Beveridgean Unemployment Gap

Pascal Michaillat, Emmanuel Saez

October 2020

This paper estimates the unemployment gap (the actual unemployment rate minus the efficient unemployment rate) using the sufficient-statistic approach from public economics. While lowering unemployment puts more people into work, it forces firms to post more vacancies and devote more resources to recruiting. This unemployment-vacancy tradeoff, governed by the Beveridge curve, determines the efficient unemployment rate. Accordingly, the unemployment gap can be measured from three sufficient statistics: the elasticity of the Beveridge curve, cost of recruiting, and social cost of unemployment. This novel formula applies to a wide range of macroeconomic models of the labor market. We implement the formula using US data. We find that the efficient unemployment rate varies between 2.6% and 4.7% since 1951, and has been stable between 3.3% and 4.0% since 1990. As a result, the unemployment gap is countercyclical, reaching 1.9–6.7 percentage points in slumps. Thus the US labor market appears inefficient—especially inefficiently slack in slumps. The unemployment gap is in turn a crucial statistic to determine optimal stabilization policies over the business cycle, including fiscal and monetary policy.

Michaillat: Brown University. Saez: University of California–Berkeley. We thank Regis Barnichon, Olivier Blanchard, Varanya Chaubey, Gabriel Chodorow-Reich, Richard Crump, Peter Diamond, Andrew Figura, Etienne Lehmann, Sephorah Mangin, and Alison Weingarden for helpful comments and discussions. This work was supported by the Institute for Advanced Study and the Berkeley Center for Equitable Growth.

Available at https://www.pascalmichaillat.org/9.html

1. Introduction

Perhaps the most important statistic in the design of stabilization policy is the unemployment gap: the difference between the actual and the socially efficient unemployment rate. In practice, many governments are mandated to reduce the unemployment gap to zero. For instance, in the United States, the 1978 Humphrey-Hawkins Full Employment Act mandates that the government maintains the economy at "full employment" using monetary and fiscal policy. Full employment should not be interpreted as no unemployment, which is physically impossible, but rather as a socially efficient amount of unemployment; therefore, the mandate of US policymakers can be interpreted as closing the unemployment gap. Macroeconomic theory also shows that many stabilization policies should move in tandem with the unemployment gap. Such policies include labor-market subsidies and taxes (Pissarides 2000, chapter 9); minimum wage (Hungerbuhler and Lehmann 2009); monetary policy (Michaillat and Saez 2019a); public expenditure (Michaillat and Saez 2019b); and short-time work (Giupponi and Landais 2018).

Two measures of the unemployment gap are commonly used (Crump et al. 2019); but neither incorporates an appropriate estimate of the efficient unemployment rate. The first is the difference between actual unemployment and its secular trend. This measure is flawed because in most models trend unemployment is not efficient (Pissarides 2000, chapter 8). The second is the difference between actual unemployment and the non-accelerating inflation rate of unemployment (NAIRU), obtained by estimating a Phillips curve. This measure is flawed because the NAIRU never indicates labor-market efficiency (Rogerson 1997). Thus, although these two measures are easy to use, they lack a theoretical foundation.

This paper proposes a new measure of the unemployment gap that builds upon the theory of efficiency in modern labor-market models (Pissarides 2000, chapter 8). These models feature both unemployed workers and job vacancies, each associated with welfare costs: more unemployment means fewer people at work so less output; more vacancies mean more work effort devoted to recruiting and also less output. Furthermore, these models feature a Beveridge curve, so unemployment and vacancies cannot be simultaneously reduced: less unemployment requires more vacancies, and fewer vacancies create more unemployment. Our analysis resolves this unemployment-vacancy tradeoff, characterizing the efficiency point on the Beveridge curve.

At the same time, we strive to develop an unemployment-gap formula that is usable for policy work. To that end, we adhere to the sufficient-statistic method from public economics (Chetty 2009). A first advantage is that our formula requires little theoretical structure. It applies to any labor market with a Beveridge curve, irrespective of the structure of the labor market, production, preferences, wage setting, and shocks; the model's relevant properties are captured

by the formula's sufficient statistics. Since the Beveridge curve appears in many models and has been observed in many countries (Elsby, Michaels, and Ratner 2015), the formula applies broadly. A second advantage is that the sufficient statistics are estimable, so the formula is simple and transparent to apply.

We begin by solving the problem of a social planner who must allocates labor between production, recruiting, and unemployment subject to the Beveridge-curve constraint. The solution gives the efficient unemployment rate, which we express as a function of actual unemployment and vacancy rates, and three sufficient statistics: the elasticity of the Beveridge curve, cost of recruiting, and social cost of unemployment.

Next, we provide estimates of the three sufficient statistics for the United States. Using standard unemployment and vacancy data, we estimate the elasticity of the Beveridge curve. Although the Beveridge curve is stable for long periods of time, it is also subjects to sudden shifts. To address this issue, we use the method proposed by Bai and Perron (1998, 2003); it allows us to estimate the elasticity while permitting multiple structural changes. We estimate that the Beveridge curve indeed experiences several structural breaks between 1951 and 2019, but the Beveridge elasticity remains fairly stable, in the 0.84–1.02 range. Next, we estimate the recruiting cost using evidence from the 1997 National Employer Survey conducted by the Census Bureau. The survey indicates that firms allocate 2.5% of labor to recruiting (Villena Roldan 2010). Last, we estimate the social cost of unemployment from the natural and field experiments analyzed by Borgschulte and Martorell (2018) and Mas and Pallais (2019). These experiments suggest that the value from home production and recreation during unemployment only replaces 25% of the marginal product of labor—implying a substantial social cost of unemployment.

Using the estimates of the statistics, we compute the efficient unemployment rate in the United States between 1951 and 2019. We find that the efficient unemployment rate averaged 3.8%. It started around 3.0% in the 1950s, climbed to reach 4.7% in 1979, fell to 4.0% in 1990, and remained in the 3.3%–4.0% range until 2019. These variations are caused by shifts of the Beveridge curve.

Since the efficient unemployment rate is slow-moving while the actual unemployment rate is countercyclical, the unemployment gap is countercyclical. We infer that the US unemployment gap is almost never zero: the US labor market does not operate efficiently. In fact the unemployment gap is generally positive, averaging 2.0 percentage points over 1951–2019, so the US labor market is generally inefficiently slack. The unemployment gap is especially high in slumps, reaching for instance 6.7 percentage points in 1982 and 2010. Unsurprisingly, inefficiencies are exacerbated in slumps.

Finally, we explore the sensitivity of the efficient unemployment rate to alternative values

of the sufficient statistics. When the Beveridge elasticity takes any value in its 95% confidence interval, the efficient unemployment rate is contained in a band whose width averages 1.0 percentage point, and is always below 2.0 percentage points (the higher the elasticity, the higher the efficient unemployment rate). And when the social cost of unemployment spans the range of plausible values given by Borgschulte and Martorell (2018) and Mas and Pallais (2019), the efficient unemployment rate is contained in a band whose width averages 1.4 percentage point, and is always below 1.8 percentage points (the lower the cost, the higher the efficient unemployment rate). Hence, for any plausible estimate of the sufficient statistics, the unemployment gap will never depart from the baseline by more than 1 percentage point. This means that our substantive conclusions—that the US labor market is almost always inefficient, generally inefficient slack, and especially inefficiently slack in slumps—are robust to alternative calibrations.

2. Beveridgean labor market

We introduce the labor-market model used to compute the unemployment gap. The main ingredient is a Beveridge curve—a negative relation between unemployment and vacancies.

2.1. Beveridge curve

We consider a labor market with both unemployed workers and job vacancies. The unemployment rate u is the number of unemployed workers divided by size of the labor force. The vacancy rate v is the number of vacancies divided by size of the labor force. Labor-market tightness is number of vacancies per unemployed worker: $\theta = v/u$. The employment rate is the number of employed workers divided by size of the labor force: $n = 1 - u$.

Unemployment and vacancy rates are related by a Beveridge curve: the vacancy rate is a strictly decreasing and convex function of the unemployment rate, $v(u)$. Many labor-market models feature a Beveridge curve and so are nested into our framework (Elsby, Michaels, and Ratner 2015). Importantly, models build around matching functions exhibit a Beveridge curve (Petrongolo and Pissarides 2001, equation (12)).[1] This category includes the canonical Diamond-Mortensen-Pissarides model (Pissarides 2000, chapter 1); but also its variants with rigid wages

[1] In many matching models, unemployment follows a law of motion, and the Beveridge curve is defined as the locus of unemployment and vacancies consistent with a steady level of unemployment. If unemployment converged slowly to steady state, unemployment and vacancies would not always be on the Beveridge curve. However, Pissarides (2009a, p. 236) notes that "Perhaps surprisingly at first, but on reflection not so surprisingly, we get a good approximation to the dynamics of unemployment if we treat unemployment as if it were always on the Beveridge curve." The reason is that labor-market flows are large so after a shock, unemployment adjusts rapidly to its new steady-state level, where inflows into unemployment equal outflows from unemployment (Pissarides 1986; Hall 2005b; Pissarides 2009b; Elsby, Michaels, and Solon 2009; Shimer 2012).

(Hall 2005a; Hall and Milgrom 2008), large firms (Cahuc, Marque, and Wasmer 2008; Elsby and Michaels 2013), and job rationing (Michaillat 2012). Even models without a matching function may feature a Beveridge curve: for instance, models of mismatch (Shimer 2007) and of stock-flow matching (Ebrahimy and Shimer 2010).

Our framework applies to any country that exhibits a Beveridge curve. There are many such countries (Jackman, Pissarides, and Savouri 1990; Nickell et al. 2002; Elsby, Michaels, and Ratner 2015), including the United States (Blanchard and Diamond 1989; Diamond and Sahin 2015; Elsby, Michaels, and Ratner 2015). As an illustration, we construct the US Beveridge curve. For the unemployment rate, we use the standard measure constructed by the Bureau of Labor Statistics (BLS) from the Current Population Survey (CPS) (figure 1, panel A). For the vacancy rate, we use two different sources because there is no continuous national vacancy series over the period. For 1951–2000, we use the vacancy proxy constructed by Barnichon (2010). Barnichon starts from the help-wanted advertising index constructed by the Conference Board—a proxy for vacancies proposed by Abraham (1987) that has become standard (Shimer 2005, p. 29). He then corrects the Conference Board index, which is based on newspaper advertisements, to take into account the shift from print advertising to online advertising after 1995. Finally, he rescales the index into vacancies, and divides the vacancy number by the size of the labor force to obtain a vacancy rate. For 2001–2019, we obtain the vacancy rate from the number of job openings measured by the BLS in the Job Opening and Labor Turnover Survey (JOLTS), divided by the civilian labor force constructed by the BLS from the CPS. We then splice the Barnichon and JOLTS series to obtain a vacancy rate for 1951–2019 (figure 1, panel B).

The Beveridge curve appears in scatterplots of the unemployment and vacancy rates (panels C–F of figure 1; for readability we separately plot the 1951–1969, 1970–1989,1990–2009, and 2010–2019 periods). The Beveridge curve is stable for long periods, during which unemployment and vacancies moved up and down along a clearly defined branch. Furthermore, until the mid-1980s, the Beveridge curve shifts outward at the end of each period of stability. After the mid-1980s, the Beveridge curve shifts back inward to positions that were typical in the 1960s and 1970s.

Plotted on a logarithmic scale, all the branches of the Beveridge curve are close to linear, so each branch is isoelastic. A central statistic to measure the unemployment gap will be the Beveridge elasticity:

DEFINITION 1. *The Beveridge elasticity is the elasticity of the vacancy rate with respect to the unemployment rate along the Beveridge curve, normalized to be positive:*

$$\epsilon = -\frac{d\ln(v(u))}{d\ln(u)} = -\frac{u}{v(u)}v'(u).$$

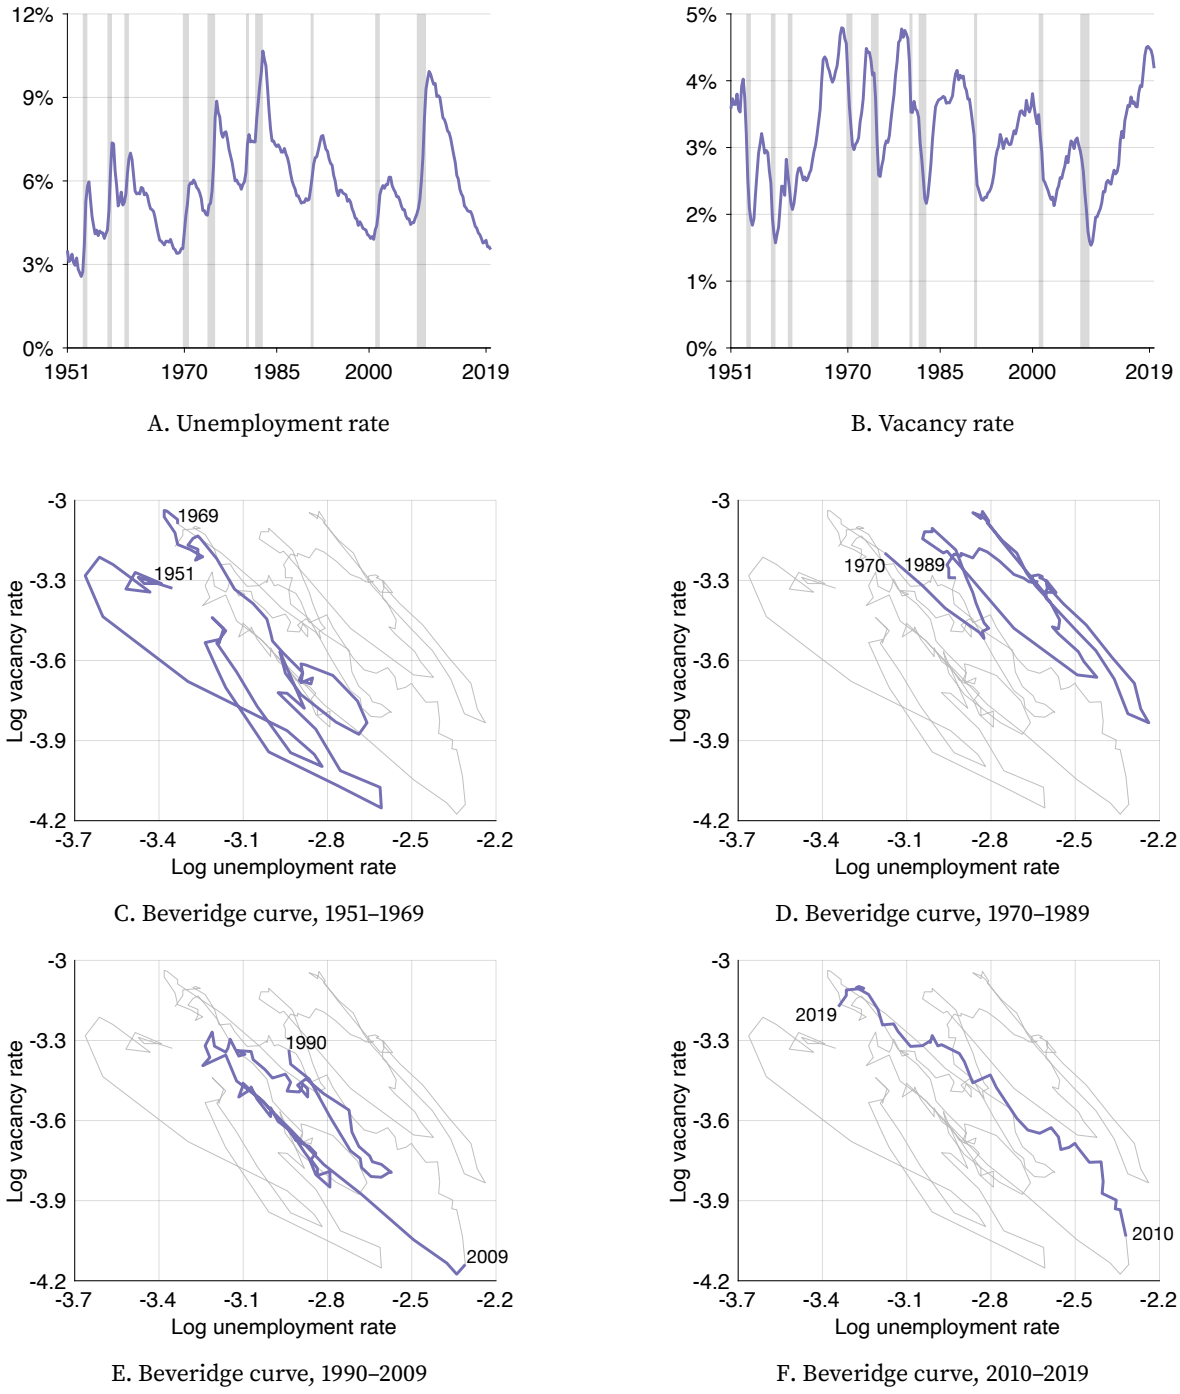


FIGURE 1. Beveridge curve in the United States, 1951–2019

Panel A: The unemployment rate is constructed by the BLS from the CPS. Panel B: For 1951–2000, the vacancy rate is constructed by Barnichon (2010) from the Conference Board help-wanted advertising index; for 2001–2019, the vacancy rate is the number of job openings measured by the BLS in JOLTS, divided by the civilian labor force constructed by the BLS from the CPS. All unemployment and vacancy rates are quarterly averages of seasonally adjusted monthly series. The shaded areas represent recessions, as identified by the National Bureau of Economic Research (NBER). Panels C–F display scatterplots of the unemployment and vacancy rates (from panels A and B) on a logarithmic scale. The four panels focus on consecutive periods: 1951–1969 in panel C, 1970–1989 in panel D, 1990–2009 in panel E, and 2010–2019 in panel F.

In a Diamond-Mortensen-Pissarides model with Cobb-Douglas matching function $m(u,v) = \mu u^{\alpha} v^{1-\alpha}$, the Beveridge elasticity is closely related to the matching elasticity α. In that model the Beveridge curve is obtained by equating inflows into unemployment with outflows from unemployment: $s \cdot (1-u) = m(u,v)$, where s is the job-separation rate, so that

$$v(u) = \left[\frac{s(1-u)}{\mu u^{\alpha}}\right]^{1/(1-\alpha)}$$

(see Shimer 2005, p. 36). Hence, the Beveridge elasticity satisfies

$$\epsilon = \frac{1}{1-\alpha}\left(\alpha + \frac{u}{1-u}\right). \tag{1}$$

2.2. Social welfare

The Beveridge curve determines the tradeoff between unemployment and vacancies. This tradeoff is central to the welfare analysis because both unemployment and vacancies enter the welfare function.

We assume that social welfare is given by a function $\mathcal{W}(n,u,v)$, where n is the employment rate, u is the unemployment rate, and v is the vacancy rate. The function $\mathcal{W}$ is differentiable and strictly increasing in n. The effects of unemployment and vacancies on welfare are given by the following two statistics, which play a key role in calculating the unemployment gap:

DEFINITION 2. *The social value of nonwork is the marginal rate of substitution between unemployment and employment in the welfare function:*

$$\zeta = \frac{\partial \mathcal{W}/\partial u}{\partial \mathcal{W}/\partial n} < 1.$$

The social cost of unemployment is $1-\zeta > 0$.

DEFINITION 3. *The recruiting cost is minus the marginal rate of substitution between vacancies and employment in the welfare function:*

$$\kappa = -\frac{\partial \mathcal{W}/\partial v}{\partial \mathcal{W}/\partial n} > 0.$$

Employed workers contribute to social welfare through market production. Unemployed workers contribute to social welfare through home production and recreation (Aguiar, Hurst, and Karabarbounis 2013); this contribution is diminished if people suffer psychic pain from being unemployed (Brand 2015). The social value of nonwork ζ measures the marginal contribution of

unemployed workers relative to that of employed workers. Since ζ is a social concept, it does not include monetary transfers received by unemployed workers from the government or others.

Unemployment is socially costly because unemployed workers' contribute less to welfare than employed workers ($\zeta < 1$). The social cost of unemployment $1 - \zeta$ measures the social loss of having a person unemployed rather than employed. Such loss comprises foregone market production and psychological pain of being unemployed, net of the value home production and recreation when unemployed.

Vacancies enter the welfare function because the recruiting activity required to fill vacancies diverts labor and other resources away from market production. The recruiting cost κ measures the resources absorbed by maintaining a vacancy, expressed in terms of labor. It is normalized to be positive.

Since the labor force is divided between employed and unemployed workers, the employment rate satisfies $n = 1 - u$, and social welfare can be written as a function of unemployment and vacancy rates:

$$(u, v) \mapsto \mathcal{W}(1 - u, u, v). \tag{2}$$

The assumptions that $\zeta < 1$ and $\kappa > 0$ imply that the function (2) is strictly decreasing in u and v—which captures the social costs of unemployment and vacancies. Moreover, to ensure that the social planner's problem is well behaved, we assume that this function is quasiconcave.

In the Diamond-Mortensen-Pissarides model the statistics ζ and κ take a simple form. Let L be the size of the labor force; p the productivity of employed workers; $z < p$ the productivity of unemployed workers; and $p \times c$ the resource cost of posting a vacancy (Pissarides 2000, chapter 1). The utility function is linear, so the welfare function is

$$\mathcal{W}(n, u, v) = (pn + zu - pcv)\, L.$$

Accordingly, the social value of nonwork and recruiting cost are

$$\zeta = \frac{z}{p} \quad \text{and} \quad \kappa = c. \tag{3}$$

3. Efficient unemployment rate and unemployment gap

In a Beveridgean labor market, we define efficiency as the solution to the problem of a social planner who is subject to the Beveridge-curve constraint:

DEFINITION 4. *The efficient unemployment and vacancy rates, denoted u^* and v^*, maximize social*

welfare (2) *subject to the Beveridge-curve constraint* $v = v(u)$. *The efficient labor-market tightness is* $\theta^* = u^*/v^*$, *and the unemployment gap is* $u - u^*$.

We now represent labor-market efficiency in a Beveridge diagram, and then derive a sufficient-statistic formula for the efficient unemployment rate and unemployment gap.

3.1. Representation in a Beveridge diagram

The Beveridge diagram, with unemployment rate on the x-axis and vacancy rate on the y-axis, is depicted in figure 2, panel A. In the diagram the Beveridge curve is downward-sloping and convex; it gives the locus of unemployment and vacancy rates that are feasible in the economy. The diagram also features an isowelfare curve: the locus of unemployment and vacancy rates such that social welfare (2) remains constant at some level (the equivalent of an indifference curve for a utility function or an isoquant for a production function). Since (2) is decreasing in unemployment and vacancies, all the points inside the isowelfare curve yield higher welfare, so the green area delineated by the isowelfare curve is an upper contour set of (2). Since the function (2) is quasiconcave, its upper contour sets are convex, which implies that the isowelfare curve must be concave.

The efficient unemployment and vacancy rates can easily be found in the Beveridge diagram. First, they have to lie on the Beveridge curve. Second, since both unemployment and vacancies impose a welfare cost, they must lie on the isowelfare curve that is as close to the origin as possible. The closest that the isowelfare curve can be while remaining in contact with the Beveridge curve is at the tangency point with the Beveridge curve. This is where the efficient unemployment and vacancy rates are found. The efficient labor-market tightness is also visible on the diagram: it is the slope of the origin line going through the tangency point.

As with indifference curves and isoquants, the slope of the isowelfare curve is minus the marginal rate of substitution between unemployment and vacancies in the welfare function (2):

$$-\frac{(\partial \mathcal{W}/\partial u) - (\partial \mathcal{W}/\partial n)}{\partial \mathcal{W}/\partial v} = -\frac{1 - (\partial \mathcal{W}/\partial u)/(\partial \mathcal{W}/\partial n)}{-(\partial \mathcal{W}/\partial v)/(\partial \mathcal{W}/\partial n)} = -\frac{1-\zeta}{\kappa}.$$

The efficient unemployment rate is found at the point where the Beveridge curve, with slope $v'(u)$, is tangent to isowelfare curve, with slope $-(1-\zeta)/\kappa$. This yields a first result:

PROPOSITION 1. *In a Beveridge diagram, efficiency is achieved at the point where the Beveridge curve is tangent to an isowelfare curve. Hence, the efficient unemployment rate is implicitly defined by*

$$v'(u) = -\frac{1-\zeta}{\kappa}, \tag{4}$$

where $\zeta < 1$ is the social value of nonwork and $\kappa > 0$ the recruiting cost.

Formula (4) simply says that when the labor market operates efficiently, welfare costs and benefits from moving one worker from employment to unemployment are equalized. When one worker moves from employment to unemployment, the reduction in welfare is the social cost of unemployment, $1 - \zeta$. Having one more unemployed worker also means having $-v'(u) > 0$ fewer vacancies. Each vacancy reduces welfare by the recruiting cost, κ, so welfare improves by $-v'(u)\kappa$ through the reduction in recruiting activity. When welfare costs and benefits are equalized, we have $1 - \zeta = -v'(u)\kappa$, which is equivalent to (4).

Of course, there is no guarantee that the labor market operates efficiently (figure 2, panel B). The labor market may be above the efficiency point, where unemployment is too low, vacancies are too high, and the unemployment gap is negative. This situation corresponds to a boom. It may also be below the efficiency point, where unemployment is too high, vacancies are too low, and the unemployment gap is positive. This situation corresponds to a slump. As slumps and booms are inefficient, they lie on a worse isowelfare curve than the efficiency point.

3.2. Formula with sufficient statistics

We rework the efficiency condition (4) to obtain an explicit expression for the unemployment gap. We begin by introducing the Beveridge elasticity: $\epsilon = -(u/v)v'(u)$ so $\epsilon\theta = -v'(u)$. With this result, we can re-express (4) as $\theta = (1 - \zeta)/(\kappa\epsilon)$.

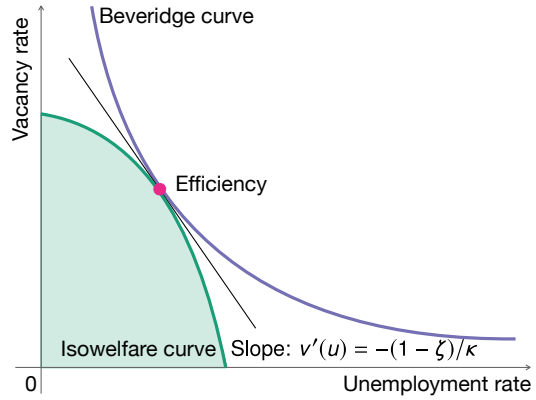
In panel B of figure 2, we also see that any point on the Beveridge curve above the efficiency point has $-v'(u) > (1 - \zeta)/\kappa$, and any point below it has $-v'(u) < (1 - \zeta)/\kappa$. Using again $\epsilon\theta = -v'(u)$, we infer that tightness is inefficiently high whenever $\theta > (1 - \zeta)/(\kappa\epsilon)$; and tightness is inefficiently low whenever $\theta < (1 - \zeta)/(\kappa\epsilon)$.

Hence we can assess the efficiency of labor-market tightness from three sufficient statistics:

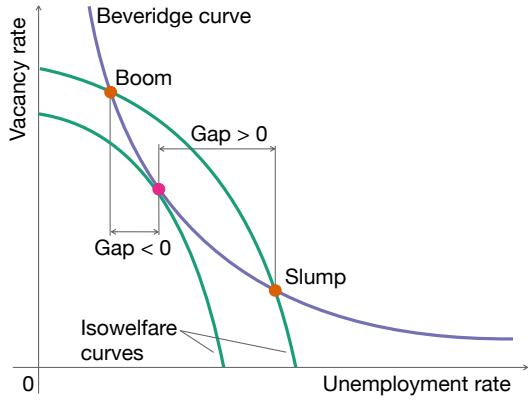
PROPOSITION 2. *Consider a point on the Beveridge curve with labor-market tightness θ, Beveridge elasticity ϵ, recruiting cost κ, and social value of nonwork ζ. Then tightness is inefficiently high if $\theta > (1 - \zeta)/(\kappa\epsilon)$, inefficiently low if $\theta < (1 - \zeta)/(\kappa\epsilon)$, and efficient if*

$$\theta = \frac{1 - \zeta}{\kappa\epsilon}. \tag{5}$$

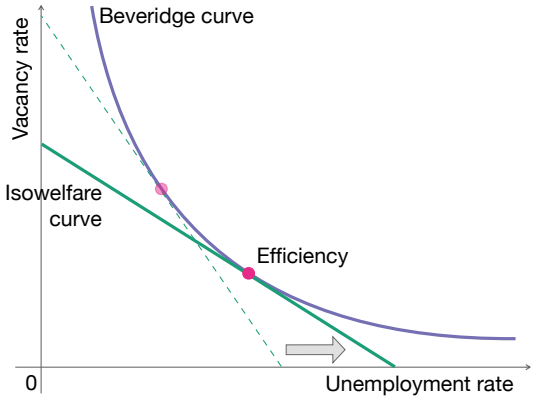
Formula (5) can be seen as a reformulation of the well-known Hosios (1990) condition. Hosios resolves the unemployment-vacancy tradeoff and derives a condition to ensure efficiency when wages are determined by Nash bargaining. The condition is that workers' bargaining power equals the matching elasticity. While the Hosios condition has had a tremendous theoretical



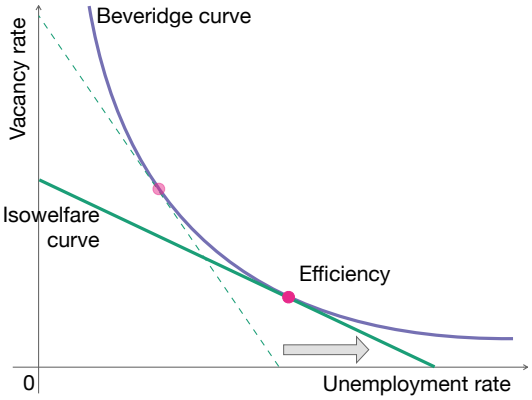
A. Efficient unemployment rate



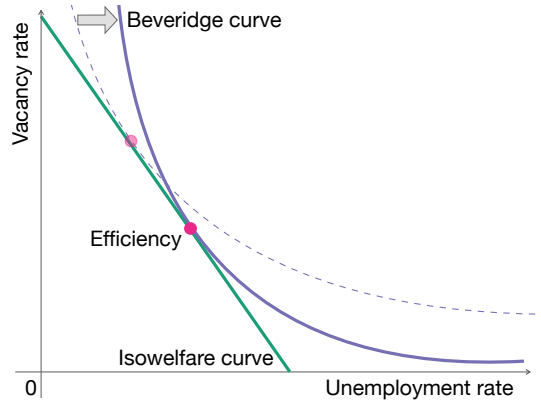
B. Unemployment gap in slumps and booms



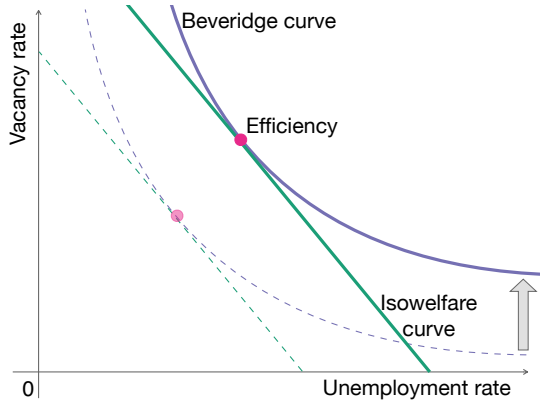
C. Increase in recruiting cost



D. Increase in social value of nonwork



E. Compensated increase in Beveridge elasticity



F. Outward shift of Beveridge curve

FIGURE 2. Efficient unemployment rate and unemployment gap in Beveridge diagram

Panel A depicts an efficient labor market in the Beveridge diagram. The statistic κ is the recruiting cost, ζ is the social value of nonwork, and $v'(u)$ is the slope of the Beveridge curve. Panel B describes inefficient labor markets. In slumps, unemployment is inefficiently high, vacancies are inefficiently low, and the unemployment gap is positive; in booms, unemployment is inefficiently low, vacancies are inefficiently high, and the unemployment gap is negative. panels C–F are special cases of panel A: assumption 1 holds, so the isowelfare curve is linear with slope $-(1-\zeta)/\kappa$, and the Beveridge curve is isoelastic with elasticity ϵ. Panel C shows that the efficient unemployment rate increases when the recruiting cost increases. Panel D shows that the efficient unemployment rate increases when the social value of nonwork increases. Panel E shows that the efficient unemployment rate increases when the Beveridge elasticity increases (keeping welfare constant). Panel F shows that the efficient unemployment rate increases when the Beveridge curve shifts outward (keeping the Beveridge elasticity constant).

impact, its practical impact has been more limited. In our view this is due to two limitations, which we aim to address. First, measuring bargaining power is challenging (Pissarides 2000, p. 229). As a result the bargaining power is usually simply calibrated to 0.5 (den Haan, Ramey, and Watson 2000; Pissarides 2000; Gertler and Trigari 2009) or to the matching elasticity (Mortensen and Pissarides 1994; Shimer 2005; Costain and Reiter 2008). In contrast, formula (5) involves labor-market statistics that are estimable. Second, Nash bargaining poorly describes wage-setting over the business cycle (Shimer 2005; Hall 2005a; Jager et al. 2018). In contrast, our formula applies to any model with a Beveridge curve, irrespective of how wages are set.

Since the statistics ϵ, κ, and ζ generally depend on tightness θ, formula (5) characterizes the efficient tightness only implicitly. This limitation is typical of the sufficient-statistic approach (Chetty 2009), but it complicates the task of computing the unemployment gap. To address it, we use a workaround developed by Kleven (2018):

ASSUMPTION 1. *The Beveridge elasticity (ϵ), recruiting cost (κ), and social value of nonwork (ζ) do not depend on the unemployment and vacancy rates.*

How realistic is this assumption? Panels C and D in figure 1 suggest that the Beveridge curve is isoelastic, so the assumption on the Beveridge elasticity seems valid in US data. We do not have comparable evidence on the recruiting cost and social value of nonwork, but at least in the Diamond-Mortensen-Pissarides model, these two statistics are independent of the unemployment and vacancy rates (equation (3)).[2]

Under assumption 1, we obtain a simple formula for the unemployment gap. The assumption implies that the Beveridge curve is isoelastic:

$$v(u) = v_0 \cdot u^{-\epsilon},$$

where the parameter $v_0 > 0$ determines the location of the curve. On the Beveridge curve, tightness is related to unemployment by

$$\theta = \frac{v(u)}{u} = v_0 \cdot u^{-(1+\epsilon)} \quad \text{and} \quad \theta^* = v_0 \cdot (u^*)^{-(1+\epsilon)}.$$

We can therefore link the unemployment gap to the tightness gap:

$$\frac{u^*}{u} = \left(\frac{\theta}{\theta^*}\right)^{1/(1+\epsilon)}.$$

[2]The Beveridge elasticity does depend on the unemployment rate in the Diamond-Mortensen-Pissarides model (equation (1)). However, as the unemployment rate is an order of magnitude smaller than the matching elasticity α, the Beveridge elasticity is approximately constant and equal to $\alpha/(1-\alpha)$.

Under assumption 1, (5) gives $\theta^* = (1-\zeta)/(\kappa\epsilon)$, which yields the following proposition:

PROPOSITION 3. *Under assumption 1, the efficient unemployment rate and unemployment gap can be measured from current unemployment rate (u), vacancy rate (v), Beveridge elasticity (ϵ), recruiting cost (κ), and social value of nonwork (ζ). The efficient unemployment rate satisfies*

$$u^* = \left(\frac{\kappa\epsilon}{1-\zeta}\cdot\frac{v}{u}\right)^{1/(1+\epsilon)} u, \tag{6}$$

from which the unemployment gap $u - u^*$ follows.

The proposition gives an explicit formula for the unemployment gap, expressed in terms of observable sufficient statistics. The formula is valid in any Beveridgean labor-market model, irrespective of the structure of the labor market, production, preferences, and wage setting. Another advantage of the formula is that we do not need to keep track of all the shocks disturbing the labor market—shocks to productivity, wages, labor-force participation, matching function, job separations, preferences, etc. The sufficient statistics are all we need to observe.[3]

A secondary benefit of assumption 1 is to simplify the shapes of the curves in the Beveridge diagram: the Beveridge curve becomes isoelastic, and the isowelfare curve becomes linear with slope $-(1-\zeta)/\kappa$. The simplification yields additional results:

PROPOSITION 4. *Under assumption 1, the following comparative statics hold:*

- *An increase in recruiting cost raises the efficient unemployment rate and lowers the efficient tightness.*
- *An increase in the social value of nonwork raises the efficient unemployment rate and lowers the efficient tightness.*
- *A compensated increase in the Beveridge elasticity (increase in ϵ keeping welfare constant) raises the efficient unemployment rate and lowers the efficient tightness.*
- *An outward shift of the Beveridge curve (increase in v_0) raises the efficient unemployment rate but does not affect the efficient tightness.*

The intuition is conveyed by the diagrams in figure 2. When recruiting is more costly or unemployment is less costly, the unemployment-vacancy tradeoff becomes more favorable to unemployment, and the efficient unemployment rate increases (panels C and D). A compensated increase in the Beveridge elasticity (analogous to a compensated price increase in the context of

[3]Without assumption 1, we could still obtain a formula for the unemployment gap, but it would require three additional statistics: the elasticities of ϵ, κ, and ζ with respect to the unemployment rate (Kleven 2018).

Hicksian demand) is an increase in ϵ compensated by a change in v_0 so that the new Beveridge curve remains tangent to the same isowelfare curve (panel E). Such a change steepens the Beveridge curve: a rise in unemployment triggers a larger drop in vacancies, so the unemployment-vacancy tradeoff is more favorable to unemployment, and the efficient unemployment rate increases. Finally, the outward shift of the Beveridge curve does not change its elasticity so the efficient tightness remains the same. Since the Beveridge curve is further out, however, the efficient unemployment rate is higher.

4. Unemployment gap in the United States, 1951–2019

We apply formula (6) to measure the unemployment gap in the United States over the 1951–2019 period. The first step is to estimate the sufficient statistics: Beveridge elasticity, recruiting cost, and social value of nonwork.

4.1. Beveridge elasticity (ϵ)

We estimate the Beveridge elasticity with linear regressions of log vacancy rate on log unemployment rate. The data are displayed in figure 1. The sample contains $T = 276$ observations, since data are quarterly from 1951Q1 to 2019Q4. Since the Beveridge curve shifts over time, we need to allow for structural breaks in the estimation. We therefore use the algorithm proposed by Bai and Perron (1998, 2003) to estimate linear models with multiple structural changes.

Statistical model. The statistical model that we estimate has m breaks, and $m + 1$ regimes. It is given by

$$\ln(v(t)) = \beta_j + \epsilon_j \cdot \ln(u(t)) + z(t) \qquad t = T_{j-1} + 1, \ldots, T_j$$

for $j = 1, \ldots, m + 1$. The observed dependent variable is the log of the vacancy rate, $\ln(v(t))$; the observed independent variable is the log of the unemployment rate, $\ln(u(t))$; $z(t)$ is the error at time t; β_j is the intercept of the linear model in regime j; ϵ_j is the Beveridge elasticity in regime j; and $T_1, \ldots, T_m$ are the m break points (we use the convention that $T_0 = 0$ and $T_{m+1} = T$).

The regression coefficients $\beta_1, \ldots, \beta_{m+1}$ and $\epsilon_1, \ldots, \epsilon_{m+1}$, and the break points $T_1, \ldots, T_m$, are unknown. [4] They are jointly estimated with the Bai-Perron algorithm. The regression coefficients are estimated by least-squares while the break points are estimated by minimizing the sum of squared residuals. The algorithm also provides confidence intervals for the break dates and regression coefficients under various hypotheses about the structure of the data and the errors

[4] This is a pure structural change model because the two regression coefficients change across regimes.

across regimes. Finally, the algorithm offers several ways to test for structural changes and to determine the number of breaks, m.

Algorithm setup. Before proceeding to the estimation, we need to make several modeling choices and to calibrate several parameters of the algorithm.

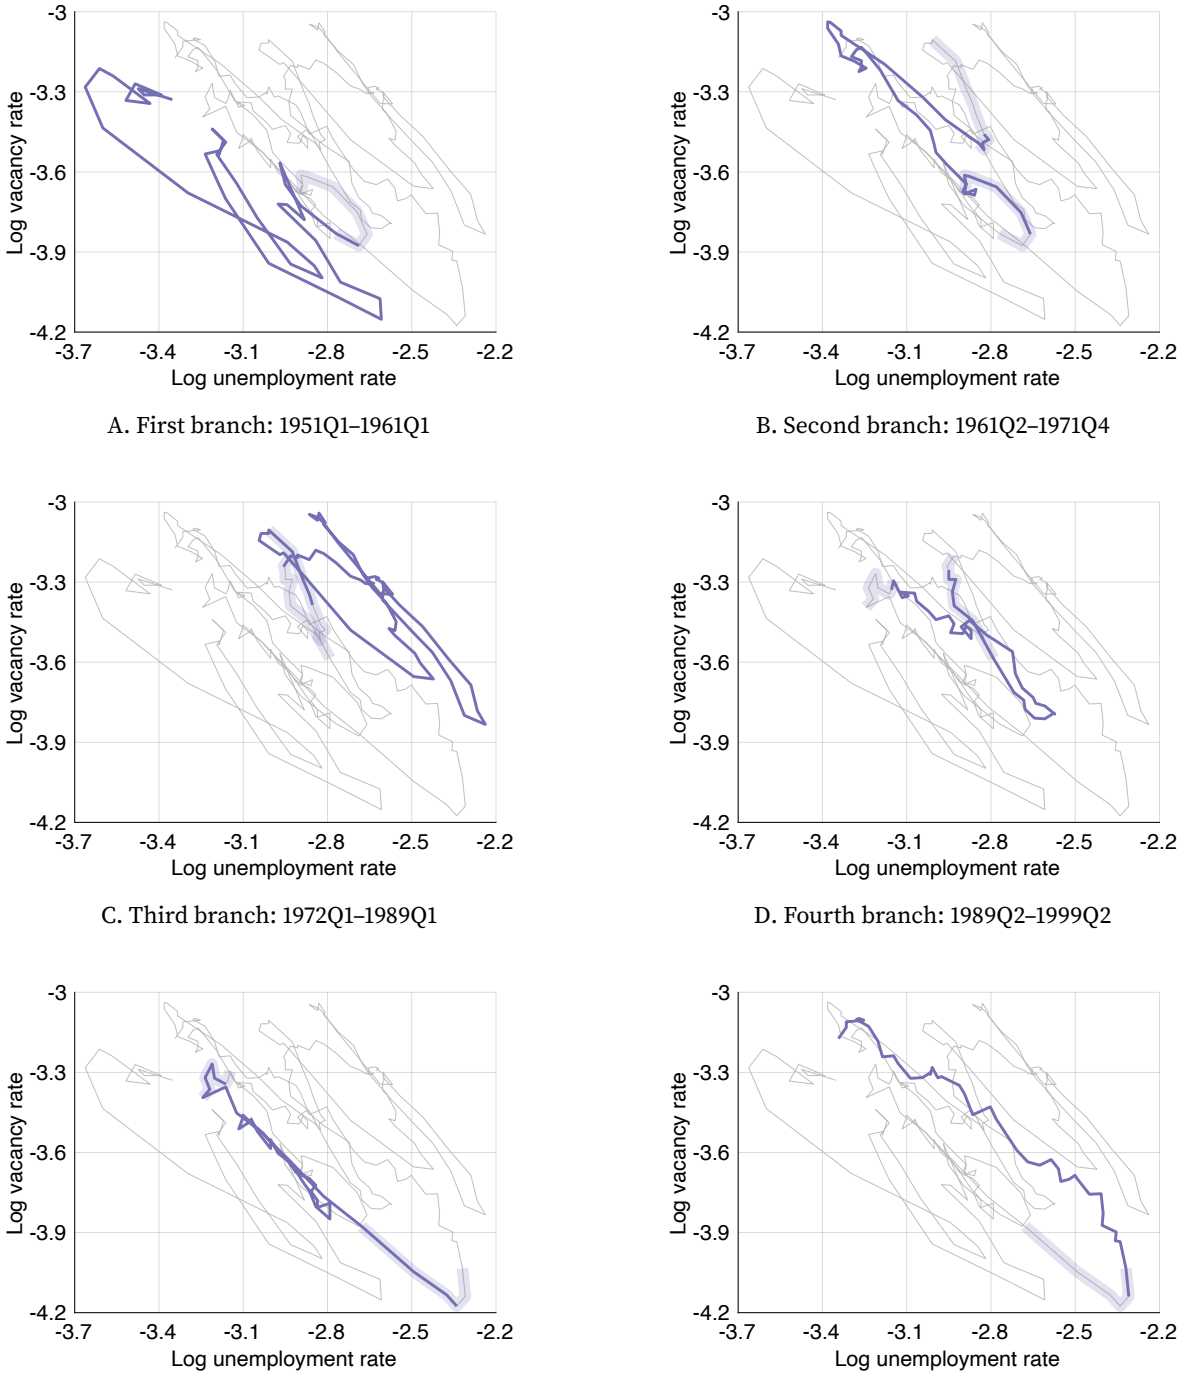
- We allow for autocorrelation in the errors, and different variances of the errors across regimes. To obtain standard errors robust to autocorrelation and heteroskedasticity, the algorithm follows the method proposed by Andrews (1991), using a quadratic kernel with automatic bandwidth selection based on an AR(1) approximation.
- We also allow different distributions of the independent and dependent variables (unemployment and vacancy rates) across regimes.
- We set the trimming parameter to $\epsilon = 0.15$, as suggested by Bai and Perron (2003, p. 15); hence each regime has at least $\epsilon \times T = 0.15 \times 276 = 41$ observations. As required by Bai and Perron (2003, p. 14), we set the maximum number of breaks to $M = 5$.

Existence and number of breaks. We first consider statistical tests for the existence of structural breaks. The algorithm provides supF tests of no structural break versus m breaks, for $m = 1, \ldots, 5$. All such tests reject the null hypothesis of no breaks at the 1% significance level. The algorithm also provides two double-maximum tests of no structural break versus an unknown number of breaks below the upper bound $M = 5$. Again, both tests reject the null hypothesis of no breaks at the 1% significance level. Given these results, it is clear that at least one break is present.

To determine the number of structural breaks, we consider two information criteria: the Bayesian Information Criterion proposed by Yao (1988), and the modified Schwarz criterion proposed by Liu, Wu, and Zidek (1997). Both information criterion select 5 breaks.

Break dates. Next we estimate the 5 break dates. The algorithm finds that the breaks in the Beveridge curve occur in 1961Q1, 1971Q4, 1989Q1, 1999Q2, and 2009Q3. The break dates are precisely estimated as all their 95% confidence intervals cover less than 2.5 years. The 6 branches of the Beveridge curve delineated by the break dates, together with the 95% confidence intervals for the dates, are represented on figure 3.

Beveridge elasticity. Finally, we estimate the Beveridge elasticities in the 6 different regimes. We find that from 1951 to 2019, the Beveridge elasticity fluctuated between 0.84 and 1.02, averaging 0.91. So overall, the elasticity remains quite stable over time. The elasticities are fairly precisely



A. First branch: 1951Q1–1961Q1

B. Second branch: 1961Q2–1971Q4

C. Third branch: 1972Q1–1989Q1

D. Fourth branch: 1989Q2–1999Q2

E. Fifth branch: 1999Q3–2009Q3

F. Sixth branch: 2009Q4–2019Q4

FIGURE 3. Beveridge-curve branches in the United States, 1951–2019

The figure is obtained by applying the methodology of Bai and Perron (1998, 2003) on log vacancy rate and log unemployment rate. The methodology is designed to estimate linear models with multiple structural changes. The figure reports the branches of the Beveridge curve between each estimated structural break (thin dark line), as well as the 95% confidence interval for each break date (wide transparent line).

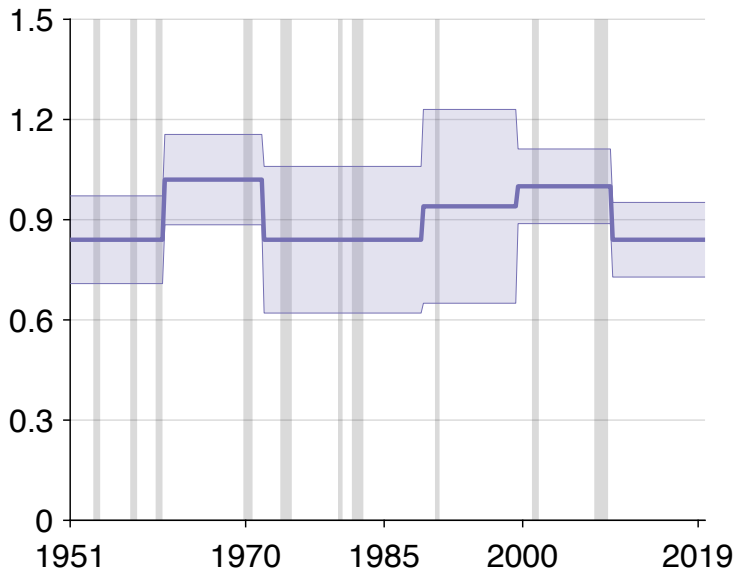


FIGURE 4. Beveridge elasticity and and 95% confidence interval in the United States, 1951–2019

The figure is obtained by applying the methodology of Bai and Perron (1998, 2003) on log vacancy rate and log unemployment rate. The methodology is designed to estimate linear models with multiple structural changes. The figure reports the estimate of the Beveridge elasticity ϵ and its 95% confidence interval. The standard errors used to compute the confidence interval are corrected for autocorrelation in the errors, as well as heterogeneity in the data and errors across regimes. The shaded areas represent recessions, as identified by the NBER.

estimated. The standard errors—corrected for autocorrelation in the error term, as well as heterogeneity in the data and error term across regimes—vary between 0.06 and 0.15.[5] The elasticities and their 95% confidence intervals are displayed in figure 4.[6]

The fit of the linear model with structural changes is good: $R^2 = 0.91$. The good fit confirms that unemployment and vacancy rates travel in the close vicinity of an isoelastic curve that shifts from time to time.

Our estimates of the Beveridge elasticity are consistent with the estimates of the matching elasticity obtained by the empirical literature studying the matching function. A midrange estimate of the matching elasticity is $\alpha = 0.5$ (Petrongolo and Pissarides 2001). The average unemployment rate over the 1951–2019 period is $u = 5.8\%$. Combining these values with the expression of the Beveridge elasticity in a Diamond-Mortensen-Pissarides model, given by (1), we obtain $\epsilon = [0.5 + 0.058/(1 - 0.058)]/(1 - 0.5) = 1.12$. This value is not far from our average

[5]The uncorrected standard errors are smaller: between 0.04 and 0.07.

[6]In our analysis, any temporary deviation from the balanced-flow assumption that occurs in the data would appear in the error term. (This is why we allow the error term to be autocorrelated and heteroskedastic.) Ahn and Crane (2020) refine the analysis by imposing more structure on labor-market flows. The extra structure allows them to quantify how much out-of-balance flows contribute to the observed movements in unemployment and vacancies. This information, in turn, allows them to estimate to estimate the Beveridge elasticity more finely.

elasticity of $\epsilon = 0.91$.

4.2. Recruiting cost (κ)

To estimate the recruiting cost we use the 1997 National Employer Survey. In the survey, the Census Bureau asked 4,500 establishments about their recruiting process, and found that firms spend 2.5% of their labor costs on recruiting (Villena Roldan 2010). Assuming that all workers are paid the same, this means that firms allocate 2.5% of their labor to recruiting: $\kappa v = 2.5\% \times (1-u)$. In 1997, the average vacancy rate is 3.3% and the average unemployment rate is 4.9% (figure 1). Hence, we find that the recruiting cost in 1997 is $\kappa = 2.5\% \times (1 - 4.9\%)/3.3\% = 0.72$. Since there is no other comprehensive measure of recruiting cost in the United States, we assume that the recruiting cost remains at its 1997 value over the entire 1951–2019 period.

It would be easy to measure the recruiting cost better in the future. It would only require the BLS to add a new question into JOLTS, asking firms to report each month the number of man-hours devoted to recruiting in addition to the number of vacancies.

4.3. Social value of nonwork (ζ)

To estimate the social value of nonwork, we rely on the work of Borgschulte and Martorell (2018) and Mas and Pallais (2019).

Raw estimates. Using military administrative data covering 1993–2004, Borgschulte and Martorell study how servicemembers' reenlistment choice is influenced by unemployment. This choice allows them to estimate the dollar value of the utility loss caused by higher unemployment during the transition to civilian life, and compare it to the earnings loss caused by higher unemployment. They find that during unemployment home production, recreation, and public benefits offset between 13% and 35% of lost earnings.

Using a large field experiment, Mas and Pallais study how unemployed job applicants choose between randomized wage-hour bundles. These choices imply that the value of home production and recreation during unemployment amounts to 58% of predicted earnings.

Translating raw estimates into social values of nonwork. Next, we translate these estimates into social values of nonwork. The first step is to express the estimates relative to the marginal product of labor rather than earnings. The marginal product of labor is higher than earnings for several reasons. First, the wage paid by firms is usually lower than the marginal product of labor. In a matching model the wedge amounts to the share of workers allocated to recruiting, so the marginal product is about 3% higher than the wage (Landais, Michaillat, and Saez 2018,

equation (1)). In a monopsony model, the wedge depends on the elasticity of the labor supply, and the marginal product may be 25% higher than the wage (Mas and Pallais 2019, p. 121). Second, the wage received by workers is lower than that paid by firms because of employer-side payroll taxes, which amount to 7.7%. Third, Mas and Pallais discount predicted earnings by 6% to capture the wage penalty incurred by workers who recently lost their jobs; we undo the discounting because the penalty does not seem to apply to the marginal product of labor (Davis and von Wachter 2011). To conclude, to obtain a marginal product of labor, Borgschulte and Martorell's earnings have to be adjusted by a factor between $1.03 \times 1.077 = 1.11$ and $1.25 \times 1.077 = 1.35$, and Mas and Pallais's earnings by a factor between $1.03 \times 1.077 \times 1.06 = 1.18$ and $1.25 \times 1.077 \times 1.06 = 1.43$. Accordingly, to obtain a value relative to the marginal product of labor, Borgschulte and Martorell's estimates must be adjusted by a factor between $1/1.35 = 0.74$ and $1/1.11 = 0.90$, and Mas and Pallais's estimates must be adjusted by a factor between $1/1.43 = 0.70$ and $1/1.18 = 0.85$.

The second step only applies to Borgschulte and Martorell's estimates, from which we subtract the value of public benefits. All servicemembers are eligible to unemployment insurance (UI). Chodorow-Reich and Karabarbounis (2016, pp. 1585–1586) find that UI benefits amount to 21.5% of the marginal product of labor. But this quantity has to be reduced for several reasons: the UI takeup rate is only 65%; UI benefits and consumption are taxed, imposing a factor of 0.83; the disutility from filing for benefits imposes a factor of 0.47; and UI benefits expire, imposing another factor of 0.83. In sum, the average value of UI benefits to servicemembers is $21.5 \times 0.65 \times 0.83 \times 0.47 \times 0.83 = 5\%$ of the marginal product of labor. Servicemembers are also eligible to other public benefits, which Chodorow-Reich and Karabarbounis quantify at 2% of the marginal product of labor. Hence, to account for benefits, we subtract $5\% + 2\% = 7\%$ of the marginal product of labor from Borgschulte and Martorell's estimates.

Combining these two steps, we find that Mas and Pallais's estimates imply a social value of nonwork between $0.58 \times 0.70 = 0.41$ and $0.58 \times 0.85 = 0.49$; and that Borgschulte and Martorell's estimates imply a social value of nonwork between $(0.13 \times 0.74) - 0.07 = 0.03$ and $(0.35 \times 0.90) - 0.07 = 0.25$. The range of plausible values for the social value of nonwork therefore is 0–0.5; we set the statistic to its midrange value, $\zeta = 0.25$.

Fluctuations of the social value of nonwork. We keep the social value of nonwork constant over the business cycle. In some models, the productivities of unemployed and employed workers do not move in tandem over the business cycle, generating fluctuations in the social value of nonwork (for instance, in the Diamond-Mortensen-Pissarides model, as shown by (3)). However, Chodorow-Reich and Karabarbounis (2016, pp. 1599–1604) find no evidence of such fluctuations in US data. Instead, they establish that the utility derived by unemployed workers from recreation and home

production moves proportionally to labor productivity—which implies that ζ is acyclical. The social value of nonwork could also exhibit medium-run fluctuations, but we omit them by lack of evidence.

4.4. Unemployment gap

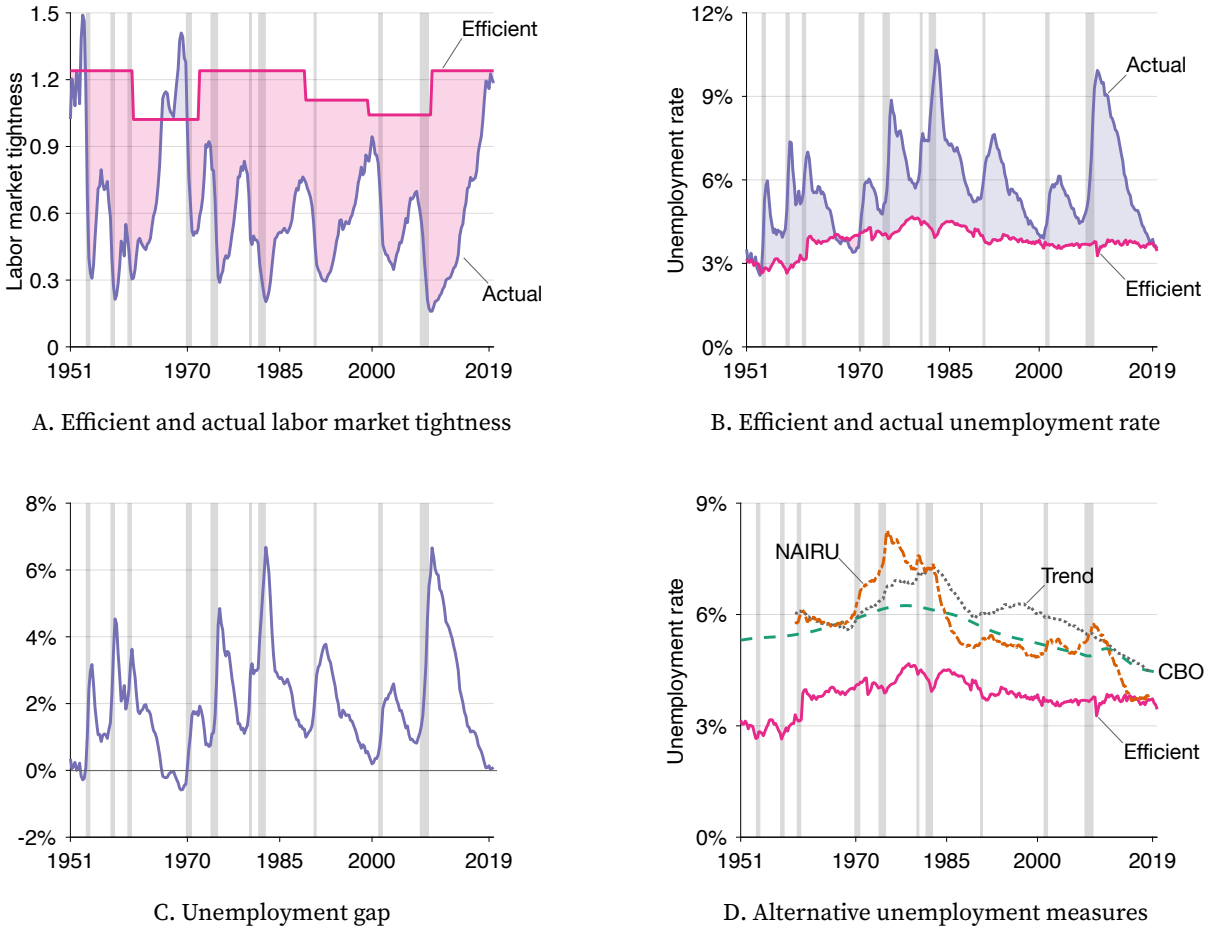
We now use our estimates of the Beveridge elasticity, recruiting cost, and social value of nonwork, as well as our unemployment and vacancy series, to measure the unemployment gap in the United States between 1951 and 2019.

Efficient labor market tightness. We begin by computing the efficient labor market tightness using formula (5) (figure 5, panel A). The efficient tightness fluctuates over time between 1.02 and 1.24, mirroring the movements of the Beveridge elasticity. The average efficient tightness is 1.16. Since actual and efficient tightness almost never coincide, the US labor market almost never operates efficiently. Compared to its efficient level, actual tightness was almost always too low during the period. The only two episodes when tightness was inefficiently high were 1952–1953, during the Korea war, and 1966–1970, at the peak of the Vietnam war; tightness was also virtually efficient in 2019. This implies that the US labor market is generally inefficiently slack.

Efficient unemployment rate. Next we compute the efficient unemployment rate from formula (6) (figure 5, panel B). The efficient unemployment rate hovered around 3% in the 1950s, rose to 4% in the 1960s, and climbed to reach 4.7% in 1979. The steady increase of the efficient unemployment rate between 1951 to 1979 was caused by a steady outward shift of the Beveridge curve (figure 3). Then, the efficient unemployment rate declined to reach 4% in 1990. The decline was caused by an inward shift of the Beveridge curve (figure 3, panels C–D). The efficient unemployment rate then remained stable through the 1990s, 2000s, and 2010s, hovering between 3.3% and 4.0%.

Interestingly the efficient unemployment rate did not increase in the aftermath of the Great Recession—despite the outward shift of the Beveridge curve (figure 3, panels E–F). This is because the Beveridge curve also became flatter after 2009: it fell from 1.0 to 0.84 (figure 4). The flattening offset the outward shift, leaving the efficient unemployment rate almost unchanged by the recession.

Unemployment gap. Measuring the distance between the actual and the efficient unemployment rate, we obtain the unemployment gap (figure 5, panel C). Between 1951 and 2019, the unemployment gap averages 2.0 percentage points. The unemployment gap is sharply countercyclical. It is close to zero at business-cycle peaks: sometimes negative (−0.3 percentage points in 1953 and −0.6



A. Efficient and actual labor market tightness

B. Efficient and actual unemployment rate

C. Unemployment gap

D. Alternative unemployment measures

FIGURE 5. Unemployment gap in the United States, 1951–2019

Panel A: Actual tightness is vacancy rate divided by unemployment rate, from figure 1). The efficient tightness is computed from (5) with $\kappa = 0.72$, $\zeta = 0.25$, ϵ from figure 4. Panel B: The actual unemployment rate comes from figure 1. The efficient unemployment rate is computed from (6) with $\kappa = 0.72$, $\zeta = 0.25$, ϵ from figure 4, and the unemployment and vacancy rates from figure 1. Panel C: The unemployment gap is the difference between the actual and efficient unemployment rates from panel B. Panel D: The efficient unemployment rate come from panel B. CBO unemployment rate is the long-term natural rate of unemployment constructed by the Congressional Budget Office. NAIRU and trend unemployment rates come from Crump et al. (2019, figure 8B). The shaded areas represent recessions, as identified by the NBER.

points in 1969); sometimes close to zero (in 2019); and sometimes positive (for instance 0.2 points in 2000 and 0.8 points in 2007). And the unemployment gap is highly positive at business-cycle troughs: for instance, 6.7 points in 1982, 3.8 points in 1992, and 6.7 points in 2009. Unsurprisingly, the largest unemployment gaps occurred in the wake of the Volcker recession and Great Recession.

Although assumption 1 is required to obtain precise values of the unemployment gap (proposition 3), it is possible to determine whether unemployment is inefficiently high or low without the assumption. Indeed, unemployment is inefficiently high whenever labor market tightness is inefficiently low. Conversely, whenever tightness is inefficiently high, we know that unemployment is inefficiently low. This in turn can be determined from the graph in panel A of figure 5, which is constructed without relying on assumption 1.

Summary and implications. In sum, the US unemployment rate is generally inefficiently high—an inefficiency exacerbated in slumps. These results have implications for policymaking and macroeconomic modeling. First, because the unemployment spikes in recessions are inefficient, it is warranted to deploy fiscal and monetary policy in bad times to reduce unemployment. Second, given that the unemployment rate is almost always inefficient, and sometimes markedly so, it might not be productive to insist on modeling the labor market as efficient—whether by assuming that the Hosios condition holds or by focusing on a competitive-search equilibrium.

Comparisons with other unemployment gaps. Finally, to provide some context, we compare our efficient unemployment rate to three other unemployment measures that are used to construct unemployment gaps: trend unemployment, the NAIRU, and the Congressional Budget Office's natural rate of unemployment (figure 5, panel D). Trend unemployment and the NAIRU are constructed by Crump et al. (2019, figure 8B) using state-of-the-art techniques. The CBO's natural rate of unemployment—which features prominently in policy discussions—is constructed by blending trend and NAIRU considerations (Shackleton 2018, appendix B). The main similarity between the four measures is that they are slow-moving. As the actual unemployment rate is sharply countercyclical, the unemployment gap constructed with either measure will be countercyclical. Another similarity is that the four unemployment measures were higher in the 1970s and 1980s, and lower after that. The main difference is that our measure is lower than the three others. On average the efficient unemployment rate is 1.7 percentage points below the CBO's natural rate of unemployment, 1.8 points below the NAIRU, and 2.1 points below trend unemployment. As a result, the unemployment gap constructed with the efficient unemployment rate will be higher than that constructed with the three other unemployment measures.

4.5. Alternative calibrations of the sufficient statistics

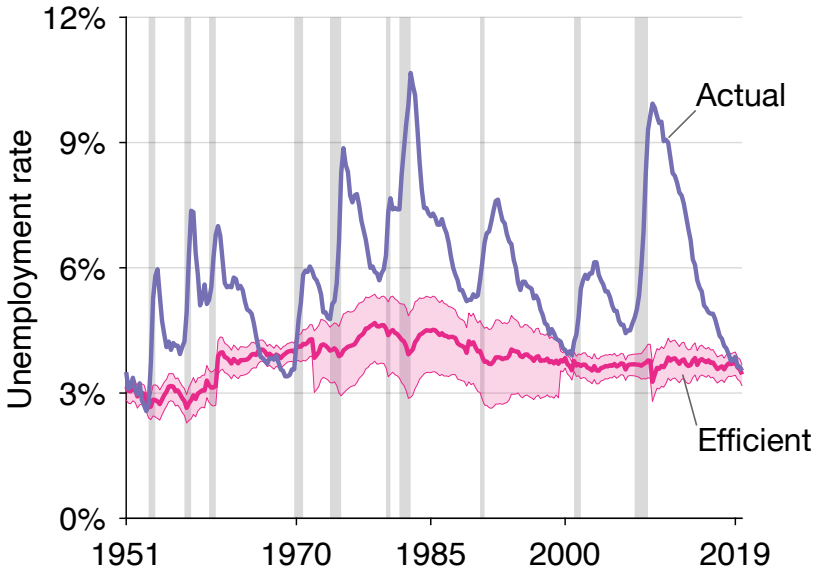
We explore the sensitivity of the efficient unemployment rate and unemployment gap to alternative calibrations of the sufficient statistics.

Beveridge elasticity. We begin by reconstructing the efficient unemployment rate when the Beveridge elasticity takes any value in its 95% confidence interval (figure 6, panel A). When the estimated elasticity is higher, the efficient unemployment rate is also higher. Hence, when the Beveridge elasticity is at the bottom end of its 95% confidence interval, the efficient unemployment rate follows the same pattern as under the baseline calibration, but is on average 0.6 percentage point lower. And when the Beveridge elasticity is at the top end of its 95% confidence interval, the efficient unemployment rate follows the same pattern as under the baseline calibration, but is on average 0.5 percentage point higher. For any Beveridge elasticity inside the confidence interval, the efficient unemployment rate is somewhere between these two extremes (pink band in the figure).

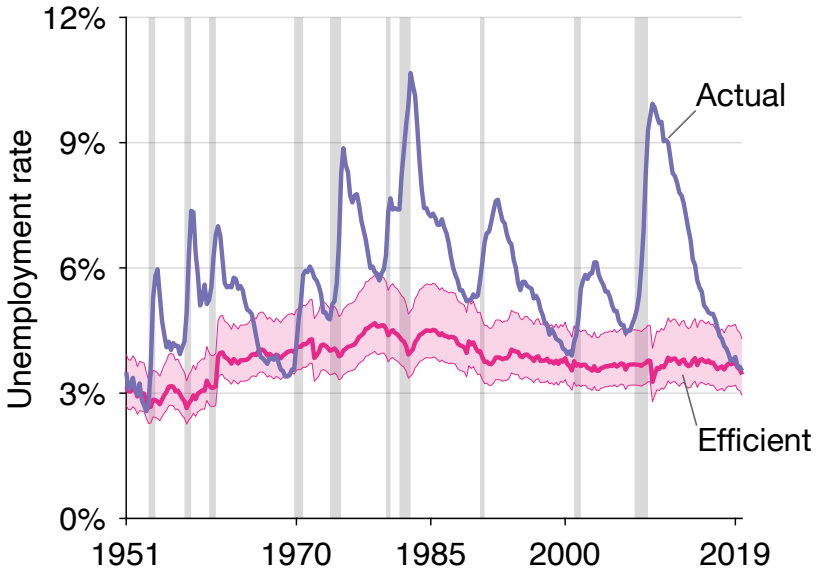
To summarize, when the estimated elasticity ϵ spans its 95% confidence interval, the efficient unemployment rate remains contained in a band whose width averages 1.0 percentage point, and is always below 2.0 percentage points. Furthermore, since vacancy data from JOLTS have become available (2001), estimates of the Beveridge elasticity have become more precise, and the 95% confidence interval for ϵ has narrowed. As a result, the band of possible efficient unemployment rates has narrowed to a width of 0.7 percent points since 2001.

Social value of nonwork. Next we reconstruct the efficient unemployment rate when the social value of nonwork spans the range of plausible values given by Borgschulte and Martorell (2018) and Mas and Pallais (2019) (figure 6, panel B). First, we consider an estimate at the low end of the range of plausible values: $\zeta = 0$. Under this calibration, the efficient unemployment rate follows the same pattern as under the baseline calibration but is on average 0.5 percentage point lower. Next, we consider an estimate at the high end of the range of plausible values: $\zeta = 0.5$, which is consistent with the estimates from the macro-labor literature (Chodorow-Reich and Karabarbounis 2016, equation (30)). The efficient unemployment rate follows again the same pattern as under the baseline calibration, but it is on average 0.9 percentage point higher.

These different calibrations show that around our baseline calibration of $\zeta = 0.25$, the efficient unemployment rate is fairly insensitive to the precise value of ζ. For any ζ in the 0–0.5 range, the efficient unemployment rate remains contained in a band whose whose width averages 1.4 percentage point, and is always below 1.8 percentage points. This is reassuring as the range of plausible values for ζ is quite broad.



A. Beveridge elasticity: $\epsilon \in$ 95% confidence interval



B. Social value of nonwork: $\zeta \in [0, 0.5]$

FIGURE 6. Efficient unemployment rate in the United States for a range of plausible calibrations of the sufficient statistics

Panel A reproduces panel B of figure 5, and adds the range of efficient unemployment rates obtained when the Beveridge elasticity spans its 95% confidence interval (displayed in figure 4). The thickest line is the efficient unemployment rate obtained with the point estimate of ϵ. The top thin line is the efficient unemployment rate obtained with the top-range value of ϵ. The bottom thin line is the efficient unemployment rate obtained with the bottom-range value of ϵ. Panel B reproduces panel B of figure 5, and adds the range of efficient unemployment rates obtained when the social value of nonwork spans the range of plausible values: $\zeta \in [0, 0.5]$. The thickest line is the efficient unemployment rate obtained with the midrange value of $\zeta = 0.25$. The top thin line is the efficient unemployment rate obtained with $\zeta = 0.5$. The bottom thin line is the efficient unemployment rate obtained with $\zeta = 0$. The shaded areas represent recessions, as identified by the NBER.

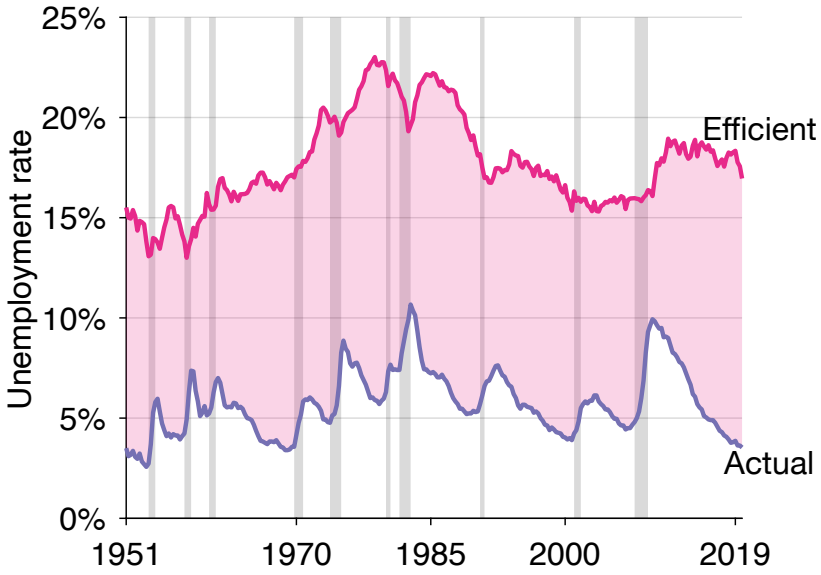


FIGURE 7. Efficient unemployment rate in the United States when the social value of nonwork is $\zeta = 0.96$, as proposed by Hagedorn and Manovskii (2008)

The figure reproduces panel B of figure 5, but uses instead the social values of nonwork proposed by Hagedorn and Manovskii (2008): $\zeta = 0.96$. The shaded areas represent recessions, as identified by the NBER.

Conclusion. For any plausible estimate of the sufficient statistics, the unemployment gap never departs from the baseline by more than 1 percentage point. This means that our substantive findings—that the US labor market is almost always inefficient, generally inefficient slack, and especially inefficiently slack in slumps—are robust to alternative calibrations.

Aside on Hagedorn-Manovskii calibration. Our baseline calibration ($\zeta = 0.25$) implies that the social value of unemployment is much lower than labor productivity; in contrast, some macro-labor studies argue that the two are very close. A well-known calibration, due to Hagedorn and Manovskii (2008), is $\zeta = 0.96$. Such a calibration has a drastic impact: it pushes the efficient unemployment rate above 13%, and sometimes as high as 23%, with an average value of 17.7% (figure 7). Under this calibration, unemployment is always inefficiently low—even at the peak of the Great Recession. This result seems implausible, suggesting that such calibration understates the social cost of unemployment.

4.6. Inverse-optimum sufficient statistics

To provide further perspective, we compute the values of the sufficient statistics that arise under the assumption that US unemployment is efficient at all time. Such inverse-optimum sufficient

statistics are commonly computed in public economics to understand the conditions under which the economy would operate efficiently (for example, Hendren 2020). Here, the distance between the inverse-optimum values of the statistics and their calibrated values is another measure of the distance between current labor market conditions and efficiency.

Beveridge elasticity. We start with the inverse-optimum Beveridge elasticity. Proposition 2 shows that the value of the Beveridge elasticity under which actual tightness θ is efficient, given the other sufficient statistics, is

$$\epsilon^* = \frac{1-\zeta}{\kappa\theta}. \tag{7}$$

To support labor market efficiency, the Beveridge elasticity would have to be strongly countercyclical, varying between 0.7 in booms and 6.5 during the Great Recession (figure 8, panel A). Furthermore, the average Beveridge elasticity would need to be much higher than estimated, with a value of 2.1 instead of 0.9. In sum, the inverse-optimum Beveridge elasticity is generally far above the 95% confidence interval for the estimated Beveridge elasticity.

Recruiting cost. We turn to the inverse-optimum recruiting cost. Again, from proposition 2, the value of the recruiting cost under which actual tightness is efficient is

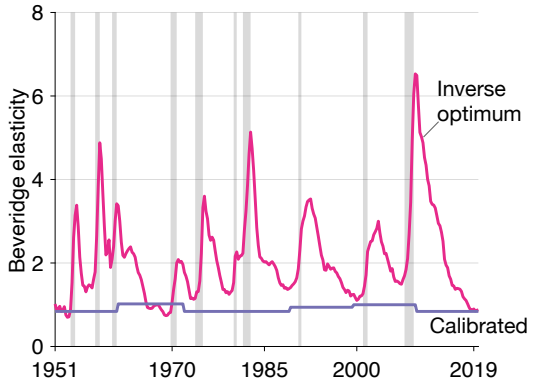
$$\kappa^* = \frac{1-\zeta}{\epsilon\theta}. \tag{8}$$

To support labor market efficiency, the recruiting cost would have to be strongly countercyclical, varying between 0.5 in booms and 5.6 during the Great Recession, with an average value of 1.7 (figure 8, panel B). Since the inverse-optimum recruiting cost is ten times higher in slumps than in booms, unemployment is continuously efficient only if filling a vacancy requires ten times more labor in bad times than in good times.

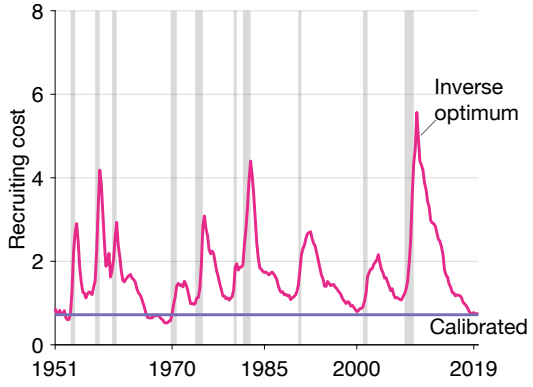
Social value of nonwork. Last, we consider the inverse-optimum social value of nonwork. Once more, proposition 2 indicates that the social value of nonwork under which actual tightness is efficient is

$$\zeta^* = 1 - \kappa\epsilon\theta. \tag{9}$$

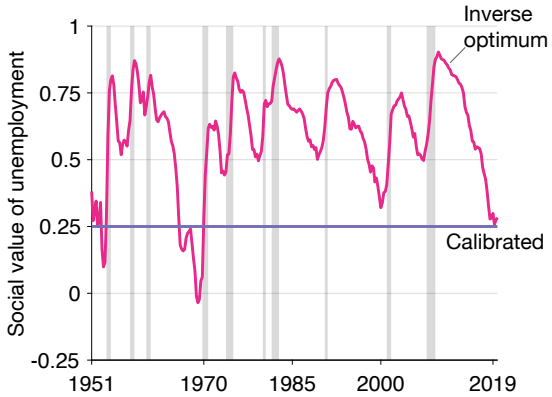
To sustain efficiency, the social value of nonwork would need to be immensely countercyclical, as low as -0.04 in booms and as high as 0.90 during the Great Recession, with an average value of



A. Beveridge elasticity



B. Recruiting cost



C. Social value of nonwork

FIGURE 8. Inverse-optimum sufficient statistics

The figure displays the Beveridge elasticity (panel A), recruiting cost (panel B), and social value of nonwork (panel C) that would ensure that the observed unemployment rate is efficient at all time. These values are obtained from (7), (8), and (9). The figure also compares these inverse-optimum values with the calibrated values of the statistics. The shaded areas represent recessions, as identified by the NBER.

0.59 (figure 8, panel C). Under the inverse-optimum social value of nonwork, recessions would merely be vacations.

5. Conclusion

This paper develops a new method to measure the unemployment gap—the difference between the actual and the socially efficient unemployment rate. We consider a labor-market model with only one structural element: a Beveridge curve relating unemployment and vacancies. This framework covers many modern labor-market models, including the Diamond-Mortensen-Pissarides model. We show that the unemployment gap can simply be measured from current unemployment and vacancy rates, and three sufficient statistics: the elasticity of the Beveridge curve, cost of recruiting, and social cost of unemployment.

We apply our unemployment-gap formula to the United States, 1951–2019. We find that the US unemployment gap is countercyclical: the gap is close to zero in booms, slightly positive or negative; it is highly positive in slumps. We infer that the US unemployment rate is generally inefficiently high, and such inefficiency worsens in slumps. Historically, it was therefore warranted to activate fiscal and monetary policy in bad times to reduce unemployment.

As the world emerges from its coronavirus quarantine and economic activity resumes, policymakers need to assess the unemployment gap. The gap indicates how much inefficient slack remains within their economies, and the extent to which stabilization policies need to be brought to bear as a result. Since our unemployment-gap measure can be calculated in real time, it could be helpful to guide such policy. Michaillat and Saez (2019a,b) have developed formulas linking optimal monetary policy and fiscal policy to the unemployment gap. Combined with these formulas, the real-time measure of the unemployment gap could help policymakers decide how much they should adjust interest rates (if the zero lower bound is not binding) or how much they should adjust public expenditure.

References

Abraham, Katharine G. 1987. "Help-Wanted Advertising, Job Vacancies, and Unemployment." *Brookings Papers on Economic Activity* 18 (1): 207–248.

Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis. 2013. "Time Use during the Great Recession." *American Economic Review* 103 (5): 1664–1696.

Ahn, Hie Joo, and Leland D. Crane. 2020. "Dynamic Beveridge Curve Accounting." arXiv:2003.00033.

Andrews, Donald W. K. 1991. "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica* 59 (3): 817–858.

Bai, Jushan, and Pierre Perron. 1998. "Estimating and Testing Linear Models with Multiple Structural Changes." *Econometrica* 66 (1): 47–78.

Bai, Jushan, and Pierre Perron. 2003. "Computation and Analysis of Multiple Structural Change Models." *Journal of Applied Econometrics* 18: 1–22.

Barnichon, Regis. 2010. "Building a composite Help-Wanted Index." *Economics Letters* 109 (3): 175–178.

Blanchard, Olivier J., and Peter Diamond. 1989. "The Beveridge Curve." *Brookings Papers on Economic Activity* 1989 (1): 1–76.

Borgschulte, Mark, and Paco Martorell. 2018. "Paying to Avoid Recession: Using Reenlistment to Estimate the Cost of Unemployment." *American Economic Journal: Applied Economics* 10 (3): 101–127.

Brand, Jennie E. 2015. "The Far-Reaching Impact of Job Loss and Unemployment." *Annual Review of Sociology* 41: 359–375.

Cahuc, Pierre, Francois Marque, and Etienne Wasmer. 2008. "A Theory of Wages and Labor Demand with Intra-Firm Bargaining and Matching Frictions." *International Economic Review* 49 (3): 943–972.

Chetty, Raj. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods." *Annual Review of Economics* 1: 451–488.

Chodorow-Reich, Gabriel, and Loukas Karabarbounis. 2016. "The Cyclicality of the Opportunity Cost of Employment." *Journal of Political Economy* 124 (6): 1563–1618.

Costain, James S., and Michael Reiter. 2008. "Business Cycles, Unemployment Insurance, and the Calibration of Matching Models." *Journal of Economic Dynamics and Control* 32 (4): 1120–1155.

Crump, Richard K., Stefano Eusepi, Marc Giannoni, and Aysegul Sahin. 2019. "A Unified Approach to Measuring u*." *Brookings Papers on Economic Activity* 50 (1): 143–238.

Davis, Steven J., and Till von Wachter. 2011. "Recessions and the Costs of Job Loss." *Brookings Papers on Economic Activity* 42 (2): 1–72.

den Haan, Wouter J., Garey Ramey, and Joel Watson. 2000. "Job Destruction and Propagation of Shocks." *American Economic Review* 90 (3): 482–498.

Diamond, Peter A., and Aysegul Sahin. 2015. "Shifts in the Beveridge curve." *Research in Economics* 69 (1): 18–25.

Ebrahimy, Ehsan, and Robert Shimer. 2010. "Stock-Flow Matching." *Journal of Economic Theory* 145 (4): 1325–1353.

Elsby, Michael W. L., and Ryan Michaels. 2013. "Marginal Jobs, Heterogeneous Firms, and Unemployment Flows." *American Economic Journal: Macroeconomics* 5 (1): 1–48.

Elsby, Michael W. L., Ryan Michaels, and David Ratner. 2015. "The Beveridge Curve: A Survey." *Journal of Economic Literature* 53 (3): 571–630.

Elsby, Michael W. L., Ryan Michaels, and Gary Solon. 2009. "The Ins and Outs of Cyclical Unemployment." *American Economic Journal: Macroeconomics* 1 (1): 84–110.

Gertler, Mark, and Antonella Trigari. 2009. "Unemployment Fluctuations with Staggered Nash Wage Bargaining." *Journal of Political Economy* 117 (1): 38–86.

Giupponi, Giulia, and Camille Landais. 2018. "Subsidizing Labor Hoarding in Recessions: The Employment and Welfare Effects of Short Time Work." CEPR Discussion Paper 13310.

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. 2005a. " Employment Fluctuations with Equilibrium Wage Stickiness." *American Economic*

Review 95 (1): 50–65.

Hall, Robert E. 2005b. "Employment Efficiency and Sticky Wages: Evidence from Flows in the Labor Market." *Review of Economics and Statistics* 87 (3): 397–407.

Hall, Robert E., and Paul R. Milgrom. 2008. "The Limited Influence of Unemployment on the Wage Bargain." *American Economic Review* 98 (4): 1653–74.

Hendren, Nathaniel. 2020. "Measuring Economic Efficiency Using Inverse-Optimum Weights." *Journal of public Economics* 187: 104198.

Hosios, Arthur J. 1990. "On the Efficiency of Matching and Related Models of Search and Unemployment." *Review of Economic Studies* 57 (2): 279–298.

Hungerbuhler, Mathias, and Etienne Lehmann. 2009. "On the Optimality of a Minimum Wage: New Insights from Optimal Tax Theory." *Journal of Public Economics* 93 (3-4): 464–481.

Jackman, Richard, Christopher Pissarides, and Savvas Savouri. 1990. "Labour Market Policies and Unemployment in the OECD." *Economic Policy* 5 (11): 449–490.

Jager, Simon, Benjamin Schoefer, Samuel G. Young, and Josef Zweimuller. 2018. "Wages and the Value of Nonemployment." NBER Working Paper 25230.

Kleven, Henrik J. 2018. "Sufficient Statistics Revisited." https://perma.cc/MCR9-BE5J.

Landais, Camille, Pascal Michaillat, and Emmanuel Saez. 2018. "A Macroeconomic Approach to Optimal Unemployment Insurance: Theory." *American Economic Journal: Economic Policy* 10 (2): 152–181.

Liu, Jian, Shiying Wu, and James V. Zidek. 1997. "On Segmented Multivariate Regression." *Statistica Sinica* 7: 497–525.

Mas, Alexandre, and Amanda Pallais. 2019. "Labor Supply and the Value of Non-Work Time: Experimental Estimates from the Field." *American Economic Review: Insights* 1 (1): 111–126.

Michaillat, Pascal. 2012. "Do Matching Frictions Explain Unemployment? Not in Bad Times." *American Economic Review* 102 (4): 1721–1750.

Michaillat, Pascal, and Emmanuel Saez. 2019a. "An Economical Business-Cycle Model."

Michaillat, Pascal, and Emmanuel Saez. 2019b. "Optimal Public Expenditure with Inefficient Unemployment." *Review of Economic Studies* 86 (3): 1301–1331.

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.

Nickell, Stephen, Luca Nunziata, Wolfgang Ochel, and Glenda Quintini. 2002. "The Beveridge Curve, Unemployment and Wages in the OECD from the 1960s to the 1990s." CEP Discussion Paper 502.

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. 1986. "Unemployment and Vacancies in Britain." *Economic Policy* 1 (3): 500–559.

Pissarides, Christopher A. 2000. *Equilibrium Unemployment Theory*. 2nd ed. Cambridge, MA: MIT Press.

Pissarides, Christopher A. 2009a. "Comments on 'A New Method for Estimating Time Variation in the NAIRU' by William T. Dickens." In *Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective*, edited by Jeff Fuhrer, Yolanda K. Kodrzycki, Jane Sneddon Little, and Giovanni P. Olivei, 235–242. Cambridge, MA: MIT Press.

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

Rogerson, Richard. 1997. "Theory Ahead of Language in the Economics of Unemployment." *Journal of*

Economic Perspectives 11 (1): 73–92.

Shackleton, Robert. 2018. "Estimating and Projecting Potential Output Using CBO's Forecasting Growth Model." CBO Working Paper 2018-03.

Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95 (1): 25–49.

Shimer, Robert. 2007. "Mismatch." *American Economic Review* 97 (4): 1074–1101.

Shimer, Robert. 2012. "Reassessing the Ins and Outs of Unemployment." *Review of Economic Dynamics* 15 (2): 127–148.

Villena Roldan, Benjamin. 2010. "Aggregate Implications of Employer Search and Recruiting Selection." CEA Working Paper 271.

Yao, Yi-Ching. 1988. "Estimating the Number of Change-Points via Schwarz' Criterion." *Statistics & Probability Letters* 6 (3): 181–189.