



## Employ America Research Report

# A Vacant Metric: Why Job Openings Are So Unreliable

Preston Mui  
Senior Economist,  
Employ America

## Executive Summary

*Many who warn of the risk to inflation of excessively tight labor markets rely on measures of job vacancies as an indication of labor market strength. This is bad practice both theoretically and empirically. Not all indicators are created equally, and in today's environment, vacancies are a poor proxy for labor market strength.*

*On the theory side, vacancies are prominent due to highly-stylized models where vacancy posting drives labor market dynamics. These models ignore many important ways that firms adjust their hiring, such as recruitment intensity. Many of the assumptions and implications of these models do not hold up in the establishment-level data.*

*Worse, the empirical sources for vacancy data are shaky at best. The threshold for reporting a vacancy in official statistics is low, making it difficult to infer search intensity from the raw number of vacancies, potentially inflating the number of observed vacancies. Changes to the cost of job posting and job searching – especially on online job boards – make comparisons across time difficult.*

*The Federal Reserve has given job vacancy data center stage in assessing the strength of the labor market. The theoretical and empirical issues with vacancies data show that this is a mistake. Worse, it is a mistake that embeds a clear directional bias today.*


*While labor market indicators in general show a strong labor market recovery from the pandemic recession, vacancies-based metrics portray a far tighter labor market than other indicators. If the Fed fixates on openings, they are setting themselves up to deliver substantially more policy tightening than is warranted by more traditional – and reliable – indicators. This is the first of a series of pieces on the hazards involved in using job vacancy data in economic policymaking.*

## Introduction

As much recent commentary has noted, many measurements of job openings and vacancies seem to paint a picture of labor markets that have gone beyond red-hot. This week's JOLTS data is no exception. However, all is not as it seems. If the Fed plans to calibrate policy to vacancy data – rather than more robust macroeconomic indicators like the level of employment, nominal wage growth, and the pace of hiring and voluntary turnover – it will likely over-tighten monetary policy relative to actual labor market conditions.

After a years-long pandemic and energy crisis, most countries are facing their highest inflation rates in decades. This high inflation coupled with a rapid recovery in the labor market during the pandemic recovery has prompted worries that today's inflation may continue, triggering a wage-price spiral (a concern which we have shown to be overblown). As a result, the Fed is raising interest rates with the goal of bringing inflation down by cooling off the labor market.

In arguing for this course of action, policymakers at the Fed have consistently pointed to an elevated level of job openings as evidence that labor markets are overheated. In his July 2022 press conference, Jay Powell repeatedly cited the high ratio of vacancies to unemployed workers in arguing that the Fed can bring down inflation without inflicting too much damage to the labor market.




Source: Bureau of Labor Statistics

**“From the standpoint of our congressional mandate to promote maximum employment and price stability, the current picture is plain to see: The labor market is extremely tight, and inflation is much too high.”**

*- Jay Powell, June 15, 2022 FOMC Press Conference*

Recently market commentators and academic researchers have made similar arguments: [Domash and Summers \(2022\)](#) argue that the recent elevated job vacancy rate is likely to lead to high inflationary pressure in the near future, while [Barnichon and Shapiro \(2022\)](#) argue that the vacancy-to-unemployment ratio (VUR) is the best measure of economic slack, based on its purported ability to forecast wage and price inflation.

This focus on vacancies points towards a far tighter labor market than other indicators suggest. While all metrics point towards a strong labor market recovery, the extent of that recovery differs across indicators. Prime-age employment, unemployment, and the quit rates have all more-or-less returned to their pre-pandemic levels. Unemployment is at its pre-pandemic low, while quits are slightly higher and employment is slightly lower. Vacancies, however, show a far hotter labor market. The vacancy rate in July 2022 was 6.9%, far above its pre-pandemic high of 4.7%.



Source: Bureau of Labor Statistics

The problem is, vacancies are simply not a very good labor market indicator. There are empirical problems in the way datasets are collected, and theoretical problems in the way they are used to make arguments. The canonical models that rely on them use a vision of the labor market that may be sufficient for academic work, but which is insufficient for the needs of real-time policymaking. In this piece, we'll walk through both theoretical and data-related reasons to be wary of arguments for tightening that rely heavily on vacancy data. In later pieces in this series, we will look at how well vacancies perform as an empirical measure of the state of the labor market.

### The Scope of the Problem

There are three key problems that make it difficult to use reported job vacancies when thinking about the state of the labor market. First, the models that give significant importance to VUR as a labor market indicator rely on highly abstracted models of job search and unemployment. These models greatly simplify many salient features of real-world job search and unemployment activity. Most notably, these models leave out any variation in job-seeking effort, movement between employers, and hiring from labor force non-participants, three dynamics that play important and empirically visible roles in labor market flows.

Second, job openings statistics are not a pure measure of labor demand. In order for a vacancy to be reported in the data, an establishment needs to take active steps to recruit for an open position, steps that have likely become easier as hiring technology has improved over time. What this means is that a measure of the number of vacancies does not account for the extent to which employers expend effort to fill a vacancy. More importantly, this methodology leaves a simple reliance on the vacancy data vulnerable to technological changes in the job posting and recruitment process. Unlike measurements of employment or quits – which have clear definitions – vacancy rates are uniquely empirically hazy.

Third, the microempirical evidence shows that actual job hiring behavior differs dramatically from existing models that give a central role to VUR. The cross-sectional behavior of vacancies and hires clearly demonstrates that differences in hiring at the firm level aren't explained by differences in job openings. There is substantial heterogeneity in both the levels and volatility of vacancy-filling rates across industries, and a large number of hires occur without corresponding vacant positions.

What, then, should policymakers pay attention to when thinking about the state of the labor market? As we have [explained before](#), “full employment” is an inherently dynamic and multi-dimensional concept. No single indicator can fully capture all aspects of the labor market. As such, the Federal Reserve should pay

attention to a variety of indicators, such as the unemployment rate, the prime-age employment-to-population ratio, and the employment cost index. Quits and hires also contain some potentially useful information unaffected by the conceptual issues that plague the vacancy data, precisely because the observations are not distorted by employer behavior. Given the conceptual and empirical issues with vacancy data, Federal Reserve officials simply put more intellectual weight on vacancy data than the data can bear.

### Why are Vacancies Supposed to Matter?

On its face, the VUR is an appealingly straightforward metric. In theory, vacancies purport to measure unmet labor demand, and unemployment represents unmet labor supply. The ratio of the two then should, in theory, provide a measure of the balance of labor demand and labor supply. But if these are proxies in theory, then what theory lies behind them? When commentators appeal to the VUR as a measure of labor market tightness, they are often appealing to some version of the Diamond-Mortensen-Pissarides (DMP) model, which has become tremendously influential in the academic macro-labor literature.


In the DMP model, firms only hire by opening vacancies, which unemployed workers encounter randomly. Since matching is random, the VUR is a key variable in the model; the higher the VUR, the lower the chance a vacancy is filled and the higher the chance an unemployed person finds a job. The VUR also drives wages in the model, since a higher VUR gives the unemployed stronger bargaining positions when wages are negotiated. An important aspect of the textbook DMP model is the assumption that all labor market adjustment can only happen through the posting and filling of vacancies.

These models have become popular because they are flexible and tractable. The DMP model provides a concise modeling environment for unemployment dynamics that is simple to work with, is easily integrated with existing macroeconomic models, and provides a lens with which to analyze a wide variety of macroeconomic questions. Within this frame, the VUR is the main determinant of wages and employment. However, the tractability of the basic model means it abstracts from a number of important features of the labor market in ways that seriously compromise its adequacy as a guide to policy making.

It is no overstatement to say that every change in labor markets must pass through a change in vacancies within this model. This simplification means that the only way for an employer to control its hiring is to change their number of vacancies; there is zero scope for companies to affect hiring by changing wages, changing hiring standards, or adjusting recruitment efforts. Worse, it models all hiring as movement of workers from unemployment to employment when, in reality, fewer than 20% of hires actually result from an unemployed person finding a job. In fact, nearly half of new hires come from workers who are

considered out of the labor force; that is, they do not have a job and do not report actively searching for a job.

Andolfatto and Birinci (2022) show that not accounting for the employed who make job-to-job transitions substantially overstates the tightness of the labor market. To account for these movements, they compare the VUR to the ratio of vacancies to unemployed plus employed workers that change jobs.



Source: Andolfatto and Birinci (2022)

What stands out about this correction is the fact that the gap between these two measures varies over the business cycle. During strong labor markets, employer-to-employer movements are more frequent, meaning that the gap between these two measures of labor market tightness widens during recoveries and tightens during recessions. If the accuracy of the measurement of labor market tightness varies with labor market tightness, it might not be the best tool for calibrating policy.

Empirically, these employment movements are an important feature of the labor market. Moscarini and Postel-Vinay (2017) find the rate of job-switching is far more related to wage growth than the rate of hires from unemployment, while Gertler, Huckfeldt and Trigari (2020) find that the wages of job switchers are more cyclical than new hires from unemployment or continuing workers. Job switching may be indicative of workers finding more productive jobs, and there is value in those new matches even if they don't result in a hire from unemployment. In any event, the model "simplifies away" certain empirical aspects of today's labor markets that policymakers would do well to retain.

## How are Vacancies Measured?

The most commonly used source of vacancy data is the Job Openings and Labor Turnover Survey (JOLTS), administered by the Bureau of Labor Statistics (BLS).

The BLS defines a “job opening” as:


- **A specific position exists and there is work available for that position. The position can be full-time or part-time, and it can be permanent, short-term, or seasonal;**
- **The job could start within 30 days, whether or not the employer can find a suitable candidate during that time;**
- **The employer is actively recruiting workers from outside the establishment to fill the position. Active recruiting means that the establishment is taking steps to fill a position. It may include advertising in newspapers, on television, or on the radio; posting Internet notices, posting “help wanted” signs, networking or making “word-of-mouth” announcements; accepting applications; interviewing candidates; contacting employment agencies; or soliciting employees at job fairs, state or local employment offices, or similar sources.**

While this definition is wordy, it is not precise. The threshold for “actively recruiting” is very low, and the vacancy data does not distinguish between low and high recruiting efforts on the part of companies. A business making a short Facebook post about a job opening is treated the same as a business that hires a recruiter and sends interviewers to job fairs. Accounting for recruitment intensity may be crucial to understanding the behavior of hires data over multiple business cycles.

For another example of the fragility of this kind of measurement, consider the Help Wanted Online (HWOL) series from the Conference Board. This data series tracks the number of online job postings, a large number of which happen on a few job boards such as Craigslist. Starting in 2004, Craigslist began to increase the price of posting a job from \$0 to \$25 across a number of geographic markets. Cajner and Ratner (2016) show that this simple change had important implications for the HWOL data. Starting in 2012, the HWOL data began to show significantly lower growth in vacancies than the JOLTS series, and about a quarter of this was attributable to the change in Craigslist prices alone. A metric designed to measure the vacancy rate can easily become a measurement of the cost of posting a vacancy on a single site, which can hardly be considered a major macroeconomic force.



## When Craigslist Raised the Price of Job Ads, Job Postings in HWOL Fell.



Source: [Cajner and Ratner \(2016\)](#).

Even though the JOLTS series does not rely on online job postings in the same way, this cautionary tale shows how difficult it is to interpret vacancies as a straightforward labor market indicator. There is now a whole industry of companies like ZipRecruiter dedicated to using sophisticated algorithms to make it easier for companies to find workers. The job board site Indeed varies the cost of posting a job based on local labor market conditions. If, over time, the costs and technologies involved in worker recruitment and job search change, so too will the interpretation of a given measurement of the vacancy or vacancy-to-unemployment ratio.

## Microevidence

The biggest problem for most theories of job-matching is simple inconsistency with the empirical behavior of firms' hiring and vacancy-behavior. [Davis, Faberman, and Haltiwanger \(2013\)](#) study hiring and vacancy behavior using establishment-level data from the JOLTS. Strikingly, they find that many assumptions in matching models just do not hold up in the establishment-level data.



As we remember from our discussion of the DMP model, hiring only occurs when firms post vacancies, and the most firms can do when hiring is hope that the unemployed encounter these vacancies. If a firm wants to hire more, it simply needs to post more vacancies. In fact, in the canonical search and matching model, posting more or fewer vacancies is the only mechanism by which employment adjusts. One would therefore expect that establishments that grow faster do so only because they have more vacancies.

However, DFW (2013) find that only a small part of the differences in employment growth rates at the establishment level can be explained by differences in vacancy posting. Firms that hire more do not primarily do so by reporting a greater number of job openings. This suggests that there is a large scope for variation in recruiting effort; establishments that hire more do so as a consequence of other actions, such as advertising more, lowering hiring standards, or increasing compensation. Since the threshold for reporting a job opening in the JOLTS data is low, none of this recruitment intensity would be reflected in the quantitative vacancy statistics.

The key mechanism that is supposed to explain aggregate movements in employment in the model plays a much smaller role in the cross-sectional evidence. This is a serious problem! The absence of this evidence should give pause to policymakers who believe that job openings are the main driver of labor market movement per se, rather than a spurious statistic that happens to generally move alongside other labor market indicators.

The presence of varying job-recruiting efforts may lead policymakers astray. For an illustrative example, take the Great Recession. During the Great Recession, the VUR fell as vacancies decreased and unemployment increased due to a generalized economic downturn. Looking at the matching function, one would expect the vacancy yield rate – that is, the rate at which a vacancy converts to a hire – to rise. If the unemployed are encountering vacancies at random, fewer vacancies and higher unemployment should lead to more frequent random encounters. However, the vacancy yield rate rose far less than one would predict based on prior statistical relationships. In their analysis, DFW find that this dynamic is partially explained by a fall in recruiting intensity during the Great Recession: employers simply put less effort into hiring workers, and thereby hired fewer. As the Federal Reserve moves to tighten policy, there is a risk that a fall in the VUR will once again understate the deterioration of the labor market, especially if there are decreases in job recruiting efforts that are not reflected in the vacancy series.

### Recruiting Intensity Fell during the Great Recession




FIGURE X

Index of Recruiting Intensity Per Vacancy, January 2001 to December 2011

Source: [Davis, Faberman and Haltiwanger \(2013\)](#)

That DFW (2013) find that 42% of hires happen at employers that report zero vacancies is a major blow to the theoretical adequacy of the DMP model. Part of this may reflect the fact that in the JOLTS, firms are asked about vacancies at the end of the month. Any vacancies that are opened and filled within a given month would show up in the hires data but not the vacancies data. The rest could be due to misreporting or hires that don't happen through a process that would be recognized as a "vacancy" under the JOLTS definition.

The empirical evidence also shows that an aggregate model of hiring behavior does a poor job accounting for sectoral and cyclical determinants of recruitment intensity. The job-filling rate in construction is four times as high as that of health and education, for example. Complicating matters, [Davis, Faberman, and Haltiwanger \(2012\)](#) find in a separate paper that the cyclical behavior of job-filling and recruiting intensity also varies across industries. Since the sectoral composition of jobs changes over time and especially during recessions, this makes comparing aggregate vacancy statistics challenging. There are simply too many ill-defined moving parts for trends to be easily disentangled from one another.

## Conclusion

Put simply, vacancies are a flawed indicator of the state of the labor market. The theoretical basis for relying on vacancies or the vacancy-to-unemployment ratio arises from highly abstract models that ignore many salient features of the labor market. Vacancy statistics truncate any variations in recruitment efforts between firms, which are a major determinant of empirical hiring.

The JOLTS vacancy data tells us precisely and only what it asks: for how many positions do firms say they are doing a non-zero amount of recruiting at a particular point in time? While the answer may contain some information about labor demand, the vacancy statistic can't tell the whole story. The potential weaknesses of this data should give pause to anyone leaning too heavily on the vacancy statistic itself as a labor market metric. Changes in recruiting technology and industry composition substantially complicate attempts to compare vacancy rates across long periods of time, or business cycles.

Policymakers have many labor market indicators at their disposal, each of them capturing different aspects of the current state of the labor market. In making policy, these indicators should be weighted proportionately to the relevance and reliability of the data. While no labor market indicator is perfect, vacancy series are fraught with data and measurement issues. However, the Federal Reserve is at risk of being led astray by the explanatory role they have given to vacancies data.

In recent speeches, Powell has expressed a desire to see the vacancy rate fall in order to soften the labor market as the Federal Reserve attempts to navigate its way to a soft landing. If other, more robust, labor market indicators deteriorate before vacancies fall significantly, the Federal Reserve should recognize the risk that chasing the vacancy statistic risks significant damage to American workers.