

CEP Discussion Paper No 1058

July 2010

**The Role of Worker Flows in the Dynamics and
Distribution of UK Unemployment**

**Michael W. L. Elsby, Jennifer C. Smith
and Jonathan Wadsworth**

Abstract

Unemployment varies substantially over time and across subgroups of the labour market. Worker flows among labour market states act as key determinants of this variation. We examine how the structure of unemployment across groups and its cyclical movements across time are shaped by changes in labour market flows. Using novel estimates of flow transition rates for the UK over the last 35 years, we decompose unemployment variation into parts accounted for by changes in rates of job loss, job finding and flows via non-participation. Close to two-thirds of the volatility of unemployment in the UK over this period can be traced to rises in rates of job loss that accompany recessions. The share of this inflow contribution has been broadly the same in each of the past three recessions. Decreased job-finding rates account for around one-quarter of unemployment cyclical variation and the remaining variation can be attributed to flows via non-participation. Digging deeper into the structure of unemployment by gender, age and education, the flow-approach is shown to provide a richer understanding of the unemployment experiences across population subgroups.

Key Words: labour market, unemployment, worker flows

JEL Classification: E24, J6

This paper was produced as part of the Centre's Labour Markets Programme. The Centre for Economic Performance is financed by the Economic and Social Research Council.

Acknowledgements

The authors thank Tim Jenkinson for comments. They also thank the Office for National Statistics and the UK Data Archive, respectively Principal Investigator and Distributor of the Labour Force Survey data. Neither is responsible for the analysis or interpretation in this paper.

Michael Elsby is Professor of Economics at the University of Edinburgh and NBER. Jennifer Smith is a Lecturer in the Department of Economics at the University of Warwick. Jonathan Wadsworth is a Senior Research Fellow at the Centre for Economic Performance, London School of Economics and Professor of Economics at Royal Holloway College, University of London.

Published by
Centre for Economic Performance
London School of Economics and Political Science
Houghton Street
London WC2A 2AE

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior permission in writing of the publisher nor be issued to the public or circulated in any form other than that in which it is published.

Requests for permission to reproduce any article or part of the Working Paper should be sent to the editor at the above address.

© M. Elsby, J. Smith and J. Wadsworth, submitted 2010

I. Introduction

A defining feature of the UK economy in recent decades has been the substantial variation in the rate of unemployment. While the unemployment rate hovered at around 5 percent throughout the 1970s, it soared to levels above 10 percent by the mid-1980s, an experience repeated in the depths of the recession of the early 1990s. There then followed a period of sustained improvement and even tranquillity, where unemployment in the UK returned to rates seen in the 1970s, only to rise again in the course of the recent recession. However, despite a larger accompanying fall in GDP than in either of the previous two recessions, the UK unemployment rate did not return to the double digit levels experienced in the two previous recessions. While welcome, the reasons for this are not yet fully understood.

In this paper, we delve into the origins of this variation in the unemployment rate over three different recessions using a dynamic approach. Based on an influential literature dating from the 1970s, it is now known that changes in the stock of unemployment are shaped by the flows of workers into and out of the unemployment pool from both employment and inactivity. Knowledge of the relative size of these flows has been used both to try to understand the main reasons underlying rises (and falls) in unemployment and to shape appropriate policy responses. Does unemployment rise as a result of increased inflows into unemployment driven by elevated rates of job loss? Or does it rise because the unemployed leave the unemployment pool at a slower rate due to declines in their ability to find jobs? Or is it some combination of the two? What are the roles of flows into and out of employment compared to the flows into and out of inactivity in shaping unemployment? Do these roles change over time?

The answers to these questions are important if policy responses are to be shaped appropriately. For example, policy that focused on encouraging outflows from unemployment may not be as relevant in an economy in which rises in unemployment were driven by changes in the rate of outflows from employment. Other studies have considered some of these issues for the UK for earlier periods (Pissarides, (1986); Layard, Nickell and Jackman (1991); Petrongolo and Pissarides (2008); Smith, (2011). Indeed the consensus that emerged from studies of the 1980s recession was that the increase in UK unemployment was prompted by an initial rise in outflow rates from employment, as firms got rid of labour, but driven subsequently by changes in unemployment outflow rates. Policy was then focused on improving search effectiveness of the unemployed and avoiding the build-up of long-term unemployment, culminating in the introduction of the various New Deals under the 1997-2010 Labour government. Recently, the debate over which flows drive recessionary increases in unemployment has re-surfaced in the United States where academics are divided over the roles of job loss or outflows from unemployment (see Shimer (2007), and Elsby, Michaels and Solon (2009)). Prior to the recent recession, a number of prominent researchers in the US suggested that the contribution of unemployment inflows to cyclical ramp-ups in the unemployment rate was negligible (Hall (2005); Shimer (2007)). If this were the case, it would be tempting to conclude that inflows into unemployment had neither a direct, nor an indirect causal impact on unemployment.

It seems therefore important to try to understand the reasons behind recent UK unemployment performance and look to see whether the explanations that were thought to hold in earlier periods still apply. Recent advances in data availability now make it possible, for the first time in the UK, to look comprehensively at these issues over three full labour market cycles, the approach that we adopt here. We exploit information on worker flows

using data on recalled labour market status available in the UK Labour Force Survey (LFS).¹ A key benefit of these data is that they provide measures of worker flows between the three labour market states, unemployment, employment and economic inactivity, from 1975 up to the end of the recent recession.²

After reviewing the evolution of labour market stocks, we document the time series of the flow transition rates between employment, unemployment and non-participation (inactivity) using these recall-based estimates of worker flows from the LFS. Inspection of the cyclical properties of these transition rates reveals a marked counter-cyclicality in the rate of job loss, mirrored by pro-cyclicalities in job-finding rates. Flows involving non-participation are also cyclical, but less so than other flows. This therefore suggests that, in each of the last 3 recessions, a combination of both *more* spells of unemployment as well as increased *duration* of those spells explains the substantial increases in unemployment witnessed.

We show how it is possible to use a formal decomposition of the variation in the unemployment rate into parts accounted for by changes in rates of job loss, job finding and flows via non-participation. Applying this decomposition suggests the following stylized account of unemployment dynamics in the UK: In contrast to the received wisdom, the leading contribution to UK unemployment cyclicalities since 1975 has in fact been the substantial rise in rates of job loss in times of recession, accounting for approximately two-thirds of the fluctuations in the unemployment rate over each cycle. Declines in unemployed workers' job-finding prospects, while undeniably important, explain just over one-quarter of the cyclical change in unemployment in each of the recessions we examine. The remaining 10 percent is attributed to flows involving non-participation. Over time, the relative roles of unemployment outflows and inflows into unemployment have been broadly constant in each of the three recessions covered by our analysis.

In the remainder of the paper, we delve further into the composition of unemployment by examining unemployment rates by gender, age and educational attainment. In terms of overall stocks, it is well-known that young, low-skilled male workers tend to face higher rates of joblessness, a fact that we confirm using LFS data. We show that an analysis of differences in worker flows provides a rich picture of the origins of these cross-sectional differences in unemployment.

While changes in flow rates involving non-participation play a small role in shaping aggregate unemployment, they are an important contributor to both the level and the dynamics of female unemployment. That men face higher unemployment rates than women can be attributed in large part to the fact that women are more likely to exit from unemployment out of the labour force.

The analysis of unemployment flows by age is similarly enlightening. The fact that younger workers are more likely to be unemployed makes it tempting to conclude that youth bear the brunt of joblessness disproportionately. Worker flows by age paint a more nuanced picture, however. We confirm the findings of other studies using different sample periods to ours that for the thirty-five years covered by our sample, while young workers are much more likely to lose their jobs, they also are more likely to find jobs. Thus, younger workers face a more fluid labour market, experiencing more jobless spells, but for shorter durations.

The breakdown of unemployment rates by educational attainment, however, points to a subgroup of the labour market that is hit harder on all margins: the less-educated. Workers who have left school prior to age 18 face significantly higher rates of job loss which are

¹ These measures of worker flows were first constructed by Gregg and Wadsworth (1995) in their analysis of the evolution of the UK labour market from the mid-1970s to the mid-1990s.

² Leading alternative sources from the British Household Panel Survey and the longitudinally-linked quarterly LFS provide measures beginning only in the late-1980s and early-1990s respectively (Smith (2011); Petrongolo and Pissarides (2008)).

aggravated further by depressed rates of job finding, and associated longer unemployment spells.

The paper is organized as follows: Section II documents the behaviour of labour market stocks over the last thirty-five years. Section III introduces worker flows, how they are defined, and their average levels over the sample period. In section IV, we introduce the law of motion for unemployment, which links variation in worker flows to variation in the stock of unemployed workers. We then describe our measures of the flows, and document their evolution since 1975. Section V then takes on the task of decomposing the variation in the unemployment rate into parts accounted for by its constituent flows, and summarizes the results. Finally, section VI analyzes unemployment rates across subgroups of the labour market. Section VII concludes.

II. A Brief History of Labour Market Stocks in the United Kingdom

The main focus of this paper is to document the evolution of the unemployment rate in the UK, and the flows that underlie it. However, unemployment is just one of the three key labour market stocks that form the basis of modern-day labour market indicators—employment, E , unemployment, U and non-participation, N . Non-participation may also be referred to as “out of the labour force” or “inactivity.” This is a rather heterogeneous group, comprising the long-term sick, discouraged workers, students, early retirees and those engaged in full-time domestic work. The headline measures of employment, unemployment and non-participation that are published regularly by the Office for National Statistics (ONS) in the UK are based on definitions developed by the International Labour Organization. These ILO definitions have been adopted by many national statistical agencies, and are summarized in Figure 1. A person who reports work for one hour or more in the survey week is classified as employed. In the event of no work, two further criteria determine labour market status: If having looked for work in the last four weeks and available to start work within the subsequent two weeks, the respondent is recorded as unemployed. Otherwise, they are classified as a non-participant.

The official source of data on these labour market stocks for the UK is the Labour Force Survey (LFS). Figure 1 summarises the magnitude of these stocks using LFS data from 1975 to 2010. On average over this period, out of a working-age population of 33.4 million, approximately 26.1 million were employed, 2.0 million were unemployed and 7.3 million were out of the labour force.³

These figures underlie the key labour market indicators that researchers, policymakers and pundits alike use to obtain a first glimpse of the overall condition of the labour market. A key indicator, and the focus of the remainder of this paper, is the unemployment rate. At a given point in time t , this is defined as the ratio of the number unemployed U_t to the number either employed, E_t or unemployed—the sum of which comprises the labour force $L_t = E_t + U_t$,

$$u_t = U_t/L_t. \quad (1)$$

The unemployment rate is intimately related to two further headline labour market measures, the employment-to-population ratio (“e-pop”) and the labour force participation rate. The former is self-explanatory. The participation rate is the fraction of the population

³ The working age population has grown by around 4.5 million over this period, a combination of earlier baby boom generations reaching adulthood and, more recently, rising net immigration.

that is in the labour force—those either working or seeking work at a point in time. To see how these three measures are intertwined, note that the unemployment rate can be related to employment and the labour force according to the identity $u_t \equiv 1 - (E_t/L_t)$. Total differentiation of this identity reveals that

$$du_t = (1 - u_t)[d \ln(L_t/Pop_t) - d \ln(E_t/Pop_t)]. \quad (2)$$

Thus, when the unemployment rate rises, it could be associated with a rise in the labour force participation rate, or a decline in the e-pop ratio, or some combination of the two. Moreover, equation (2) further reveals that it is *logarithmic* changes in the labour force participation rate and e-pop ratio that shape changes in the level of the unemployment rate.⁴

Figure 2 plots the working age unemployment rate, the e-pop ratio, and the participation rate from the early 1970s to the latest available data. Based on equation (2), the e-pop ratio and the participation rate are plotted on a logarithmic scale, while the unemployment rate is on a standard scale so that equal-sized variation in each of these series places them approximately on an equal footing with respect to changes in the unemployment rate.⁵

The unemployment rate in the UK has varied substantially over the sample period, ranging from a low of 4% in 1975 to a high of 12% in 1986 with three notable cycles. A notable feature of Figure 2 is that the unemployment rate is strongly countercyclical. The recessions starting in 1973, 1979, 1990 and the most recent recession beginning in 2008 all have been accompanied by sharp rises in the unemployment rate, as shown in Table 1. The trough to peak rise in unemployment in the latest recession is noticeably lower than in earlier downturns, which is remarkable given the large fall in GDP this time round—the contrast in the unemployment and output changes across recessions apparently violating the stable relationship predicted by Okun’s Law (1962).

Figure 2 provides a perspective on the origins of these cyclical rises in joblessness. The rise in the unemployment rate in each recession is accompanied by commensurate declines in the e-pop ratio. In contrast, the aggregate labour force participation rate is only mildly procyclical, falling modestly in the 1980 and 1990s downturns, and much less so in the latest downturn. Thus, the majority of the rise in aggregate unemployment in times of recession can be traced to near-symmetric declines in employment.⁶

A further characteristic of cyclical movements in unemployment evident from Figure 2 is the persistence of high unemployment rates following an initial recessionary shock. Table 1 confirms that the duration of elevated unemployment rates far exceeds the duration of declining output in recessions.

In the remainder of this paper, we look at what accounts for the substantial cyclical variation in the UK unemployment rate. Cyclical variation, however, is not our only focus. A particular benefit of using microdata on individual workers is that one can analyse the unemployment experiences of particular subgroups of the labour market. Figures 3, 4 and 5 plot unemployment rates by gender, age and education groups from 1975 on.⁷ These figures confirm that the experience of unemployment is not allocated uniformly across the

⁴ This effectively means that percentage changes in participation and employment rates drive changes in unemployment rates.

⁵ Note that the tick marks on the Figures are reported in levels to help the exposition.

⁶ This is not true, however, when the stocks are disaggregated—for example, by gender. Contrasting trends in inactivity rates by gender observed over our sample period are offset when aggregated.

⁷ To examine unemployment rates by age and level of education, we need the LFS micro data. Official statistics from the ONS on unemployment disaggregated by age are only available from 1992, and no breakdowns are published by education.

population. Rather, some groups are much more likely to be in want of work than others. Specifically, young, male, less-educated workers face significantly higher unemployment rates than average. In addition, these same workers are more likely to experience larger rises in unemployment in times of recession. In the light of this, an important question that arises is what accounts for the fact that particular subgroups are hit harder than others in the labour market.

In what follows, we also attempt to provide an account of the proximate determinants of variation in unemployment in the UK across groups and across time. The next section introduces these flows, and explains how they shape the evolution of unemployment.

III. An Introduction to Worker Flows

The stocks illustrated in Figures 1 to 5 provide important information on the state of the aggregate labour market over time. But, they miss a feature of labour market dynamics that modern research on unemployment views as crucial, namely the *fluidity* of the labour market. At all points in time, many individuals flow in and out of the three labour market states. Unemployed workers flow into employment as they find new jobs. Employed workers flow into unemployment when they lose their jobs. Likewise, employed and unemployed workers flow out of and into non-participation as they enter and exit the labour force.

A long and distinguished line of research has identified the existence of substantial *worker flows* as a defining characteristic of labour markets. Much of this research evolved in the United States, where readily available micro-data on worker flows began to be exploited in the early 1970s.⁸ Seminal contributions by Kaitz (1970), Perry (1972) and Marston (1976) were among the first to describe how existing data sources could be used to estimate worker flows, and how these flows shape the evolution of unemployment. Figure 6 adapts a diagram introduced by Blanchard and Diamond (1990) that has become a staple in any analysis of worker flows. The three labour market states form the three corners of a triangle, and the arrows between them represent worker flows between the three states.

Figure 6 reports the results of two approaches to measuring worker flows. The first is to measure the *number* of workers who move between *E*, *U* and *N*. These *gross flows* are reported next to each of the arrows in Figure 6, and are the average annual flows estimated from LFS data back to 1975 for the UK.

A unifying theme in research on worker flows in the US is that the magnitude of gross worker flows dwarfs that of the net change in the respective labour market stocks. Figure 6 reveals that the same is true for the UK. While the annual *net* inflow into unemployment averaged 362 thousand since 1975, annual *gross* flows in and out of unemployment are much larger: On average over the last thirty-five years, 707 thousand individuals initially in work were measured as unemployed one year later, and 597 thousand unemployed workers were in a job one year later. Flows between unemployment and non-participation also dominate the net changes in their respective stocks. Notably, the numbers flowing between non-participation and employment are between 1.5 to 2 times as large as the flows between unemployment and employment.

Thus, while the variation in the aggregate unemployment rate illustrated in Figure 2 is substantial, it belies a teeming mass of individuals who are continually losing and finding jobs, and entering and exiting the labour force at all points in time.

⁸ In fact, the US Bureau of Labor Statistics (BLS) first published measures of worker flows starting in 1949. Publication was discontinued in 1953, however, due to concerns with the quality of the data (National Commission on Employment and Unemployment Statistics, 1979).

Although numbers of individuals flowing between labour market states are instructive, from the perspective of an individual worker, what really matters is the probability they face of losing a job, or finding a job, or entering non-participation, and so on. These probabilities are reflected in *flow transition rates*, and are the second main approach to the measurement of worker flows.

Of course, the numbers of workers flowing between labour force states and the associated transition rates are closely related concepts. To see this, consider the job loss flow, the flow from employment to unemployment, E to U . If we denote the number of employed workers in a given year t who are unemployed in the subsequent year by EU_t , the associated transition rate is simply equal to

$$\lambda_t^{EU} = EU_t/E_t. \quad (3)$$

More generally, the probability of a transition from an origin state A to a destination state B is given by the number of workers making that transition over a given period, divided by the stock of individuals in the origin state at the start of the period, $\lambda_t^{AB} = AB_t/A_t$.

Average annual flow transition rates between the three labour market states for the period 1975 to 2010 based on LFS data are reported in parentheses in Figure 6. Looking at transitions between employment and unemployment, it can be seen that, although the gross flows in each direction are comparable in size, the fact that the pool of the unemployed is less than one twelfth the size of the employment stock means that the probability of an unemployed worker finding a job over the course of a year (39.4%) far exceeds the probability of an employed worker losing theirs and becoming unemployed (3.0%).

In the following sections, we show how these worker flows and their corresponding flow transition rates determine changes in labour market stocks. In addition, we demonstrate how these worker flows can be used to help understand why unemployment rises in recessions, and why certain groups are hit harder than others.

IV. How Flows Shape Stocks

The previous sections have documented two important sources of variation in the labour market. On the one hand, labour market stocks such as the unemployment rate have varied substantially over time with expansions and contractions that buffet the aggregate economy. On the other hand, at any given point in time, many workers flow between the employment, unemployment and non-participation. In this section, we show how these two sources of variation are intertwined.

(i) The law of motion for unemployment

The link between variation in unemployment and its constituent flows is formalized by the *law of motion for unemployment*. This states that any change in the unemployment stock is due either to people joining the unemployment pool, or to people leaving the pool. In turn, these inflows and outflows can be traced to flows in and out of unemployment from employment and non-participation. Formally, the change in unemployment across two periods t and $t + 1$ is equal to the difference between inflows and outflows, and may be expressed as

$$\Delta U_{t+1} = EU_t + NU_t - UE_t - UN_t. \quad (4)$$

Unemployment rises when inflows exceed outflows. Inflows may originate from employment as workers lose their jobs (EU), or from non-participation as individuals begin to search for a job (NU). Likewise, outflows from unemployment occur when unemployed workers find new jobs (UE), or when individuals cease searching for a job (UN).

The law of motion for unemployment may also be expressed in terms of flow transition rates. Recall from equation (3) that the gross flow between an origin state A and destination state B is associated with the flow transition rate by the relation $AB = \lambda^{AB} \cdot A$. It follows that we can re-express equation (4) as

$$\Delta U_{t+1} = \lambda_t^{EU} E_t + \lambda_t^{NU} N_t - (\lambda_t^{UE} + \lambda_t^{UN}) U_t. \quad (5)$$

Before examining the empirical behaviour of these flow transition rates, and how they have contributed to unemployment changes in the UK, we first describe how these flows can be measured.⁹

(ii) Measuring worker flows

In the majority of research that estimates worker flows, the approach has been to use *longitudinal data*—data that includes repeated observations on the same individuals over time. Using such data, it is straightforward to compute gross flows and their associated transition rates. For example, the E to U transition rate is simply the fraction of those who report that they are employed in a given survey who subsequently report that they are unemployed in the next survey.

Many US studies have exploited the fact that the major source of data on labour force status in the US, the Current Population Survey (CPS), has a longitudinal element to it. Households in the CPS are surveyed for four consecutive months, rotated out of the survey for eight months, and then return for a final four months. This “rotating-panel” structure has allowed researchers in the US to compute worker flows back to 1967.

The UK analogue to the CPS in the US, the Labour Force Survey, has incorporated a rotating-panel element since 1992. Individuals who remain at the same address are surveyed for five consecutive quarters before rotating out of the survey. Estimates of worker flows in the UK based on these data have been studied by Gomes (2010), Petrongolo and Pissarides (2008) and Elsby and Smith (2010).

For the purposes of understanding the cyclical dynamics of the UK labour market, estimates of worker flows based on longitudinal data from the LFS are subject to another important drawback: They are available only from 1992 onwards. Consequently, such estimates cover just one full recession, limiting their ability to inform us on the propagation of recessions through the labour market.

For this reason, we explore an alternative and relatively under-studied set of measures of labour market flows that extend back to 1975. The LFS asks individuals about their labour force status a year prior to the interview date. This information on *recalled status* may be combined with the individual’s reported current status to infer measures of annual worker flows, and thereby the accompanying transition rates.

⁹ The estimates we report are derived from data in which an individual’s labour force status is observed at discrete points in time, specifically each year. In reality, however, individuals may make multiple transitions within a year that we do not observe in discrete-time data—there is a time aggregation problem. Our estimates of the number of individuals making any particular transition will tend to miss some transitions and wrongly add others. Authors such as Shimer (2007) and Elsby, Michaels and Solon (2009) have provided empirical methods for correcting estimates using data for the US. The latter show that, for monthly data, while correcting for time aggregation does influence the levels of the estimated flows, their cyclicity is affected only modestly. However, it is possible that the effects of annual time aggregation might be more severe.

To see how, consider transitions from employment to unemployment. The gross *EU* flow is simply the sum of respondents who report that their current status is unemployed, while their recalled status one year prior to the survey was employed. The associated transition rate λ^{EU} is just the gross *EU* flow divided by the number whose recalled status was employed.

These measures are not published officially, and so must be computed using the microdata that underlie the LFS. These microdata files are available for every other year from 1975 to 1983, and on an annual basis thereafter. The frequency of the estimates that we study in what follows mirrors the frequency of the available data.

As mentioned above, the information in the LFS on recalled status has the invaluable benefit of being asked of all individuals, not just those who remain at the same address, unlike the Quarterly LFS longitudinal data available from 1992. It is also straightforward to calculate flows, since the current and recalled status of a particular individual are simple to match. The use of recalled data does raise issues about the accuracy of remembered status, however. Studies investigating recall accuracy indicate that over short periods—up to about three years—recall bias is not severe (Paull (2002); Elias, (1996)). If individuals are asked to remember over longer periods, unemployment tends to be underreported; this does not appear to be simply short spells being forgotten, but is rather a general tendency to underreport. The one-year recall required of respondents in this paper falls well within the horizon where results should not be adversely affected by recall bias. However, Bell and Smith (2002) find recalled stocks accurate, and transitions between employment and unemployment correctly recalled, but where spells are short, transitions between unemployment and non-participation estimated from recalled data tally less well with contemporaneous reports. As Akerlof and Yellen (1985) suggest, this might be because individuals tend to remember better the most personally-important or salient events. Moves between the two non-employment states are unlikely to be as psychologically ‘painful’ or ‘enjoyable’ as losing or gaining a job.¹⁰

(iii) A brief history of worker flows in the United Kingdom

With our estimates of worker flows in hand, we can now begin to document the evolution of these flows in the UK over the last thirty-five years. Figure 7 plots the respective time series for all six transition rates in equation (5) above. Panels A and B depict flow transition rates describing the probabilities of joining the pool of the unemployed (by either losing a job or entering from non-participation) and leaving unemployment (through either finding a job or exiting out of the labour force). Flow transition rates between employment and non-participation are presented in panel C. In all panels, times of recession are indicated by shaded regions that correspond to sustained periods of falling GDP. For reasons that will become clear in section V, the series are plotted on logarithmic scales.¹¹ As we noted in section II, the unemployment rate is markedly countercyclical, rising in recessions and subsiding in booms. Figure 7 allows one to tell a heuristic story of how changes in flow transition rates correspond to the historical behaviour of the unemployment rate.

Panels A and B reveal that unemployment rises in recessions because rates of inflow into unemployment tend to rise in downturns, and rates of outflow tend to fall. Specifically, the *E* to *U* transition rate—the job loss probability—is strongly countercyclical, rising sharply in all recessions. These are accompanied by more modest rises in the inflow rate from non-participation into unemployment. Symmetrically, the *U* to *E* transition rate—the job

¹⁰ It is also worth noting that recalled status is subjective, and does not necessarily accord with ILO definitions.

¹¹ The values on the vertical axes are simply the fraction of workers in a particular state making the relevant transition. Plotting these on a logarithmic scale has the effect of stretching the distance between lower transition rates in such a way that an equal vertical distance on the scale represents a similar percentage change in transition rates.

finding probability—is clearly procyclical, falling systematically in every recession since 1975. Again, these also are accompanied, with an approximate one year lag, by modest reductions in the outflow rate from unemployment to non-participation.¹² So, casual observation of the flow transition rates underlying the aggregate unemployment rate in Figure 2 would suggest that *both* a rising rate of inflow and a declining rate of outflow are proximate causes for increased unemployment in times of recession.

An additional feature of the behaviour of the unemployment rate in the UK is that, even after the economy (GDP) starts to recover, the unemployment rate often continues to rise and remains persistently high for some time. How can this be related to the evolution of the flows in Figure 7?

Again, we see that both inflows and outflows play an important role in driving the persistence of UK unemployment. Workers continue to lose jobs at an elevated rate for some time after output begins to recover. A particularly prominent example of this is the recession of the early 1980s: Even eight years after the end of the downturn, the job loss rate had not returned to its pre-recession level.

The job-finding probability also displays persistence. After both the recessions of the early 1980s and early 1990s, the U to E transition rate continues to fall and remains stubbornly low for many years after the recession ends. Falling rates of job finding mean rising durations of unemployment. A simple way to think about this is to note that the overall exit rate from unemployment is just the sum of the U to E and U to N transition rates, $x = \lambda^{UE} + \lambda^{UN}$. It follows that the probability that an unemployment spell lasts T periods is simply $(1 - x)^{T-1}x$, the probability one fails to exit unemployment for $T - 1$ periods, times the probability of exiting in the T^{th} period. In this environment, then, unemployment duration is geometrically distributed,¹³ and so average duration is simply equal to $\sum_{T=0}^{\infty} T (1 - x)^{T-1}x = 1/x$. So, increased unemployment duration, and declining rates of outflow from unemployment are just two sides of the same coin.

An important literature starting in the 1980s pointed to a European unemployment problem of persistently high rates of long-term unemployment, from which the UK also suffered, (see Layard, Nickell and Jackman (1991)). These trends are lucidly surveyed by Machin and Manning (1999), who show the importance of rising unemployment duration in driving increased unemployment in Europe in the 1980s and 1990s.¹⁴

What of flows between employment and inactivity? Their evolution is depicted in Panel C of Figure 7. Flows out of employment to non-participation rise in recessions, as some individuals who lose or quit their jobs choose to leave the labour force. Thus, both sets of outflow rates from employment are countercyclical—those to unemployment more so than those to non-participation. The timing of changes in these inflows and outflows is such that, on average, both series tend to shift at the onset of a downturn rather than one leading the other systematically.

Cyclical movements in flows out of non-participation to employment likewise mirror those of U to E transitions. The smaller scale of flow transition rates from non-participation to employment reflects the fact that a smaller proportion of those out of the labour force is in

¹² This lag is probably due to the build-up of long-term unemployment in the stock of unemployed following a negative labour demand shock. A higher share of long-term unemployment is associated with higher subsequent rates of labour force withdrawal (OECD (2002)).

¹³ Empirically, unemployment duration is not geometrically distributed. We have assumed that the rate of exit from unemployment x is the same for all unemployed workers. In reality, of course, this is not the case. Nonetheless, the inverse relationship between exit rates and duration serves as a useful rule of thumb in practice.

¹⁴ As Figure 7B shows, however, the 2000s were notable for a marked rise in the U to E transition rate, which reached a thirty year high in 2007, prior to falling back somewhat the latest downturn, though not to the levels experienced in previous recessions.

a position to, or desires to, gain employment. However, in terms of absolute size of gross flows, N to E flows dominate those from U to E , for the simple reason that that stock of non-participants is so much larger than the stock of unemployed workers. The procyclicality of N to E flows indicates that job finding by non-participants is slowed by recessions in a similar manner to job finding by unemployed workers.¹⁵

V. Decompositions of Unemployment Variation

The previous section showed that increased unemployment in times of recession can be traced both to increased rates of inflow to, as well as reduced rates of outflow from, unemployment. In this section we show how one can be more formal about the relative roles of flow transition rates in shaping unemployment variation. Specifically, we pose the question of what fraction of the overall variance in the unemployment rate across time can be attributed to each of the flows.¹⁶

The recent literature has explored two avenues. The first, which we consider in section V(i), is referred to as the “two-state” approach in the literature. Its name derives from the fact that this approach does not explicitly take into account all three labour market states. It abstracts from flows between unemployment and non-participation, focusing instead on flows between the two states of employment and unemployment. We will see that this approach is a useful benchmark by which to get one’s bearings for the role of worker flows in unemployment dynamics.

Of course, this two-state abstraction does not provide the full picture of unemployment dynamics. In section V(ii), we take on the more complicated task of showing how the two-state model may be extended to an analysis of flows among employment, unemployment and non-participation—the “three-state” approach. As we will see, this approach turns out to be a very natural generalization of the two-state framework.¹⁷

(i) Two-state approach

The building block of the two-state model of unemployment dynamics is the following law of motion for the stock of unemployed workers,

$$\Delta U_{t+1} = s_t E_t - f_t U_t, \quad (6)$$

where s_t is the *inflow rate* into unemployment, and f_t the *outflow rate*. The literature sometimes refers to these flow rates respectively as separation and job-finding rates. We do

¹⁵ Using annual data, it is hard to determine whether changes in inflows lead changes in outflows (or vice versa) around unemployment turning points.

¹⁶ The decompositions we use provide a breakdown of flow steady-state unemployment, and not the realised unemployment rate. We shall see in Figure 8 that the steady-state unemployment rate is a leading indicator of the actual unemployment rate. Consequently, in order to understand the contributions of the flow transition rates to the evolution of actual unemployment, it is necessary to take into account these dynamic effects. Elsby, Hobijn and Şahin (2009), and Smith (2011) show how it is possible to take into account these dynamics in decomposing the variation in the realised unemployment rate.

¹⁷ It is important to note that the results of these decompositions are the outcome of an *accounting* exercise, and are not necessarily indicative of the degree to which these changes in flow rates *cause* changes in unemployment. For example, it is quite possible to construct stories for why changes in the flows might be interrelated: If workers who lose their jobs in a recession experience a loss of skill, a rise in job loss could retard the job-finding rate. If that were true, one could argue that the real “causal” contribution of the job-loss rate would be larger than the 65 percent figure suggested in Table 2. These possible inter-linkages between flows have been emphasised by Burgess and Turon (2005), who showed empirically that allowing for this endogeneity did indeed increase the role of the inflow rate.

not follow this practice for two reasons. First, as we know from sections III and IV, these flows are really a combination of flows between U and E and flows between U and N . Second, not all separations of workers from employers result in an inflow into the unemployment pool—some workers may line up a new job to start immediately after they separate from their old one.

A growing literature has sought to relate variation in the unemployment rate to variation in the flow transition rates s_t and f_t (see, among others, Elsby, Michaels and Solon (2009); Fujita and Ramey (2009); Shimer (2007)). To see how this might be done, a useful starting point is to define the unemployment rate that would prevail in the long run if the inflow and outflow rates in equation (6) did not change from their current level. This *steady-state unemployment rate* is found by setting $\Delta U_{t+1} = 0$ in equation (6) and solving to obtain

$$u_t^* \equiv U_t^*/L_t = s_t/(s_t + f_t). \quad (7)$$

In practice, of course, the flow transition rates s_t and f_t do move over time, as we have seen in Figure 7, and therefore so does the steady-state unemployment rate. Thus, the actual unemployment rate that we observe in the data u_t is in fact continually converging toward a moving target u_t^* .

Another way to see this is to note that one may rewrite equation (6) above as

$$\Delta U_{t+1} = -(s_t + f_t)(U_t - U_t^*). \quad (8)$$

This makes it clear that actual unemployment U rises whenever steady-state unemployment U^* lies above current unemployment, and vice versa. In this way, steady-state unemployment acts as a *leading indicator* of the future path of realized unemployment.

The prognostic nature of the steady-state unemployment rate can also be seen clearly in the data. Figure 8 graphs the steady state unemployment rate implied by the flow transition rates, together with the actual unemployment rate from 1975 to 2010.¹⁸ Over the cycle, movements in steady state and actual unemployment rates appear similar. However, it is clear that steady-state unemployment acts as a leading indicator for actual unemployment. At times when unemployment is rising—in recessions—the steady-state unemployment rate rises and peaks before the actual unemployment rate.¹⁹

The importance of the steady-state unemployment rate for us is that it provides a link from variation in the flow transition rates s_t and f_t to variation in the unemployment rate. This link can be used to inform a decomposition of unemployment variation into the relative contributions of the two flows in driving cyclical unemployment. In particular, Elsby, Michaels and Solon (2009) pointed out that simple log differentiation implies that a Taylor-series approximation to changes in the steady-state unemployment rate can be broken down as follows:

$$\Delta \ln u_t^* \approx \alpha_t [\Delta \ln s_t - \Delta \ln f_t], \text{ where } \alpha_t = (1 - u_{t-1}^*). \quad (9)$$

A useful implication of this is that, in order to ascertain the relative roles of the inflow and outflow rates in driving fluctuations in the *steady-state* unemployment rate, just compare

¹⁸ Figure 8 is based on a measure of s_t capturing the *overall* inflow rate to unemployment, and f_t the overall outflow rate. s_t and f_t are calculated on the basis of equation (13) below.

¹⁹ To see this formally, note that equation (8) can be rearranged to show that actual unemployment will lag further behind steady state unemployment, the faster unemployment is changing—and the fastest changes in unemployment tend to occur at the start of recessions (see Smith (2011)).

the *logarithmic* variation in the two flows. It is for this reason that the flow transition rates displayed in Figure 7 are presented on log scales.

Equation (9) is used by Elsby, Michaels and Solon (2009) to zoom in on each cyclical ramp-up in the US unemployment rate over time to trace out the cumulative log rise in the inflow rate, and the cumulative log decline in the outflow rate, since unemployment began rising.

Figure 9A summarizes the results of this approach for the UK. This confirms the informal story implied by Figure 7—that *both* the ins and the outs of unemployment play an important role in driving cyclical unemployment in the UK. Figure 9A also enables us to infer the quantitative contributions of the two margins. In each cyclical upswing in the unemployment rate, Figure 9 suggests something like a two-thirds inflows to one-third outflows split of the increase in unemployment. Thus, both flows matter, with the inflows being relatively more dominant. Figure 9 also suggests that this decomposition of unemployment variation has remained quite stable across recessions in the UK.

This basic impression is also substantiated by an alternative summary measure of the two contributions suggested by Fujita and Ramey (2009). They note that equation (9) above can be used to derive a useful *decomposition of variance* for the steady-state unemployment rate. Specifically, they note that the variance of log changes in the unemployment rate can be written as

$$var(\Delta \ln u_t^*) \approx cov(\alpha_t \Delta \ln s_t, \Delta \ln u_t^*) + cov(-\alpha_t \Delta \ln f_t, \Delta \ln u_t^*). \quad (10)$$

This is useful because the variance of changes in the steady-state unemployment rate is a single-statistic measure of fluctuations in unemployment over time. This decomposition of variance in turn implies very natural summary measures of the contributions of the two flows to changes in the steady-state unemployment rate, namely the ratio of their variance contribution to the total variance of the log change in steady-state unemployment,

$$\beta_s = \frac{cov(\alpha_t \Delta \ln s_t, \Delta \ln u_t^*)}{var(\Delta \ln u_t^*)}, \text{ and } \beta_f = \frac{cov(-\alpha_t \Delta \ln f_t, \Delta \ln u_t^*)}{var(\Delta \ln u_t^*)}. \quad (11)$$

Because the decomposition in equation (9) holds only approximately for discrete changes in steady-state unemployment, the contributions β_s and β_f will approximately sum to unity. In practice, however, we shall see that the approximation in fact works very well.

The first column of Table 2 summarizes the results of this decomposition of variance. These reiterate the message of Figure 9. While reductions in the outflow rate in times of recession account for around 30 percent of the rise in unemployment, elevated rates of inflow account for around 70 percent. This basic result – the dominance of inflows – holds for each of the three, very different, recessions observed over the sample period.²⁰ So, this suggests that, for the UK economy, both flows matter for shaping unemployment fluctuations, with the inflow rate being relatively more dominant.

(ii) Three-state approach

As we noted earlier, the second approach to decomposing changes in the unemployment rate takes seriously the possibility that flows in and out of non-participation can matter for fluctuations in joblessness. Indeed, we saw in Figure 7 that there did appear to be some cyclicity in the flow transition rates to and from non-participation, albeit much smaller than that seen in flows between employment and unemployment.

²⁰ Equally if we restrict the sample period to that used by Petrongolo and Pissarides (2008), similar relative contributions of inflows and outflows appear to hold, (67% and 33% respectively).

That the participation margin might have a bearing on cyclical unemployment fluctuations is a possibility only recently taken seriously in the literature on worker flows (see, in particular, Smith, forthcoming). Much of the previous literature has instead tended to ignore the potential role of non-participation flows. Often, this is justified with reference to the comparatively acyclical profile of the labour force participation rate, relative to the unemployment rate, which is evident in Figure 2.

This line of argument, however, is an important example of a stock-flow fallacy. While the *stock* of labour force participants might move little over the business cycle, small changes in the *transition rates* between unemployment and non-participation can nevertheless have a large impact on unemployment, for the simple reason that non-participation is so much larger than unemployment as a stock, as we saw in Figure 1.

Once non-participation is reintroduced, we revert to the full law of motion for unemployment, stating that changes in unemployment depend on inflows from and outflows to employment and non-participation. Each flow can be expressed in terms of the relevant transition rate multiplied by the relevant stock. The full law of motion, given in equation (5), is reproduced here for convenience:

$$\Delta U_{t+1} = \lambda_t^{EU} E_t + \lambda_t^{NU} N_t - (\lambda_t^{UE} + \lambda_t^{UN}) U_t. \quad (12)$$

Combining the law of motion for unemployment with similar laws of motion relating changes in employment and non-participation to their respective inflows and outflows, one can re-express the components of the steady-state unemployment rate in equation (7) above as follows (see Shimer (2007)):

$$u_t^* = \frac{s_t}{s_t + f_t}, \text{ where } s_t = \lambda_t^{EU} + \lambda_t^{EN} \frac{\lambda_t^{NU}}{\lambda_t^{NU} + \lambda_t^{NE}}, \text{ and } f_t = \lambda_t^{UE} + \lambda_t^{UN} \frac{\lambda_t^{NE}}{\lambda_t^{NU} + \lambda_t^{NE}}. \quad (13)$$

The overall inflow rate s_t is now split into two parts. The first term is straightforward: it is just the direct inflow rate from employment to unemployment—the job loss rate. The second term, $\lambda_t^{EN} \lambda_t^{NU} / (\lambda_t^{NU} + \lambda_t^{NE})$, can be interpreted as the indirect inflow rate from employment to unemployment via nonparticipation. It multiplies the flow transition rate between employment and nonparticipation, λ_t^{EN} , by the proportion of outflows from nonparticipation that transition to unemployment, $\lambda_t^{NU} / (\lambda_t^{NU} + \lambda_t^{NE})$, and so captures the probability that an individual transitions from E to N and also subsequently moves N to U . The overall outflow rate f_t is similarly split into two components: the direct rate of outflow from unemployment to employment (the job-finding rate λ_t^{UE}) and the indirect rate of outflow from unemployment to employment via non-participation (Petrongolo and Pissarides (2008); Smith (2011)).

In particular, note that the log change in the inflow rate can be approximated as follows

$$\Delta \ln s_t \approx \omega_t^s \Delta \ln \lambda_t^{EU} + (1 - \omega_t^s) \Delta \ln \lambda_t^{ENU}, \text{ where } \omega_t^s = \lambda_t^{EU} / s_t. \quad (14)$$

That is, the log change in the inflow rate is just a share-weighted sum of the log changes in the job-loss rate λ_t^{EU} and λ_t^{ENU} . Following a similar logic to decompose the contribution of the outflow rate, one can re-write the two-state decomposition above as

$$\Delta \ln u_t^* \approx \alpha_t [\omega_t^s \Delta \ln \lambda_t^{EU} + (1 - \omega_t^s) \Delta \ln \lambda_t^{ENU} - \omega_t^f \Delta \ln \lambda_t^{UE} - (1 - \omega_t^f) \Delta \ln \lambda_t^{UNE}], \quad (15)$$

where $\omega_t^f = \lambda_t^{UE} / f_t$.

Mirroring the two-state analysis of each recession in Panel A of Figure 9, Panel B breaks down the contributions of the inflow and outflow rates into parts associated with flows between employment and unemployment, and indirect flows via non-participation. Figure 9B reveals that variation in job-loss and job-finding rates has accounted for the vast majority of increases in unemployment in each recession since the early 1980s. In contrast, the participation margin has accounted only for a modest fraction of the variation, perhaps 10 percent of each upswing. Of the two, inflows into unemployment via nonparticipation appear to be the more important

Equation (15) above also allows one to compute analogous “beta” contributions in the three-state case, one for each of the four EU , ENU , UE , and UNE transitions. As before, these four betas will sum to unity only approximately, since both equations (14) and (15) hold only approximately. Again, however, we shall see that the above decomposition holds with a high degree of precision in practice.

The full three-state decomposition of steady-state unemployment changes, summarising the relative influence of the four flow transition rate components, is reported in Table 2. Again, the decomposition of variance reiterates the message of Figure 9B. In the UK, variation in the rate of job loss has been the dominant driver of unemployment dynamics, accounting for nearly 65 percent of overall variation in steady-state unemployment. Changes in the job-finding rate have also been influential: nearly 25 percent of unemployment variance can be attributed to these. In terms of non-participation flows, together these account for approximately 10 percent of the variance of unemployment, with inflow rates involving non-participation being more influential than non-participation outflow rates (7.5 percent versus 2 percent, respectively).

Taken together, then, our analysis of the flow-based origins of the cyclicity of unemployment in the UK has identified job loss as a leading determinant of the variation in joblessness since the mid-1970s. However, it is important to note that job loss does not account for all of the variation. In addition, we will see below in section VI that it can be especially important to recognize the roles of job-finding and flows via non-participation in accounting for the unemployment profiles of particular subgroups of the labour market.

(iii) Some caveats to the analysis

In order to keep the matters simple and transparent, the analysis above has required a certain degree of simplification. For the reader who wishes to delve further into these matters, in this subsection we discuss some of the more important issues that have been omitted up until now, and provide a guide to the relevant literature.

Time aggregation

Throughout the paper we have expressed the law of motion for unemployment in *discrete time*, whereby an individual’s labour force status is observed at discrete points in time, t and $t + 1$. In the LFS data we have been using throughout the paper, this period of time corresponds to a year.

In reality, however, some individuals may make multiple transitions within a year that we do not observe given the discrete-time nature of the available data. Consider a person we currently observe to be unemployed, who reports that they were employed a year ago. In the data, we simply see a single transition. But that person could have reached her present labour market state through an infinity of different paths. She could have made just one transition—from employment to unemployment—over the course of the year, in which case our annual observation is accurate. Or, she could have flowed back and forth between employment and unemployment several times. She also may have flowed through nonparticipation. Although

each of these paths will not be equally probable, the principle still holds that multiple transitions will be missed in annual data on worker flows.

It follows that the annual flow transition rates that we report may not be an accurate representation of the flow rates at a higher frequency, say monthly or weekly. There is a *time aggregation* problem: Our estimates of the number of individuals making any particular transition will tend to miss some transitions and wrongly add others.

Authors such as Shimer (2007) and Elsby, Michaels and Solon (2009) have provided empirical methods for correcting estimates using data for the US. The latter show that, for monthly data, while correcting for time aggregation does influence the *levels* of the estimated flows, their cyclicalities is affected only modestly. However, it is possible that the effects of annual time aggregation might be more severe.

Dynamic decompositions

A second caveat relates to the decompositions of unemployment variation that we have used repeatedly in the sections above. It is important to note that all of these decompositions have provided a breakdown of *steady-state* unemployment, and not the realized unemployment rate. While we saw in Figure 8 that the steady-state unemployment rate exhibits cyclical swings that are similar to those of the actual unemployment rate, they are not the same, and it seems reasonable to ask how one might decompose the variation in the realized unemployment rate.

As we have mentioned, the steady-state unemployment rate is in fact a leading indicator of the actual unemployment rate. That is, there is a *dynamic* relationship between the actual and steady-state unemployment whereby the former is continually evolving toward the latter. Consequently, in order to understand the contributions of the flow transition rates to the evolution of actual unemployment, it is necessary to take into account these dynamic effects—the notion that *past* changes in the flow transition rates can impact the observed unemployment rate in the present.

The literature has begun to think about these issues very recently. Elsby, Hobijn and Şahin (2009), and Smith (2011) show how it is possible to take into account these lead-lag dynamics in decomposing the variation in the realized unemployment rate. Intuitively, it is possible to show that the actual unemployment rate is a *distributed-lag* of current and past steady-state unemployment rates—a kind of weighted-average, with weights that decline the further into the past one goes. As a result, a chain of events links changes in worker flows to changes in actual unemployment: Changes in transition rates lead to changes in steady-state unemployment, which in turn shape the evolution of the realized unemployment rate in present and future periods. Elsby, Hobijn and Şahin (2009), and Smith (2011) show how this logic can be used to motivate a dynamic decomposition of the actual unemployment rate.

Accounting and causality

All of the analysis so far has summarized an *accounting* exercise of inferring the roles of flow transition rates as proximate determinants of unemployment variation. It is tempting to interpret the results of this analysis as indicative of the degree to which these changes in flow rates “cause” changes in unemployment. However, such an interpretation would be premature. It is quite possible to construct stories for why changes in the flows might be interrelated. Imagine, for example, that workers who lose their jobs in a recession experience a loss of skill, or human capital, and consequently find it harder to find new jobs.²¹ In such a

²¹ There is suggestive evidence for this hypothesis in the United States. Since the work of Jacobson, LaLonde and Sullivan (1993), a wealth of evidence has surfaced for the existence of substantial and persistent earnings losses that workers face following a job loss. Gregg, Knight and Wadsworth (2002) estimate a wage penalty in

world, a rise in the job loss rate could retard the job-finding rate. If that were true, one could argue that the real “causal” contribution of the job-loss rate would be larger than the 65 percent figure suggested in Table 2. These possible inter-linkages between flows have been emphasised by Burgess and Turon (2005), who showed empirically that allowing for this endogeneity did indeed increase the role of the inflow rate.

However, these accounting exercises are nonetheless an important and useful guide for the progression of research on the determination and evolution of unemployment. Prior to the recent recession, a number of prominent researchers in the US suggested that the contribution of unemployment inflows to cyclical ramp-ups in the unemployment rate was negligible (Hall (2005); Shimer (2007)). If this were the case, it would be tempting to conclude that unemployment inflows had neither a direct, nor an indirect causal impact on unemployment. Just as we have seen for the UK, based on the decompositions introduced above, the evidence does in fact point to the conclusion that both rates of inflow and outflow matter for unemployment dynamics in the US (Elsby, Michaels and Solon (2009); Fujita and Ramey (2009)).

The fact that both flows have a substantial accounting contribution implies that students of the labour market should apply their energies to understanding both margins, as well as their potential interactions, in developing their understanding of unemployment dynamics.

VI. Worker Flows and the Structure of Unemployment across Groups

Up to now, we have focused on changes in labour force stocks over time, and the role of worker flows as proximate determinants of these. As we saw in section II above, temporal—cyclical—changes are not the only source of variation in the labour market. In Figures 3, 4 and 5, we saw that there is substantial heterogeneity in unemployment across different subgroups in the labour market, specifically by gender, age and educational attainment. In this section, we look more closely at this heterogeneity. We review the flow-origins of the different experiences of unemployment across groups. We provide two sets of complementary estimates that summarise key features of worker flows for the different groups and the relationship between these flows and the groups’ varying unemployment experiences.

Our starting point is estimates of the differences in average flow transition rates across groups using recall-based LFS data from 1975 to 2010. Table 3 summarizes these results, and provides a rich picture of the average unemployment propensities and durations of the different groups.

These provide a sense for how differences in average flow transition rates across groups map into differences in their respective group-specific unemployment rates. In addition, however, we shall see that they provide an important perspective on the differences in the nature of the typical unemployment experience of different groups. For example, the participation margin will be seen to play a more significant role for some groups, while the job-finding margin is crucial for other groups.

The second summary measures we document are a direct extension of the decomposition of the time-series variation in the aggregate unemployment rate captured by equation (15) above. In particular, we apply analogous decompositions to the variation over time in group-specific unemployment rates noted in Figures 3, 4 and 5. Doing so provides a

the order of 10% for a sample of displaced UK workers. Machin and Manning (1999) summarise the earlier evidence on unemployment duration dependence in UK.

sense of the degree of heterogeneity in which flows matter for different groups. We will see that, mirroring the heterogeneity in average flow transition rates, there are also considerable differences in the origins of variation in unemployment over time within groups. Table 4 summarizes these results.

(i) Gender

We noted in Figure 3 that the unemployment rate for men has tended to be higher than that faced by women in the last thirty-five years. What might account for this? Inspection of Table 3 indicates that, although men are more likely to lose their job on average, this effect is offset in large part by the fact that men are less likely to enter unemployment via non-participation than women. Taken together, the difference in their overall rates of inflow into the unemployment pool is in fact rather small.

Instead, Table 3 reveals that it is a reduced rate of outflow from unemployment relative to women that appears to account for much of the higher unemployment rate faced by men. In turn, the majority of this difference in rates of outflow can be traced to the fact that women are much more likely to exit unemployment via non-participation than men.

A common theme in the flow transition rates of women is the relative importance of flow transitions via non-participation. This feature of female unemployment dynamics accords well with the observation that women are more likely to move in and out of the labour force with the demands of childcare responsibilities, an activity more than proportionately allocated to women. It suggests that an understanding of the female labour market requires an understanding of the participation decisions women face, including issues of possible gender differences in contributory-based benefit eligibility.

This basic conclusion is reaffirmed in the analysis of the variation over time of gender-specific unemployment rates in Table 4. There we see that the overall inflow/outflow decomposition of unemployment variation is about the same for men and women, with around 70 percent of unemployment variation accounted for by inflows, and the remaining 30 percent contributed by outflows. However, the *composition* of these effects is very different across men and women. The role of flows via participation in driving fluctuations in male unemployment over time is negligible, accounting for around 6 percent of its variance. Thus, job-loss and job-finding flows between unemployment and employment are the crucial determinants of male joblessness.

In contrast, for women the participation margin contributes over 20 percent of the fluctuations in the female unemployment rate. Thus, in addition to being an important determinant of differences in the average levels of unemployment between men and women, changes over time in flows via non-participation are also a distinguishing characteristic of female unemployment dynamics.

(ii) Age

Figure 4 documented substantial differences in unemployment rates across workers of different ages. Most notably, the unemployment rate faced by young workers is substantially higher than older age groups'. In recent years, unemployment rates among those aged 16 to 24 have risen to more than double those faced by all other age groups.

Table 3 provides a unique perspective on the unemployment experiences of young workers. It reveals substantial heterogeneity in rates of inflow to and outflow from the unemployment pool across age groups. Younger workers are much more likely to flow into the unemployment pool, with inflow rates approximately double those faced by workers aged 25 to 49. Table 3 reveals that higher rates of youth unemployment are more than entirely explained by this phenomenon, and in particular by markedly higher rates of job loss among younger workers. In direct contrast, we see that differences in outflow rates work *against*

higher youth unemployment: Younger workers aged 16 to 24 exit unemployment *faster* than their older counterparts. Thus, youth face a much more fluid labour market than older workers, flowing in and out of the unemployment pool frequently.

The message of Table 3 is therefore much more nuanced than the impression presented by inspecting the overall levels of the unemployment rates by age. Casual observation of Figure 4 might suggest that the brunt of unemployment is borne by younger workers. Table 3 reveals, however, that the unemployment spells faced by youth tend to be substantially shorter than those experienced among older workers. Younger workers tend to “churn” through the labour market more frequently, a message that is missed in a simple reading of the overall unemployment rates.²²

Turning now to an analysis of what drives cyclical movements in unemployment rates by age, Table 4 again reveals quite clear distinctions between younger and older workers. The main stylized fact that emerges is that the role of the participation margin in driving unemployment fluctuations appears to be U-shaped in age. That is, flows via non-participation play an important role in the unemployment dynamics of workers aged 16 to 24 and 50 plus, while being relatively unimportant for workers aged 25 to 49, presumably reflecting the greater salience of the timing of labour force entry for younger workers and retirement for older workers.

(iii) Education

Unemployment varies considerably by education attainment. Figure 5 showed that less-well-educated workers are much more likely to be in want of work. Unemployment rates of individuals who left school before age 18 hover at around double those faced by university-educated workers.

The message portrayed by labour market stocks, then, is that less-skilled workers are hit harder in the labour market. However, we saw in our analysis of unemployment flows by age group that it can be the case that inspection of labour market stocks can miss an important part of the picture of unemployment experiences of different groups.

So, do less-educated workers really bear the brunt of unemployment? The measures of average flows in Table 3 suggest that they do, from the perspective of both the inflow and outflow margins. Workers who left school prior to age 18 not only face significantly higher rates of entry into unemployment, they also experience substantially longer jobless spells relative to their more-educated counterparts. Thus, higher rates of unemployment among the low-skilled appear to be a consequence of both increased incidence and increased duration of unemployment spells.

While the unemployment experiences of the less-educated are conspicuously more severe at all points in time, their dynamics over time are not much different from the aggregate picture presented in Table 2. Just as is the case for the overall dynamics of unemployment, there is something like a 70:30 inflow/outflow split of unemployment over time among the less-educated.

VII. Conclusion

The UK unemployment rate has recently stabilised following the end of the third recession experienced over the last thirty years. Economists have long realised that a better

²² It is worth noting that increased participation in tertiary education by young adults observed in the UK over the last twenty years will have reduced the size of the youth labour force and so raised the unemployment *rate* for a given unemployed stock. The unemployment-population ratio shows a much lower rate of increase over the most recent downturn.

understanding of what drives changes in the unemployment rate—and hence an appropriate policy response—can be gleaned from an examination of the numbers of workers moving into unemployment relative to the numbers moving out of unemployment. In a downturn many individuals lose their jobs and others fail to find work immediately after job loss. Yet, equally, some people are able to find work even in the depths of a recession.

This analysis has, for the first time in the UK, used individual micro data to show that recessionary ramp-ups during the last thirty-five years can be accounted for by rises in the unemployment inflow rate and falls in the outflow rate—with changes in inflows over the cycle accounting for around 70% of unemployment variation, and outflows for the remaining 30%. This result holds, broadly, in each of the last 3 recessions, in good times as well as bad. It also appears true when the data are disaggregated by gender, age and education. The analysis is not designed to not reveal why the flow rates were different this time round compared to earlier recessions, but does have important implications for the unemployment policy debate. Prior emphasis on unemployment duration and job finding led to a flurry of active labour market policy prescriptions. While our results reaffirm the potential importance of these policies (including the possibility that the reforms of the late 1990s may have mitigated the rise in unemployment in the recent recession), they also highlight the need to be aware of the importance of job loss in shaping unemployment. The appropriate policy responses to job loss rest on the nature of these job losses: Are they the outcome of a mutual agreement on behalf of firms and workers to go their separate ways? Or, do they represent the loss of otherwise profitable relationships that are severed inefficiently? While our analysis does not answer this question, it revives as an important point of discussion for future policy debates and academic research.

Several other novel aspects of the analysis in this paper are worth highlighting. First, a three-state decomposition of log unemployment variation was developed and used—for the first time. Second, we have applied decomposition methods to disaggregated UK population sub-groups, which has allowed us to investigate heterogeneity in unemployment levels and cyclicity. Thirdly, the data we have drawn on also distinguish this paper from previous UK research. Their micro, individual-level, nature has allowed us to focus on flows between the three states, which we have shown to be important particularly when analysing differences across genders and age groups. Furthermore, the unusually large time dimension of these micro data has meant we can look back over three complete business cycles.

The failure of the unemployment rate to rise as far in the latest downturn as many people feared, and indeed relative to past downturns, appears to be in part because the outflow rate from unemployment stayed comparatively high this time round, and in part because of a lower than expected rise in the job loss rate. An explanation of why these differences arose goes beyond the scope of this paper, but is an area ripe for further study. It raises the tantalising prospect for policymakers that the relatively small decline in the job finding rate in the recent recession might have been due to successful active labour market policies—the various New Deal policies targeted at different vulnerable labour market subgroups—and perhaps also the revamped unemployment benefit regime—the Job Seeker’s Allowance, with its associated sticks and carrots, now being stress-tested for the first time in the recent recession. Similarly, the rapid fall-back in the job loss rate might be a response to the expansionary government policy and depreciation pursued over the recession as well as a greater prevalence of real wage cuts among workforces under threat. The importance of changes in the job loss rate imply that demand-management policies designed to counteract the impact of adverse shocks on employment can be effective in alleviating rises in unemployment.

References

- Akerlof, G. A. and Yellen, J. L. (1985), 'Unemployment through the Filter of Memory', Quarterly Journal of Economics, **100**, 747-773.
- Bell, B. and Smith, J. (2002), 'On Gross Worker Flows in the United Kingdom: Evidence from the Labour Force Survey', Bank of England Working Paper.
- Blanchard, O. J. and Diamond, P. (1990), 'The Cyclical Behaviour of the Gross Flows of US Workers', Brookings Papers on Economic Activity, 2, 85-143.
- Burgess, S. and Turon, H. (2005), 'Unemployment Dynamics in Britain', Economic Journal, **115**, 423-448.
- Elias, P. (1996), 'Who Forgot They Were Unemployed?', University of Warwick Institute for Employment Research Working Paper.
- Elsby, M. W. L., Michaels, R. and Solon, G. (2009), 'The Ins and Outs of Cyclical Unemployment', American Economic Journal: Macroeconomics, **1**, 84-110.
- Elsby, M. W. L., Hobijn, B. and Şahin, A. (2009), 'Unemployment Dynamics in the OECD', NBER Working Paper 14617.
- Elsby, M. W. L., Hobijn, B. and Şahin, A. (forthcoming), 'The Labor Market in the Great Recession', Brookings Papers on Economic Activity.
- Elsby, M. W. L. and Smith, J. C. (2010), 'The Great Recession in the UK Labour Market: A Transatlantic Perspective', National Institute Economic Review, **214**, R26-R37.
- Fujita, S. and Ramey, G. (2009), 'The Cyclicity of Separation and Job Finding Rates', International Economic Review, **50**, 415-430.
- Gomes, P. (2010), 'Labour Market Flows: Facts from the United Kingdom', mimeo, London School of Economics.
- Gregg, P., Knight, G. and Wadsworth, J. (2002), 'They Get Knocked Down. Do They Get Up Again? Displaced Workers in Britain', in P. Kuhn (ed.), *Losing Work, Moving On: International Perspectives on Worker Displacement*, Kalamazoo Press.
- Gregg, P. and Wadsworth, J. (1995), 'A Short History of Labour Turnover, Job Tenure and Job Security', *Oxford Review of Economic Policy*, **11**, 73-90.
- Gregg, P. and Wadsworth, J. (2010), 'The UK Labour Market and the 2008-9 Recession', *National Institute Economic Review*, **212**, R61-R72.
- Hall, R. E. (2005), 'Job Loss, Job Finding, and Unemployment in the US Economy over the Past Fifty Years', in M. Gertler and K. Rogoff (eds), *NBER Macroeconomics Annual 2005*, Cambridge, MA, MIT Press.
- Jacobson, L. S., LaLonde, R. J. and Sullivan, D. G. (1993), 'Earnings Losses of Displaced Workers', American Economic Review, **83**, 685-709.

- Layard, R. Nickell, S. and Jackman, R., (1991), *Unemployment: Macroeconomic Performance and the Labour Market*, Oxford University Press.
- Kaitz, H. B. (1970), 'Analyzing the Length of Spells of Unemployment', Monthly Labor Review, **93**, 11-20.
- Machin, S. and Manning, A. (1999), 'The Causes and Consequences of Longterm Unemployment in Europe', in O. C. Ashenfelter and D. Card (eds), *Handbook of Labor Economics*, **3C**, Amsterdam, Elsevier.
- Marston, S. T. (1976), 'Employment Instability and High Unemployment Rates', Brookings Papers on Economic Activity, 1, 169–203.
- National Commission on Employment and Unemployment Statistics. (1979), *Counting the Labor Force*, Washington, DC: U.S. Government Printing Office.
- Okun, A. M. (1962), 'Potential GNP: its Measurement and Significance', American Statistical Association, *Proceedings of the Business and Economics Section*, 98-103. Reprinted as Cowles Foundation Paper 190
<http://cowles.econ.yale.edu/P/cp/p01b/p0190.pdf>
- Paull, G. (2002), 'Biases in the Reporting of Labour Market Dynamics', Institute for Fiscal Studies Working Paper WP02/10.
- Perry, G. L. (1972), 'Unemployment Flows in the US Labor Market', Brookings Papers on Economic Activity, 2, 245–78.
- Petrongolo, B. and Pissarides, C. A. (2008), 'The Ins and Outs of European Unemployment', American Economic Review Papers and Proceedings, **98**, 256-262.
- Pissarides C. A., (1986), 'Unemployment and Vacancies in Britain', Economic Policy, No. 3, October, pp. 499-540
- OECD (2002), 'The Ins and Outs of Long-Term Unemployment', Employment Outlook, Chapter 4, OECD, Paris.
- Shimer, R. (2007), 'Reassessing the Ins and Outs of Unemployment', NBER Working Paper 13421.
- Smith, J. C. (2011), 'The Ins and Outs of UK Unemployment', Economic Journal, Vol. 121, May, pp. 402-444.

Table 1: Unemployment rate in recessions

Recession start	1973	1979	1990	2008
Unemployment rate at start of recession	3.4%	5.3%	7.0%	5.3%
Total rise in unemployment rate	2.3 percentage points (pp)	6.7pp	3.7pp	2.8pp
Period of unemployment rate rise (duration)	1974q1–1977q3 (15 quarters)	1979q3–1984q1 (19 quarters)	1990q3–1993q1 (11 quarters)	2008q2–2010q1 (8 quarters)
Total fall in GDP	-3.3%	-6.1%	-2.6%	-6.6%
Period of GDP fall (duration)	1973q3–1975q3 (9 quarters)	1979q3–1981q3 (7 quarters)	1990q3–1991q3 (5 quarters)	2008q2–2009q3 (6 quarters)

Source: ONS.

Table 2: Role of inflows and outflows in unemployment dynamics, 1975-2010

% Contribution to unemployment variance of changes in:			
1975-2010			
		<i>Of which</i>	
U Inflow rate	71	<i>Job loss rate (E to U)</i>	64
		<i>Inflow rate via nonparticipation</i>	7
U Outflow rate	30	<i>Job finding rate (U to E)</i>	28
		<i>Outflow rate via nonparticipation</i>	2
1979-1984			
Inflow rate	62	<i>Job loss rate (E to U)</i>	55
		<i>Inflow rate via nonparticipation</i>	8
Outflow rate	39	<i>Job finding rate (U to E)</i>	29
		<i>Outflow rate via nonparticipation</i>	10
1990-1993			
Inflow rate	85	<i>Job loss rate (E to U)</i>	80
		<i>Inflow rate via nonparticipation</i>	5
Outflow rate	16	<i>Job finding rate (U to E)</i>	17
		<i>Outflow rate via nonparticipation</i>	-1
2008-2010			
Inflow rate	80	<i>Job loss rate (E to U)</i>	76
		<i>Inflow rate via nonparticipation</i>	4
Outflow rate	21	<i>Job finding rate (U to E)</i>	26
		<i>Outflow rate via nonparticipation</i>	-5
1990-2007			
Inflow rate	74	<i>Job loss rate (E to U)</i>	66
		<i>Inflow rate via nonparticipation</i>	8
Outflow rate	26	<i>Job finding rate (U to E)</i>	27
		<i>Outflow rate via nonparticipation</i>	-1

Source: Authors' calculations using Labour Force Survey microdata (using recalled labour force status one year ago). Biennial data between 1975 and 1983 are linearly interpolated so the decomposition relates to annual changes in unemployment throughout the sample.

Notes: The table shows the proportion of the variance of steady state unemployment accounted for by changes in the relevant transition rate, using the log decomposition described in the text. Components might not sum to 100% due to approximation error.

Table 3: Average unemployment and flow transition rates by gender, age and education

	Unemp- loyment rate	Inflow rate	<i>Of which:</i>		Outflow rate	<i>Of which:</i>	
			<i>Job loss rate</i>	<i>via non- partici- pation</i>		<i>Job finding rate</i>	<i>via non- partici- pation</i>
(a) Gender:							
Men	8.0	4.0	3.5	0.5	46.6	37.5	9.1
Women	6.9	3.8	2.6	1.3	58.9	43.2	15.8
(b) Age:							
16-24	13.6	7.5	6.2	1.2	55.1	44.7	10.4
25-34	7.6	4.2	3.3	0.9	51.8	41.2	10.5
35-49	5.3	3.7	2.2	1.6	40.5	38.6	1.8
50 plus	5.9	3.8	2.3	1.4	43.4	26.8	16.6
(c) Age left full-time education:							
< 16	8.2	4.2	3.0	1.2	43.9	31.9	12.1
16-17	8.4	4.3	3.5	0.8	51.1	40.8	10.3
18-20	5.6	3.0	2.4	0.6	61.0	50.1	11.0
> 20	3.9	2.1	1.7	0.4	66.4	56.8	9.6

Source: Authors' calculations using Labour Force Survey microdata 1975-2010 (using recalled labour force status one year ago). Biennial data between 1975 and 1983 are linearly interpolated.


Table 4: Steady state unemployment variance decomposition by gender, age and education (%)

	Contribution			Contribution		
	of	Of which:		of	Of which:	
	inflow	<i>Job</i>	<i>via non-</i>	outflow	<i>Job</i>	<i>via non-</i>
	rate	<i>loss</i>	<i>partici-</i>	rate	<i>finding</i>	<i>partici-</i>
		<i>rate</i>	<i>pation</i>		<i>rate</i>	<i>pation</i>
(a) Gender:						
Men	71	66	5	30	30	1
Women	69	55	14	31	24	8
(b) Age:						
16-24	66	57	8	36	33	3
25-34	76	71	6	25	25	1
35-49	60	58	5	41	41	0
50 plus	69	61	11	31	26	4
(c) Age left full-time education:						
< 16	70	64	7	32	33	-2
16-17	68	64	5	33	32	1
18-20	84	74	5	16	10	7
> 20	74	67	8	26	25	2

Source: Authors' calculations using Labour Force Survey microdata 1975-2010 (using recalled labour force status one year ago). Biennial data between 1975 and 1983 are linearly interpolated so the decomposition relates to annual changes in unemployment throughout the sample.


Notes: The table shows the proportion of the variance of steady state unemployment accounted for by changes in the relevant transition rate, using the log decomposition described in the text. Components might not sum to 100% due to approximation error.

Figure 1: Labour force stocks



Source: Labour Force Survey microdata (every other year between 1975 and 1983, and every year thereafter).

Notes: Numbers are average stocks (in millions) during 1975 to 2010. All numbers relate to the working age population (men aged 16-64 and women aged 16-59). Data are not weighted to account for changing frequency of observation.

Figure 2: Key labour force ratios

Source: ONS.

Figure 3: Unemployment by gender




Figure 4: Unemployment by age






Figure 5: Unemployment by age left full-time education




Sources: ONS (Figure 3). Labour Force Survey microdata (every other year between 1975 and 1983, and every year thereafter) (Figures 4 and 5).

Figure 6: Gross flows and flow transition rates


Source: Labour Force Survey microdata (every other year between 1975 and 1983, and annually thereafter).

Notes: *E*, *U* and *N* represent employment, unemployment and nonparticipation, respectively, for the working age population (men 16–64 and women 16–59). Numbers in boxes are average stocks and average annual net inflows (thousands) between 1975 and 2010. Numbers next to arrows are the relevant average annual gross flows (thousands). Annual flow transition rates are in parentheses. Flows and flow transition rates are based on recalled status one year ago.

Figure 7: Aggregate flow transition rates




Source: Labour Force Survey microdata (every other year between 1975 and 1983, and every year thereafter).

Figure 8: Actual and steady state unemployment rates

Source: Labour Force Survey microdata (every other year between 1975 and 1983, and every year thereafter).

Figure 9: Changes in log inflow and outflow rates by recession



Source: Labour Force Survey microdata (every other year between 1975 and 1983, and every year thereafter).

CENTRE FOR ECONOMIC PERFORMANCE
Recent Discussion Papers

1057	Fabrice Defever	Incomplete Contracts and the Impact of Globalization on Consumer Welfare
1056	Fadi Hassan	The Penn-Belassa-Samuelson Effect in Developing Countries: Price and Income Revisited
1055	Albert Marcet Ramon Marimon	Recursive Contracts
1054	Olivier Cadot Leonardo Iacovone Denisse Pierola Ferdinand Rauch	Success and Failure of African Exporters
1053	Björn Eriksson Tobias Karlsson Tim Leunig Maria Stanfors	Gender, Productivity and the Nature of Work and Pay: Evidence from the Late Nineteenth-Century Tobacco Industry
1052	Hartmut Lehmann Jonathan Wadsworth	The Impact of Chernobyl on Health and Labour Market Performance
1051	Jörn-Steffen Pischke	Money and Happiness: Evidence from the Industry Wage Structure
1050	Tim Leunig Joachim Voth	Spinning Welfare: the Gains from Process Innovation in Cotton and Car Production
1049	Francesca Cornaglia Andrew Leigh	Crime and Mental Wellbeing
1048	Gianluca Benigno Hande Küçük-Tüger	Portfolio Allocation and International Risk Sharing
1047	Guy Mayraz	Priors and Desires: A Model of Payoff-Dependent Beliefs
1046	Petri Böckerman Alex Bryson Pekka Ilmakunnas	Does High Involvement Management Lead to Higher Pay?
1045	Christos Genakos Tommaso Valletti	Seesaw in the Air: Interconnection Regulation and the Structure of Mobile Tariffs
1044	Giordano Mion Luca David Opmolla	Managers' Mobility, Trade Status and Wages
1043	Javier Ortega Gregory Verdugo	Immigration and the Occupational Choice of Natives: A Factor Proportions Approach

- | | | |
|------|---|---|
| 1042 | Nicholas Bloom
Benn Eifert
Aprajit Mahajan
David McKenzie
John Roberts | Does Management Matter? Evidence from India |
| 1041 | Joshua D. Angrist
Stacey H. Chen
Brigham R. Frandsen | Did Vietnam Veterans Get Sicker in the 1990s? The Complicated Effects of Military Service on Self-Reported Health |
| 1040 | Tanvi Desai
Felix Ritchie | Effective Researcher Management |
| 1039 | Ralf Martin
Mirabelle Muûls
Laure B. de Preux
Ulrich J. Wagner | Anatomy of a Paradox: Management Practices, Organisational Structure and Energy Efficiency |
| 1038 | Giordano Mion
Linke Zhu | Import Competition from and Outsourcing to China: A Curse or Blessing for Firms? |
| 1037 | William Brown
David Marsden | Individualisation and Growing Diversity of Employment Relationships |
| 1036 | John Van Reenen | Does Competition Raise Productivity through Improving Management Quality? |
| 1035 | Florence Kondylis
Marco Manacorda | School Proximity and Child Labor Evidence from Rural Tanzania |
| 1034 | Lars Boerner
Albrecht Ritschl | Communal Responsibility and the Coexistence of Money and Credit under Anonymous Matching |
| 1033 | Gianluca Benigno
Pierpaolo Benigno
Salvatore Nisticó | Second-Order Approximation of Dynamic Models with Time-Varying Risk |
| 1032 | Gianluca Benigno
Chuigang Chen
Christopher Otrok
Alessandro Rebucci
Eric R. Young | Financial Crises and Macro-Prudential Policies |
| 1031 | Dennis Novy | International Trade without CES: Estimating Translog Gravity |
| 1030 | Alex Bryson
John Forth | The Evolution of the Modern Worker: Attitudes to Work |