Do Unemployment Benefits Increase Unemployment? New Evidence on an Old Question[#]

By

Peter Fredriksson[*] and Martin Söderström[**]

25 September 2006

Abstract

We examine the relationship between unemployment benefits and unemployment using Swedish regional data. To estimate the effect of an increase in unemployment insurance (UI) on unemployment we exploit the fact the generosity of UI varies regionally because there is a ceiling on benefits. The *actual generosity of UI* varies within region over time due to, e.g., differences in expected regional wage growth and variations in the benefit ceiling. We find fairly robust evidence suggesting that the actual generosity of UI does matter for regional unemployment. Increases in the actual replacement rate contribute to higher unemployment as suggested by theory. We also show that removing the wage cap in UI benefit receipt would reduce the dispersion of regional unemployment. This result is due to the fact that low unemployment regions tend to be high wage regions where the benefit ceiling has a greater bite. Removing the benefit ceiling thus implies that the actual generosity of UI increases more in low unemployment regions.

[#] We thank Bertil Holmlund, Oskar Nordström Skans, Peter Skogman Thoursie, as well as seminar participants at the University of Stockholm and Uppsala university for very helpful comments. Financial support from the Swedish council for Working Life and Social Research is gratefully acknowledged.

[*] Institute for Labour Market Policy Evaluation (IFAU) and Department of Economics, Uppsala University. Address for correspondence: IFAU, P.O. Box 513, SE-751 20 Uppsala, Sweden. Email: peter.fredriksson@ifau.uu.se.

[**] Department of Economics, Uppsala University. Address for correspondence: Department of Economics, Uppsala University, P.O. Box 513, SE-751 20 Uppsala, Sweden. Email: martin.soderstrom@nek.uu.se.

1. Introduction

Whether the provision and generosity of unemployment insurance (UI) increase unemployment has been the subject of much research.[1] Theory generally predicts that unemployment will rise in response to an increase in UI generosity. However, the empirical evidence is not as unequivocal as the theory suggests.

There are a number of studies using micro data to identify the effects of UI generosity for those already unemployed; Meyer (1995) surveys the most convincing experimental evidence. But the provision of UI affects other margins as well. In addition to affecting search behavior, UI may affect, e.g., wage-setting and quitting behavior. In other words, we are most interested in the general equilibrium effects of variations in the generosity of UI.

Aggregate time series data have the potential of capturing general equilibrium effects of benefit generosity. However, the use of aggregate data creates severe identification problems. This may be part of the explanation for the fact that the estimated effects are much smaller than one would think based on theory. Now, what "one would think based on theory" is usually based on models where UI is equivalent to the "wage" during unemployment. Most empirical specifications are also derived from this simple model. Of course, real-world UI systems are much more complex and modeling their institutional features may yield different conclusions – a point forcefully made by Atkinson & Micklewright (1991).

The use of data over countries or regions, observed at different points in time, is presumably a more promising way to estimate the equilibrium effects of variations in UI benefit generosity. The prototypical US study in this vein (e.g. Katz & Meyer, 1990) uses policy changes at the state level to identify the effects. However, this approach can be criticized because policy changes at the state level are endogenous with respect to the local cycle; see Card & Levine (2000), and Lalive & Zweimüller (2004).

We also use regional panel data. However, the approach to identification is different and, to our knowledge, novel. The source of variation comes from a nationally determined policy. We exploit the fact that in most real-world UI systems there is a ceiling on the amount of

[1] See Holmlund (1998), Krueger & Meyer (2002), and Fredriksson & Holmlund (2005) for recent reviews of the literature.

benefits received.[2] This ceiling comes from the fact that there is a cap on income which is used to calculate the actual benefit received; increases in income above the cap produce no increase in the actual benefit.[3] Coupled with the fact that there are well-known regional wage differentials within countries, this implies that the *actual generosity* of UI varies regionally. More importantly, it will vary within region over time because changes in the ceiling produce regional variations in generosity depending on whether the region is above and/or below the ceiling before and after the policy change; moreover, differences in regional wage growth yield regional variation in actual generosity for a given national ceiling.

The fact that the level and changes in the regional wage may produce changes in the actual generosity of UI is, as such, not that useful. Regional wages and wage growth are endogenous with respect to regional unemployment. The challenge is therefore to find a strategy for constructing measures of predicted wages which are plausibly exogenous to local unemployment. Given an exogenous predicted wage, variations in the ceiling will produce differential changes in the actual generosity of UI depending on whether the region is predicted to be above or below the wage cap.

This empirical strategy is implemented using Swedish data during 1974-2002. To generate predicted wages we exploit individual data. For each individual and time point we estimate what the wage would be if his or her characteristics were priced on the national labor market. We then calculate the UI benefit and the actual replacement rate (given the estimated wage) should this individual become unemployed. Finally, the measures of UI generosity are aggregated to the regional level and related to regional unemployment. Notice that the non-linearity of the benefit schedule – induced by the benefit ceiling – implies that the unemployment effect of changes in the actual generosity of UI is identified even if we hold predicted wages and other labor force characteristics constant.[4]

Whether unemployment responds to changes in UI benefit generosity is one of the classic questions in labor economics that dates back to, e.g., Pigou (1932). The policy relevance of

[2] In the US, the maximum benefit amount even varies by state (Krueger & Meyer, 2002).

[3] Carling *et al* (2001) use a similar approach, albeit applied to micro data, when examining whether unemployment duration is affected by variations in UI generosity. They use the fact that because of the benefit ceiling some benefit recipients are treated with a policy change while others are not.

[4] Later on we will illustrate that the aggregate movements in the benefit ceiling are more or less idiosyncratic.

this question should thus be clear. But there is an additional reason to re-examine the issue: the design of the national unemployment insurance system has implications for the regional unemployment distribution because of the ceiling in benefit receipt.

It is an empirical fact that regional unemployment differentials are very stable in Europe. Figure 1 illustrates this for regional labor markets in Sweden. It is clear that regions which were high unemployment regions in the mid 1970s are also high unemployment regions in the beginning of the 2000s, and vice versa; the regression line has a slope of 0.91 with at t-value of 5.3. The benefit ceiling implies that UI is more generous in high unemployment/low wage regions, a fact that may further increase the spread of regional unemployment differentials. Therefore, it is interesting to examine whether (and how much) a more “neutral” design of the UI system – one that has no benefit ceiling – would reduce the dispersion of unemployment across regions.

Figure 1: Regional unemployment persistence

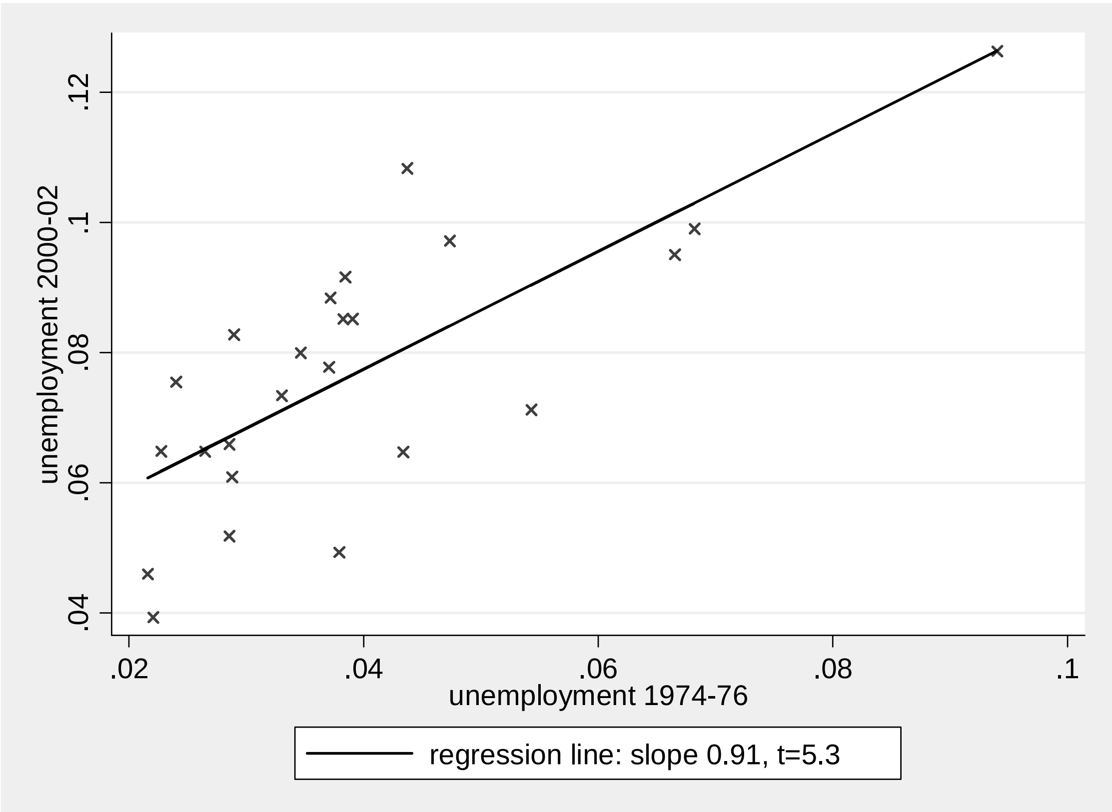

Notes:“Unemployment” is defined as the sum of the openly unemployed and participants in active labor market programs as a share of the labor force.
Sources: Labor Force Surveys and National Labor Market Board.

The remainder of the paper is outlined as follows. Section 2 presents a simple model of regional unemployment that we use for specification and interpretation purposes. Section 3 describes the Swedish institutional setting. Section 4 presents the data and our empirical strategy. We use individual data to calculate measures of the composition of the regional labor force. We also use the individual data to estimate earnings regressions which are used to generate individual expected wages and measures of UI generosity at the individual level. These measures are then used to generate a measure of the actual generosity of UI which is independent of the regional state of the labor market. Section 5 illustrates the identification strategy further. In particular we ask the question: What variation identifies the actual replacement rate? Section 6 presents the estimation results. In Section 7 we conduct two policy experiments to simulate the effects of UI policies on aggregate unemployment and the distribution of unemployment across regions. First we remove the benefit ceiling while holding the nominal replacement rate fixed. Then we raise the nominal replacement rate with the wage cap still in place. Section 8 concludes.

2. A simple model

We want to use this model as a guide for thinking about how a national UI policy may affect regional unemployment and how this is useful for identification purposes.

To model local wage determination we opt for a model involving search frictions and individualistic wage bargaining.[5] Assuming risk neutrality on the part of workers and firms, most kinds of decentralized bargaining models yield a wage equation of the following kind:

$$w_{ijt}(x) = \beta y_{jt}(x) + (1-\beta)O_{jt}(x) \tag{1}$$

where x denotes the (exogenous) characteristics of the worker involved in the bargain, y labor productivity, and O the outside option, i.e., the flow value of unemployment; see Pissarides (2000) for instance. The weighting parameter, β, reflects worker bargaining power, i indexes the bargaining unit, j the regional labor market, and t time. Thus, according to eq. (1), the

[5] We could equally well have modeled local wage determination as the outcome between a local union and a firm, but it is more convenient to have a model where we can think of firms as having only one job slot.

bargained wage is a weighted average of inside (y) and outside opportunities (O). We take outside opportunities to be given by:

$$O_{jt}(x) = \mu[n_{jt}w_{jt}(x) + u_{jt}b_{jt}(x)] + (1-\mu)[n_{-j,t}w_{-j,t} + u_{-j,t}b_{-j,t}(x)] \tag{2}$$

where u (n) denotes the un(employment) rate, and b the unemployment benefit; notice, that the unemployment benefit depends on the characteristics of the worker because there is a ceiling on benefit receipt. The index "*-j*" denotes aggregates over regions excluding j and μ is a measure of the allocation of search across regions. In equation (2) we have invoked the simplifying assumption that worker search intensity is fixed; coupled with an assumption that the separation rate is independent of x this implies that the un(employment) rate is independent of x.[6] Thus, the opportunities outside the firm are given by a weighted average of the opportunities inside and outside the region respectively. If the regions are small we can equally well write equation (2) as:

$$O_{jt}(x) = \mu[n_{jt}w_{jt}(x) + u_{jt}b_{jt}(x)] + (1-\mu)[n_t w_t(x) + u_t b_t(x)] \tag{3}$$

where the absence of a regional subscript signifies national aggregates. Conditional on the characteristics of the bargaining pair, the outcome of the wage-bargain is symmetric. Hence, $w_{ijt}(x) = w_{jt}(x)$. Inserting (3) into (1) and imposing symmetry we get

$$w_{jt}(x) = \frac{\beta y_{jt}(x) + \mu(1-\beta)u_{jt}b_{jt}(x) + (1-\beta)(1-\mu)O_t(x)}{1-\mu(1-\beta)(1-u_{jt})} \tag{4}$$

where we have introduced the notation $O_t(x) = n_t w_t(x) + u_t b_t(x)$ and used $n_{jt} \equiv 1 - u_{jt}$. This is the regional wage equation for a worker-firm pair where workers have characteristics x.

There is undirected search on the part of workers and firms. Hence, when posting a vacancy firms do not know what kind of worker they will meet – the decision to enter the market is based on the average productivity and the average wage in the region. Firms must make an up-front capital investment in order to open up a vacancy. This capital investment

[6] We acknowledge that having search intensities fixed, μ constant, and separation rates independent of x, may be short-cuts for which there is little justification except tractability. Notice that our empirical work does not rely on these assumptions.

commands a flow cost of *K*. Upon making this capital investment the firm opens a vacancy, which is filled with probability $q_{jt} = q(\theta_{jt})$ where $\theta_{jt} \equiv v_{jt} / u_{jt}$ and $q'(\cdot) < 0$. Thus, vacancy posting behavior satisfies a zero-profit condition

$$\frac{y_{jt} - w_{jt}}{\varphi} = \frac{K}{q(\theta_{jt})} \tag{5}$$

where $1/\varphi$ is the expected duration of the match. In equation (5), we do not index, e.g., the wage by x since the relevant quantity for a firm's entry decision is the wage averaged over the distribution of observed characteristics in the region: $w_{jt} = \int w_{jt}(x) dG_{jt}(x)$, where $G_{jt}(x)$ denotes the distribution of x in region j at time t. Equation (5) then says that the expected present value of the match (the left-hand side) equals the expected present value of the set-up cost (the right-hand side).[7]

To proceed, let us assume that realized productivity (i.e. the productivity realized after the match) is given by

$$y_{jt}(x) = \gamma_t x + \lambda_j + \lambda_t + \varepsilon_{jt} \tag{6}$$

where λ_j is a region-specific effect, λ_t a time-specific effect, and ε_{jt} a region-specific shock. Ex ante productivity (i.e. prior to the match) is given by

$$y_{jt} = \gamma_t \int x dG_{jt}(x) + \lambda_j + \lambda_t + \varepsilon_{jt} \tag{7}$$

Finally, we note that the unemployment benefit for a worker of type x is given by

$$b_{jt}(x) = \rho_t^n \left[w_{jt}(x) I(w_{jt}(x) \le w_t^{cap}) + w_t^{cap} I(w_{jt}(x) > w_t^{cap}) \right] \tag{8}$$

where ρ_t^n is the nominal replacement rate, $I(\cdot)$ the indicator function, and w_t^{cap} the cap on earnings used to calculate UI benefits.

Let us consider the average unemployment benefit received by workers in region j at time t upon unemployment. This equals

[7] Notice that we have implicitly assumed that there is no discounting when specifying equation (5). This assumption can also help rationalizing equation (2).

$$b_{jt} = \rho_t^n[(1-\psi_{jt}) + \psi_{jt} \frac{w_t^{cap}}{w_{jt}}]w_{jt} = \rho_{jt}^a(\rho_t^n, \psi_{jt}, w_t^{cap}, w_{jt})w_{jt} \tag{9}$$

where ψ_{jt} is the fraction of workers above the earnings cap and $\rho_{jt}^a(\cdot)$ is *the actual replacement rate*.

Notice that here we define the actual replacement rate as the ratio between average benefits and average regional wages. From an institutional point of view, it would have been more accurate to calculate the actual replacement rate for an individual worker, since unemployment benefits are usually tied to the wage of the individual worker.[8] However, this added realism would have come at substantial loss of tractability. This approach is simpler and we do not think it affects any qualitative conclusion.

Before solving the model, it is instructive to consider the determinants of the actual replacement rate. Since

$$\rho_{jt}^a = \rho_t^n[(1-\psi_{jt}) + \psi_{jt} \frac{w_t^{cap}}{w_{jt}}]$$

it is straightforward to verify that

$$\frac{\partial \rho_{jt}^a}{\partial \rho_t^n} = [(1-\psi_{jt}) + \psi_{jt} \frac{w_t^{cap}}{w_{jt}}] > 0$$

$$\frac{\partial \rho_{jt}^a}{\partial w_t^{cap}} = \rho_t^n \{ \frac{\psi_{jt}}{w_{jt}} + \frac{\partial \psi_{jt}}{\partial w_t^{cap}} [(w_t^{cap} / w_{jt}) - 1] \} \geq 0$$

where $\partial \psi_{jt} / \partial w_t^{cap} < 0$, i.e., if the wage cap increases the share above the ceiling is reduced. Consider comparing the magnitudes of these derivatives in two extreme regions: one where everyone has wages below the wage cap - $\psi = 0$ - and another where every wage is above the ceiling, i.e., $\psi = 1$. Evaluating the derivatives at these two extreme points we have

$$\left. \frac{\partial \rho_{jt}^a}{\partial \rho_t^n} \right|_{\psi=0} = 1 > \left. \frac{\partial \rho_{jt}^a}{\partial \rho_t^n} \right|_{\psi=1} = \frac{w_t^{cap}}{w_{jt}}$$

[8] In our empirical work we will calculate the actual replacement rate at the individual level and aggregate this measure to the regional level.

$$\left.\frac{\partial \rho_{jt}^{a}}{\partial w_{t}^{cap}}\right|_{\psi=1} = \frac{\rho_{t}^{n}}{w_{jt}} > \left.\frac{\partial \rho_{jt}^{a}}{\partial w_{t}^{cap}}\right|_{\psi=0} = 0$$

In other words, there are interaction effects in the model. Changes in the ceiling will increase generosity more in high-wage regions than in low-wage regions, while a change in nominal replacement rate will have the opposite effect.

Now, let us solve the model. Equations (6)-(9) imply that we can write the average regional wage as

$$w_{jt} = \frac{\beta y_{jt} + (1-\beta)(1-\mu)O_t}{1-\mu(1-\beta)(1-u_{jt}(1-\rho_{jt}^{a}(\cdot)))} \tag{10}$$

To complete the model we need an equation characterizing the flow equilibrium in the regional labor market. Equating the outflow from employment ($\varphi(1-u_{jt})$) with the inflow into employment ($\mu\alpha(\theta_{jt})u_{jt} + (1-\mu)\alpha(\theta_t)u_{jt}$) yields the relationship $u_{jt} = u(\theta_{jt})$.[9] Unemployment is decreasing in market tightness (θ) since the job offer arrival rate (α) increases in tightness (i.e. $\alpha'(\cdot) > 0$). It is convenient to invert the flow equilibrium condition (i.e. $\theta_{jt} = \theta(u_{jt})$) and use it to eliminate θ in equation (5). We get

$$\frac{y_{jt} - w_{jt}}{\varphi} = \frac{K}{q(\theta(u_{jt}))} \tag{5'}$$

Conditional on the state of the national market, equation (5') and equation (10) yield two equations in two unknowns: w_{jt} and u_{jt}.

The comparative statics with respect to the parameters of the UI system are fairly straightforward. An increase in the generosity of UI raises regional wage pressure (holding unemployment constant) and eventually increases unemployment by virtue of the zero-profit condition. Hence, we have $(\partial u_{jt} \big/ \partial \rho_t^{n}) \geq 0$ and $(\partial u_{jt} \big/ \partial w_t^{cap}) \geq 0$. From an empirical point of view, however, these predictions are not that helpful. If we control flexibly for time (by

[9] This flow equilibrium is consistent with the assumption that mobility occurs only when a job has been found. In flow equilibrium, the inflow into employment from other regions must be exactly balanced by an outflow from the region under consideration (which motivates the second term in the expression for the inflow into employment).

introducing time dummies in the empirical specification) it will not be possible to identify these effects. For empirical work, it is more useful to note the sign of two interaction effects. First of all, the effect of an increase in the nominal replacement rate will be greater in a low-wage region than in a high-wage region; in particular $(\partial u_{jt} / \partial \rho_t^n)\big|_{\psi=0} > (\partial u_{jt} / \partial \rho_t^n)\big|_{\psi=1} > 0$. In other words, the variation in the statutory replacement rate is less relevant in a region where the wage is higher (i.e. the share above the ceiling is higher). Second of all, the effect of increase in the benefit ceiling will be greater in a high-wage region than in a low-wage region; in particular $(\partial u_{jt} / \partial w_t^{cap})\big|_{\psi=1} > (\partial u_{jt} / \partial w_t^{cap})\big|_{\psi=0} = 0$. The sign of these two interaction effects follows from the properties of the actual replacement rate derived above.

To make full use of these predictions we must, of course, take account of the fact that wages (and hence the share above the benefit ceiling) are endogenous to unemployment. More specifically, the concern is that the region-specific shock in labor productivity (ε_{jt}) will spill-over onto unemployment as well as wages. Since the shock has an effect on the regional wage, it will have an effect on the actual replacement rate. In section 4, we outline how we try to eliminate this simultaneity problem.

3. The Swedish institutional setting

The "Swedish model" is a frequently used term for describing institutions in the Swedish labor market. The Swedish model featured centralized collective wage bargaining and extensive use of active labor market policy.

Given the (historical) reliance on centralized bargaining one might ask if the preceding model is a relevant characterization of the Swedish labor market. However, even during the heydays of the Swedish model, there was bargaining at different layers. There has always been additional wage drift at the local level, which constitutes a substantial fraction of the aggregate wage increase. Historically, wage drift at the local level accounted for 45 percent of total wage increases (Nilsson, 1993); between 1997 and 2002 wage drift amounted to 31 percent of the total increase.[10]

[10] This figure comes from the business cycle statistics reported by Statistics Sweden. Incidentally, it is not obvious how one should define wage drift since the early 1990s. During the 1990s, decentralized or individualistic bargaining has become increasingly common; see below.

Wage-setting institutions have changed rather drastically over the past couple of decades.[11] Centralized bargaining started to crumble in the beginning of the 1980s (Edin & Holmlund, 1995). During the 1990s, there was also a substantial move towards decentralization of wage negotiations. This started in the beginning of the 1990s, when some central agreements for white-collar workers in the private sector neither contained total wage increases nor minimum wage increases. It was entirely up to the employer and the employee to determine the wage; see Lindgren (2005). This trend towards decentralization has resulted in only 7 percent of the employed having their wages completely determined by the central industry bargain in 2004; moreover, the norm in the public sector is individualistic wage determination (Fredriksson & Topel, 2006).[12]

3.1 Unemployment benefits in Sweden

Receipt of unemployment insurance benefits requires the fulfilment of an employment requirement and a membership requirement.[13] The duration of UI receipt is formally 60 weeks. As explained above, UI benefits replace a fixed fraction (currently 80 %) of previous earnings up to a ceiling.

For those who do not fulfil the membership requirement there is an Unemployment Assistance (UA) system. Compensation on UA is unrelated to previous earnings and the generosity of UA is much lower than UI; on average it replaces roughly 40 percent of previous earnings.

Since the key aspect of our model is the effect of unemployment insurance on wage-setting, we will simply ignore the UA-system in the sequel. To us, this seems like an

[11] Despite these changes, unions figure as prominently in the Swedish labor market as they did during the beginning of the 1980s. The unionization rate in Sweden has hovered around 80 percent over the past couple of decades (OECD, 2004).

[12] At the same time as there has been decentralization of the wage bargain, a new coordination regime has emerged. In 1997, the so-called Industrial Agreement (IA) was struck between unions and employers in the manufacturing sector. This agreement involves a set of procedural rules, similar in many ways to the laws governing collective bargaining in the US. It stipulates, inter alia, time-tables for negotiations, rules for conflict resolution, and gives a prominent role for mediators. The IA-model may have delivered incentives for wage restraint at the aggregate level. But it is reasonable to think that it has had a minor influence on the regional wage structure, since the main function of the IA is to establish a set of procedural rules of the game.

[13] The information on the UI rules in this section comes from Olli Segendorf (2003). The employment requirement stipulates that the individual must have worked for a certain number of days during the year immediately preceding unemployment. Sweden is one of the few countries were UI is voluntary, hence the receipt of UI also requires the membership in a UI fund for at least 12 months and the payment of a small fee.

innocuous omission: the relevant issue in the wage bargaining framework is the level of benefit entitlement for an average *employed worker* upon unemployment entry.[14]

Another feature of unemployment benefits in Sweden is more relevant in that respect. All collective agreements provide additional compensation for (some) workers in the case of redundancies. Despite their relevance, it is very hard to get the full picture of the conditions and payments involved (Sebardt, 2005, provides very useful information, however).[15]

The redundancy payments regulated by collective agreement may come in two forms: either as a lump-sum severance payment or as a supplementary unemployment benefit. Although lump-sum severance payments may be non-negligible and should affect incentives in the wage bargain, we choose to ignore them here. The main reason for this omission is that eligibility is a function of tenure – which is information that we do not have. Furthermore, for the biggest group having a lump-sum severance payment – public sector workers – the lump-sum is proportional to the previous wage with no ceiling imposed. With this construction, the severance payment does not contribute to identification.[16]

For our whole study period, 1974-2002, there are no supplementary unemployment benefits for the vast majority of workers. Thus, in terms of the periodic unemployment benefit payments, the rules of the public unemployment insurance system apply. There are some notable exceptions, however. Starting in 1990, all central government employees got additional insurance via a collective agreement. Given that the employee has an open-ended contract, or has been on fixed-term contracts for at least three years, there is no benefit

[14] This is partly the reason for also ignoring the duration of benefit receipt. More importantly, however, benefit duration is unrelated to previous wages and hence do not contribute to identification.

[15] Indeed, Wadensjö (1993) adequately refers to the additional compensation provided by collective agreement as the “unknown part of the social insurance system”. The information in the rest of this section relies heavily on Sebardt (2005).

[16] The main agreements providing lump-sum severance pay concern public sector employees and private sector blue-collar workers. For local public sector employees, such constructions have existed since 1984. The severance pay is proportional to the previous wage (with no ceiling). At most the employee can be paid half of their annual earnings. This happens in the case of employment for 18 years in the local public sector. For each year of “tenure” less than 18 years there is a proportional reduction in the lump-sum payment. For blue-collar workers, the severance payment is only a function of tenure and age. A rough description is that only individuals above age 50 qualify; in addition, the worker should have at least 10 years of tenure. The payment is proportional to tenure, but increasing with age for given tenure; see Sebardt (2005). Of course, the existence of severance payments raises the nominal replacement rates for the workers affected by them. Notice that we can to some extent control for the incidence of severance payments by controlling for age and industry composition.

ceiling. That is, the employee gets the statutory replacement rate independent of the previous wage.

The oldest collective agreement offering supplementary benefits applies to white-collar workers in the private sector. This has been in place for the entirety of our study period. The supplementary benefit structure is more complex than for government employees. Supplementary benefits are only offered for workers above age 40 who have at least 5 years of tenure. Their basic structure is that workers should be offered an actual replacement rate which is no less than 70 percent. A simple way to think about these payments is thus that they kick-in at a wage equaling the benefit ceiling divided by 0.7.[17]

The final collective agreement offering supplementary unemployment benefits refers to local public sector employees. This agreement was struck in 1984. It features a strict eligibility requirement. It is given only to redundant employees over 45 satisfying a "tenure" requirement. For 45 year-olds the tenure requirement is that they should have worked in the public sector for 17.5 years.[18] Should they qualify for supplementary benefits, they are given a benefit equaling the nominal replacement rate times the previous wage with no ceiling imposed.[19]

4. Data and construction of key variables

We use three principal data sources: (i) LINDA – an individual (register) data base (see Edin & Fredriksson, 2000, for a description); (ii) regional (open) unemployment from the Labor Force Surveys; and (iii) regional active labor market program rates from the Labor Market Board.

The LINDA data set is based on a combination of income tax registers, population censuses, wage data, and other sources. Unfortunately, the wage data are not available for our

[17] This is almost how the system works at present; the complication that we have not mentioned is that for wages above 20 price base amounts the slope of the benefit-wage schedule becomes 0.25. Further, relative to the system that existed during 1974-2002, it is a slight simplification at the bottom end. Those below the wage cap implied by a replacement rate of 70 percent were given a relatively small nominal amount as well; this nominal amount raises the nominal wage replacement rate for those below the wage cap in the public UI system.

[18] The tenure requirement decreases with age: at age 60, 10 years of tenure is required.

[19] On top of all this, some UI funds offer their members the option of purchasing private unemployment insurance. However, this possibility is very recent and hence does not concern us.

entire study period. We only have access to wages for a representative sample of workers from 1998 and onwards. Apart from wages, the individual variables we use in our analysis are based on register information. Earnings and some other characteristics (gender, age, education, marital status, and industry affiliation) are obtained from the income tax registers, which also contain information on region of residence and country of birth from the population registers. The earnings information and most of the other individual characteristics are available throughout the time period; see appendix for more information on data availability.

The individual data are used to calculate measures of the composition of the regional labor force and to run individual earnings regressions. The estimated parameters from the earnings regressions are used to generate expected wages had the characteristics of the individual worker been priced at the national labor market. We use this strategy to free the estimates from the simultaneity bias caused by local shocks affecting both regional unemployment and wages. Having generated these expected wages we calculate the average of these wages at the regional level and the actual generosity of UI at the regional level.

4.1 Construction of key independent variables

We start by estimating individual earnings regressions separately by year. These equations have the following structure

$$\ln y_{ijt} = \alpha_t + \alpha_{jt} + \beta_t X_{it} + \varepsilon_{ijt} \tag{11}$$

where *i* indexes individuals, *j* regions, and *t* time. In equation (11), *y* denotes earnings, α_{jt} is a region-fixed effect – normalized such that $\sum_j \alpha_{jt} = 0$ – and *X* denotes the vector of covariates. The covariate vector includes information on gender, age (separate dummies for each five-year age category) educational attainment, marital status, country of birth, and industry. We run these equations for each year between 1970 and 1998 including only individuals who are 16-59 years of age.[20] We control flexibly for region at the estimation stage to avoid sorting

[20] The upper age limit is due to the fact that the information on education is only consistently available for individuals less than 60 years-of-age. See appendix for more details.

bias in the coefficient vector, β_t; such a bias might arise if high-skilled individuals cluster in regions hit by positive wage shocks. When estimating these equations we exclude the lowest quintile of the earnings distribution. The rationale for this is that we want the parameter estimates to resemble what one gets when estimating traditional wage equations; see Antelius & Björklund (2000).

Using the estimates of the parameters in (11) we want to generate an expected wage – the wage that each individual would obtain if his/her characteristics were priced on the national labor market. Our main strategy to compute such a wage is as follows

$$w_{it}^{e,1} = \delta_{it} - \bar{\delta} + \bar{w}_t \tag{12}$$

where $\delta_{it} = \exp(\hat{\beta}_{t-4} X_{it})$, $\bar{\delta}$ denotes the mean of δ_{it}, and $\bar{w}_t$ denotes the average wage in the country.[21] Thus, the individual gets assigned the same wage independently of where s(he) is located. We lag the "national price vector", $\hat{\beta}$, four years in order to ensure that the expected wage is independent of any region-specific shocks. Big regions, such has Stockholm, are likely to be very influential in the estimation of $\hat{\beta}_t$. If we would have used $\hat{\beta}_t$ rather than $\hat{\beta}_{t-4}$ a potential worry is that the wage predictions would not have been independent of shocks to unemployment in Stockholm. Another reason for not using $\hat{\beta}_t$ concerns skilled-biased technical change. Suppose there is skilled-bias technical change. This will presumably raise the return to education and will represent a favorable employment shock in regions rich on observed and unobserved human capital. This scenario will induce a negative correlation between the wage prediction and the error-term in the unemployment equation.

Given a measure of the expected wage, we proceed to define an individual indicator variable for having predicted wages above the wage cap. Moreover, we calculate the actual replacement rate at the individual level as

$$\rho_{it}^{a} = \rho_t^{n}\left[I(w_{it}^{e} \leq w_t^{cap}) + \left(w_t^{cap} / w_{it}^{e}\right) I(w_{it}^{e} > w_t^{cap})\right] \tag{13}$$

We then average over all individuals residing in the region which gives us

[21] According to equation (12) we adjust the predictions such that they are mean zero and center them on the mean national wage.

$$\rho_{jt}^{a}=\frac{\sum_{i\in j}\rho_{it}^{a}}{N_{jt}},\ w_{jt}^{e}=\frac{\sum_{i\in j}w_{it}^{e}}{N_{jt}} \tag{14}$$

where N_{jt} is the number of individuals residing in region j at time t. ρ_{jt}^{a} is the key independent variable in the empirical analysis.

In equations (13) and (14) we have calculated the actual replacement rates as if only the public UI system is relevant. Obviously, we would also like to take the existence of supplementary unemployment benefits into account. In the next section, we outline how we try to accommodate this feature.

4.2 Supplementary unemployment benefits

Since we do not have adequate information in the data, taking supplementary unemployment benefits into account is bound to involve some approximations. In the data, we observe in what sector the individual works but we do not observe whether the individual is a blue-collar or a white-collar worker. Further, we do not observe tenure for the individual worker.

The supplementary unemployment benefit in the central government sector is fairly straightforward to approximate. Historically, the vast majority of workers in the public sector were on open-ended contracts. Therefore, we simply assume that all workers are eligible for this system from 1990 and onwards. Since this agreement implies that there is no benefit ceiling, we set the actual replacement rate equal to the nominal one from 1990 and onwards for central government workers.

White-collar worker status in the private sector is proxied with workers in the private sector having at least three years of (theoretical) upper-secondary education. The supplementary benefit was paid to individuals who were at least 40 years-of-age with at least 5 years of tenure in the firm. The question then is: What does the tenure structure look like for white-collar workers in the private sector above 40? To examine this question we used survey data from the Swedish Level of Livings Survey (LNU) in 2000; Erikson & Åberg (1987) describe the LNU data. It turned out that 75 percent of workers in the private sector with at least 3 years of upper-secondary education had tenure of at least 5 years. Therefore, as an approximation we assume that all workers that we classify as private sector white-collar workers are eligible for supplementary benefits if they satisfy the age constraint. The workers that qualify for this supplementary benefit are given the benefit structure outlined in section 3.1, i.e. the actual replacement rate never falls below 0.7.

The final supplementary benefit agreement concerns local government employees. In this case the age constraint is 45 and the "tenure" requirement is almost 18 years. Since this requirement appears very stringent, we have chosen to ignore this agreement altogether.[22]

4.3 Data

There are many steps involved in creating these regional panel data. The full detail of our data collection effort is presented in Appendix A. Here we describe the main steps and present the main characteristics of the data.

We begin by creating a data set involving individual characteristics and earnings from 1970 to 2002. The included individual characteristics are fairly standard. We have information on gender, age, marital status, region of residence (at the county level), educational attainment, industry affiliation (2-digit ISIC), and country of birth. With respect to country of birth we distinguish between individuals of native, Nordic, OECD, or non-OECD origin.[23] In terms of education, we distinguish between compulsory school (or less), upper secondary school, and tertiary education.

We first utilize these data for estimating individual earnings and wage regressions. On the basis of the estimated equations we generate an expected "wage" for each individual as described above. The mean of the predictions is adjusted such that it corresponds to the national average wage for each point in time. We are implicitly assuming that the estimates of the slope parameters in the earnings regressions are the same as they would be in the wage

[22] Also, in this case we had a brief look at the LNU data. The "tenure" requirement in the agreement pertains to the total number of years worked in the local public sector. This is not observed in the LNU data. If we look at tenure with the current employer – a reasonable approximation of the number of years of continuous employment in the local public sector – we find that a quarter of those aged 45 are eligible. Eligibility increases with age. At age 50, half of the relevant population is eligible and at age 59 around three quarters are eligible. At any rate, a small share of the population is eligible for this supplementary unemployment benefit and, therefore, we ignore it.

[23] Individuals are generally classified as being of OECD origin if they were born in a country which was a member of the OECD in 1985. The only exceptions from this rule are Turkey – which is included among the non-OECD countries – and the Nordic countries.

regressions. This may be a questionable assumption since earnings variations are also due to variations in hours worked. But notice that we trim the lower tail of the earnings distribution to minimize this problem.

Then we also need information on the relevant parameters of the UI system: the benefit ceiling and the nominal replacement rate. The benefit ceiling is specified in nominal terms, so it comes as no surprise that it has been changed frequently. On 20 occasions the ceiling was changed during the time period. One would expect the ceiling to be adjusted according to the rate of wage inflation such that the “insurance value” is left unchanged. However, during most of the time period, the ceiling is changed on the discretion of the legislator and, as we illustrate later, there is a good deal of hap-hazardness introduced by these discretionary changes. The nominal replacement rates have been changed more infrequently. There have been four changes in the nominal replacement rate between 1974 and 2002.

Figure 2: The nominal and actual replacement rates and the wage cap, 1974-2002.

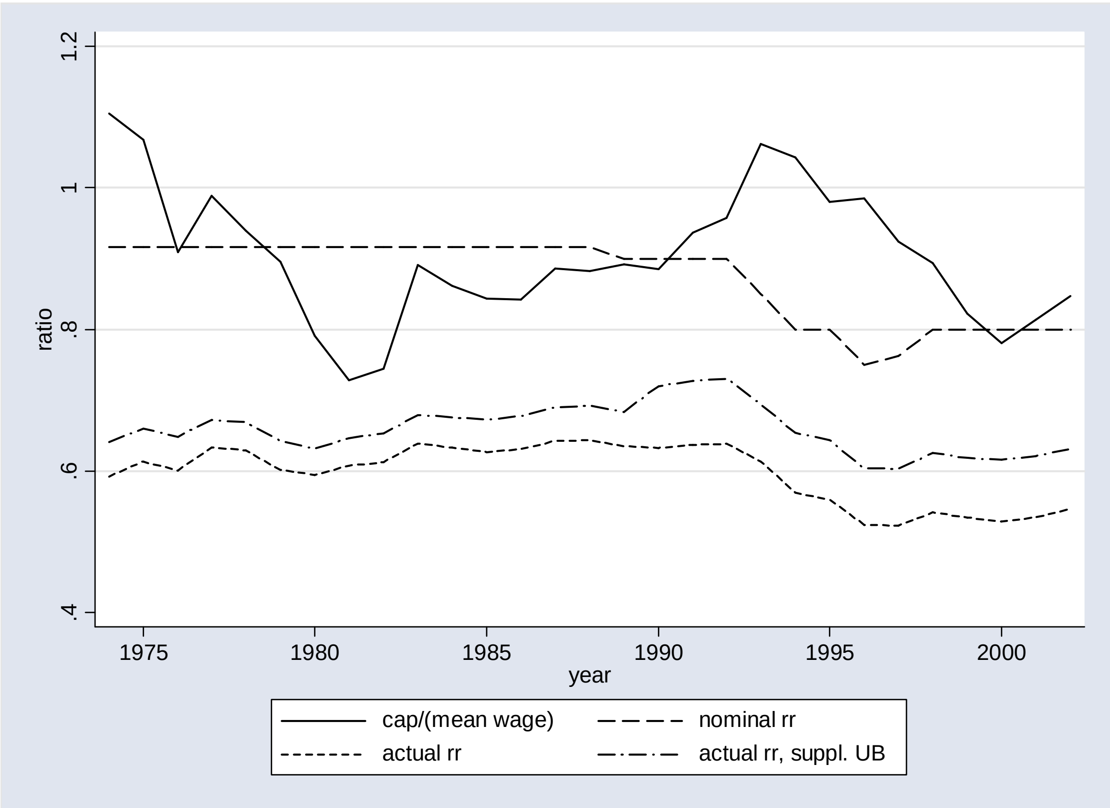

Sources: See data appendix

In Figure 2 we plot the evolution of the nominal replacement rate and the wage cap (divided by mean wages) over time at the national level. Along with these two series, we also plot the evolution of the actual replacement rate – unadjusted as well as adjusted for the incidence of supplementary unemployment benefits.

Figure 2 shows that there is a good deal of idiosyncratic variation in the wage cap and that this variation contributes to most of the variation in the actual replacement rate (we substantiate this claim more in the next section). Figure 2 also shows that benefit generosity was scaled back following the unemployment crisis in the beginning of the 1990s.

Our key outcome measure is defined as the sum of open unemployment and participants in labor market programs as a share of the labor force. With some abuse of language we refer to this sum as "unemployment" in the sequel.

Figure 3. Unemployment, mean and spread, 1974-2002.

Notes: Unemployment is defined as the sum of the openly unemployment and participants in active labor market programs as a share of the labor force.
Sources: Labor Force Surveys and the Labor Market Board

Figure 3 shows the development of mean unemployment along with the evolution of the extremes in the distribution (the min. and max. values) to give a sense about the regional variation in the data. The most striking event in this figure is the adverse shock that hit Sweden in the beginning of the 1990s. In just three years unemployment shot up from around three percent in 1990 to roughly 13 percent in 1993. The aggregate unemployment rate was stable at this high level until 1997. In some regions, however, unemployment continued to rise to reach 22 percent in 1997. The period since then has seen substantial fall in unemployment.

Figure 4. The actual replacement rate, mean and spread, 1974-2002.

Notes: The actual replacement rate have been generated using coefficients estimated on earnings data from 1970-1998. The actual replacement rate takes supplementary unemployment benefits into account.

Figure 4 gives a sense about the regional variation in our key measure of the generosity of the UI system. It shows the variation in the actual replacement rates over time and across regions when supplementary unemployment benefits have been taken into account. The actual replacement rate stood at a high in the early 1990s when it equalled 73 percent. Since then it has fallen quite rapidly to 63 percent in 2002. The variation across regions was particularly high around 2000. It is evident that there is a good deal of variation across regions as well as time, which we can potentially utilize in the following sections.

5. What variation identifies the actual replacement rate?

This section is devoted to illustrating in more detail where the identifying variation in the actual replacement rate comes from. To fix ideas we begin, in section 5.1, with a simple graphical example where we ignore the existence of supplementary unemployment benefits. In section 5.2 we turn to the data and examine the empirical importance of the determinants of the actual replacement rate.

5.1 A simple graphical example

Figure 5 provides a simple graphical illustration, where supplementary unemployment benefits are ignored. The bold (solid) line depicts the benefit schedule. According to this schedule, benefits increase linearly with wages for all wages below the cap (w^{cap}); the rate of increase in benefits is given by the nominal replacement rate (nominal rr). For wages above the cap there is no increase in benefits as indicated by the flat segment of the benefit schedule.

Suppose, for simplicity, that there is no dispersion of wages within region. Then the actual replacement rate (actual rr) in the high-wage (w_h) region is given by the slope of the dashed line, while the nominal and actual replacement rates coincide in the low-wage region (w_l). Now, suppose that the benefit ceiling increases. Then such a change has no effect in the low-wage region. But it has a substantial effect in the high-wage region, as indicated by the thinner dashed line in Figure 5.

It is easy to see that if there is wage growth in the high-wage region – i.e., w_h is pushed further to the right in the figure – then this will lower the actual replacement rate. It is also straightforward to verify that if the nominal replacement rate increases this will have the biggest effect on the generosity of UI in the low-wage region.

Figure 5: The effects of variations in the benefit ceiling

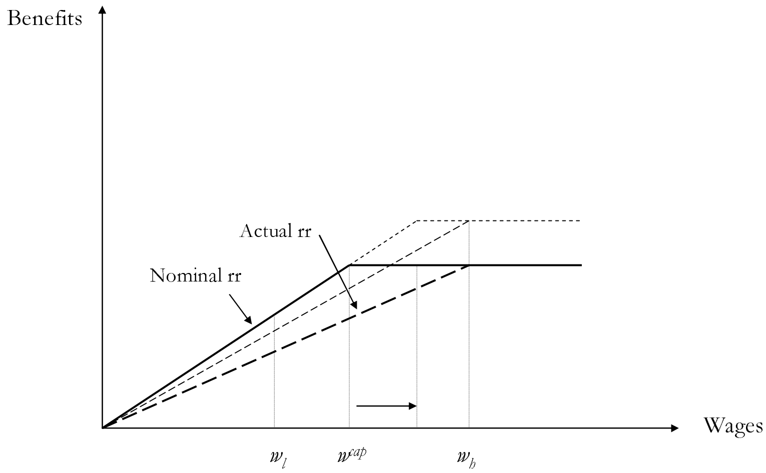

Other possible sources of variation are more subtle, however. Consider wage dispersion within regions. Let us focus on two regions where mean wages are the same and coinciding with w^{cap}. Suppose, further, that in the two regions the wages are symmetrically distributed around the mean. Then in the region with the greater variation in wages, the top end of the distribution will have a lower actual replacement rate on average. Thus, the standard deviation of the wage distribution should be negatively associated with the actual replacement rate.

5.2 A look at the data

Above we argued that the variations in expected wages, the benefit ceiling, the nominal replacement rate, as well as the variation in the spread of the expected wage distribution all contribute to the variation in the actual replacement rate. Here we illustrate the importance of each source of variation.

To facilitate the interpretation of the independent variables we standardize these variables with their standard deviations. Table 1 presents the results. In panel A) we show the results when not taking the existence of supplementary unemployment benefits into account. All the estimates have signs which are consistent with the discussion above. So, for instance, if expected wages increases by a standard deviation this yields a reduction of the actual replacement rate by half a percentage point; see column (2). It is also interesting to note that the variation in the wage cap is such a powerful predictor of the actual replacement rate; this confirms the impression given already in Figure 2. An increase in the cap has the effect of increasing the actual generosity of UI more in regions which are expected to be high-wage.[24]

Table 1. What explains the variation in the actual replacement rate?

	(1)	(2)	(3)
A) No account for supplementary UB			
Expected wage	-0.392** (0.033)	-0.502** (0.177)	-0.701** (0.178)
Expected wage interacted with wage cap		0.466** (0.056)	0.550** (0.064)
Expected wage interacted with nominal replacement rate		-0.321** (0.132)	-0.197 (0.130)
Standard deviation of expected wage			-0.062** (0.014)
# observations	696	696	696
Within R^2	0.63	0.80	0.81
B) With account for supplementary UB			
Expected wage	-0.281** (0.042)	-1.30** (0.437)	-1.32** (0.442)
Expected wage interacted with cap		0.601** (0.103)	0.611** (0.113)
Expected wage interacted with nominal replacement rate		0.463 (0.362)	0.477 (0.361)
Standard deviation of expected wage			-0.007 (0.022)
# observations	696	696	696
Within R^2	0.29	0.43	0.43

Note: Dependent variable in percent. The table reports standardized coefficients. An individual coefficient has the interpretation of percentage point change in response to a standard deviation increase in one of the independent variables. The regressions also control for regional fixed effects, region-specific trends, fixed time effects, and exogenous labor force characteristics. Within R^2 reports the share of the variance explained by the four variables in the table after having controlled for other covariates, region-specific FEs and trends, as well as time effects. Other covariates include age, education, immigrant status, gender, industry affiliation, and the share covered by supplementary UB (only panel B). Regressions are weighted by population. Standard errors, reported in parentheses, allow for clustering at the county level. Significance levels: * = 10%, ** = 5%

[24] Notice that it is only the interaction effect which is identified. The main effect of the wage cap is "swamped" by the time fixed effects.

In panel B) we consider the variation in the generosity of UI when supplementary unemployment benefits are taken into account. The evidence presented in panel B) is not as clean as the estimates presented in the previous panel. For instance, the statutory replacement rate no longer has a greater effect in regions that are predicted to be low-wage (which should be the case according to the simple benefit formula). And the standard deviation of the expected wage distribution ceases to be a significant predictor of the actual replacement rate. Nevertheless, the estimates again suggest that the wage cap is the most significant contributor to the explained variance of the actual replacement rate.

Another aspect of the results in Table 1 is also worth noting. With the four variables we do not account fully for the variation in the actual replacement rate. In other words, there is residual variation, since the explained variance does not equal unity. There are a number of reasons for this. At the individual level, the benefit schedule depicted in Figure 5 is deterministic. This is not the case at the aggregate regional level. To explain the variation in the actual replacement rate fully at the regional level, we would have to include all moments of the expected wage distribution; obviously, this is not feasible. Further, supplementary unemployment benefits introduce additional noise, which is evidenced by the fact that explained variance is lower in panel B) than in panel A).[25]

In summary, the most important finding in this section is that a substantial fraction of the variation in the actual replacement rate at the regional level is due to variations in the national wage cap. An increase in the wage cap has a greater positive effect on UI generosity in regions which are expected to be high-wage. Thus it should be possible to identify the effect of the actual replacement rate on regional unemployment using only the variation in the wage cap. This identification strategy is the one that we will mainly pursue in the next section.

6. What is the effect of increases in the actual replacement rate?

With the exercise in section 5 as a background we now proceed to examine the relevance of the UI system for regional unemployment. We begin with a very basic question. Do the

[25] A final reason is that we are not using the functional form implied by Figure 5. Since this is not the right function at the regional level, we have no reason to impose it.

parameters of the UI system have any impact on regional unemployment? This is a relevant question given that many collective agreements supplement unemployment benefits. To investigate this issue, we first estimate the equation

$$\ln u_{jt} = \kappa^{w} w_{jt}^{e} + \kappa^{cap}(w_{t}^{cap} \times w_{jt}^{e}) + \kappa^{\rho}(\rho_{t}^{n} \times w_{jt}^{e}) + \kappa^{\sigma}\sigma_{jt} + \kappa^{X} X_{jt} + \mu_{j}^{1} + \mu_{t}^{1} + \delta_{j}^{1} t + \omega_{jt}^{1} \qquad (15)$$

where w_{jt}^{e} denotes the expected wage and σ_{jt} the standard deviation of the expected wage distribution. The vector of characteristics, X, includes the same components as in the individual earnings regressions since any exclusion restriction with respect to the components of X is bound to be arbitrary. Furthermore, X includes a control for supplementary unemployment benefits. The specification of equation (15) also takes region-specific effects, time fixed effects, as well as region-specific trends into account.

The idea behind equation (15) is that the first four components conceptually drive the variation in the actual replacement rate at the regional level. One can potentially make the argument that the expected wage and the standard deviation capture omitted variables in the unemployment equation. But it is very hard to see that this is a relevant argument for the interaction terms. In particular, if we find that $\kappa^{cap} > 0$ this strongly suggests that the design of the national UI system has implications for regional unemployment. The same line of argument goes for the interaction with the nominal replacement rate where we would expect $\kappa^{\rho} < 0$.

Table 2 reports the results. We mainly focus on the specification where the dependent variable is the log of unemployment; see column (1). But in column (2) we also report the results of a specification where we use the unemployment rate as the dependent variable. Again, we standardize the key independent variables to facilitate the interpretation of the coefficients of these variables.

In column (1) the interaction between the expected wage and the wage cap enters significantly with a positive sign. Thus, changes in the wage cap produce a greater increase in unemployment in regions which are expected to have a high wage. A standard deviation increase in this interaction term raises unemployment by almost 4 percent. The remaining interaction variable is not significant and does not have the predicted negative sign; this result is consistent with the estimates reported in panel B) of Table 1.

Table 2. Basic estimates

	(1) ln(unemployment)	(2) Unemployment (percent)
Expected wage	-0.069	1.41^{**}
	(0.054)	(0.628)
Expected wage interacted with wage cap	0.036^{**}	-0.090
	(0.015)	(0.158)
Expected wage interacted with nominal replacement rate	0.044	-1.30^{**}
	(0.046)	(0.514)
Standard deviation of expected wage distribution	0.008^{*}	-0.021
	(0.005)	(0.031)
Other covariates	Yes	Yes
Region-specific FEs	Yes	Yes
Region-specific trends	Yes	Yes
Time effects	Yes	Yes
Overall R^2	0.981	0.982
Within R^2	0.045	0.045
# observations	696	696

Note: Key independent variables are standardized and have the interpretation of the effect on the dependent variable in response to a standard deviation increase in the independent variable. All regressions are estimated using a within-estimator and include controls for gender, age, marital status, educational attainment, immigrant status, industry affiliation, and the share of individuals covered by collective agreements with supplementary unemployment benefits. Within R^2 reports the share of the variance explained by the four variables in the table after having controlled for other covariates, region-specific FEs and trends, as well as time effects. Regressions are weighted by population. Standard errors, reported in parentheses, are clustered by county. Significance levels: * = 10%, ** = 5%.

The estimates in column (2) tell a slightly different story. In this case, the interaction with the wage cap is not significant; but the interaction with the nominal replacement rate enters significantly with the predicted negative sign.

Theory provides little guidance to the question of whether the dependent variable should be specified in logs or as the rate of unemployment. In the sequel, however, we focus on the estimates where log unemployment is the dependent variable. Given the results presented in Table 1, the estimates in column (1) is consistent with the hypothesis that the actual replacement rate drives the evolution of log unemployment. The estimates in column (2) are consistent with this hypothesis to a lesser extent.

The most important result contained in Table 2, however, is that parameters of the national UI system do affect regional unemployment. Having established this we proceed to estimating equations imposing more structure.

The specification in equation (16) imposes more structure. In this case we relate unemployment directly to the actual replacement rate. Thus

$$\ln u_{jt} = \varphi^{\rho}\rho_{jt}^{a} + \varphi^{w}w_{jt}^{e} + \varphi^{\sigma}\sigma_{jt} + \varphi^{X}X_{jt} + \mu_{j}^{2} + \mu_{t}^{2} + \delta_{j}^{2}t + \omega_{jt} \tag{16}$$

There are two potential ways to estimate (16). The first is akin to a control function approach. Controlling for expected wages and the standard deviation of the expected wage distribution, the remaining variation in the actual replacement rate has two components: one source of variation is due to the interaction terms between the expected wage and the wage cap as well as the nominal replacement rate respectively; the other source of variation is the residual variation in the actual replacement rate. Given the substantial difficulties involved in measuring the actual generosity of UI, the residual variation is likely to contain a lot of noise; this implies that the control function approach will generate estimates that are biased downwards due to attenuation.

The other approach to estimating (16) is to just utilize the predicted variation in the actual generosity of UI stemming from the variation in the interaction terms – the most important of these interactions being the variation stemming from the changes in the wage cap. Implementing this strategy is straightforward; it amounts to estimating equation (16) using standard IV methods.

Table 3 reports estimates of equation (16). The dependent variable is the log of unemployment. In columns (1) and (2), the actual replacement rate does not account for supplementary unemployment benefits; in column (3) it does. In column (1) we use the control function approach; columns (2)-(3) are based on the IV-approach. The equations are all estimated using a traditional within-estimator.

Table 3 suggests that the estimation approach matters a great deal for the results. Column (1) – which is based on the control function approach – reports an insignificant estimate on the actual replacement rate. However, if we use only the variation induced by the interactions terms, the estimate is significant. We are inclined to interpret the divergence in the results as being due to bias because of measurement error. As illustrated in Table 1, the noise is substantial, particularly when the measure of benefit generosity attempts to account for supplementary unemployment benefits.

Table 3. The effect of the actual replacement rate on unemployment

	ln(unemployment)		
	(1)	(2)	(3)
Actual replacement rate (percent) (No account for supplementary UB)	0.018 (0.018)	0.045** (0.023)	
Actual replacement rate (percent) (Account for supplementary UB)			0.054** (0.023)
Expected wage/1000	0.023* (0.012)	0.040** (0.016)	0.036** (0.013)
Standard deviation of expected wage/1000	0.032** (0.012)	0.030** (0.012)	0.023* (0.012)
Other covariates	Yes	Yes	Yes
Region-specific FEs	Yes	Yes	Yes
Region-specific trends	Yes	Yes	Yes
Time effects	Yes	Yes	Yes
Overall R^2	0.981	0.981	0.980
Estimation approach	Control function	IV	IV
# observations	696	696	696

Notes: Columns (2)-(3) use the interactions between the expected wage and the wage cap as well as the nominal replacement rate respectively to identify the coefficient on the actual replacement rate. All regressions are estimated using a within-estimator and include controls for gender, age, marital status, educational attainment, immigrant status, industry affiliation, and the share of individuals covered by collective agreements with supplementary unemployment benefits. Table B1 in Appendix B reports the coefficient estimates on the majority of the remaining covariates for the specification reported in column (3). Regressions are weighted by population. Standard errors, reported in parentheses, are clustered by county. Significance levels: * = 10%, ** = 5%.

Our preferred estimate is shown in column (3). The coefficient estimate suggests that unemployment rises by 5 percent (i.e. the unemployment rate increases from, say, 6 to 6.3 percent) in response to increase in the actual replacement rate of 1 percentage point.[26] The elasticity of unemployment with respect to benefit generosity implied by this estimate is remarkably high. Evaluated at the mean actual replacement rate in 2002 (63 %), the elasticity equals 3.4.

[26] This change in benefit generosity roughly corresponds to the weighted standard deviation of the actual replacement rate within regions and time.

The estimates reported in Table 3 are higher than we have found elsewhere in the literature. The estimates in columns (2) and (3) are roughly four times higher than Nickell (1998) obtained in his study of a cross-section of OECD countries. Krueger & Meyer (2002) report a benefit elasticity of one when taking the effect on the incidence as well as duration of unemployment into account.

Of course, it is hard to pinpoint why we get higher estimates than those available elsewhere in the literature. Relative to Nickell (1998), we would argue that effects that we estimate are more credibly identified than in his cross-country regression. The estimate reported in the Krueger & Meyer (2002) is obtained by surveying micro studies mostly pertaining to the US. Here we cannot argue that our estimate is more credibly identified. But clearly the parameter we estimate is different in the sense that it takes equilibrium adjustments into account to a greater extent than in micro studies. Also, we obtain this estimate using Swedish data where unemployment benefits are substantially more generous than in the US. This is a relevant issue since, e.g., standard matching models suggest that the general equilibrium effect on unemployment of a given variation in UI generosity is greater the higher is UI benefits from the outset; some illustrative simulations on this theme are reported in Holmlund (1998), and Hornstein *et al* (2005).

We have subjected the specification in column (3) to some specification checks. First we used the unemployment rate as the dependent variable. The estimate is substantially weaker. A percentage point increase in UI generosity causes unemployment to rise by 0.090 percentage points; the standard error of this estimate is 0.062. Second, we introduced a lag of the actual replacement rate. This virtually had no effect on the estimate and the coefficient on the lag was not significant. Third, we transformed the model by taking first differences. This reduced the size as well as the precision of the estimate. The coefficient on the actual replacement rate was reduced to 0.024 with a standard error of 0.015. Despite the fact that the estimates sometimes become less precise, we view them as fairly robust to specification changes.

7. Policy interventions

The purpose of this section is to conduct two policy simulations. In particular we are interested in the effect of removing the benefit ceiling and the effect of increasing the nominal replacement rate. These two policy changes have obvious implications for aggregate unemployment – i.e. aggregate unemployment increases. The more interesting effects are those on the regional distribution of unemployment. In almost all countries, regional unemployment differentials are very stable over time; see Figure 1 and, e.g., Fredriksson (1999) for a collection of evidence. Perhaps the design of the social insurance system contributes to this feature?

Figure 6: The correlation between unemployment and the actual replacement rate

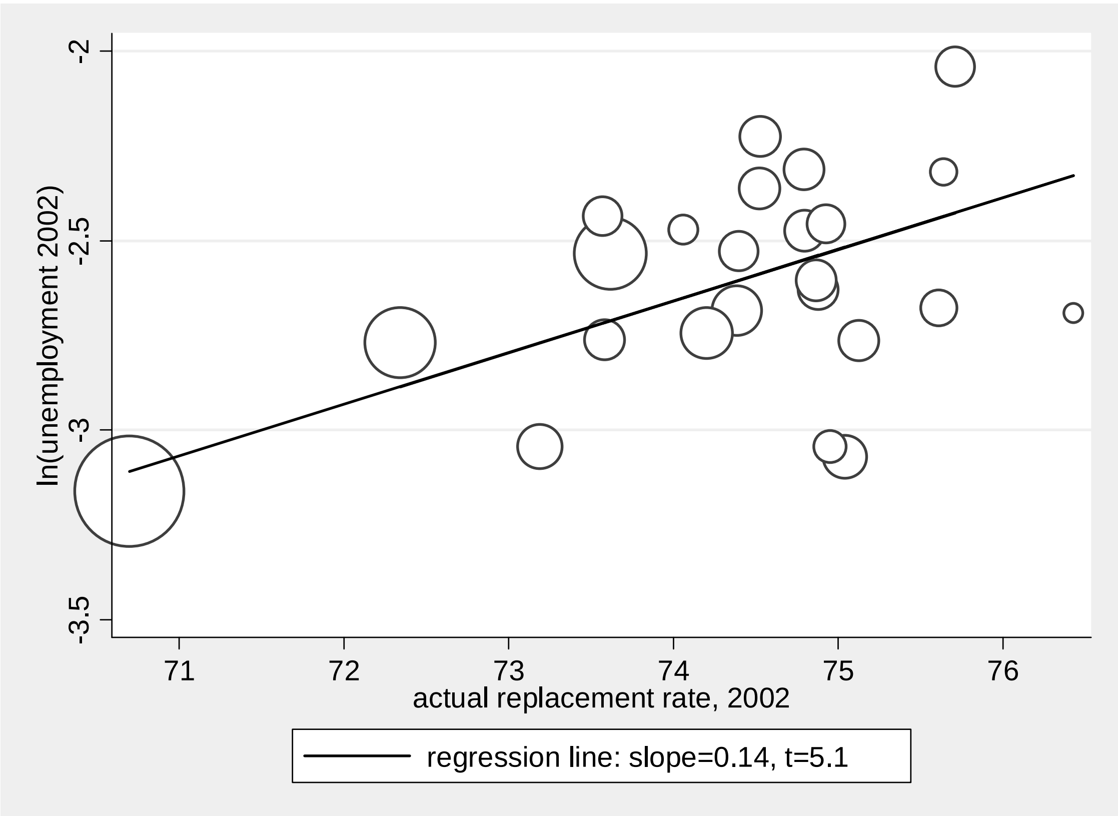

The starting point for these experiments is the regional distribution of unemployment and actual replacement rates in 2002. In contrast to the previous analysis we actually have wage data for 2002 and hence we have a better estimate of the "true" actual replacement rate. In Figure 6 we show the correlation between regional unemployment and the actual replacement rate accounting for supplementary unemployment benefits. As shown by the slope of the regression line, a percentage point increase in the actual replacement rate is associated with 14 percent higher unemployment. This just illustrates that high-wage regions tend to be low-unemployment regions. Since unemployment benefits replace a lower fraction of previous wages in high-wage regions they also tend to have a lower actual replacement rate. This simultaneity bias thus inflates the estimate of the relationship between benefit generosity and unemployment.

In 2002, the aggregate unemployment rate stood at 6.8 percent. To generate the situation after a policy change we use the estimate on the actual replacement rate reported in column (3) of Table 3. We set the coefficient on the actual replacement rate to 0.05.

Figure 7. Policy simulation – Removing the cap

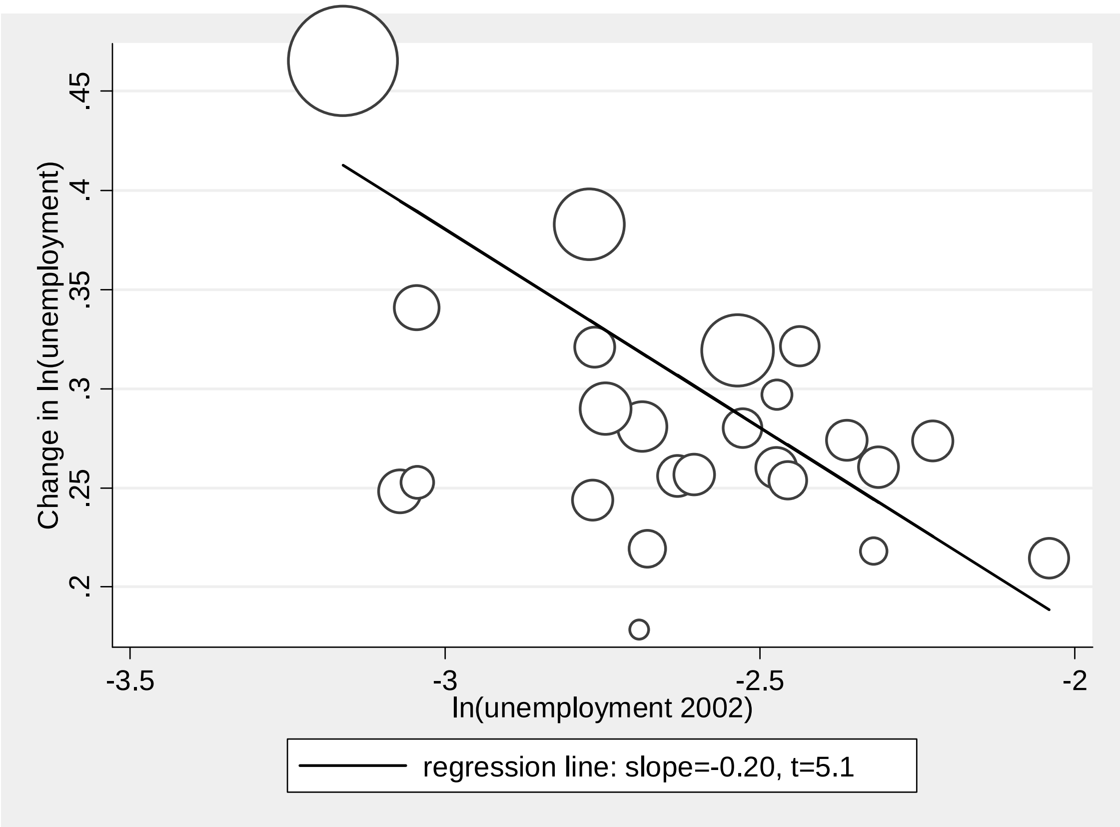

Notes: This graph is based on a hypothetical policy experiment where the benefit ceiling is removed. The implied change in unemployment is calculated using a coefficient on the actual replacement of 0.05. The regression line is based on a weighted regression using regional population as weights.

What happens if we remove the benefit ceiling? Obviously this has the effect of making the system more generous – the actual replacement rate rises by 6.5 percentage points, on average. As a consequence, there is an increase in overall unemployment from 6.8 percent to 9.1 percent. What is more the spread of the regional unemployment distribution is reduced. This is illustrated in Figure 7, which relates the change in log unemployment – induced by the reform – to the log of the unemployment rate prior to the change. As the graph shows, there is a greater change in regions where unemployment was low initially; the slope of the regression line is negative with a *t*-ratio of 5. The intuition for this result is that the proposed policy change has a bigger effect in high-wage regions, which also tend to be low-unemployment regions. Thus, the policy change contributes to reducing unemployment differentials.

Now, what about raising the nominal replacement rate to 85 %? Again, this makes the system more generous and the actual replacement rate rises by 3.8 percentage points. Consequently, the aggregate unemployment rate rises to almost 8.3 percent. What is more, this policy change has the opposite effect on the regional distribution of unemployment in comparison to the change in the benefit ceiling. As the figure shows, the unemployment rate increases more in regions which were high-unemployment locations initially. The intuition is analogous to the previous case. High-unemployment regions tend to be low-wage regions. Consequently, changes in the statutory replacement rate have a bigger impact on the actual generosity of the UI system in these locations. Comparing the slope of the regression lines in Figures 7 and 8, we see that removing the benefit ceiling has a stronger differential impact across regions.

To sum up, the results of these policy simulations show that the design of the national UI system has repercussions on the regional labor market. Moreover, they concur with the simple model in section 2. The impact on the regional distribution of unemployment differs across the policy experiments. If UI is made more generous by raising the benefit ceiling this will compress unemployment differentials, while if generosity increases because of an increase in the statutory replacement rate this will exacerbate regional unemployment differences.

Figure 8: Policy simulation – Increasing the nominal replacement rate.

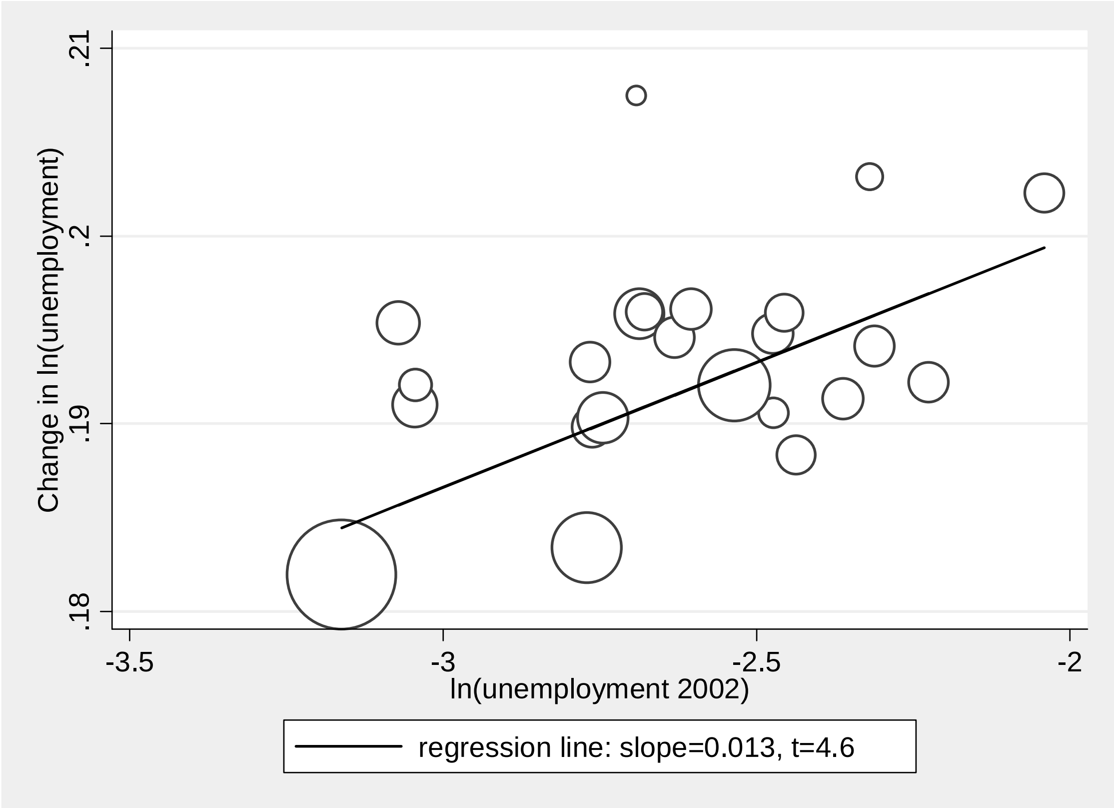

Notes: This graph is based on a hypothetical policy experiment where the nominal replacement rate is raised to 85 %. The implied change in unemployment is calculated using a coefficient on the actual replacement of 0.05. The regression line is based on a weighted regression using regional population as weights.

8. Conclusions

We have presented new evidence on the unemployment effects of increasing UI benefit generosity. The empirical strategy has been to utilize the fact that the nationally imposed benefit ceiling causes actual UI generosity to vary regionally. This paper has thus used variations in the national UI rules to estimate the effects at the regional level. Hence, the estimates should thus not suffer from the potential policy endogeneity hampering studies using regional policy changes for identification.

The evidence suggests that benefit generosity increases unemployment. We view this evidence as fairly robust since the estimates are similar across alternative specifications. The magnitudes involved are rather substantial and appear to be relatively high compared to estimates available elsewhere in the literature. The estimates suggest that an increase in the (actual) replacement rate of 5 percentage points contributes to increasing unemployment by 25 percent.

We have also shown that the benefit ceiling may contribute to exacerbating regional unemployment differentials. Lowering the ceiling reduces benefit generosity more in high-wage regions. Since high-wage regions also tend to be low-unemployment regions, the result follows. Moreover, a reduction in the statutory replacement rate has the opposite effect. Given that a benefit ceiling exists, a reduction in the statutory rate will reduce benefit generosity more in high-unemployment regions. Thus, these simple policy experiments illustrate that national rules in social protection systems can have (perhaps unintended) repercussions at the regional level.

References

Antelius, J and A Björklund (2000), "How Reliable are Register Data for Studies of the Return on Schooling? An examination of Swedish Data", Scandinavian Journal of Educational Research, Vol. 44.

Atkinson, A and J Micklewright (1991), "Unemployment Compensation and Labor Market Transitions: A Critical Review", Journal of Economic Literature, Vol. 29.

Card, D and P Levine (2000), "Extended Benefits and the Duration of UI Spells: Evidence from the New Jersey Extended Benefit Program", Journal of Public Economics, Vol. 78.

Carling, K, B Holmlund, and A Vejsiu (2001), "Do Benefit Cuts Boost Job Finding? Swedish Evidence from the 1990s", Economic Journal, Vol. 111.

Edin, P-A and P Fredriksson (2000), "LINDA – Longitudinal INdividual DAta for Sweden", Working Paper 2000:19, Department of Economics, Uppsala University.

Edin, P-A and B Holmlund (1995), "The Rise and Fall of Solidarity Wage Policy?" in R Freeman och L Katz (eds) Differences and Changes in Wage Structures, University of Chicago Press.

Erikson, R and R Åberg (1987), Welfare in Transition, Clarendon Press, Oxford.

Fredriksson, P (1999), "The Dynamics of Regional Labor Markets and Active Labor Market Policy: Swedish Evidence", Oxford Economic Papers, Vol. 51.

Fredriksson, P and B Holmlund (2005), "Improving Incentives in Unemployment Insurance: A Review of Recent Research", forthcoming in Journal of Economic Surveys.

Fredriksson, P and R Topel (2006), "Wage Determination and Employment in Sweden since the early 1990s", manuscript, Department of Economics, Uppsala University.

Holmlund, B (1998), "Unemployment Insurance in Theory and in Practice", Scandinavian Journal of Economics, Vol. 100.

Hornstein, A, P Krusell, and G Violante (2005), "Unemployment and Vacancy Fluctuations in the Matching Model: Inspecting the Mechanism", Federal Reserve Bank of Richmond Economic Quarterly, Vol. 91.

Katz, L and B Meyer (1990), "The Impact of the Potential Duration of Unemployment Benefits on the Duration of Unemployment", Journal of Public Economics, Vol. 41.

Krueger, A and B Meyer (2002), "The Labor Supply Effects of Social Insurance", in A J Auerbach and M Feldstein (eds) Handbook of Public Economics, Vol. 4, North-Holland, Amsterdam.

Lalive, R and J Zweimüller (2004), "Benefit Entitlement and Unemployment Duration: The Role of Policy Endogeneity", forthcoming in Journal of Public Economics.

Lindgren, B (2005), "Lönespridning i näringslivet 2004", Svenskt Näringsliv

Meyer, B (1995), "Lessons from the U.S. Unemployment Insurance Experiments", Journal of Economic Literature, Vol. 33.

Nickell, S (1998), "Unemployment: Questions and some Answers", Economic Journal, Vol. 108.

Nilsson, C (1993), "The Swedish Model: Labour Market Institutions and Contracts", in J Hartoog and J Theeuwes (eds) Labor Market Contracts and Institutions, Elsevier, Amsterdam.

OECD (2004), Employment Outlook, Organization of Economic Cooperation and Development, Paris.

Olli Segendorf, Å (2003), "Arbetsmarknadspolitiskt kalendarium II", Rapport 2003:9, Institute for Labour Market Policy Evaluation, Uppsala.

Pigou, A (1932), The Theory of Unemployment, MacMillan, London.

Pissarides, C (2000), Equilibrium Unemployment Theory, MIT Press.

Sebardt, G (2005), Redundancy and the Swedish Model – Swedish Collective Agreements on Employment Security in a National and International Context, Iustus Förlag, Uppsala (PhD Thesis).

Wadensjö, E (1993), "Socialförsäkringssystemets okända delar", in A Lindbeck, T Persson, A Sandmo, and B Swedenborg (eds) Nya villkor för ekonomi och politik, SOU 1993:16, Ministry of Finance, Stockholm.

Appendix A: Creating regional panel data

This appendix describes the construction of the regional panel data. Regional labor force composition, predicted wages and replacement rates are calculated from individual data. We use LINDA, a 3.35% representative sample of the Swedish population; see Edin & Fredriksson (2000). From this register, we select all individuals between the ages of 16-59, from 1970 to 2002.[27] In the early 1970s data contain roughly 130 000 individuals per year; in 2002 about 150 000. LINDA has a panel dimension which is very useful when constructing the data. If information is missing in one year, we can check if this information is available at another time point. This panel structure of the data is extremely valuable when comes to impute missing information on educational attainment as discussed below.

The regions correspond to the counties of Sweden. Between 1970 and 1996 there were 24 counties in Sweden. In 1996 two counties were merged, and in 1997 another three counties were merged.[28] Hence, from 1997 and onwards, there are 21 counties in Sweden. Since we also have data at the municipality level we can reconstruct the original 24 counties. We have thus used the municipality data to split the merging counties – thus creating 24 regions for the full time period.

The individual characteristics used in this paper are standard. Gender is identified by a female dummy. We define a set of age-group dummies for each five-year interval; the youngest group thus contains individuals aged 16 to 20, and the oldest group contains the ages 56 to 59. Marital status identifies married individuals. Immigrants are divided into three groups depending on the country of birth. We identify three groups: Nordic, Oecd, and non-Oecd immigrants. The definition of Nordic ancestry is obvious; the categorization into Oecd and non-Oecd immigrants is perhaps less obvious. We have used the following rule: individuals are classified as being of OECD origin if they were born in a country which was a member of the OECD in 1985. The only exceptions from this rule are Turkey – which is included among the non-OECD countries – and, of course, the Nordic countries.

Industry affiliation is defined by two-digit ISIC-codes, generating 33 industry dummies. The coding changed in 1993, but at the two-digit level it is possible to link the two coding systems. However, information on industry affiliation is missing for four years: 1974, 1976, 1977, and 1979. To deal with this issue we use the following simple rule: the information

[27] We have to restrict the analysis to individuals younger than 60, since educational information is not available for those older than 59 for the full time period.

[28] In 1996, the county of Skåne was created by merging the counties of Malmöhus and Kristianstad. In 1997, the county of Västra Götaland was created by a merger of the counties of Älvsborg, Göteborg och Bohuslän, and Skaraborg.

observed in 1975 is used for the individual also in 1974 and 1976; analogously, the information observed in 1978 is used also in 1977 and 1979.

Educational attainment is divided into three categories: compulsory, secondary, and tertiary schooling. Starting in 1991 educational information is available each year. Prior to 1991 we only observe educational attainment at two time points: in 1970 and 1990. We have used the following procedure to attach educational information to the individuals during 1971-1989. If an individual is at least 25 years-of-age in 1970, education is assumed to be completed and the observation from 1970 is used to fill out the missing information during 1971-1989. If the individual is younger than 25 in 1970, we use data from 1990. Different rules are used depending on educational attainment in 1990 and age at the time point of observation. For an individual who has completed tertiary education, we assign the level of attainment should this individual turn 25 during 1971-1989. Should this individual turn 21 during this time period he or she assigned secondary schooling and when the individual is below age 21 he or she is assigned compulsory schooling. For an individual who has completed secondary schooling in 1990, we use this attainment level from the point when the individual turns 21 and onwards. Prior to turning 21, compulsory schooling is used has the highest attainment level. An individual who had completed compulsory schooling in 1990 is classified as having attained compulsory schooling from the time point when he or she enters our data.

Our key measures (expected wages and actual replacement rates) are constructed using earnings and wage information as described in the text, as well as the UI rules described below.

Finally, the regional panel is constructed by averaging over all individuals residing in a particular region. This gives us annual information on the composition of the regional population as well as the key explanatory variables of interest. Ideally, we would have liked to calculate the characteristics of the regional labor force. But this was not possible since there was no indicator of labor force status in our data. However, the measurement error involved is likely to be small.

To these regional panel data we match information on unemployment. Regional unemployment data are defined for the age-category 16-64; they are collected from the Labor Force Surveys and the Labor Market Board. As the measure of unemployment we use the sum of open unemployment and participants in labor market programs as a share of the labor force.

Unemployment insurance

The design of the public unemployment insurance system has varied somewhat over time. There are two distinct time-periods – the first covers the period from 1974 to 1988, and the second the period 1989-2002.

Between 1974 and 1988, individuals were sorted into different benefit levels depending primarily on how much they earned. The various UI funds used different benefit ceilings. There was a national benefit ceiling, however, and the replacement rate could never exceed 91.7 % of previous income. This implies that the maximum benefit level varied between individuals, depending on which particular UI fund the individuals were members of. Since we cannot observe membership in a particular UI fund, we use the "average maximum benefit level" as a proxy for the maximum level. This measure is reported in the Annual Financial Report of the Labor Market Board; it is calculated as a weighted average over individuals, where the weights are based on the number of members in a particular UI fund.

From 1989 and onwards, the construction of the UI system is more straightforward. An unemployed individual then receives a certain amount (in percent) of the previous wage, up to a maximum level.

Even though the design of the system has varied somewhat over time, we implement the rules in essentially the same way. An individual receives a benefit equal to the nominal replacement rate multiplied by foregone income, but the benefit can never exceed the ceiling. The ceiling is here defined from 1974-1988 by the "average maximum benefit level" and from 1989-2002 as the "maximum benefit".

Table A1 displays the benefit levels (in SEK per day), and the nominal replacement rates from 1974 to 2002 as observed on December 31st each year.[29] Column (1) displays the national benefit ceiling. Remember that this variable is only used as a measure of the benefit ceiling from 1989 to 2002.[30] As described above, the average maximum benefit level in column (2) is used between 1974 and 1988. Note that prior to 1977, the difference between the average maximum benefit level and the national benefit ceiling is substantial. But from then on the differences across UI funds become smaller, to eventually disappear completely.

Column (3) reports the date when the benefit ceiling was changed. Typically, this was not at the beginning of a calendar year. Hence, the annual benefit ceiling used in this paper is calculated in column (4) using the information in columns (1) to (3). For example, there is a

[29] Information is taken from the annual reports of the Labor Market Board.

[30] From 2001 and onwards, the benefit ceiling decreases after 100 days of unemployment, as shown by the figures in the brackets. In this paper, we use only the maximum during the first 100 days.

change in the maximum benefit on July 1st 1979, and then the calendar year benefit ceiling (179.39), is simply calculated as the mean of the average benefit values observed on December 31st 1978 (171.16) and December 31st 1979 (187.62).

Column (5) shows the nominal replacement rate, column (6) reports when it was changed, and column (7) the annual average of the nominal replacement rate used in the analysis. Hence, the bold figures in column (4) and (7) are the primary input in our analysis. From these data, we calculate the wage cap by dividing the benefit ceiling with the nominal replacement rate. The wage cap and the nominal replacement rate are shown in Figure 2.

Table A1. The unemployment insurance system, 1974-2002.

Year	National benefit ceiling SEK/day (Dec 31) (1)	Average benefit ceiling SEK/day (Dec 31) (2)	Date of change in ceiling (3)	**Benefit ceiling SEK/day (calendar year) (4)**	Nominal repl. rate (Dec 31) (5)	Date of change in nominal repl. rate (6)	**Nominal repl. rate (calendar year) (7)**
1974	130	98,07		**98.07**	0.917		**0.917**
1975	130	116,58		**116.58**	0.917		**0.917**
1976	160	122,22	July 1st	**119.4**	0.917		**0.917**
1977	160	151,76		**151.76**	0.917		**0.917**
1978	180	171,16	July 1st	**161.46**	0.917		**0.917**
1979	195	187,62	July 1st	**179.39**	0.917		**0.917**
1980	195	192,19		**192.19**	0.917		**0.917**
1981	210	206,80	April 1st	**203.1475**	0.917		**0.917**
1982	230	227,66	July 1st	**217.23**	0.917		**0.917**
1983	280	278,80	Jan 1st	**278.8**	0.917		**0.917**
1984	300	298,87	July 1st	**288.835**	0.917		**0.917**
1985	315	314,48	July 1st	**306.675**	0.917		**0.917**
1986	360	359,20	July 1st	**336.84**	0.917		**0.917**
1987	400	400	July 1st	**379.6**	0.917		**0.917**
1988	425	425	July 4th	**412.5**	0.917		**0.917**
1989	450		Jan 2nd	**450**	0.9	Jan 2nd ; 90%	**0.9**
1990	495		Jan 1st	**495**	0.9		**0.9**
1991	543		Jan 7th	**543**	0.9		**0.9**
1992	564		Jan 6th	**564**	0.9		**0.9**
1993	564		Jan 4th ; 598 July 5th ; 564	**581**	0.8	July 5th ; 80%	**0.85**
1994	564			**564**	0.8		**0.8**
1995	564			**564**	0.8		**0.8**
1996	564			**564**	0.75	Jan 1st ; 75%	**0.75**
1997	580		Dec 29th	**564**	0.8	Sep 29th ; 80%	**0.7625**
1998	580			**580**	0.8		**0.8**
1999	580			**580**	0.8		**0.8**
2000	580			**580**	0.8		**0.8**
2001	680 (580)		July 2nd	**630**	0.8		**0.8**
2002	730 (680)		July 1st	**705**	0.8		**0.8**

Note: Column (4) is based on cols. (2) and (3) during 1974-88, and cols. (1) and (3) during 1989-2002. Column (7) is based on columns (5) and (6) throughout the time period.

Appendix B: Coefficient estimates from baseline specification

Table B1 shows coefficient estimates corresponding to our preferred specification, that is, column (3) in Table 3.

Table B1. Estimates on a selection of observed regional control variables.

	Coefficient estimate (standard error)
Covered by collective agreement with supplementary UB	1.74** (0.637)
Female	2.58 (2.01)
Married	-1.99* (1.12)
Nordic	-0.784 (3.45)
Oecd	-3.57 (7.51)
non-Oecd	-4.79 (3.20)
Secondary schooling	-5.06** (1.84)
Tertiary schooling	-3.85* (2.13)
Age 21-25	3.26 (2.24)
Age 26-30	4.26 (2.84)
Age 31-35	6.89** (3.16)
Age 36-40	8.03** (3.16)
Age 41-45	8.68** (2.58)
Age 46-50	9.05** (3.02)
Age 51-55	11.63* (3.31)
Age 56-59	10.51* (3.04)
Region-specific fixed effects	Yes
Region-specific trends	Yes
Time effects	Yes
Overall R^2	0.981
# observations	696

Note: The regressions also include a constant and industry employment shares. Regressions are weighted by population. Standard errors, reported in parentheses, are clustered by county. Significance levels: * = 10%, and ** = 5%.