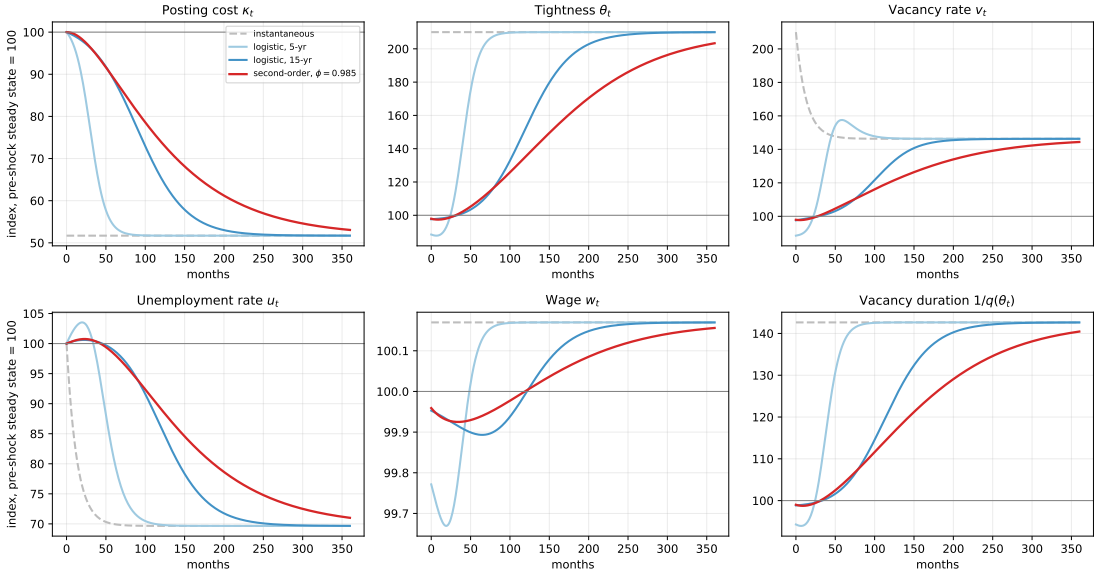
C.8 Impulse responses to a vacancy posting-cost shock

The decline in vacancy posting costs documented in Section 4 is a permanent, technology-driven change that was phased in over two decades, not a one-off jump. This subsection therefore characterises the model's dynamics for *permanent* reductions in κ that arrive *gradually*, which is the empirically relevant experiment. The economy starts in the high-cost steady state, and κ_t falls to κ_L along three alternative paths, all ending at exactly the same terminal value so that the long-run effects are identical by construction:

- a *logistic* (S-shaped) diffusion path completing over roughly five years;
- the same shape completing over roughly fifteen years;
- a damped *second-order* path, $\kappa_t - \kappa_L = 2\phi(\kappa_{t-1} - \kappa_L) - \phi^2(\kappa_{t-2} - \kappa_L)$, which has zero initial slope and so starts flat, accelerates, and then settles – a hump in the growth rate rather than in the level.

An instantaneous fall is retained only as a benchmark. Firms have perfect foresight over the whole path.

Figure C.3: Gradual permanent fall in the vacancy posting cost



Notes: All paths end at the same terminal posting cost κ_L, so the long-run effects are identical; only the timing differs. Panels show percentage deviations from the high-cost steady state of Table C.1. Tightness is a jump variable and unemployment adjusts through the law of motion (C.1).

Three features stand out.

First, anticipation makes firms wait. When the decline is gradual and foreseen, tightness initially falls *below* its old level before rising to the new steady state. This is not a numerical artefact: it follows from the two forward-looking terms in (C.7), which both work in the same direction. Anticipating cheaper posting raises future tightness and hence, by Proposition 1, raises future *wages*, which lowers the value of a job started today; and a vacancy carried unfilled into next period is worth κ'/q', which is itself lower when κ falls. The result is an option value of delay – why hire now at today's cost when posting is about to get cheaper? – and the effect is largest for the paths that are flattest at the start.

This has a direct empirical implication. A gradual fall in posting costs should show up in tightness with a *lag*, and possibly after a small initial dip, rather than contemporaneously. Regressing tightness on the contemporaneous cost measure would therefore understate the effect, which is one reason the error-correction specification of Section 5, with its lag structure and long-run relationship, is the appropriate econometric framework.

Second, the vacancy rate overshoots. After a permanent fall, v rises by more on impact than in the long run, because $v = \theta u$ and unemployment is still at its old, higher level when tightness jumps. As unemployment falls – with a half-life of about six months at the new job-finding rate – vacancies decline towards the new steady state. The model therefore traces an anti-clockwise loop in the Beveridge plane, the same pattern the raw data display around recessions. Gradual phasing damps

the overshoot substantially.

Third, the speed of the transition is governed by the cost path, not by labour-market adjustment. Unemployment converges quickly relative to the decadal horizon over which posting costs actually fell, so the shape of θ_t essentially inherits the shape of κ_t. This is why the sequence-of-steady-states approximation used in Appendix C.7 performs well, and it explains why the model attributes the secular rise in the V/U ratio to the *trend* in posting costs rather than to transitional dynamics.

Which other shocks would the model need? Posting-cost shocks move the economy along the Beveridge curve at low frequency, but they are one column of a fuller shock menu, and each of the standard DMP shocks targets a distinct feature of the data that posting costs cannot generate:

- *Aggregate productivity shocks* (x_t) are the canonical cyclical driver of job creation: like κ_t they move the economy along the Beveridge curve, and under the small-surplus calibration used here ($b = 0.95$) they are strongly amplified (Hagedorn and Manovskii, 2008). They would be needed to fit the business-cycle variation in tightness that the posting-cost path deliberately leaves unexplained.

- *Separation-rate shocks* (λ_t) are needed for the sharp recession spikes in unemployment inflows – the exogenous-λ model adjusts u only gradually through a lower job-finding rate – and they generate the pronounced counter-clockwise Beveridge loops observed in the early 1990s, 2008–09 and 2020.

- *Matching-efficiency shocks* (m_t) *shift* the Beveridge curve itself, as – through the separation margin – do λ_t shocks; what distinguishes m_t is that it moves the curve at unchanged separation rates, generating persistent BC displacement without a spike in inflows. Such shocks would be required to capture outward BC movements such as the post-pandemic mismatch episode, and they matter for the interpretation of the secular trend: if online recruitment raised m as well as lowering κ, part of the tightness drift attributed here to posting costs would instead reflect matching efficiency (the identification caveat of Appendix C.2).

- *Discount-rate or financial shocks* (β_t) operate through the same job-creation channel as productivity, providing a financial-cycle channel into hiring (Hall, 2017); like κ_t and x_t shocks they move the economy along the Beveridge curve and are observationally similar to x_t shocks in this model.

A full business-cycle version of the model would thus combine the measured κ_t path with productivity and separation shocks (and, for Beveridge-curve shifts, matching-efficiency shocks), with the posting-cost series doing the work it is uniquely suited to: pinning down the low-frequency movement of the job creation curve from direct measurement rather than from a residual.

D Price evidence on vacancy posting costs

D.1 Why the WARC series cannot validate itself

The posting-cost index of Section 4 is constructed as $\kappa_t = (A_t/P_t)/V_t$, so that in logs

$$\ln \kappa_t = \ln(A_t/P_t) - \ln V_t. \tag{D.1}$$

The vacancy stock therefore enters the cost measure with a negative sign, and it is also the dependent variable of the long-run relationship (8). This is a property of the construction, not of the data, and it has three consequences that must be confronted before the series can be used as evidence.

First, the sign of the long-run coefficient is close to guaranteed. If real recruitment advertising expenditure A_t/P_t were constant, (D.1) would reduce to $\ln \kappa_t = c - \ln V_t$ and a regression of $\ln V_t$ on $\ln \kappa_t$ would return exactly -1. We verify this numerically: holding A_t/P_t at its sample mean and letting only V_t vary delivers a coefficient of -1.000. The estimated coefficient is well away from that bound, so expenditure carries real information, but the null against which the estimate should be judged is not zero.

Second, the same identity contaminates the descriptive correlation between the cost measure and tightness. Since $\ln \theta_t = \ln V_t - \ln U_t$, the covariance $\mathrm{Cov}(\ln \kappa_t, \ln \theta_t)$ contains $-\mathrm{Var}(\ln V_t)$ by construction. Over 1982 Q2–2024 Q3 the raw correlation is -0.61, but the term in $-\ln V_t$ accounts for 114% of the covariance: the numerator on its own, $\mathrm{corr}(\ln(A_t/P_t), \ln \theta_t)$, is $+0.11$. Real UK recruitment advertising expenditure has essentially no unconditional relationship with labour market tightness.

Third, and decisively for what follows, *no measure built from the WARC data can settle the question*, because every such measure divides by the vacancy stock. Nor does dividing by hires rather than vacancies solve it: cost per hire removes the exact identity, but hires and vacancies correlate at 0.90 in our sample, so the mechanical benchmark falls only from -1.000 to -0.983. Establishing that the measured secular decline reflects a genuine fall in the cost of posting a vacancy requires a cost measure containing no equilibrium quantity at all.

Three such measures exist, and none is a perfect measure of κ. The UK has a producer price index for television and print advertising space from 2000 (Section D.2); Germany has a survey measure of advertising cost per hire from 2000 (Section D.3); and the United States has a producer price index for newspaper advertising space from 1980 (Section D.4). The two price indices do not cover the online channel and the German survey does not separate price from intensity, so each is partial in a different way. What they share is that the vacancy stock does not appear in any of them, so their agreement with the WARC series is informative in a way that no internal check on the WARC series can be. That – rather than replication for its own sake – is why we assemble them.

It is worth being precise about what each source can and cannot do, because they address different problems. Measured advertising spend per vacancy confounds three things: the identity above; the *price* of a posting against the *quantity* of advertising bought per vacancy (recruiting intensity); and

secular movement against cyclical movement. Table D.1 summarises which source separates which.

Table D.1: What each source can separate

	UK spend (WARC)	UK price (HUFK)	Germany (BIBB)	US price (WPU)
Free of the $1/V_t$ identity	no	**yes**	**yes**	**yes**
Separates price from intensity	no	partial	no	partial
Distinguishes secular from cyclical	phases only	**yes**	**yes**	**yes**
Spans the print-to-online transition	**yes**	**yes**	**yes**	**yes**
Long sample (pre-1990)	**yes**	no	no	**yes**
Prices the online channel	**yes**	no	**yes**	no
Recruitment-specific	**yes**	no	**yes**	no
Domestic to the paper's question	**yes**	**yes**	no	no

Notes: "Free of the $1/V_t$ identity" means the cost measure is not constructed by dividing by the vacancy stock. HUFK is ONS SPPI M7312, *Media representation services*; WPU is BLS WPU361102, *Advertising space sales in newspapers*. Both are producer prices for traditional advertising space – HUFK covers television and print only and WPU newspapers only – so neither prices the online channel, and both separate price from intensity only within the channels they cover. The German measure is a survey-reported advertising cost per hire divided by a survey-reported vacancy duration; neither input uses the vacancy stock. No single source satisfies every criterion, which is why they are used together.

D.2 A UK price index for television and print advertising space

A UK price series exists, but its coverage must be stated precisely before it is used, because it does not price the channel that recruitment advertising moved to.

The ONS Services Producer Price Index includes series HUFK, *M7312 Media representation services*, quarterly from 1996 Q1. The methodology article accompanying its introduction (Luke, 2005) states that it "measures the changing cost of the space an advertisement takes up, be it air-time on a television station or column inches in a newspaper". For printed media it is built from British Rate and Data rate-card prices divided by circulation, sampled across twelve strata, with weights from Advertising Association revenue fixed in the base year 2000.

What it covers. The index spans *television and printed media only*, which together the article puts at "75 per cent of the revenue generated by the industry". On the base-year industry weights, newspapers and magazines are 47.7 and television 26.8; direct mail (8.0), directories (5.9), outdoor (5.6), radio (4.0), cinema (0.9) and *internet* (1.1) are outside the index. Two implications follow, and both cut against a naive reading. Roughly 36% of the index is television, a medium recruitment advertising essentially never uses. And internet advertising, which carried a weight of 1.1 in 2000 and is the channel to which recruitment subsequently migrated, is excluded altogether. We have not been able to establish from published ONS documentation whether coverage was extended after 2005, and we do not assume that it was.

HUFK is therefore *not* a price index for the advertising that firms doing the hiring actually bought. It

is a price index for the channels they were leaving. The part of it that is recruitment-relevant is the classified strata – national newspapers classified (weight 3.7), regional newspapers classified (12.5), business magazines classified (3.5) and consumer magazines classified (1.1), about 28% of the index – which the article notes are "used primarily for advertising jobs, properties and educational courses". Those strata are constructed by ONS but only the M7312 aggregate is published.

What it can nonetheless establish. Two things, both worth having.

First, it is a *price*, containing neither the vacancy stock nor hires, so it is free of the identity (D.1). That property is unaffected by its coverage.

Second, it bounds how much of the measured decline can be attributed to the old channel simply becoming cheaper. Writing measured cost per vacancy as

$$\kappa_t = p_t \, x_t, \tag{D.2}$$

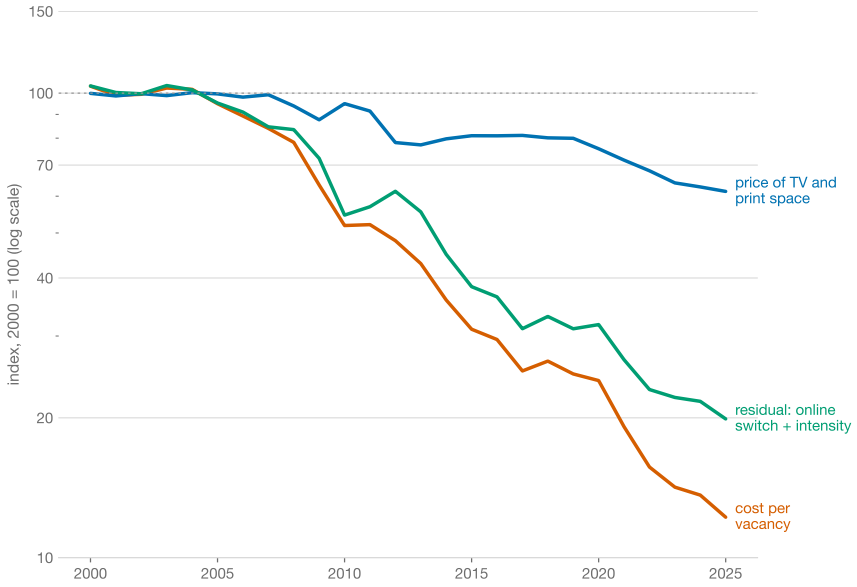
with p_t the real price of traditional advertising space, we report the factorisation in Figure D.1. Between 2000 and 2024 real cost per vacancy falls 84% while p_t falls 37%, so the old channel's own repricing accounts for about a quarter of the decline in logs even if one attributes the entire fall in p_t to recruitment. The remaining three quarters is the residual x_t, and this decomposition cannot say how much of it is substitution towards a cheaper unpriced channel and how much is a fall in advertising bought per vacancy. What it does establish is that the decline in measured posting costs is *not* mostly the old channel discounting: something other than the price of print and television space is doing the work, and the two candidates are precisely the technology substitution the paper proposes and the recruiting-intensity story that competes with it.

The price series also delivers the statistic the WARC-based measure cannot. Over 2000–2024, $\mathrm{corr}(\ln \kappa_t, \ln \theta_t) = -0.44$ but contains $-\ln V_t$; the corresponding correlation using the price, $\mathrm{corr}(\ln p_t, \ln \theta_t) = -0.34$, contains no equilibrium quantity at all. Estimating the long-run relationship with the price in place of the constructed cost index gives a coefficient of -1.21 (HAC standard error 0.69). This is imprecise – the sample is 99 quarters rather than 170 – but it is the correct sign, and it is measured against a null of *zero* rather than against the mechanical null of -1 that applies to (D.1).

The coverage problem has a known direction here. If $\ln p_t$ is treated as a noisy proxy for the true price of posting a vacancy, the resulting attenuation biases the coefficient *towards zero*, so -1.21 understates rather than overstates the relationship. That argument requires the measurement error to be classical, which partial coverage need not deliver, so we present it as a direction of bias rather than a correction. We note also that -1.21 sits close to the elasticity of tightness with respect to the posting cost of about -1.1 implied by the calibrated model of Section 6, which is derived from the free-entry condition and the two regime means without using the regression at all.

Three limitations remain, and the first is the binding one. `HUFK` prices television and print rather than the online channel that came to dominate recruitment advertising, so it can bound the contribution

Figure D.1: Decomposing the UK posting-cost decline into price and quantity



Sources: Advertising Association and WARC Expenditure Report; ONS SPPI series HUFK; ONS GDP deflator; authors' calculations.
Notes: Log scale, 2000 = 100. The series begins in 2000 because HUFK is carried back at a constant 90.5 for all sixteen quarters from 1996 Q1 to 1999 Q4, reflecting the CSPI base year rather than observed prices.

of the old technology's repricing but cannot price the new technology. Obtaining the four classified strata from ONS would give a genuinely recruitment-relevant UK price index and is the single most valuable extension of this exercise. Second, the series does not extend back to 1982, so it cannot be used across the paper's full sample. Third, discounts are applied to display advertising only; the article notes that "discounting factors for classified advertisements are not available", so the classified component is closer to a list-price basis – the same valuation convention as the US series, which helps comparability between the two but means neither is net of the discounts large advertisers negotiate.

D.3 Germany: advertising cost per hire from the BIBB survey

The German *Kosten-Nutzen-Erhebung*, run by the Bundesinstitut für Berufsbildung (BIBB), asks firms directly for their *Inserierungskosten* – advertising and posting outlays per newly hired skilled worker – separately from the personnel cost of the interview process and separately from external consultant and agency fees. The 2007 questionnaire (Q138) defines the item as the cost of press advertisements, enquiries to the employment agency, internal job postings and notices, per skilled worker to be hired. This is the closest counterpart to the UK recruitment advertising numerator that exists in any European survey, with the decisive advantage that it is reported per *hire* rather than per unit of the vacancy stock.

Converting to the model's flow cost uses cost per hire = $\kappa/q(\theta)$, so that

$$\kappa_t^{DE} = \frac{\text{Inserierungskosten}_t}{\text{vacancy duration in months}_t}, \tag{D.3}$$

with the denominator taken from the IAB-Stellenerhebung's *tatsächliche Besetzungsdauer*. Both inputs are survey measurements. Neither is constructed from the vacancy stock, so κ^{DE} is free of the identity (D.1).

The survey has five waves. BIBB publishes nominal euros of the survey year, and the German GDP deflator rises by roughly a half over the sample, so the series must be deflated before the waves are compared; Table D.2 reports both.

Table D.2: German advertising cost per filled vacancy, BIBB *Inserierungskosten*

Wave	Nominal (€)	2023 prices (€)	Duration (days)	κ^{DE} (€/month)
2000	684	1,060	78.0	414
2007	432	617	73.0	257
2012/13	275	365	83.7	133
2017/18	358	438	82.2	162
2022/23	656	679	96.7	214

Sources: BIBB *Kosten-Nutzen-Erhebung*, weighted published figures (Walden and Herget 2002; BIBB REPORT 8/2009; Schönfeld et al. 2016; BIBB REPORT 1/2020; BIBB REPORT 2/2025); IAB-Stellenerhebung; authors' calculations.
Notes: Deflated by the German GDP deflator, 2023 = 1; half-year waves use the average of the two calendar-year deflators. Figures are weighted and representative, conditional on having hired a skilled worker externally, and include firms reporting zero spend.

Deflation matters for the reading, not merely for the levels. In nominal terms the final wave sits only 4% below the first, which makes the series look like a round trip. In constant prices the fall to 2012/13 is −66% and the 2022/23 wave is still 36% below 2000: the secular decline dominates and the post-2012 rise is a partial retracement.

Three caveats, all from BIBB's own documentation. The series covers the *mid-skill segment only – Fachkraft* denotes occupations normally requiring a completed dual apprenticeship, so graduate, professional and managerial vacancies are excluded, and the measure is narrower than the UK economy-wide series. The basis *breaks at 2017/18*: from that wave costs refer to the last skilled worker hired, whereas in 2007 and 2012/13 firms estimated an average over the preceding three years, so part of the 2012/13–2017/18 rise may be definitional. And there are only five unevenly spaced waves; the annual series used below interpolates between them in logs with a monotone spline, which manufactures no information and whose departure from log-linear interpolation never exceeds 7%. Weighting convention also matters: BIBB's weighted 2007 figure is €432 where an unweighted version of the same wave and variable is €563, and we use the weighted figures throughout.

D.4 The United States: expenditure and price

The United States supplies the opposite pairing – a long expenditure series that fails, and a long price series that does not.

Expenditure. The Newspaper Association of America published annual newspaper classified advertising expenditure by category, including an *Employment* line, from 1975 to 2011. It is the only US recruitment advertising expenditure series in existence and the direct counterpart of the UK numerator. Employment classified spend peaks at $8.7bn in 2000 and falls to $0.74bn by 2011. Constructing $\kappa^{US} = (A/P)/V$ on the Barnichon composite vacancy series reproduces the UK exercise – and it fails. Tightness is *lower* in the low-cost regime than in the high-cost regime, the low-cost cloud sits inside the high-cost one in the Beveridge plane, and the two-point identification returns an inadmissible matching elasticity. Two features of the US data explain much of this: the NAA series is drawn from a panel of around seventeen companies rather than a census, and it carries a documented methodology break over 1991–94 with the category split unpublished for 1992 and 1993.

Price. The Bureau of Labor Statistics publishes a producer price index for *advertising space sales in newspapers* (series `WPU361102`), monthly from December 1980. This is a price collected from sellers with the advertising unit held fixed, not an expenditure total, and it is therefore much closer to the model's κ than any spend measure can be. It contains no vacancy stock.

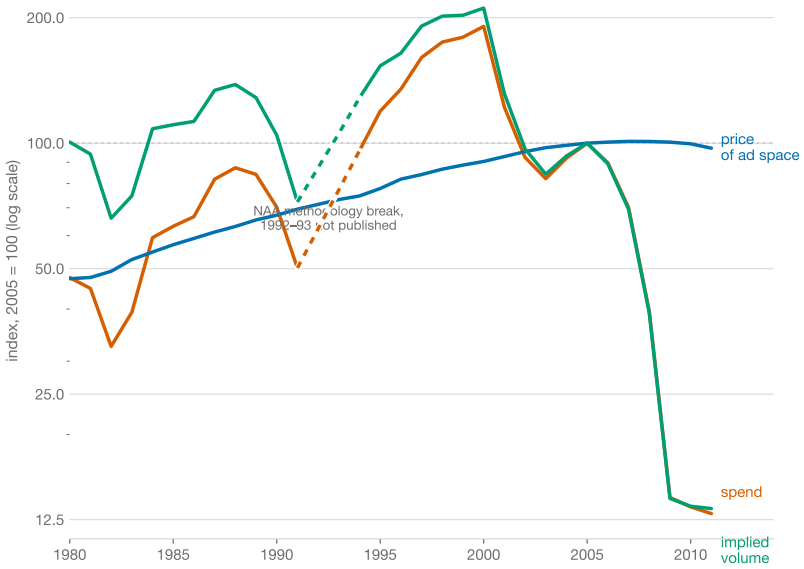
Deflating by the GDP deflator gives the result in Figure D.2. The real price of newspaper advertising space *rose* by about 75% between 1985 and 2005 and was still near its peak in 2010 – precisely while US recruitment advertising expenditure was collapsing by 93% from its 2000 level. Implied real advertising volume falls by 94% over 2000–2011 while the real price is roughly flat. Essentially the entire fall in expenditure is quantity.

This matters for the interpretation of the UK decline. The fall in recruitment advertising expenditure was not incumbents cutting prices into a slack market – print advertising became *more* expensive in real terms while employers left it. That is the signature of substitution towards a cheaper technology, which is the mechanism the paper proposes, and it is a mechanism whose timing is governed by the diffusion of online platforms rather than by the labour market cycle.

The price series also reverses the sign of the US result. Using the expenditure-based measure, $\mathrm{corr}(\ln \kappa^{US}, \ln \theta) = +0.56$; using the real price, corr $= -0.28$, the same sign as the UK. In first differences the price series is essentially acyclical (+0.11) while the expenditure measure remains strongly procyclical (+0.50), which is what one expects if expenditure per vacancy is tracking recruiting intensity and the price is moving for technological reasons.

Two limitations, stated plainly. `WPU361102` covers *all* newspaper advertising space, not help-wanted specifically; no recruitment-specific price index exists in the United States, and the Census Bureau's classified advertising revenue series carries no help-wanted split. And it prices the medium being *abandoned*, so it says nothing directly about the price of the online substitute asserted to

Figure D.2: US recruitment advertising decomposed into price and quantity



Sources: Newspaper Association of America, *Annual Newspaper Classified Advertising Expenditures* (Employment); BLS PPI `WPU361102`; BEA GDP deflator; authors' calculations.
Notes: Log scale, 2005 = 100. Implied volume is real expenditure divided by the real unit price. The dashed segment spans the NAA methodology break, for which the category split was not published in 1992 and 1993.

be cheaper. It bounds the argument rather than completing it. A rising real price for a shrinking product is also consistent with a declining high-fixed-cost industry repricing as volume falls, and we do not claim to separate that from a pure substitution story.

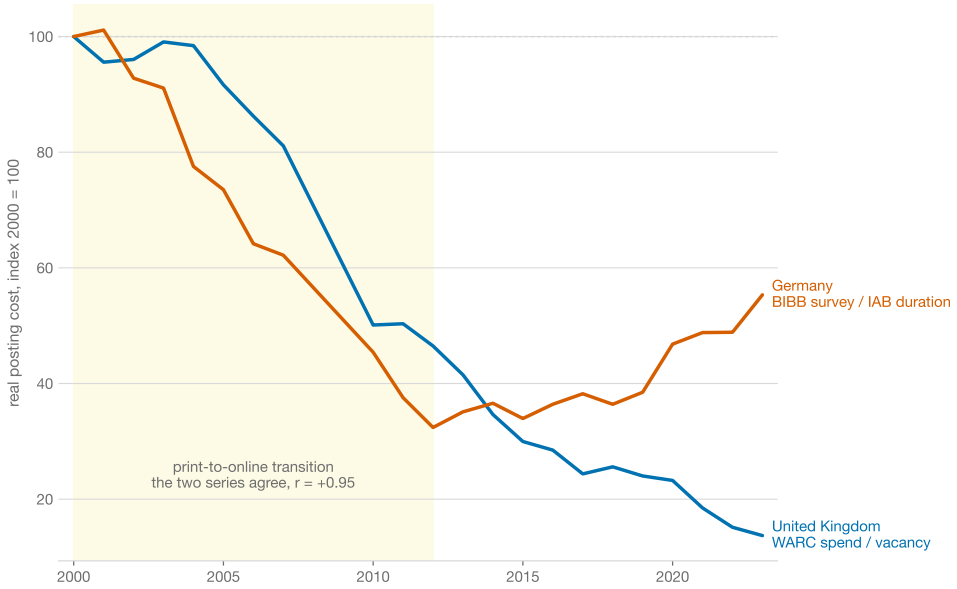
We note for completeness that no US employer survey has ever collected advertising spend as a separable dollar figure. The EOPP-lineage surveys (EOPP 1980 and 1982, SBA 1992, Upjohn 1993) measure recruiting as a single bundled hours count. The contrast within the 1992 SBA instrument is instructive: Q48 asks for "the total number of hours spent by your company personnel recruiting, screening, and interviewing all applicants", while Q63A asks "how much did the company spend" – but only for off-site training. The designers knew how to ask an expenditure question and did not ask it of hiring. There is consequently no US analogue to the BIBB item.

D.5 What the international evidence establishes

The decisive comparison is between the UK series and the German one, because the German measure is free of the identity (D.1) while the UK measure is not. Figure D.3 plots them together over their 2000–2023 overlap.

Over 2000–2012 the two series move together closely: the correlation of their logs is +0.95, and both fall by roughly half. This is the useful result. A measure that never divides by the vacancy stock, built from two independent German surveys, traces the same secular decline as the UK measure that does. The UK decline over the print-to-online transition is therefore not an artefact of the construction (D.1).

Figure D.3: UK and German real posting costs over the print-to-online transition



Sources: Advertising Association and WARC Expenditure Report; BIBB *Kosten-Nutzen-Erhebung*; IAB-Stellenerhebung; authors' calculations.
Notes: Both series in constant prices, indexed to 2000 = 100. The German series interpolates between five survey waves in logs. The shaded region marks the period over which the two series are compared.

After 2012 the two series separate. The UK continues to fall, reaching 86% below its 2000 level by 2023; Germany troughs in 2012 and then rises, ending 45% below. Over the full overlap the correlation of logs falls to +0.68 and in log differences to +0.29. The German reversal coincides with German tightness reaching record highs, and within the German series the two phases carry opposite signs: $\text{corr}(\ln \kappa^{DE}, \ln \theta^{DE}) = -0.28$ over 2000–2012 and +0.70 over 2013–2023.

We read this as a genuine and important limitation rather than a puzzle. Advertising spend per hire is the product of the price of a posting and the intensity with which firms advertise, and when hiring becomes difficult firms advertise more, in more places, and for longer. The German data show that the intensity term can dominate the price term over a decade. The identifying assumption behind any advertising-based cost measure – the UK series included – is therefore that the *secular* component reflects price while the cyclical component reflects intensity. The German evidence supports that assumption over the transition and cautions against it thereafter; the US price index supports it directly, being acyclical in differences.

Two further points of interpretation follow.

The German data cannot deliver an independent estimate of the matching elasticity, and the reason is instructive. Splitting the German sample at the median cost gives regime means of $u = 9.25\%$, $v = 2.72\%$ in the high-cost regime and $u = 4.16\%$, $v = 2.64\%$ in the low-cost regime. The move is almost purely horizontal: unemployment more than halves while the vacancy rate is essentially unchanged. A movement *along* a Beveridge curve must trade vacancies against

unemployment; this is a curve that has shifted inward, consistent with the Hartz reforms of 2003–05 and the subsequent fall in German structural unemployment. The two-point estimator attributes the whole move to the matching technology and returns $\xi = -0.11$, outside $(0, 1)$. Germany violates the common-matching-technology assumption on which the UK identification rests, by construction, and so corroborates the mechanism without being able to quantify it.

Finally, the German level provides an untargeted check on the UK calibration. German measured cost per hire in 2021 is 0.17 months of output per worker, against the UK calibration's 0.58 in the high-cost regime and 0.42 in the low-cost regime. Since the UK level is pinned by the choice of the flow value of non-employment b, this suggests $b = 0.95$ places UK hiring costs on the high side. The comparison is indicative rather than decisive: the IAB question covers own labour time plus direct outlays for a single *Stellenbesetzung*, whereas the UK figure is constructed from advertising expenditure on a different basis, and the two are not like for like.