



IFS

WAGE RISK AND EMPLOYMENT RISK OVER THE LIFE CYCLE

*Hamish Low
Costas Meghir
Luigi Pistaferri*



THE INSTITUTE FOR FISCAL STUDIES
WP08/06

Wage Risk and Employment Risk over the Life Cycle*

Hamish Low[†], Costas Meghir[‡] and Luigi Pistaferri.[§]

September 2008

Abstract

We specify a structural life-cycle model of consumption, labour supply and job mobility in an economy with search frictions that allows us to distinguish between different sources of risk and to estimate their effects. The sources of risk are shocks to productivity, job destruction, the process of job arrival when employed and unemployed and match level heterogeneity. Our model allows for four main social insurance programmes. In contrast to simpler models that attribute all income fluctuations to shocks, our framework allows us to disentangle the effects of the shocks from the responses to these shocks. Estimates of productivity risk, once we control for employment risk and for individual labour supply choices, are substantially lower than estimates that attribute all wage variation to productivity risk. Increases in productivity risk impose a considerable welfare loss on individuals and induce substantial precautionary saving. Increases in employment risk have large effects on output and, primarily through this channel, affect welfare. The welfare value of government programs such as food stamps which partially insure productivity risk is greater than the value of unemployment insurance which provides (partial) insurance against employment risk and no insurance against persistent shocks.

Keywords: uncertainty, life-cycle models, unemployment, precautionary savings

JEL Classification: D91, H31, J64

*We are grateful to the editor Richard Rogerson, two anonymous referees, Joe Altonji, Gadi Barlevy, Marco Bassetto, Audra Bowlus, David Card, Thomas Crossley, James Heckman, Bruce Meyer, Dale Mortensen, Giuseppe Moscarini, Chris Pissarides, Victor Rios-Rull, Chris Taber, Robert Townsend, Gianluca Violante, Randy Wright, conference participants at the NBER Summer Institute, Society of Economic Dynamics, Stanford Institute for Theoretical Economics, the 2006 Cowles Macro-Labor Conference, the “Macroeconomics of Imperfect Risk Sharing” conference in Santa Barbara, as well as participants in numerous seminars and workshops. Pistaferri acknowledges funding from NSF through grant 0453119. We thank the ESRC Centre for Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies. Costas Meghir also thanks the ESRC for funding under the ESRC Professorial Fellowship RES-051-27-0204.

[†]University of Cambridge and IFS.

[‡]University College London and IFS.

[§]Stanford University, SIEPR and CEPR.

1 Introduction

There is an extensive literature analyzing individuals' precautionary response to income risk under incomplete markets. The theoretical literature has clarified the circumstances under which precautionary behavior arises. Empirical analysis has concentrated on assessing the levels of income risk and the persistence of shocks;¹ on showing that insurance markets are indeed incomplete;² and on measuring the effects of uninsurable idiosyncratic risk on life-cycle consumption profiles and wealth accumulation.³ In most studies idiosyncratic risk is identified as the variance of an appropriately defined residual in a panel data model of income but the underlying sources of risk are not distinguished and exogenous shocks are not disentangled from the effects of actions (such as changes in labour supply and job mobility) taken in response to such shocks. While we have learned a lot by this first generation of models, to obtain a better understanding of individual responses to risk and to carry out policy analysis, it is necessary to go deeper and to understand the sources of risk and to recognize that many of the observed fluctuations are the result of endogenous choices.

In order to disentangle shocks from the responses to the shocks, we specify a structural life-cycle model of consumption, labour supply and job mobility. We then specify the underlying sources of shocks that are the key drivers of observed fluctuations in earnings. These include shocks to individual productivity that persist across different jobs and across time; firm level shocks that lead to job destruction; the stochastic process of job offers when employed and unemployed; and variation in the quality of the match offered. Our model captures how these basic underlying shocks transmit into observed behavior, welfare and earnings fluctuations. Without the labour supply and job mobility choices, we would obtain a misleading picture on the extent to which individuals can self-insure and the extent to which observed earnings fluctuations reflect risk.⁴

Within our framework, we can distinguish between employment risk and productivity risk. Productivity risk is individual-specific uncertainty which exists independently of the employer's characteristics. Employment risk captures the uncertainty about having a job and also about the firm type. This includes the possibility of firm closure or job destruction, the difficulty of finding a new

¹For example, see MaCurdy (1982), Abowd and Card (1989), and Meghir and Pistaferri (2004), Guvenen (2007).

²See Cochrane (1991) and Attanasio and Davies (1996).

³See amongst others Zeldes (1989), Deaton (1991), Carroll (1992), Carroll and Samwick (1997), Banks, Blundell and Brugiavini (2001), Gourinchas and Parker (2002), and Attanasio, Banks, Meghir and Weber (1999).

⁴Krussell et al. (2008) highlight the importance of modelling labour market frictions alongside labour supply choices in understanding the aggregate implications of incomplete markets.

job match while unemployed, and the extent of unobserved heterogeneity across firms. In a fully competitive labor market with no worker-firm match heterogeneity and no search costs, the distinction between employment and productivity risk would be meaningless because unemployment would arise only due to low productivity resulting in the individual's market wage being below the reservation wage. Unemployment itself would not be a source of risk.⁵

Shocks differ in their available insurance opportunities. For example, layoffs are usually partially insured by the unemployment insurance system, while individual productivity shocks, other than major observable health shocks, are rarely insured in any formal way because of moral hazard and limited enforcement and commitment reasons. It is precisely this lack of formal insurance that prompts prudent individuals to engage in precautionary behavior. Furthermore, the individual's response to earnings risk will depend partly on the availability of outside insurance - private or public. With few exceptions (Hubbard, Skinner and Zeldes, 1995), the literature on precautionary savings has assumed that only self-insurance is available. In this paper, we propose a model in which people can self-insure, but may also be eligible for government provided insurance mirroring three popular programs in the US: Unemployment Insurance (UI), Disability Insurance (DI), and Food Stamps. These systems provide partial insurance only.

The parameters of our model are obtained partly from estimating the characteristics of the wage process with endogenous participation and mobility choices, and partly from calibrating our life cycle model to fit observed participation profiles and unemployment durations. We use longitudinal wage and job mobility data from the Survey of Income and Program Participation (SIPP) and unemployment duration data from the Panel Study of Income Dynamics (PSID).

The empirical results we report relate to the nature of the income process and the basic implications of the model. First, we show that our preferred stochastic process for income (the sum of a random walk, an i.i.d. component, and a firm-worker match fixed effect) provides a good fit of the key facts in the data. Second, we find that if mobility is ignored the estimated variance of the permanent innovation to wages doubles, leading to an impression of much greater risk in the earnings process. This is because much of the wage fluctuations are due to individuals moving to

⁵Some recent papers have analysed the joint precautionary saving-labour supply decision. Low (2005), Pijoan-Mas (2006), and Domeij and Floden (2006) analyze the joint saving and labour supply decision, but in a context without exogenous job destruction, search frictions or job mobility. French (2004) analyses labour supply and savings behaviour, but focuses on older workers. Heathcote et al. (2007) consider a joint saving-labor supply decision, again without frictions, and focusing on the degree of partial insurance.

jobs with better match specific effects; ignoring this biases measured uncertainty upwards.

Turning to counterfactual experiments, we assess the effects of different types of risk by varying some key parameters one at a time (including the variance of productivity risk, and the job destruction rate) and reporting the change in labour supply, output and savings. We also compute the willingness to pay to avoid the various changes in risks.

When productivity risk increases, individuals are worse-off because of the increased risk and also because output declines. However, individuals are willing to pay substantially more than the output loss to compensate for the increased risk. When job destruction increases, output also declines and unemployment increases, as we would expect. The environment becomes riskier as highly valued jobs can be lost at a faster rate, but the welfare effects of this risk are mitigated by the utility value of leisure (which in our model is a substitute for consumption). Overall, although welfare falls as job destruction increases, the willingness to pay to return to the original lower rate of job destruction is *less* than the loss in output.

Finally, we measure the value that people assign to an increase in the various government provided insurance programs in our model, and compare this to the value of a revenue equivalent cut in proportional taxes. The welfare value of programs such as food stamps, which partially insure productivity risk, is greater than the value of unemployment insurance which provides (partial) insurance against employment risk and no insurance against persistent shocks. The relatively low value of unemployment insurance is in line with the results of Hansen and Imrohoroglu (1992).⁶

The layout of the paper is as follows. Section 2 presents the model and discusses the distinction between employment and productivity risk. Section 3 describes the data. Section 4 describes the estimation strategy and results for estimating the wage process. Section 5 presents the calibration process for the remaining parameters. Section 6 discusses possible alternatives to our wage process. Section 7 presents our calculations of the behavioral effects and the welfare costs of uncertainty and the welfare benefit of government insurance programs. Section 8 concludes.

⁶Lentz (2003) also analyzes the value of unemployment insurance, allowing for the interaction between search frictions and saving. Costain (1999) proposes an equilibrium search model with precautionary savings that attempts to measure the welfare effects of unemployment insurance. Rendon (2006) examines the relationship between wealth accumulation and job search dynamics in a model where the motivation for accumulating wealth is to finance voluntary quits in order to search for a better job. However, all these papers, along with Hansen and Imrohoroglu (1992), ignore the risk to individuals' own productivity which is independent of any particular match.

2 Model

2.1 Overview

We specify a model where individuals choose consumption and make work decisions so as to maximize an intertemporal utility function, in an environment with search frictions. We view the key sources of shocks underlying earnings fluctuations as being shocks to individual productivity, firm level shocks leading to job destruction, the process of job offers when unemployed and when employed, and the quality of the match offered. Thus individuals face multiple sources of uncertainty: in each period employed individuals may be laid off or may receive offers of alternative employment; unemployed workers may or may not be offered a job; all individuals face uninsurable shocks to their productivity.

The economy offers partial social insurance in the form of a number of programs. These are Food Stamps, Unemployment Insurance, Disability Insurance and Social Security (pensions). When simulating the model, changes to these programs are funded by proportional taxation; thus individuals are linked through the government budget constraint. The model has numerous sources of dynamics. These include asset accumulation, the fact that job offer probabilities are state dependent and that current actions affect future eligibility to the various programs. We consider two types of individuals separately: the lower educated individuals, which include all those with a high school diploma or less, and the higher educated individuals with at least some college.

In this section we start by describing the stochastic process of wages. Then we describe the process of job arrival and job destruction. With the sources of shocks specified we then describe the individual optimization problem and the distinction between employment and productivity risk. The empirical analysis follows in the next sections.

2.2 Structure of Wages and Shocks

We begin the model specification by outlining the process for wages. We assume that wages w_{it} in the data are governed by the process:

$$\ln w_{it} = d_t + x'_{it}\psi + u_{it} + e_{it} + a_{ij}(t_0) \quad (1)$$

where w_{it} is the real hourly wage, d_t represents the log price of human capital at time t , x_{it} a vector of regressors including age, u_{it} the permanent component of wages, and e_{it} measurement error.⁷ All

⁷More generally e_{it} represents a mix between a transitory shock and measurement error. In the usual decomposition of shocks into transitory and permanent components, researchers work with annual earnings data where transitory

parameters of the wage process are education specific.

The term $a_{ij(t_0)}$ denotes a firm-worker match specific component where $j(t_0)$ indexes the firm that the worker joined in period $t_0 \leq t$.⁸ It is drawn from a normal distribution with mean zero and variance σ_a^2 . We model the match effect as constant over the life of the worker-employer relationship, and so if the worker does not change employer between t and $t + 1$, there is no wage growth due to the match effect. If the worker switches to a different employer between t and $t + 1$, however, there will be some wage growth which we can term a mobility premium. In this case we define the random variable $\xi_{it+1} = a_{ij(t+1)} - a_{ij(t_0)}$ as the wage growth due to inter-firm mobility between t and $t + 1$. Since offers can be rejected when received, only a censored distribution of ξ_{it+1} is observed. The match effect $a_{ij(\cdot)}$ is complementary to individual productivity. It is constant over time but it will be assumed uncertain across firms.⁹ Both the match effect and the idiosyncratic shock have education-specific distributions. The information structure is such that workers and firms are completely informed about u_{it} and $a_{ij(\cdot)}$ when they meet (jobs are “search goods”). The importance of match effects in explaining wages has been stressed by Topel and Ward (1992) and Abowd, Kramarz and Margolis (1999). Postel-Vinay and Robin (2002) show in an equilibrium setting how firm and individual heterogeneity translate into a match effect.

Finally, we assume that there are constant returns to scale in labor implying that the firm is willing to hire anyone who can produce non-negative rents. However, we assume the firm does not respond to outside offers. If firms did respond, this would imply that the match specific effect would increase each period with some probability and would manifest itself as a return to job tenure. However, returns to tenure are thought to be small, once one controls for endogeneity of job mobility.¹⁰ This provides some evidence that outside offers are not an important source of wage

shocks may well be important because of unemployment spells. In our framework, this source of transitory shocks is modelled explicitly through the unemployment process, through the choice about whether or not to work and through job mobility. For this reason, attributing the transitory shock entirely to measurement error seems appropriate. Further, in the empirical section we find that the variance of e_{it} is low, and indeed lower than the variance of measurement error obtained on *annual earnings* by validation studies on the SIPP data we use (see Abowd and Stinson, 2005).

⁸We should formally have a j subscript on wages but since it does not add clarity we have dropped it. Note also that in the absence of firm data one cannot distinguish between a pure firm effect and a pure match effect. In the latter case, one can imagine $\alpha_{ij(t_0)}$ as being the part of the matching rent that accrues to the worker. We take the bargaining process that produces this sharing outcome as given.

⁹Ideally we would like to allow also for shocks to the match effect. These will act as within firm aggregate shocks. Restricting match effects to be constant is forced upon us by the lack of matched firm and individual data. In section 6, we consider the alternative assumption of modelling individual productivity as a fixed effect and the match component as stochastic.

¹⁰Altonji and Williams (2005) assess this literature and conclude that their preferred estimate for the US is a return to tenure of 1.1% a year.

growth on the job. While dealing with the effect of outside offers may be interesting, we leave this for future research.

We assume that the permanent component of wages follows a random walk process:

$$u_{it} = u_{it-1} + \zeta_{it} \tag{2}$$

where ζ_{it} is a normally distributed random shock with mean zero and variance σ_{ζ}^2 . We assume this shock reflects uncertainty.¹¹

Given a particular level of unobserved productivity, the worker will be willing to work for some firms but not for others, depending on the value of the match. We assume that the measurement error e_{it} is normally distributed with variance σ_e^2 . As far as the policy implications of the model are concerned, we are interested in estimating σ_a^2 and σ_{ζ}^2 . We describe later how these are identified and estimated.

The specification we presented, while consistent with much of the evidence and in line with a number of papers,¹² is not uncontroversial. Two alternatives might be a model with a stationary AR(1) process with a fixed effect in wage growth or a model where the match specific effect is itself stochastic. We discuss these alternatives in Section 6 and justify our choice.

2.3 Job destruction and job arrival rates

In each period workers receive an alternative job offer with probability λ^e . Those who are currently unemployed receive an offer with probability λ^n . Individuals become unemployed either because they choose to quit following particular wage realizations or because of exogenous job destruction, which happens each period with probability δ . The friction parameters (as well as the variance parameters discussed earlier) are all assumed to be specific to an education group.

The composition of those becoming unemployed is not random in our model, despite the fact that the job destruction rate acts as a random shock independently of individual skill levels. First, people with bad productivity shocks will quit their jobs and the extent to which this happens depends on the variance of the wage innovations. Second the job destruction rates can differ by education group.

¹¹An issue is how much of the year to year variability of wages reflects uncertainty. A large component of this variability is measurement error, which we control for. Beyond that, primarily for lack of adequate data, we abstract from the important issues that have to do with consumers having superior information vis-à-vis the econometrician (For discussions and empirical analysis see Blundell and Preston (1998), Manski (2004) Pistaferri (2001, 2003) and Cuhna, Heckman and Navarro (2004).

¹²See MaCurdy (1982), Topel (1991), Abowd and Card (1989), Moffitt and Gottschalk (2001), and Meghir and Pistaferri (2004) .

Thus there is selection into the unemployment pool both in terms of observable and unobservable skill characteristics, and this selection means those becoming unemployed are less productive on average than the employed.

We assume there is no exogenous depreciation of skills following job loss. Instead, the loss of the particular match on entering unemployment may lead to wages on re-entry being lower because the new firm will on average have a lower match value. This is the case because individuals in work will have improved over the average offer through job mobility, before a job in which they are employed is destroyed.¹³ Thus firm heterogeneity implies that exogenous job destruction will lead to wage losses and appear as scarring, which we document in the empirical analysis below.

We assume that job destruction and job offer arrival rates are constant over time and so we ignore business cycle effects. We focus instead on the implications of idiosyncratic risk for behavior and for welfare. By contrast, Lucas (1987) and others focus on the welfare benefit of smoothing out the aggregate business cycle risk, and Storesletten et al. (2001) focus on smoothing out the variation in idiosyncratic risk. In our comparative static analysis, however, we show the effects of different values of job destruction and job offer arrival rates across a range consistent with the variation observed over the business cycle.

2.4 Individual Optimization

We consider an individual with a period utility function

$$U_t = U(c_{it}, P_{it})$$

where P is a discrete $\{0, 1\}$ labor supply participation variable and c is consumption. The individual is assumed to maximize lifetime expected utility,

$$\max_{c, P} V_{it} = E_t \sum_{s=t}^L \beta^{s-t} U(c_{is}, P_{is})$$

where β is the discount factor and E_t the expectations operator conditional on information available in period t (a period being a quarter of a year). Individuals live for L periods, may work from age 22 to 62, and face an exogenous mandatory spell of retirement of 10 years at the end of life. The date of death is known with certainty.

¹³Indeed, as stated by Jacobson, LaLonde and Sullivan (1993), “workers possessing skills that were especially suited to their old positions are likely to be less productive, at least initially, in their subsequent jobs. Such a fit between workers’ skills and the requirements of their old jobs could have resulted from on-the-job investment in firm-specific human capital or from costly search resulting in particularly good match with their old firms”.

The labour supply choice in our model is a discrete choice. However, since one period is one quarter, this discrete choice can generate substantial variation in annual hours of work.¹⁴ The worker's problem is to decide whether to work or not and, if the opportunity arises, whether to switch firm. When unemployed he has to decide whether to accept a job that may have been offered or wait longer. If eligible, the unemployed person will have the option to apply for disability insurance. There is a fixed and known probability of being successful, conditional on applying. Whether employed or not, the individual has to decide how much to save and consume. Accumulated savings can be used to finance spells out of work and early retirement.

We use a utility function of the form

$$U(c, P) = \frac{(c \cdot \exp\{\eta P\})^{1-\gamma}}{1-\gamma}$$

We consider cases where $\gamma > 1$ and $\eta < 0$, implying that individuals are reasonably risk averse, participation reduces utility and that consumption and participation are Frisch complements (i.e. the marginal utility of consumption is higher when participating).

The intertemporal budget constraint during the working life has the form

$$A_{it+1} = R [A_{it} + (w_{it}h(1 - \tau_w) - F_{it}) P_{it} + (B_{it}E_{it}^{UI} (1 - E_{it}^{DI}) + D_{it}E_{it}^{DI}) (1 - P_{it}) + T_{it}E_{it}^T - c_{it}] \quad (3)$$

where A are beginning of period assets, R is the interest factor, w the hourly wage rate, h a fixed number of hours (corresponding to 500 hours per quarter), τ_w a proportional tax rate that is used to finance social insurance programs, F the fixed cost of work, B_{it} unemployment benefits, T_{it} the monetary value of food stamps received, D_{it} the amount of disability insurance payments obtained, and E_{it}^{UI} , E_{it}^{DI} , and E_{it}^T are reciprocity $\{0, 1\}$ indicators for unemployment insurance, disability insurance, and the means-tested transfer program, respectively. Note also that there are costs of applying for disability insurance which we discuss below.

We assume that individuals are unable to borrow either against the social insurance programs

¹⁴In the data, the variation in annual hours is predominantly due to changes in participation status during the year. By using a quarter as the decision time we generate quite a lot of potential variation over the year. Hours elasticities for workers are found to be very small in most empirical microeconomic studies for men; see MaCurdy et al. (1990), Pencavel (2002) and Meghir and Phillips (2008) as examples.

or against future earnings:

$$A_{it} \geq 0$$

In practice, this constraint has bite because it precludes borrowing against unemployment insurance, against disability insurance, against social security and against the means-tested program.

At retirement, people collect social security benefits which are paid according to a formula similar to the one we observe in reality (see below). These benefits, along with assets that people have voluntarily accumulated over their working years, are used to finance consumption during retirement.

Unemployment Insurance We assume that unemployment benefits are paid only for the quarter immediately following job destruction. We define eligibility for unemployment insurance E_{it}^{UI} to mirror current legislation: benefits are paid only to people who have worked in the previous period, and only to those who had their job destroyed (job quitters are therefore ineligible for UI payments, and we assume this can be perfectly monitored).¹⁵ We assume $B_{it} = b \times w_{it-1} \bar{h}$, subject to a cap, and we set the replacement ratio $b = 75\%$. The replacement ratio is set at this high value because the payment that is made is intended to be of a similar magnitude to the maximum available to someone becoming unemployed. The cap is set according to the median state (Meyer, 2002).

In the US, unemployment benefit provides insurance against job loss and insurance against not finding a new job. However, under current legislation benefits are only provided up to 26 weeks (corresponding to two periods of our model) and so insurance against not finding a new job is limited. Our assumption is that there is no insurance against the possibility of not receiving a job offer after job loss. This simplifying assumption means that, since the period of choice is one quarter, unemployment benefit is like a lump-sum payment to those who exogenously lose their job and so does not distort the choice about whether or not to accept a new job offer. The only distortion is introduced by the tax on wages, used to finance UI.

¹⁵We have simplified considerably the actual eligibility rules observed in the US. A majority of states have eligibility rules which are tougher than the rule we impose, both in terms of the number of quarters necessary to be eligible for any UI and in terms of the number of quarters of work necessary to be eligible for the maximum duration (Meyer, 2002). However, making eligibility more stringent in our model is numerically difficult because the history of employment would become a state variable. Our assumption on eligibility shows UI in its most generous light.

Universal Means-Tested Program In modelling the universal means-tested program, our intention was to mirror partially the actual food stamps program but with three simplifying differences. First, the means-testing is only on income rather than on income and assets;¹⁶ second, the program provides a cash benefit rather than a benefit in kind;¹⁷ and third, we assume there is 100% take-up. These assumptions mean that in our model there is no disincentive for poor individuals to hold existing assets (as in Hubbard, Skinner and Zeldes, 1995); there is still however the disincentive to accumulate caused by the programs, as the public insurance will lead to a lower need for precautionary savings.

For the purposes of the program gross income is defined as

$$y_{it}^{gross} = w_{it}hP_{it} + (B_{it}E_{it}^{UI}(1 - E_{it}^{DI}) + D_{it}E_{it}^{DI})(1 - P_{it}) \quad (4)$$

giving net income of $y = (1 - \tau_w)y^{gross} - d$, where d is the standard deduction that people are entitled to when computing net income for the purpose of determining food stamp allowances. The value of the program is then given by

$$T_{it} = \begin{cases} \bar{T} - 0.3 \times y_{it} & \text{if } y_{it} \leq \underline{y} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

where \bar{T} is the maximum payment and where \underline{y} should be interpreted as a poverty line. In the actual food stamp program, only people with net earnings below the poverty line are eligible for benefits (which we denote by $E_{it}^T = 1$). The maximum value of the payment, \bar{T} , is set assuming a household with two adults and two children, although in our model there is only one earner.

Disability Benefits and Social Security Workers may find themselves in circumstances that would lead them to apply for disability insurance, the final element of the budget constraint. First, we allow only individuals who face a negative productivity shock to apply for disability. The requirement of a negative shock to wages is meant to mimic a health shock, on the basis of which an individual could claim to be eligible. Second, we require people to remain unemployed for at least one quarter before being able to apply for disability insurance and then they must remain

¹⁶The difficulty with allowing for an asset test in our model is that there is only one sort of asset which individuals use for retirement saving as well as for short-term smoothing. In reality, the asset test applies only to liquid wealth and thus excludes pension wealth (as well as real estate wealth and other durables).

¹⁷We assume that the means-tested transfer is paid in cash rather than in the form of coupons (as with Food Stamps). While this is in contrast with the reality, it would be of little practical importance if stamps were inframarginal or if there were “trafficking”. Moffitt (1989) finds evidence for both phenomena.

unemployed in the quarter that the application is made. Again, this is meant to reflect the actual rules of the system: there is a waiting period of 5 months between application and receipt of benefits, and during this period the individual must be unemployed. Third, we assume that only workers above the age of 50 are eligible to apply for disability benefits.¹⁸

Conditional on applying, individuals have a fixed probability of obtaining the benefit, which we obtain from actual data (50%, see Bound et al., 2004). If successful, the individual remains eligible for the rest of their working life and disability insurance becomes an absorbing state. If not successful, the individual has to remain unemployed another quarter before taking up a job. Individuals can only re-apply in a subsequent unemployment spell. The combination of disability and the means-tested program turns out to be very important in fitting the declining labor force participation profiles with age. Disability payments can provide a high replacement rate which is not affected by the duration of unemployment. However, the requirement that individuals spend two quarters unemployed before the disability application is resolved would discourage a large proportion of applicants were it not for the means-tested (food stamps) program which provides a floor to income during this application process.

The value of disability insurance is given by

$$D_{it} = \begin{cases} 0.9 \times \bar{w}_i & \text{if } \bar{w}_i \leq a_1 \\ 0.9 \times a_1 + 0.32 \times (\bar{w}_i - a_1) & \text{if } a_1 < \bar{w}_i \leq a_2 \\ 0.9 \times a_1 + 0.32 \times (a_2 - a_1) + 0.15 \times (\bar{w}_i - a_2) & \text{if } a_2 < \bar{w}_i \leq a_3 \\ 0.9 \times a_1 + 0.32 \times (a_2 - a_1) + 0.15 \times (a_3 - a_2) & \text{if } \bar{w}_i > a_3 \end{cases} \quad (6)$$

where \bar{w}_i is average earnings computed before the time of the application and a_1 , a_2 , and a_3 are thresholds we take from the legislation. We assume \bar{w}_i can be approximated by the value of the permanent wage at the time of the application. Whether an individual is eligible (i.e., $E_{it}^{DI} = 1$) depends on the decision to apply ($DI_{it} = 1$) while being out of work and on having received a large negative productivity shock. We assume that the probability of success is independent of age. Eligibility does not depend on whether an individual quits or the job is destroyed.

By contrast with our assumption of a 50% probability of success for DI is our assumption of 100% take-up for our universal means-tested program and for unemployment insurance. We assume that this difference arises because of the difficulty of verifying disability compared to the income test

¹⁸Interestingly, this was an actual requirement of the program at the time of inception (1956). In our model, it reflects the fact that health shocks triggering disability are rare before this age.

and the unemployment test.

In retirement, all individuals receive social security calculated using the same formula used for disability insurance.

2.5 Employment Risk and Wage Risk

We allow for different types of shocks that constitute risk an individual is facing and we distinguish earnings and employment fluctuations driven by endogenous decisions versus unexpected shocks. The direct shocks to wages are interpreted as productivity risk. The job destruction process, the rate at which job offers are sampled in and out of work and the heterogeneity of firms constitute employment risk.

The distinction between employment and wage risk becomes relevant in the presence of search frictions and is further reinforced by the probability of job destruction. Firm heterogeneity adds another dimension to this risk: it means that some jobs may be available with a match value that would lead to a wage worth taking for an unemployed individual, even following a very bad productivity shock. Search frictions however, make it hard to find such a job and create uncertainty in both the length of unemployment and in prospective earnings. Moreover firm heterogeneity generates an option value to waiting in the unemployment state if the job arrival rate when on the job is lower than the job arrival rate when unemployed. The model allows us to identify the effects of changes in each of these risks from the behavioral reactions to their presence/change.

The productivity shocks that we observe are assumed to be uninsurable uncertainty. These productivity shocks may, for example, reflect health shocks or demographic shocks but we do not specify their source in this model. We assume that there is no commitment from the side of the firm (or the worker), so Harris-Holmstrom (1983) type contracts are not implementable. Further, we assume there is no private insurance market against employment risk. This incomplete markets set-up is consistent with results from Attanasio and Davis (1996) and others.¹⁹

¹⁹It is possible that observed wages may have already been smoothed out relative to productivity by implicit agreements within the firm. This means that productivity risk may be greater than observed wage movements within a firm, which implies that the process for productivity shocks is not properly identified for the unemployed. In other words, productivity shocks are a combination of actual shocks plus insurance, but this insurance is only present if the individual is working. If the unemployed experience greater productivity risk than estimated, this will impact on the reservation wage and on job search. For the time being we ignore this issue as far as permanent shocks are concerned.

2.6 Value Function and Model Solution

The solution of the model consists of policy functions for consumption, participation, and realizations of earnings, career paths, assets etc. There is no analytical expressions for these. Instead, the model must be solved numerically, beginning with the terminal condition on assets, and iterating backwards, solving at each age for the value functions conditional on work status. In this section, we discuss the key features of the solution, more details on the method are provided in Appendix A.

When employed, the state variables are $\{A_{it}, u_{it}, a_{ij(t_0)}\}$, corresponding to current assets, individual productivity and the match effect. The match effect is indexed by t_0 , which is the date the job began.²⁰ When unemployed and not on disability, the state variables are $\{A_{it}, u_{it}, DI_t^{Elig}\}$, corresponding to current assets, individual productivity and an indicator of whether the individual is eligible to apply for disability in that period. When unemployed and receiving disability, the state variables are $\{A_{it}, D_{it}\}$ where D_{it} is the amount of disability benefit received defined by equation (6). Consumption is chosen to maximize each value function conditional on all other decisions. Once consumption is substituted out of each value function the discrete labour supply and mobility decisions can be made.

The value function for an employed individual incorporates the fact that in the next period he will have the choice of quitting into unemployment, moving to a new job if he gets an alternative offer or staying with the firm. However if the job is destroyed the individual will have to move to unemployment. Thus the value function for an individual i who is working in period t is

$$V_t^e(A_{it}, u_{it}, a_{ij(t_0)}) = \tag{7}$$

$$\max_c \left\{ \begin{array}{l} U(c_{it}, P_{it} = 1) + \\ \beta \delta E_t \left[V_{t+1}^n(A_{it+1}, u_{it+1}, DI_{it+1}^{Elig} = 1) \right] \\ + \beta (1 - \delta) (1 - \lambda^e) E_t \left[\max \left\{ \begin{array}{l} V_{t+1}^n(A_{it+1}, u_{it+1}, DI_{it+1}^{Elig} = 1), \\ V_{t+1}^e(A_{it+1}, u_{it+1}, a_{ij(t_0)}), \end{array} \right\} \right] \\ + \beta (1 - \delta) \lambda^e E_t \left[\max \left\{ \begin{array}{l} V_{t+1}^n(A_{it+1}, u_{it+1}, DI_{it+1}^{Elig} = 1), \\ V_{t+1}^e(A_{it+1}, u_{it+1}, a_{ij(t_0)}), \\ V_{t+1}^e(A_{it+1}, u_{it+1}, a_{ij(t+1)}) \end{array} \right\} \right] \end{array} \right\}$$

²⁰Ideally we should model the behaviour of the firm. If the firm has a fixed number of positions, and if there are firing costs, a firm with characteristic $a_{ij(\cdot)}$ may not make an offer to any worker. High $a_{ij(\cdot)}$ firms may wish to wait to locate high u_{it} workers, in the same way that high u_{it} workers may wish to wait for high $a_{ij(\cdot)}$ firms. At present we ignore this issue.

The expectation operator is conditional on information at time t . If there is no offer available in $t + 1$, the expectation operator is over the productivity shock only; if an offer is an offer in $t + 1$, the expectation taken in t is also over the type of the firm making the offer.

Among the unemployed, we distinguish between those who have the option of applying for disability and those who are ineligible to apply (either because the individual is under 50 or because he has not had a negative productivity shock or because he has had an application turned down in the current unemployment spell).

For an individual who is eligible to apply for disability, the value function is given by

$$V_t^n(A_{it}, u_{it}, DI^{Elig} = 1) = \max_{c, Apply} \left\{ u(c_{it}, P_{it} = 0) + \beta \begin{cases} V_{t+1}^A & \text{if } Apply = 1 \\ V_{t+1}^{NA} & \text{if } Apply = 0 \end{cases} \right\} \quad (8)$$


where

$$V_{t+1}^{NA} = \lambda^n E_t \left[\max \left\{ \begin{array}{l} V_{t+1}^n(A_{it+1}, u_{it+1}, DI^{Elig} = 1), \\ V_{t+1}^e(A_{it+1}, u_{it+1}, a_{ij(t+1)}) \end{array} \right\} \right] \\ + (1 - \lambda^n) E_t [V_{t+1}^n(A_{it+1}, u_{it+1}, DI^{Elig} = 1)] \\ V_{t+1}^A = S \times V_{t+1}^{DI}(A_{it+1}, D_{it+1}) + (1 - S) \times E_t [V_{t+1}^n(A_{it+1}, u_{it+1}, DI^{Elig} = 0)]$$

and S is the exogenous probability of a successful application. When deciding whether or not to apply, the individual already knows if he has a job offer in that period. If the disability application is successful, we can calculate the resulting value function, V_{t+1}^{DI} , analytically: the amount of the disability insurance payment, D_{it} , depends on the permanent wage only and not on the particular firm that the individual has most recently been working for. This amount is earned each year until retirement.

Based on a comparison of the value functions, in each period the individual decides whether or not to work; and if working, whether or not to move to another job if the opportunity arises; and if not working, whether or not to apply for disability benefit. The decision about whether or not to move to another job if an outside offer is received is, in practice, more straightforward than the other decisions because we assume that there is no cost of switching firm. This means that the decision to switch firm involves a simple comparison of the $a_{ij(\cdot)}$ and the individual will move if the new offer

Figure 1: Consumption as a function of current assets conditional on current period work status



is from a higher $a_{ij(.)}$ -firm than the current one.²¹

Because of the discrete nature of labour supply, consumption may not be continuous in assets and value functions may not be necessarily differentiable, which complicates the optimization problem. To give a clear example of this, we show the solution without retirement and so the life-cycle ends at age 62. The same qualitative pictures are observed with retirement. Figure 1 shows consumption as a function of assets in the period preceding the end of life, $T - 1$, for participants and non-participants, and for different firm types, conditioning on individual productivity. The sharp decline in consumption when participating at a given firm in $T - 1$ arises at the asset stock which induces the individual not to work in the *next* period, T . Because the individual is not working in period T , lifetime income is lower and consumption falls in both periods. On the other hand, since leisure is higher in the next period, overall welfare is higher: the value function is monotonically increasing in assets. The extent of the fall depends on the degree to which consumption and leisure are substitutes. If we look at the solution in earlier time periods or the solution with retirement included, these sharp drops are smoothed out. This is partly because the fall in income associated with a change in participation in one period in the future can be smoothed out over several periods. It is also partly because uncertainty smooths the discreteness: a marginal increase in asset holdings in period t will

²¹If we were to allow for a cost of switching firm in the numerical solution, then the decision about whether or not to switch would depend on a comparison of the value function at the existing firm and the value function at the new firm. This difference will depend on the expected duration of the new job, the worker's horizon and all elements of the dynamic programming problem.

only change participation in $t + 1$ in particular states and so has less of an impact on consumption in period t than if participation in $t + 1$ changed in all states.

3 Data

We use the 1993 panel of the Survey of Income and Program Participation (SIPP) to estimate our wage dynamics parameters because it records all job to job transitions and the resulting new wage each time. However, the SIPP follows individuals only for 3 years and this means that it is less useful for duration analysis. We use the 1988-1996 Panel Study of Income Dynamics (PSID) to construct participation and unemployment duration profiles. In both data sets, we stratify the sample by education, low (those with a high school diploma or less), and high (those with some college or more).

3.1 The SIPP

The main objective of the Survey of Income and Program Participation (SIPP), conducted by the US Census Bureau, is to provide accurate and comprehensive information about the income and welfare program participation of individuals and households in the United States. The SIPP offers detailed information on cash and non-cash income on a sub-annual basis. The survey also collects data on taxes, assets, liabilities, and participation in government transfer programs.

The SIPP is a nationally representative sample of individuals 15 years of age and older living in households in the civilian non-institutionalized population. Those individuals, along with others who subsequently come to live with them, are interviewed once every 4 months for a certain number of times (from a minimum of 3 to a maximum of 13 times). Each year, a new panel starts, so some overlapping is expected. The first sample, the 1984 Panel, began interviews in October 1983 and surveyed individuals 9 times. The second sample, the 1985 Panel, began in February 1985 and surveyed individuals for 8 times. We use the 1993 panel, which has 9 interviews in total (or three years of data for those completing all interviews).²²

The Census Bureau randomly assigns people in each panel to four rotation groups. Each rotation group is interviewed in a separate month. Four rotation groups thus constitutes one cycle, called a wave, of interviewing for the entire panel. At each interview, respondents are asked to provide

²²The raw data can be obtained at <http://www.nber.org/data/sipp.html>.

information covering the 4 months since the previous interview. The 4-month span is the reference period for the interview.

Our sample selection is as follows. The raw data has 62,721 records, one for each individual, corresponding to 1,767,748 month/person observations (note that, due to attrition, not all individuals complete 9 interviews). We drop females, those aged below 25 or above 60, those completing less than 9 interviews, the self-employed, those who are recalled by their previous employer after a separation, those with missing information about the state of residence, and some outliers in earnings.²³ Our final sample includes 6,226 individuals corresponding to 224,136 month-person observations, or 3 years of data per individual. We report some sample statistics in Table 8 in the appendix.

Our measure of (firm-specific) hourly wage is obtained by dividing annual earnings earned at the firm by annual hours worked at the firm. Individuals may have multiple hourly wage observations within a year if they work for multiple firms (concurrently or not). We use only the main job (the one that pays the highest proportion of annual earnings). In the SIPP, each firm an individual is working for is assigned an ID.²⁴ We set $M_{it} = 1$ if the employer the individual is working for at time t is different than the one he was working for at time $t - 1$. We allocate individuals to the low and high education groups based on response to a question about the highest grade of school attended. An important advantage of the SIPP over the PSID when it comes to estimating the wage process allowing for job mobility is that the SIPP does not average pay over different employers. Thus the full effect of a move from one employer to another is observed.

3.2 The PSID

The PSID data are drawn from the 1988-1996 family and individual-merged files. The PSID started in 1968 collecting information on a sample of roughly 5,000 households. Of these, about 3,000 were representative of the US population as a whole (the core sample), and about 2,000 were low-income families (the Census Bureau's SEO sample). Thereafter, both the original families and their split-offs (children of the original family forming a family of their own) have been followed. In the empirical analysis we use the core sample after 1988 because detailed data on monthly employment status and other variables of interest are available only after that year.

²³An outlier is defined as one whose (annualized) earnings fall by more than 75% or grow by more than 250%.

²⁴We use corrected firm IDs (see Stinton, 2003).

Our sample selection is as follows. We focus on males with no missing records on race, education, or state of residence. We drop those with topcoded wages, the self-employed, those with less than three years of data, and those with missing records on the monthly employment status question. Education level is computed using the PSID variable with the same name.

The PSID asked individuals to report their employment status in each month of the previous calendar year and their year of retirement (if any). We use these questions to construct a quarterly participation indicator for each individual and unemployment durations. We classify as not employed in a given month those who report to be unemployed/temporarily laid off, out of the labor force, or both, in that month. We treat unemployment and out-of-labor force as the same state; this tallies with the definition of unemployment that we use in the simulations (see Flinn and Heckman, 1991, for a discussion of the difference between these two reported states).²⁵ In principle, the durations are both left- and right-censored. Some spells begin before the time of the first interview, while some spells are still in progress at the time of the last interview. To avoid problems of left censoring we only use spells that begin in the sample. In calculating durations, we take our sample to be individuals who exit between 1988 and 1992. However, we use more recent years of PSID data (1993-1996) to calculate durations for those whose spells are right-censored by the 1988-1992 window. This reduces the censoring from 13.09% of all spells to 5.52%.

4 Estimating the Wage Process

Wages are observed conditional on individuals working; within-firm wage growth, which identifies the variance of permanent productivity shocks, is only observed if the individual does not change job; between firm wage growth, which helps identify heterogeneity across firms is observed only for job movers. Further, employment and mobility decisions are all endogenous and if this is ignored we risk biasing the estimates of the variances to wages and of firm heterogeneity.

To address this problem our approach is as follows: First we model the selection process into and out of work and between firms. We then construct sample selection terms and estimate wage growth equations conditioning on these terms. We finally obtain the estimates of the variances of interest by modelling the first and second moments of unexplained wage growth for various subgroups. We

²⁵If the distinction in the data between out-of-labor force and unemployment reflects a difference in search intensity, we could make a meaningful distinction in our model only if we introduced a search decision with a cost attached.

simplify the problem by assuming normality of all error terms.

Define the latent utility from labor market participation as $P_{it}^* = z'_{it}\varphi + \pi_{it}$. The associated labor market participation index is $P_{it} = 1 \{P_{it}^* > 0\}$, which is unity for participants. Workers separate from their current employer voluntarily (quits) or involuntarily (layoffs). As argued by Borjas and Rosen (1980), job turnover, regardless of who initiates it, represents the same underlying phenomenon, that of workers' marginal product being higher elsewhere. Let $M_{it}^* = k'_{it}\theta + \mu_{it}$ denote the latent utility from moving in year t to an employer that is different from the one in $t - 1$. The indicator $M_{it} = 1 \{M_{it}^* > 0\}$ singles out the “movers”. We assume: $(\pi_{it} \quad \pi_{it-1} \quad \mu_{it})' \sim N(\mathbf{0}, \mathbf{I})$.

Taking first differences of the wage equation (1), using the process for permanent shocks (2) and recalling that $\xi_{it} = (a_{ij(t)} - a_{ij(t_0)})$, we obtain:

$$\Delta \ln w_{it} = \Delta d_t + \Delta x'_{it}\psi + \zeta_{it} + \Delta e_{it} + \xi_{it}M_{it}$$

Wage growth is only observed for those who work in both years. To achieve identification of the relevant parameters, we make the following assumptions:

1. $E(a_{ij(t)}a_{ij(s)}) = \sigma_a^2$ if $j(s) = j(t)$ and zero otherwise.
2. We denote $\sigma_\zeta^2 = E(\zeta_{it}^2)$ and $\sigma_e^2 = E(e_{it}^2)$ (for all i, t) the variances of the permanent productivity shock and measurement error, respectively. We denote $E(\zeta_{it}\pi_{is}) = \sigma_\zeta\rho_{\zeta\pi}$ if $s = t$ and assume it to be zero otherwise.²⁶
3. Given the definition of the mobility premium ξ_{it} , we assume $E(\xi_{it}\pi_{is}) = \sigma_\xi\rho_{\xi\pi}$ if $s = t$, $E(\xi_{it}\pi_{is}) = \sigma_\xi\rho_{\xi\pi-1}$ if $s = t - 1$, and zero otherwise.
4. We allow for contemporaneous correlation between the unobservable of the job mobility decisions (μ) and the shocks to the permanent productivity component and the match effect: $E(\zeta_{it}\mu_{is}) = \sigma_\zeta\rho_{\zeta\mu}$, and $E(\xi_{it}\mu_{is}) = \sigma_\xi\rho_{\xi\mu}$ for all $s = t$ and zero otherwise.
5. We assume that the distribution of innovations to the match effect ξ_{it} and the productivity shock are uncorrelated ($E(\xi_{it}\zeta_{is}) = 0 \forall t, s$), and that there is no selection on measurement error ($E(e_{it}\pi_{is}) = E(e_{it}\mu_{is}) = 0 \forall t, s$).

²⁶We denote with ρ_{ab} the correlation coefficient between a and b , and with σ_a the standard deviation of a .

Suppose now that we select only those who work at t and $t - 1$ ($P_{it} = 1, P_{it-1} = 1$). It is easy to show that:

$$\begin{aligned}
E(\Delta \ln w_{it} | P_{it} = 1, P_{it-1} = 1) &= E(\Delta \ln w_{it} | P_{it} = 1, P_{it-1} = 1, M_{it} = 1) \Pr(M_{it} = 1) \\
&\quad + E(\Delta \ln w_{it} | P_{it} = 1, P_{it-1} = 1, M_{it} = 0) (1 - \Pr(M_{it} = 1)) \\
&= \Delta d_t + \Delta x'_{it} \psi + G_{it}
\end{aligned} \tag{9}$$

where G_{it} is a “selection” term induced by labor market participation and inter-firm mobility (see Appendix B for details).²⁷ We estimate the components of this selection term in a first stage by running separate probit regressions,²⁸ and use these to then estimate the parameters of (9) consistently in a second stage using only participants.

We now need to estimate the variance of the permanent shocks the variance of the firm level heterogeneity and the variance of the measurement error. Estimation is based on the moments of unexplained wage growth (observed only for participants in both periods):

$$g_{it} = \Delta(\ln w_{it} - d_t - x'_{it} \psi) = \zeta_{it} + \Delta e_{it} + \xi_{it} M_{it} \tag{10}$$

We use the first and second moments of 10 for movers ($M_{it} = 1$) and for stayers ($M_{it} = 0$), as well as the first-order autocovariance, always correcting for selection due to participation and mobility. In addition to the two variances of interest we also estimate the relevant correlations that drive selection. The estimation process takes also into account that the wage growth we model is annual, while the participation decision is quarterly, in accordance with the model. The details of the entire estimation process are given in Appendix B.

Standard errors are computed using the block-bootstrap procedure suggested by Horowitz (2002). In this way we account for serial correlation of arbitrary form, heteroskedasticity, as well as for the fact that we use a multi-step estimation procedure, pre-estimated residuals and selection terms. We should point out that this procedure is likely to be conservative, since it allows for more serial correlation than that implied by the moment conditions we use.

²⁷In estimation we do not use the restrictions on the parameters of interest imposed by (9). This only results in a loss of efficiency, but it does not affect consistency. We estimate the standard errors by the block bootstrap.

²⁸The assumed orthogonality assumption between π_{it} and μ_{it} allows us to do this.

4.1 Results

4.1.1 Participation and mobility

We start by estimating quarterly participation probits using the SIPP data. These include a quadratic in age, a dummy for whites, region dummies, a dummy for married, year dummies as well as unearned household income, and an index of generosity of the welfare system, which here we proxy with the generosity of the state-level UI system.²⁹ The latter two are excluded from the wage equation and are the instruments that identify selection into work - the unearned income as a pure income effect and UI as a fixed cost of work while eligible.³⁰ The participation equation for each quarter is reported in the appendix in Table 9. The main point is that unearned income has a strong and significantly negative effect on the probability of working. UI generosity is also a significant factor discouraging work, but only for the lower education group and not for the college graduates.

We also estimate a mobility probit, which will allow us to control for the censoring of the between firm wage growth. The dependent variable is whether an individual who was working in period τ is in a different job in period $\tau + 1$. Thus for the purposes of this estimation, mobility may include those moving jobs via unemployment.³¹ The mobility probit includes the same variables as the participation equation, as well as industry dummies and an indicator as to whether the person was working for a non-profit organization, in both cases for period τ . Unearned income influences positively mobility for both education groups; UI generosity influences positively mobility for the lower education group but not the college graduates. The effect of UI on mobility is theoretically ambiguous. On the one hand, it increases the reservation wage leading to individuals quitting employment following negative wage shocks and increasing mobility through this mechanism. On the other hand, when UI is low, durations of unemployment will be shorter and wage increases will occur through job-to-job mobility. Our results indicate that the former effect dominates. Our

²⁹To obtain a measure of the generosity of the UI program in the state where the worker lives, we rank states according to the ratio between maximum weekly UI benefit (which we take from current legislation) and average weekly wages (which we calculate from the CPS- using males only). Our measure of generosity is the rank variable, which varies over time and across states. We obtain similar results if we rank states pooling data for all years. Ideally, one would like to use an index of generosity of the Food Stamps program, but this is a federal program and its time-series variability is negligible.

³⁰In practice we exploit variation over states and time. For the exclusion restrictions to be valid the US labour market should be sufficiently integrated and sufficient trade should be taking place, so that variability in benefits in one state does not affect the price of human capital in that state.

³¹We also tried to distinguish between “voluntary” movers (with no spell of unemployment in between two employment spells) and “involuntary” movers (those who move jobs via unemployment). We modified the selection process outlines in the Appendix, but find that the estimates of the variances of interest (σ_c^2 and σ_a^2) change very little (although the correlation coefficients in Table 1 change considerably). Thus here we consider the simpler model with a single mobility index.

results also show that mobility declines with age for both groups. As people age, they tend to locate in better firms, and thus it becomes increasingly unlikely that an outside offer is sufficiently good to trigger mobility. Job destruction is an important force disrupting this age effect. Table 10 in Appendix C presents the results.³²

4.1.2 Variance Estimates

Armed with these results, we move on to estimate the parameters of the wage process by the method of moments, imposing constraints across equations. The moments we fit, together with the corrections for selection are reported in appendix B. The results are reported in Table 1. The σ parameters refer to the standard deviations of the various stochastic components of wages. The ρ parameters are the correlations between the various stochastic shocks and the shocks driving selection. They are defined in appendix B. We estimate the model for the whole sample to have a comparison with previous work (column 1) and separately by the two education groups (columns 2 and 3).

When we control for selection into employment and for job mobility, we find that in the whole sample the standard deviation of the permanent shock, σ_ζ , is about 0.10, the standard deviation of the transitory shock (measurement error), σ_e , 0.09, and the standard deviation of the match specific effect, σ_a , 0.23. These parameters are all very precisely estimated. They imply that matching heterogeneity is a very important component of wage dispersion; wages can fluctuate $\pm 46\%$ between firms for the same individual.

Columns (2) and (3) report the results of estimating the model separately for our two education groups. The stochastic process of wages is very similar across the two education groups.

What happens if we ignore the fact that mobility is endogenous and attribute all wage fluctuations to the permanent and transitory shocks (σ_ζ and σ_e)? This, implicitly, has been the assumption made in papers estimating the covariance structure of earnings (MaCurdy, 1982; Abowd and Card, 1989; Meghir and Pistaferri, 2004) and in the precautionary savings papers estimating risk via the standard transitory/permanent shock decomposition (Carroll and Samwick, 2001; Gourinchas and Parker, 2002). In column (4) we report the results of this experiment for the whole sample. The estimated standard deviation of the permanent shock σ_ζ increases by about 50%: a large proportion

³²Unearned income and UI generosity as well as the industry dummies act as identifying instruments because they are excluded from the wage *growth* equation.

Table 1: Wage variance estimates

	Whole sample	Low education	High education	Neglect mobility (All)
	(1)	(2)	(3)	(4)
Standard Deviations				
σ_ζ	0.104 (0.014) [0%]	0.095 (0.021) [1%]	0.106 (0.018) [0%]	0.144 (0.009) [0%]
σ_e	0.088 (0.014) [0%]	0.085 (0.033) [0%]	0.090 (0.022) [0%]	0.088 (0.006) [0%]
σ_a	0.226 (0.011) [0%]	0.222 (0.018) [0%]	0.227 (0.016) [0%]	
Correlations				
$\rho_{\zeta\pi}$	0.090 (0.234) [80%]	0.072 (0.273) [92%]	-0.195 (0.284) [31%]	0.377 (0.163) [2%]
$\rho_{\zeta\mu}$	-0.626 (0.404) [0%]	-0.875 (0.472) [0%]	-0.651 (0.470) [0%]	
$\rho_{\xi\pi}$	0.192 (0.288) [49%]	0.327 (0.351) [27%]	0.038 (0.617) [87%]	
$\rho_{\xi\pi_{-1}}$	-0.046 (0.227) [87%]	-0.048 (0.235) [95%]	-0.110 (0.581) [76%]	
$\rho_{\xi\mu}$	0.231 (0.169) [13%]	0.274 (0.201) [10%]	0.2650 (0.2151) [16%]	

Note: σ_ζ , σ_e , and σ_a are the standard deviations of the permanent shock, measurement error, and the match component. $\rho_{\zeta\pi}$ ($\rho_{\xi\pi}$) is the correlation between the permanent shock (mobility premium $\xi = a_j - a_{j-1}$) and unobserved heterogeneity in the participation equation. $\rho_{\xi\pi_{-1}}$ is the correlation between the mobility premium and unobserved heterogeneity in the participation equation in the previous period. $\rho_{\zeta\mu}$ ($\rho_{\xi\mu}$) is the correlation between the permanent shock (mobility premium) and unobserved heterogeneity in the mobility equation. Standard errors (in parenthesis) are computed using the block bootstrap. Bootstrap p-values in square brackets. We constrain all the correlation coefficients to lie between -1 and 1.

of wage fluctuations usually attributed to unexpected shocks is in fact a result of endogenous mobility choices. This is likely to be important for the welfare costs of risk because individuals change jobs quite frequently and because they do not have to accept worse paying jobs than the one they have. However match dispersion does itself introduce risk: first because individuals with good matches who are displaced can expect to be hired at a lower rate (on average); and second, because individuals face uncertainty about the quality of offers they are likely to receive. On the other hand, match dispersion also offers the possibility of job improvements, which as we shall see, is a dominant factor in the effect of such dispersion on welfare when the arrival rate is sufficiently high.

The estimated correlations that drive selection are reported on the lower panel of Table 1 and conform to the expected signs. The most significant is the negative correlation between the permanent shock and mobility ($\rho_{\zeta\mu}$). For the purposes of correcting for selection in estimation, mobility is defined as any job change, including those taking place through unemployment. Since a good productivity shock will encourage people to work, it will also result in less job changes than otherwise - hence the negative and significant correlation. The next most important one, with an overall p-value of 13% is the correlation between a good alternative offer and mobility, which is positive and quite large, as expected.

5 Calibrated Parameters

We now need to set the remaining parameters required to complete the model. We set the coefficient of relative risk aversion γ equal to 1.5, taken from Attanasio and Weber (1995), whose model of consumption also allows for nonseparable labor supply. The real interest rate is set equal to the real return on 3 month treasury bills, at an annual rate $r = 0.015$, and this is set equal to the discount rate ($\frac{1}{\beta} - 1$). The remaining parameters we obtain through calibration using the structural model outlined in section 2.

Given the estimated parameters of the wage process and those set above, we now set the remaining parameters to fit the life-cycle participation profile and unemployment duration profile for men, by education group. Our approach is to choose the parameters for each education group to minimize the weighted sum of squares of the distance between statistics calculated in the data and corresponding simulated statistics. The weights are the inverse of the variances of the sample statistics.

Table 2: Parameters Obtained through Calibration

Parameter	High Education	Low Education
Job destruction rate δ	0.030	0.049
Job arrival rate - Unemployed λ^n	0.84	0.76
Job arrival rate - Employed λ^e	0.72	0.67
Fixed cost of work F	0.34	0.34
Disutility of participation η	-0.59	-0.53

Note: The values of δ , λ^n and λ^e are given as quarterly rates. The value of the fixed cost F for each education group is expressed as a ratio to average earnings of that group at age 22.

The statistics we use are the average participation rate in four ten year age bands (22-31, 32-41, 42-51 and 52-6) and the mean duration of unemployment in eight five-year age bands. In Table 11 in the appendix we show the fit of the moments we have targeted. In Table 2, we present the calibrated parameter values, with job destruction and arrival rates given at quarterly rates. In Figures 2 and 3 we show the calibrated profiles.

The job destruction rate is about 50% higher for the lower educated individuals than for the higher educated ones. The contact rates are higher for the more educated and they are higher for those out of work than when in work, possibly reflecting increased costs of search when working. The value of η for high education individuals implies that work is equivalent to a 45% loss of consumption. For those of low education the equivalent consumption loss is 41%.³³ These values also imply consumption and leisure are substitutes, and thus it is consistent with the observed fall of consumption upon retirement (or unemployment). Finally, the fixed costs of work for both education groups correspond to about 34% of average earnings for a 22 year old in that group

Figure 2 shows participation profiles for the high educated and low educated. Each figure compares the profile in the data with the calibrated profile (labelled “employed”).³⁴ For both education groups, participation rates are constant or display a slow decline until the age of 45, followed by a sharp decline to age 62. Part of this fall reflects early retirement, rather than temporary periods

³³The consumption equivalent is calculated as $(1 - \exp[\eta])$.

³⁴The profiles from the data are calculated controlling for cohort effects and assuming that time effects average out to zero and are orthogonal to the time trend. The estimated participation rate is equal to the actual rate at age 40.

out of the labor force. Since early retirement is an endogenous labor supply response, we treat this in the same way as we treat unemployment. There is a level difference between the two groups: the high educated participate more than the low educated up to age 45 (participation rates around 96%, compared to 90% for the low educated), and the subsequent decline is less marked. Our match to participation is fairly good for both skill groups.

We also plot the participation rate that would be obtained if all offers received (including those offers from an existing employer) were accepted. This highlights the extent that the downturn in participation with age is due to more offers being rejected, rather than to fewer offers being made.

Figure 2: Participation over the Life-Cycle




Figure 3: Mean Duration over the Life-Cycle




Figure 3 compares mean unemployment durations over the life-cycle in the simulations and in the data for both education groups. Durations have a maximum length determined by the number of

quarters until age 62. In the data, durations are measured in months and are expressed as fractions of a quarter. We are able to match very well the durations for the high education individuals. For the low education ones we get a very similar profile but shifted a bit to earlier ages.

In the model $\lambda^n, \lambda^e, \delta, F$ and η are independent of age and so the age effects that we find in all the simulated profiles can be explained only by endogenous saving and labor supply behavior in response to the budget constraint and the welfare benefit structure: the match in the slope of profiles over the life-cycle is not an artefact of age varying parameters and is a demonstration of the strength of the model.

5.1 Implications of the model

We have calibrated the model using only participation and unemployment duration data, given the pre-estimated wage process and given an intertemporal substitution parameter from the literature. However, the model has implications for a range of different variables. In particular, we use the model to predict the wage loss associated with a spell of unemployment, the duration of employment (“tenure”), the extent of consumption loss on unemployment, the arrival rate of accepted offers, and the ratio of mean wealth over stages of the life cycle to mean life-cycle income. Table 3 reports the model predictions and corresponding statistics, obtained from various different sources of data, for a number of statistics that are not used in the calibration.

The Cost of Displacement There is empirical evidence that displaced workers experience earnings losses following job loss. Some authors impute this to exogenous skill depreciation during periods of unemployment (Rogerson and Schindler, 2001; Ljunqvist and Sargent, 2002). An alternative that is consistent with our model is that wages on re-entry may be lower than before job loss because of the loss of a particular good match on entering unemployment. We report in Table 3 the extent of the wage fall on re-entry.

For the high educated, wages on re-entry are, on average, 19% lower than before displacement. For the low educated, the loss is 14%. These figures are similar to those found in the literature. In particular, we compare these figures with those reported by Jacobson, LaLonde, and Sullivan (1993) for their non-mass layoff sample (after controlling for time trends). They report that 1 quarter after displacement, earnings of displaced workers are 19% less than before displacement. Finally,

Table 3: Model Implications

Statistic	Data		Model	
	High Educ	Low Educ	High Educ	Low Educ
Mean wage loss ⁺				
$\ln w_\tau - \ln w_t$	-0.19		-0.19	-0.14
$\ln w_{\tau+4} - \ln w_t$	-0.076		-0.086	-0.058
Median duration of employment (Age 22-46) ⁺⁺	12		21	15
Mean consumption loss (Age 25-60)*				
$\Delta \ln c_{t+1} / \Delta \ln y_{t+1}$	0.56		0.42	0.57
Arrival rate of accepted job offers (Age <40):				
On-the-job			0.036	0.041
From unemployment			0.50	0.46
Fraction of job offers accepted (Age <40):				
On-the-job			0.049	0.067
From unemployment			0.60	0.60
Mean wealth / mean lifetime income**				
Age 30-35	1.58	1.71	0.89	1.36
Age 50-55	7.27	5.13	5.66	5.34

⁺The data numbers for wage loss are taken from Jacobson, Lalonde and Sullivan (1993). w_t is the wage at displacement, w_τ is the wage at reentry, $w_{\tau+4}$ is the wage one year later.

⁺⁺ Employment durations are reported as number of quarters until employment with a particular employer ends.

*The data numbers for consumption loss are taken from Browning and Crossley (2001).

**The wealth data come from the 1994 PSID wealth supplement and include housing wealth and private pension holdings. Mean lifetime income is defined as average annual household income for heads of household aged 22-62.

one implication of our model is that the displacement costs are likely to be relatively short lived. Indeed, we calculate that 1 year after returning to work, wages of the low educated are only 5.8% below their pre-displacement wages; for high educated individuals the figure is 8.6%. These figures are very close to the ones we extrapolate from Jacobson, LaLonde, and Sullivan (1993) over a similar time horizon.

Employment Durations The PSID has data on job tenure which could be used to pin down the arrival rate of offers while on the job, λ^e . Heads of household are asked how many years they have been working with their current employer. There are three main difficulties with these data. First, the spells are right-censored. Second, there is a substantial initial conditions problem with a very large proportion of spells being left censored. Third, a number of authors have questioned the reliability of these measures of reported tenure (Brown and Light, 1992). Because of these difficulties we do not use these data in the calibration. Instead, we report statistics on employment duration from the data and from the simulations in Table 3. We expect the simulated employment durations to be longer than those obtained from the PSID. For the low educated, median employment duration in the simulations is 15 quarters for those younger than 46. For the high educated the median simulated spell length is 21 quarters. In the data, median employment duration is 12 quarters but this is likely to be downward biased because of censoring.

Consumption Fall at Unemployment Gruber (1997) and Browning and Crossley (2001) have explored empirically the consumption loss associated with unemployment. Consumption will be lower in unemployment if the job loss followed a permanent loss in productivity, which implies a life cycle wealth effect and a lower incentive to work. In addition if there are non-separabilities between consumption and leisure, as in our model, and leisure is a substitute for consumption, individuals will cutback on consumption as leisure increases. Finally if individuals are unable to smooth consumption through borrowing against future income they will also have to cut back more than they otherwise would. Our model contains these three effects, with the source of the income loss being the loss of the match as well as possibly the wealth effect implied by a negative permanent shock; the non-separability being built into the structure of the utility function; the liquidity constraint being the restriction that assets have to be non-negative.

In Table 3 we report average consumption loss by education group and compare it to Browning and Crossley (2001).³⁵ They use data from a Canadian sample which includes many two earner households; moreover the Canadian welfare system implies different replacement rates. To control for these differences we compare the percentage consumption lost relative to the percentage income lost following unemployment. Our figures are remarkably close to the Browning and Crossley number, whose comparable figure is 56%.³⁶ We calculate a 57% relative loss for the low educated and a 42% for the high educated. The Browning-Crossley sample contains 70% low educated individuals. So on this score the model fits the facts very well indeed.

Arrival rate of offers on-the-job Table 3 reports the arrival rate of accepted job offers among workers and among the unemployed. For workers, the arrival rate of accepted offers is low because workers only choose to move if they receive a better offer than the wage at their existing firm. Among the unemployed a much higher proportion of offers are accepted: the table shows that, among the low educated, 60% of offers are accepted by the unemployed, whereas only 7% of offers are accepted by the employed. The fast movement out of unemployment is not surprising because the offer arrival rate when employed is not much lower than it is for the unemployed making the option value of unemployment low.

Wealth Accumulation The final row in Table 3 reports, for individuals aged 30-35 and 50-55, the ratio of average wealth holdings to average lifetime income. Our model captures fairly accurately the level of wealth holdings of the low educated at both stages of life. For the high educated, the model under-predicts slightly the wealth holdings, but the rate of wealth accumulation is similar.

6 Evidence on the Wage Process and Alternatives

Our choice of specification for the stochastic process of wages is based on a long and well established literature. However this does not make it uncontroversial. In this section, we address three issues regarding the specification: first, we consider whether controlling for mobility affects estimates of the persistence of shocks; second, we consider modelling the match component as being stochastic;

³⁵Our numbers are not comparable to the Gruber calculation because he only uses food. One would need to inflate his number by dividing it by the marginal budget share for food to get back to a total consumption figure.

³⁶Their figures are a 14% consumption loss and a 25% loss in income, which imply the number we report (56%).

finally, we consider introducing heterogeneity in (deterministic) income growth rates, as in Guvenen (2006).

Mobility and Persistence of Shocks One difficulty with relying on the existing literature to support our specification of a unit root in wages is that the papers that provide supporting evidence do not control for worker mobility. In particular, job mobility decisions might create the impression of greater persistence in income: for example, if mobility is ignored, wage increases due to moving firm will be seen as permanent shocks, whereas when we control for mobility, shocks are identified from within firm wage movements which may be less persistent. Our specification imposes a unit root on within firm wage movements, but the concern is that wages are not as persistent as they seemed to be when mobility was ignored. To address this, we estimate the autocovariance properties of wage growth residuals for workers who do not change job, controlling for selection. These autocovariances still conform to the random walk process, rather than to one with less than unit root persistence. In particular the autocovariances are statistically and economically insignificant after the first two lags and there is no evidence of a gradually declining pattern which would have been observed if a simple AR(1) process had generated the data.³⁷ Thus allowing for mobility cannot account for the earlier finding of a random walk.

Stochastic Match Component The analysis in this paper has been based on the assumption that the match specific effect is constant for the duration of the match and that the shocks to individual productivity persist beyond the current job. In fact, since these shocks are permanent, they persist forever. If we had matched employer-employee data, we could have allowed for an even richer specification with a stochastic match component. Given our data restrictions, however, the alternative assumption that we explore is that the match effect is subject to permanent stochastic shocks, while individual heterogeneity is captured by a fixed effect.

$$\ln w_{it} = d_t + x'_{it}\psi + f_i + e_{it} + a_{ij(t_0)t} \quad (11)$$

³⁷We construct an estimate of wage growth residual in the PSID using the estimates from the SIPP discussed in Section 4. We use the PSID because it allows us to look at longer autocovariances. In the PSID whether one is a mover or stayer is identified correctly (up to measurement error). For stayers, the autocovariances at lags 0,1, 2, 3, and 4 are, respectively 0.0854 (s.e. 0.0038), -0.0213 (0.0015), -0.0005 (0.0013), 0.0021 (0.0018), and -0.0029 (0.0030).

where f_i is an individual specific unobserved fixed effect, e_{it} reflects measurement error and transitory shocks and is taken to be an independently and identically distributed normal random variable, and $\alpha_{ij(t_0)t}$ is the value in period t of the match component for an individual who joined firm j in period $t_0 \leq t$. The evolution of the match component is described by

$$a_{ij(t_0)t} = \begin{cases} a_{ij(t_0)t} = a_{ij(t_0)} & \text{in the first period working for firm } j \text{ (} t = t_0 \text{)} \\ a_{ij(t_0)t} = a_{ij(t_0)t-1} + \omega_{ij(t_0)t} & \text{in subsequent periods in firm } j \end{cases} \quad (12)$$

In equation (12), $\omega_{ij(t_0)t}$ is an i.i.d. normal random variable. However, this is now a permanent shock to the match specific effect. This model has fundamentally different implications from our preferred specification. For example, with this specification, the shocks $\omega_{ij(t_0)t}$ are in effect transitory: the individual can change jobs following a bad realizations of the match effect, thus wiping the slate clean from past shocks even if these are permanent at the match level. At the individual level, there are no permanent shocks. We now show why our choice fits the facts better. To achieve this we estimate the alternative specification and we calibrate a version of the model using the alternative wage process, focusing on the low education group.

First, we show realized pay in the simulations of the preferred specification and of the alternative specification. In the left-hand graph in Figure 4 we show how the simulated cross-sectional variance of pay evolves for a cohort under our assumption of permanent shocks to productivity allowing for all the endogenous decisions; this is the upward sloping line marked “Preferred”. The fall in the cross-section variance that occurs at older ages and observed in the simulations arises because of the sharp decline in participation among men over 50 and can also be seen in the data graph on the right hand side panel. The flat line, marked “Alternative” depicts the simulated cross-sectional variance with the alternative model. The right-hand graph in Figure 4 shows the increase in the variance in the data. The line labelled “estimated age effect” reports the actual cross-sectional variance at each age, while the line labelled “predicted age effect” reports the predicted cross-sectional variance using the specification assumed in our preferred model.³⁸ This predicted age effect implicitly assumes that there is full participation: the slope is the conditional variance of the permanent shock. This shows that, for most of the life cycle, the cross-sectional variance of wages is increasing with age of the cohort. Our preferred model describes the wage process far better than the alternative model.

³⁸The “estimated age effect” is obtained controlling for cohort effects and assuming that time effects average out to zero and are orthogonal to a time trend.

Figure 4: Cross-Section Variance of Wages by Age





Figure 5: Standard Deviation of Initial Conditions



The second piece of evidence in support of our preferred model comes from the estimates of the variance of initial wages required in the two specifications. Figure 5 shows that the standard deviation of the initial wage decreases with year of birth using the alternative specification. This contradicts evidence from MaCurdy and Mroz (1995) and others who show that more recent cohorts face a higher variance on entering the labour market. However, with the alternative specification, the only way the model can fit the increasing variance of wages by age that exists in the data is to attribute the increase to cohort effects, i.e. older cohorts have a higher variance on entering the labour market. By contrast, using our preferred specification, we estimate that the variance of initial conditions is indeed higher for more recent cohorts.

The third bit of evidence we provide in support of our preferred model arises from calibrating the model using the alternative specification for wages, and then comparing in Table 4 the resulting participation profile to the data and to the profile of our preferred specification.

The model with the alternative specification is incapable of capturing the downturn in participation from about aged 45 onwards and the increased duration of unemployment spells for older people that is present in the data. This contrasts with the success of our preferred specification. The reason for this is that in the alternative model anyone quitting due to a negative shock to the match specific effect can start again with a new job and “wash away” the past. By contrast, when these negative shocks are individual specific and so persist across matches, an individual is more likely to remain unemployed, particularly with the availability of means-tested benefits and disability insurance.

Table 4: Observed and Matched Participation Rates, Low Education

Age	Data	Preferred Model	Alternative Model
22 – 31	0.90	0.91	0.90
32 – 41	0.89	0.89	0.88
42 – 51	0.83	0.84	0.88
52 – 61	0.63	0.66	0.89

A final point is that our model also implies an increase in the cross sectional variance of consumption, very much as in the data (see Deaton and Paxson, 1994). Without effective persistence in the shocks this cannot happen with the alternative model.

Thus overall, the alternative stochastic specification with the random walk in the match component is not consistent with the data and does not allow the model to match key moments of the data anywhere near as well as we do with our preferred model. This is not to say that a richer model that combined aspects of both stochastic specifications for wages could not do even better. However, given the data limitations, we offer a parsimonious model, that is capable of replicating basic features of the data.

Heterogeneous Income Growth Model Guvenen (2008) argues that the increase in the cross-sectional variance with age can also be achieved by allowing for heterogeneity in income growth rates (as in Baker, 1997), alongside a lower degree of persistence of income shocks (an AR(1) parameter of 0.82). This specification will produce an increase in the cross-section variance of income with age, as in the data and in our preferred specification. If there were less persistence in income shocks but without the income growth heterogeneity, then the variance would be concave in age and so both

components are needed to match the data.

Distinguishing this specification from one with a permanent shock is not straightforward: in the data, the autocovariance of earnings growth is zero for observations more than two or three periods apart which is consistent with the permanent shock model and not quite in line with the deterministic individual specific growth, which would imply nonzero autocovariances at all lags. Guvenen points out that with a high enough positive autoregressive coefficient the effects of the random growth on the autocovariance structure can be obscured and its presence can only be identified by considering the covariance of income growth many years apart when the persistent shock has little effect (12 lags or so, see Guvenen (2008)). However, long panel data sets have too much attrition over such long time periods to provide a reliable test of this view. A further difficulty with this specification, as pointed out by Guvenen (2007), is that it would not match the growth in the variance of consumption because of the limited innovations to income over the life-cycle. In order to match the growth in the variance of consumption, it is necessary to assume that individuals do not know their own growth term and have to learn about it over time. This gives rise to innovations over the life-cycle, although it is difficult to distinguish statistically between learning about a random growth model and a model with a permanent shock.³⁹ Further, Haider and Solon (2006) suggest that such heterogeneity in trends may be most important early in the life-cycle, but that there is little evidence for its importance beyond age 30. Guvenen and Smith (2008) use consumption and income data to try to separate out the two models. Further, we do not preclude that introducing labour supply choices into a Guvenen framework might generate interesting alternative implications. We have chosen what is a parsimonious specification that fits the data well both in terms of income and in terms of consumption behavior.

7 The Implications of Risk

Our model and characterization of shocks has important implications for the impact of risk on behavior and welfare. Understanding these is relevant particularly when designing and evaluating policies such as unemployment insurance, food stamps or other transfers (e.g. tax credits), which effectively

³⁹Farber and Gibbons (1996) assume that individual productivity is unknown to the firm, but it is learned over time through observation of output, and so wages are updated in a Bayesian sense. They prove that this will result in the wage residual being a martingale. Thus our unit root characterization can also be consistent with a less than complete information case, but we have not considered the implications of the learning case as yet.

insure part of the risk individuals face. In this model, we have exogenous, uninsured idiosyncratic shocks and so welfare will increase if insurance is provided. We also have behavioral responses to insurance built in both through changes in participation and through changes in savings. This means we can quantify the risk sharing benefits of different sorts of insurance as well as identifying the behavioral effects induced by the insurance programs. In this section, we show first the effects of varying productivity risk, looking at the effects on participation, output and asset accumulation, as well as welfare. We then show the effects of varying the various aspects of employment risk, including job destruction and firm heterogeneity. There is no common metric for the different types of risk and as such we cannot quantify their *relative* importance; we can however quantify their effects individually.

In the model the actions of individuals are linked to each other because we require the government budget to balance over the life cycle of a cohort, which is assumed to have N members. Thus we impose

$$\sum_{i=1}^N \sum_{t=1}^T \frac{1}{R^t} [(B_{it}E_{it}^{UI} (1 - E_{it}^{DI}) + D_{it}E_{it}^{DI}) (1 - P_{it}) + E_{it}^T T_{it}] = \sum_{i=1}^N \sum_{t=1}^T \frac{1}{R^t} \tau_w w_{it} h P_{it} + Deficit \quad (13)$$

where B_{it} is unemployment insurance, D_{it} is disability insurance and T_{it} are food stamps; E_{it}^{UI} , E_{it}^{DI} and E_{it}^T are 1/0 indicators of eligibility for each of the programs respectively and $P_{it} = 1$ denotes employment. On the right hand side $\tau_w w_{it} h P_{it}$ represents tax revenue from a working individual. The deficit term represents unaccounted expenditures or revenues and will be kept constant across all simulation experiments. Following a simulated policy change we select the tax rate τ_w to satisfy this government budget constraint; individuals take τ_w as given.⁴⁰ Budget balance is imposed *within* a particular education group. We therefore abstract from the insurance *between* groups that Attanasio and Davis (1996) found to be important. Allowing the budget to balance over all education groups would confound the issue we are considering with distributional questions.

To define the welfare cost of risk write the life time expected utility of an individual as

$$E_0 U_k = E_0 \sum_t \beta \frac{(c_{kt} \exp \{\eta P_{kt}\})^{1-\gamma}}{1-\gamma}$$

where the subscript k refers to the implied consumption and labour supply stream in the baseline

⁴⁰We assume that unemployment insurance and disability insurance are financed by the tax on wages, even though in reality the financing is partly imposed upon the firms. However, if the incidence of the tax falls on the workers, as most empirical studies find, our assumption is inconsequential.

economy ($k = 1$) or an alternative economy with different risk characteristics ($k = 2$) and E_0 is the expectation at the beginning of working life. Now define π as the proportion of consumption an individual is willing to pay to be indifferent between environment $k = 2$ and $k = 1$. This is implicitly defined by

$$E_0U_2|\pi \equiv E_0 \sum_t \beta \frac{((1 - \pi) c_{2t} \exp \{\eta P_{2t}\})^{1-\gamma}}{1 - \gamma} = E_0U_1$$

which implies that $\pi = 1 - \left[\frac{E_0U_1}{E_0U_2|\pi=0} \right]^{\frac{1}{1-\gamma}}$.

Since there are no aggregate shocks in the economy and no business cycle fluctuations, we do not consider the value of, for example, smoothing the effect of the business cycle (as in Lucas, 1987) or the value of removing variation in the extent of idiosyncratic risk over the life-cycle (as in Storesletten et al., 2001). Such insurance removes heteroskedasticity but the average level of risk remains. Thus we focus entirely on the cost to the individual of idiosyncratic risk, which would be insured in a first best setting.

7.1 Wage Risk

We start by considering the impact of the permanent shock to wages. We have already shown that allowing for job mobility substantially reduces the amount of risk that we attribute to unexpected changes in productivity. Indeed this reduction is likely to be very important.

In Figure 6 we report π , the willingness to pay to avoid changes in risk relative to the estimated baseline. This willingness to pay arises because individuals are averse to the greater risk associated with increases in σ_ζ . Note that when we change the risk faced by the individual many aspects of behavior will change, including labour supply as well as unemployment and employment durations. These will result in output changes, which we also show on the graph, labelled as $\Delta \ln y$. The willingness to pay parameter has factored in all these aspects. While changes in wage risk end up implying relatively low changes in output (particularly for the higher educated), they imply large welfare losses. Thus a 50% increase of σ_ζ to 0.159 for the high educated individuals, implies a 3.7% loss in output but a willingness to pay to avoid this increase of 19.2% of consumption (the numbers are in Table 5). It is not straightforward to compare across education groups because the baseline is different. However, we note that if we increase the variance for the low educated to the same level (0.159) welfare goes down by less (16.4%). This is partly due to the effect of the welfare programs,

Figure 6: Welfare Costs and Output Effects of Varying σ_ζ





Figure 7: Effect on Participation of Varying σ_ζ



which are more important for this lower wage group. Output declines by more for the low educated, as is visible from the graph, driven by the decrease in participation shown in Figure 7. However, the overwhelming impression here is that wage risk is a major determinant of welfare, well beyond its impact on output, making insurance for such risk potentially very valuable.

More detail on the effects of varying wage risk is provided in Table 5: productivity shocks have substantial effects on unemployment durations, on consumption growth and on rates of asset accumulation particularly among the young.⁴¹ In particular, for the high educated group, as σ_ζ increases by 50%, the median rate of wealth accumulation for individuals aged 25-35 almost doubles. This faster accumulation results in an earlier and faster rate of decumulation when old, alongside lower participation rates and longer unemployment durations.

⁴¹If, in addition to the participation decision, hours of work were flexible, individuals would be able to self-insure to a greater extent than in our model.

Table 5: Comparative Statics: Varying σ_ζ

σ_ζ	π	$\Delta \ln y$	Mean	Mean		Median ($\Delta A/y$)			Age at max
			Duration	$\Delta \ln c_t$					
			47-52	25-44	45-62	25-35	36-50	51-62	
<i>High Education</i>									
0.053	0.128	0.020	2.0	0.033	0.005	0.026	0.16	0.18	62
0.095	0.032	0.006	4.6	0.029	0.012	0.11	0.24	0.10	62
0.106	0.0	0.0	5.7	0.029	0.013	0.14	0.25	0.076	62
0.159	-0.192	-0.037	12.5	0.030	0.015	0.27	0.24	-0.068	58
<i>Low Education</i>									
0.053	0.087	0.047	2.8	0.018	0.005	0.014	0.19	0.085	62
0.095	0.0	0.0	7.8	0.017	0.010	0.10	0.17	-0.010	57
0.106	-0.026	-0.011	9.1	0.017	0.011	0.14	0.17	-0.036	57
0.159	-0.164	-0.055	13.9	0.019	0.014	0.24	0.16	-0.16	55

Note: For the columns concerning the amount of assets, the denominator is average realized earnings (net of the fixed cost of work) in the education-specific baseline. Duration is measured in quarters.

7.2 Employment Risk

The two important parameters associated directly with employment risk are job destruction and the variance of the match specific effect. We now consider the implications of varying each of these parameters.

Job Destruction Figure 8 shows the impact on welfare and output of varying job destruction, δ . Increases in job destruction have large effects on output partly through increasing unemployment and partly through limiting the time that individuals are matched with the best firms. Individuals are willing to pay to avoid the increase in job destruction. There are three aspects to this willingness to pay: there is a loss of income reflected in the overall loss in output; there is an increase in employment risk; and there is an offsetting increase in leisure time. For both education groups, the effect of the increased riskiness, which otherwise would have raised the welfare loss above the loss of output, is offset by the value of increased leisure resulting from the fall in employment shown in Figure 9. More details of the effects of varying δ on behavior is provided in Table 12 in the appendix. Durations of unemployment are shorter when δ is higher because of a composition effect: more of the unemployed are out of work because of job destruction rather than because of low productivity, and so are more likely to receive job offers above their reservation wage. The rate of wealth accumulation

Figure 8: Welfare Costs and Output Effects of Varying Job Destruction





Figure 9: Effect on Participation of Varying Job Destruction



decreases with a higher value of δ because the increase in time unemployed reduces opportunities for accumulation.

Firm Heterogeneity Figure 10 shows the impact on welfare and output of varying firm heterogeneity, σ_a . Increasing firm heterogeneity implies a greater variety of firms in terms of their productivity. However, because of on-the-job search the best firms tend to be over-represented in terms of accepted offers, which pushes output up, as clearly seen in the figure. This effect of selection into the best firms leads to greater participation among older individuals when heterogeneity is increased, as shown in Figure 11. Increased heterogeneity therefore pushes welfare up, but not as much as output: increasing heterogeneity implies an increase in the cost of job loss because workers are in danger of losing a more coveted job. Interestingly the effect is not symmetric around our baseline estimates: decreasing firm heterogeneity decreases welfare as much as it does output.

More details of the effects of varying σ_a on behavior are provided in Table 13 in appendix C. For those age 25-35, greater firm heterogeneity leads to faster wealth accumulation for the high education group, but to slower wealth accumulation for the low education group. This reflects the

Figure 10: Welfare Costs and Output Effects of Varying Firm Heterogeneity





Figure 11: Effect on Participation of Varying Firm Heterogeneity



offsetting incentives caused on the one hand by greater risk inducing faster accumulation, while on the other hand, the greater expected future income induces slower accumulation.

Offer Arrival Rates The impact of varying the arrival rate of job offers when unemployed is shown in Table 14, and better described through text. Varying the quarterly arrival rate, λ^n , between 0.66 and 0.96 increases output for the low educated by 2%, and for the high educated who experience less unemployment by 1%. Further, average unemployment durations decrease by about 1 month. These effects translate into welfare gains of 1.6% and 1.2% respectively. Compared to the output change, this suggests the high educated value the increased arrival rate more highly. This is because the opportunity cost of being unemployed is greater for the high educated: the high educated receive higher wage offers and are not as well insured as the low educated by the unemployment insurance and food stamps programs. For both groups, rates of wealth accumulation are hardly affected by the changing arrival rate.

7.3 Implications of Government Insurance

Our framework is well suited for evaluating the welfare effects of the various programs. Such an evaluation requires a life-cycle model where risk plays an important role, and where labour supply is endogenous in order to capture the key source of moral hazard and a further mechanism of self-insurance over and above savings. The Food Stamps program tends to provide partial insurance for income loss whatever the source of the loss, while UI offers compensation when income loss is associated with job destruction. As such each program can be thought of as targeting different risks and it is important to quantify how these are valued, although they are unlikely to provide anything close to full insurance. We now turn to a brief examination of the welfare effects of these two social insurance programs.⁴²

We consider a small (1%) increase in the government spending on social insurance and compare the welfare effects of channelling this change, in turn, into UI and into the Food Stamps-type program. This calculation focuses on the insurance benefit of these programs because there is no cross-group redistribution. The results are presented in rows 1 and 2 of Table 6. Row 3 considers the

⁴²There are two caveats to these comments: first, these calculations ignore the interactions that may arise between increases in the tax rate needed to fund increased generosity of a program and the take-up of that program: the increased tax rate will make programs more valuable by reducing the benefit of being at work. Second, in practice these programs are funded by taxing the general population and consequently involve a large component of cross group insurance.

Table 6: Welfare effects of government programmes

	High Education	Low Education
<i>Scenario</i>	<i>Willingness to pay %</i> ($\pi \times 100$)	<i>Willingness to pay %</i> ($\pi \times 100$)
Unemp. Insurance	0.19	0.24
Food stamps	0.25	0.30
Tax Change	0.08	0.15

welfare benefit of using the extra spending to reduce the proportional tax rate. For both education groups the most valuable program is the means-tested program because it provides some insurance against large negative (and permanent) shocks. In considering the tax cut, the two groups are willing to pay 0.08% and 0.15% of consumption, respectively, to see the 1% increase in expenditure going to a tax cut. This implies both groups prefer the money to be spent on UI or the means-tested program, rather than on a decrease in taxation within their own group.⁴³

8 Conclusions

In this paper we have specified a life-cycle model of consumption, labour supply and job mobility when match effects are important determinants of wages. Standard lifecycle models have one source of risk, which is captured by fluctuations in incomes, and is assumed to be exogenous. However, this confounds earnings fluctuations due to exogenous reasons with those resulting from responses to shocks, such as changes in employment or job mobility. In our model we have introduced explicitly the key underlying sources of risk, and further by endogenizing employment and job mobility we are able to provide a framework where the effects of these different sources of risk can be investigated. In our model the underlying sources of risk are from shocks to individual wages, shocks to the firm that lead to job destruction, the process of job offers when the individual is employed and when he is out of work, and the value of the match specific effect that is associated with a job offer. Finally, an important aspect of our model is that we allow for the presence of social insurance programs, mirroring the Food Stamps program, the Unemployment Insurance, Disability Insurance and social

⁴³We do not present numbers for increasing the spending on the disability insurance programme. In our model, spending on the Disability Insurance and Social Security program is valued less than the tax cut. However, our model of DI in effect abstracts from the impact of large or catastrophic health shocks against which insurance is likely to be very valuable and which is provided mainly by DI in reality.

security for the elderly. This means that we can measure the effects of shocks allowing for the institutional environment that offers some partial insurance.

First we estimate the wage process using data from the SIPP, allowing for job to job mobility and controlling for the endogeneity of participation and job changes. We show that representing the stochastic process of wage rates as a permanent shock (unit root) and a transitory component (measurement error), alongside a match-specific fixed effect, we can fit the wage data very well. We show that ignoring job mobility overstates productivity risk, as measured by the standard deviation of a permanent shock, by 50%: much of the earnings fluctuations often attributed to income risk is due to individuals accepting better wage offers. Some of these wage fluctuations do not reflect risk; they are the response to obtaining a better offer in a world with heterogeneous matches. However, match heterogeneity implies new sources of risk: individuals may lose a valuable match when the job is destroyed, and they are uncertain about the value of the offer they will receive.

The remaining parameters are obtained by calibrating the model to fit the observed life-cycle participation profiles and unemployment durations. We show that an alternative specification of the wage process, where there are no permanent productivity shocks to the individual, but only shocks to the match specific effect, cannot reproduce the participation profiles. Our conclusion is that permanent shocks are an important ingredient for fitting the observed data on life-cycle work patterns. After showing that our specification does fit a number of key facts we use it to carry out comparative static analysis with respect to changes in the different types of risk. We show that increases in productivity risk decrease output and welfare: indeed individuals are willing to pay more than the lost output to return to the pre-increase state of affairs. Increases in job destruction also cause declines in welfare and output. However, individuals are willing to pay less than the decline in output to avoid the increase in this risk. On the other hand, small increases in the variance of match heterogeneity, although leading to increased risk, increases welfare (as well as output) because it offers enhanced opportunities for wage growth through job mobility. This makes it all the more important not to attribute wage fluctuations due to job mobility to productivity risk. In our comparative static exercises we also consider the effect of risk changes on participation and savings.

In the final section we quantify the value of two of the social insurance programs: Food Stamps and Unemployment insurance. Both the low and the high educated prefer an increase in the gen-

erosity of either of these programs to a decrease in proportional taxation costing the same amount. Both education groups prefer the extra expenditure spent on Food Stamps rather than UI. This is because Food Stamps is a broader insurance program, offering some protection against declines in productivity, while UI is just insuring a transitory loss of work.

A Appendix: Numerical Solution

Households have a finite horizon and so the model is solved numerically by backward recursion from the terminal period. At each age we solve the value function and optimal policy rule, given the current state variables and the solution to the value function in the next period. This approach is standard. The complication in our model arises from the combination of a discrete choice (to participate or not) and a continuous choice (over saving). This combination means that the value function will not necessarily be concave. The discrete choice about whether to move or not is less problematic because we assume that there is no cost of moving. This means that the decision to move depends only on the relative size of the match effect in the current and new firm.

There are five state variables in this problem: age, employment status, the asset stock, the permanent component of earnings, u_{it} , and the match component, $a_{ij}(t_0)$. Age and employment status are both discrete. We also discretize both the permanent component of earnings and the distribution of possible matches, leaving the asset stock as the only continuous state variable. Since the permanent component of earnings is non-stationary, we are able to approximate this by a stationary, discrete process only because of the finite horizon of the process. We select the discrete nodes in this process to match the paths of the mean shock and the unconditional variance over the life-cycle. In particular, the unconditional variance of the permanent component must increase linearly with age, with the slope given by the conditional variance of the permanent shock. Our estimates of the wage variance are for annual shocks, but the model period is one quarter. We reconcile this difference by imposing that each quarter an individual receives a productivity shock with probability 0.25, and this implies that productivity shocks occur on average once a year. This timing means individuals who stay with the same firm expect pay to be constant over a year.

Value functions are increasing in assets A_t but they are not necessarily concave, even if we condition on labor market status in t . The non-concavity arises because of changes in labor market status in future periods: the slope of the value function is given by the marginal utility of consumption, but this is not monotonic in the asset stock because consumption can decline as assets increase and expected labor market status in future periods changes. This problem is also discussed in Lentz and Tranaes (2001). By contrast, in Danforth (1979) employment is an absorbing state and so the conditional value function will be concave. Under certainty, the number of kinks in the conditional

value function is given by the number of periods of life remaining. If there is enough uncertainty, then changes in work status in the future will be smoothed out leaving the expected value function concave: whether or not an individual will work in $t + 1$ at a given A_t depends on the realization of shocks in $t + 1$. Using uncertainty to avoid non-concavities is analogous to the use of lotteries elsewhere in the literature. In the value functions (7) and (8), the choice of participation status in $t + 1$ is determined by the maximum of the conditional value functions in $t + 1$.

In solving the maximization problem at a given point in the state space, we use a simple golden search method. We solve the model and do the calibration assuming this process is appropriate. We then check that the results in our baseline case are unaffected when we use a global optimizing routine, simulated annealing. It is worth stressing that there are parameter values for which the techniques we used do not work. In particular, as the variance of shocks gets sufficiently low, the non-concavities in the expected value functions become problematic.

B Appendix: Deriving Moments for the Variance of Wages

B.1 The Preferred model

In our preferred model, wages are given by

$$\ln w_{it} = d_t + x'_{it}\psi + u_{it} + e_{it} + a_{ij(t_0)}$$

where $u_{it} = u_{it-1} + \zeta_{it}$ is the permanent component, e_{it} the measurement error, and $a_{ij(t_0)}$ is the match effect. Thus wage growth is

$$\Delta \ln w_{it} = \Delta d_t + \Delta x'_{it}\psi + \zeta_{it} + \Delta u_{it} + \xi_{it}M_{it}$$

where $\xi_{it} = (a_{ij(t)} - a_{ij(t_0)})$ denotes the change in the match effect for those who switched employment. The latent indices associated to working and moving are:

$$\begin{aligned} P_{it}^* &= z'_{it}\gamma + \pi_{it} \\ M_{it}^* &= k'_{it}\theta + \mu_{it} \end{aligned}$$

for all t . Define participants in period t by $(P_{it} = 1) \equiv (P_{it}^* > 0)$ and similarly those who have changed workplace since the previous year by $(M_{it} = 1) \equiv (M_{it}^* > 0)$.

Wage growth is measured annually. However, all decisions by the individual are made quarterly in the model. To make these two consistent we assume that the individual receives a wage shock each quarter with 0.25 probability. Then when we observe a wage over a year we assume it is the result of aggregating wages over the quarters that the individual worked in this firm. Thus the participation Mills ratio we use is the average Mills ratio over the number of quarters the individual worked with the firm he is currently observed. This effectively assumes that the participation and wage shocks are both independent over time.⁴⁴ Denoting the number of quarters the individual worked by Q and noting that the Mills ratio for a working quarter is $\lambda_{i(q)}^P = \frac{\phi(z'_{it(q)}\gamma_q)}{\Phi(z'_{it(q)}\gamma_q)}$ (with ϕ and Φ being the standard normal density and probability functions respectively), for annual wages the Mills ratio we use is the average of the quarterly Mills ratios, i.e. $\overline{\lambda_{it}^P} = \frac{1}{Q} \sum_{q=1}^Q \frac{\phi(z'_{it(q)}\gamma_q)}{\Phi(z'_{it(q)}\gamma_q)}$. When computing higher order moments we also construct $\overline{z'_{it}\gamma\lambda_{it}^P} = \frac{1}{Q} \sum_{q=1}^Q \left\{ \frac{\phi(z'_{it(q)}\gamma_q)}{\Phi(z'_{it(q)}\gamma_q)} z'_{it(q)}\gamma_q \right\}$. This procedure is not exact aggregation, but an approximation because the dependent variable is the log of the average wage and the right hand side is a model for the average log wage (both within the year). However, the approximation error is likely to be negligible for two reasons: first, it is proportional to the difference between the arithmetic and geometric mean of wages within a year that depends on the within year/within firm variance of wages, which is small. Second, any systematic error in the levels is removed by taking the growth of wages and the remaining part will be absorbed by the measurement error, leaving the estimate of the variance of the permanent effect σ_ζ unaffected.

In all calculations the mobility equation is kept constant across quarters and measures the probability that an individual changes jobs between any two years, which is what we need to correct for mobility selection.

Conditioning on participation in periods t and $t - 1$, we obtain:

⁴⁴We make this assumption for tractability: otherwise we would have to jointly condition on an eight dimensional selection vector, i.e. the sequence of participation decisions in each quarter over two years. As we do it we still condition on each of the participation outcomes but not on the joint event.

$$\begin{aligned}
E(\Delta \ln w_{it} | P_{it} = P_{it-1} = 1) &= E(\Delta \ln w_{it} | M_{it} = 0, P_{it} = P_{it-1} = 1)(1 - \Pr(M_{it} = 1)) \\
&\quad + E(\Delta \ln w_{it} | M_{it} = 1, P_{it} = P_{it-1} = 1) \Pr(M_{it} = 1) \\
&= \Delta d_t + \Delta x'_{it} \psi + G_{it}
\end{aligned}$$

where

$$G_{it} = \rho_{\zeta\pi} \sigma_{\zeta} \overline{\lambda_{it}^P} + \rho_{\xi\pi} \sigma_{\xi} \overline{\lambda_{it}^P} \Phi(k'_{it} \theta) + \rho_{\xi\mu} \sigma_{\xi} \phi(k'_{it} \theta) + \rho_{\xi\pi-1} \sigma_{\xi} \overline{\lambda_{it-1}^P} \Phi(k'_{it} \theta)$$

The ρ_{ls} are correlation coefficients between stochastic terms l and s . Thus, G_{it} is a “selection term” accounting for conditioning on multiple indices. The estimation of the equation above is standard (Heckman 2-step method).

The variances of the wage shocks are identified by the restrictions imposed on the moments of residual wage growth $g_{it} \equiv \zeta_{it} + \Delta u_{it} + \xi_{it} M_{it}$. Using formulae from Tallis (1961), the first moment for job stayers and movers respectively is:

$$E(g_{it} | P_{it} = P_{it-1} = 1, M_{it} = 0) = -\rho_{\zeta\mu} \sigma_{\zeta} \tilde{\lambda}_{it}^M + \rho_{\zeta\pi} \sigma_{\zeta} \overline{\lambda_{it}^P} \quad (14)$$

$$E(g_{it} | P_{it} = P_{it-1} = 1, M_{it} = 1) = (\rho_{\zeta\mu} \sigma_{\zeta} + \rho_{\xi\mu} \sigma_{\xi}) \lambda_{it}^M + (\rho_{\zeta\pi} \sigma_{\zeta} + \rho_{\xi\pi} \sigma_{\xi}) \overline{\lambda_{it}^P} + \rho_{\xi\pi-1} \sigma_{\xi} \overline{\lambda_{it-1}^P} \quad (15)$$

where $\lambda_{it}^M = \frac{\phi(k'_{it} \theta)}{\Phi(k'_{it} \theta)}$ and $\tilde{\lambda}_{it}^M = \frac{\phi(k'_{it} \theta)}{1 - \Phi(k'_{it} \theta)}$. The parameters of the model are clearly not identified from the first moments alone. Consider then the second moment for workers that either stay or move:

$$E(g_{it}^2 | P_{it} = P_{it-1} = 1, M_{it} = 0) = \sigma_{\zeta}^2 \left(1 - \rho_{\zeta\pi}^2 \overline{z'_{it} \gamma \lambda_{it}^P} + \rho_{\zeta\mu}^2 k'_{it} \theta \tilde{\lambda}_{it}^M - 2\rho_{\zeta\pi} \rho_{\zeta\mu} \overline{\lambda_{it}^P} \tilde{\lambda}_{it}^M \right) + 2\sigma_e^2 \quad (16)$$

and

$$\begin{aligned}
E(g_{it}^2 | P_{it} = P_{it-1} = 1, M_{it} = 1) &= \sigma_{\xi}^2 \left(1 - \frac{1}{2} \rho_{a\pi}^2 \overline{z'_{it} \gamma \lambda_{it}^P} - \frac{1}{2} \rho_{a\pi}^2 \overline{z'_{it-1} \gamma \lambda_{it-1}^P} - 2\rho_{a\mu}^2 k'_{it} \theta \lambda_{it}^M \right. \\
&\quad \left. + 2\rho_{a\mu} \rho_{a\pi} \lambda_{it}^M \overline{\lambda_{it}^P} - 2\rho_{a\mu} \rho_{a\pi} \lambda_{it}^M \overline{\lambda_{it-1}^P} - \rho_{a\pi}^2 \overline{\lambda_{it}^P} \overline{\lambda_{it-1}^P} \right) \\
&\quad + \sigma_{\zeta}^2 \left(1 - \rho_{\zeta\pi}^2 \overline{z'_{it} \gamma \lambda_{it}^P} - \rho_{\zeta\mu}^2 k'_{it} \theta \lambda_{it}^M \right. \\
&\quad \left. + 2\rho_{\zeta\mu} \rho_{\zeta\pi} \lambda_{it}^M \overline{\lambda_{it}^P} \right) + 2\sigma_e^2 \quad (17)
\end{aligned}$$

Finally, we consider the first order autocovariance $E(g_{it} g_{it-1} | \cdot)$. At least in principle, we could use information on those who work for three periods in a row and classify them on the basis of their

mobility decisions. In practice, there are too few observations in the relevant categories to be able to get structural identification in this case. We thus assume $\Pr(M_t = 1, M_{t-1} = 1) \approx 0$ and consider only the restrictions on the unconditional autocovariance, namely

$$E(g_{it}g_{it-1}) = -\sigma_e^2$$

The only set of parameters that are left to identify are the cohort-specific variances of the initial draw of the permanent component u_{i0} . In general, at time t :

$$\ln w_{it} = d_t + x'_{it}\psi + \alpha_{ij(t_0)} + u_{i0} + \sum_{s=1}^t \zeta_{is} + \varepsilon_{it}$$

Conditioning on participation in that period:

$$E(\ln w_{it} | P_{it} = 1) = d_t + x'_{it}\psi + (\sigma_\alpha \rho_{\alpha\pi} + \sigma_{u(b)} \rho_{u\pi} + t\sigma_\zeta \rho_{\zeta\pi}) \overline{\lambda_{it}^P}$$

where $\sigma_{u(b)}$ varies with cohort. We thus regress $\ln w_{it}$ onto $d_t, x'_{it}, \lambda_{it}^P$, the interaction of λ_{it}^P with year of birth dummies, and the interaction of λ_{it}^P with experience, using only data on labor market participants. Next, we define the residual in levels as

$$e_{it} = (\ln w_{it} - d_t - x'_{it}\psi) = \alpha_{ij(t_0)} + u_{i0} + \sum_{s=1}^t \zeta_{is} + \varepsilon_{it}$$

The parameter $\sigma_{u(b)}$ can be identified using the following restrictions:

$$E(u_{it} | P_{it} = 1) = \left(\sqrt{2}\sigma_\alpha \rho_{\xi\pi} + \sigma_{u(b)} \rho_{u\pi} + t\sigma_\zeta \rho_{\zeta\pi} \right) \overline{\lambda_{it}^P} \quad (18)$$

$$\begin{aligned} E(u_{it}^2 | P_{it} = 1) &= \sigma_\alpha^2 \left(1 - 2\rho_{\xi\pi}^2 \overline{z'_{it}\gamma\lambda_{it}^P} \right) + \sigma_{u(b)}^2 \left(1 - \rho_{u\pi}^2 \overline{z'_{it}\gamma\lambda_{it}^P} \right) \\ &\quad + t\sigma_\zeta^2 \left(1 - \rho_{\zeta\pi}^2 \overline{z'_{it}\gamma\lambda_{it}^P} \right) + \sigma_\varepsilon^2 \end{aligned} \quad (19)$$

These moments depend partly on some of the parameters characterizing unexplained wage growth (such as $\sigma_\varepsilon, \sigma_\zeta$, etc.). Since estimating all the moments jointly is complicated, we adopt a two-step procedure. We first estimate $\sigma_\varepsilon, \sigma_\zeta, \sigma_\alpha, \rho_{\zeta\pi}, \rho_{\xi\pi}$ using the restrictions on unexplained wage growth (14)-(17). We then use (18) and (19) (conditioning on the estimated values of $\sigma_\varepsilon, \sigma_\zeta, \sigma_\alpha, \rho_{\zeta\pi}, \rho_{\xi\pi}$) to estimate $\sigma_{u(b)}$ separately for each cohort. Due to convergence problems, we fix $\rho_{u\pi}$ to a common value. The estimates of $\sigma_{u(b)}$ for the low education group are plotted against year of birth in Figure 5.

B.2 The Alternative model

Consider a model in which the matching effect follows a random walk over the life of a job:

$$\begin{aligned}\alpha_{ij(t_0)t} &= \alpha_{ij(t_0)t-1} + \omega_{ij(t_0)t} \\ &= \underbrace{\alpha_{ij(t_0)}}_{\text{initial "fixed" matching component}} + \underbrace{\sum_{s=t_0+1}^t \omega_{ij(t_0)s}}_{\text{cumulative shocks to the initial match component}}\end{aligned}$$

The notation is as follows: $j(t_0)$ indexes a job that started in period t_0 . In the first year in a new job (say t_0), the match component is $\alpha_{ij(t_0)}$. In the second year, it's $\alpha_{ij(t_0)} + \omega_{ij(t_0)t_0+1}$, and so forth.

Assume wages are now given by:

$$\ln w_{it} = d_t + x'_{it}\psi + f_i + e_{it} + \alpha_{ij(t_0)t}$$

where f_i is an individual time-invariant random effect. Define the growth residual as:

$$g_{it} = \Delta(\ln w_{it} - d_t - x'_{it}\psi)$$

Then

$$\begin{aligned}g_{it} &= \left[(\alpha_{ij(t_0)t} - \alpha_{ij(t_0)t_0}) - \sum_{s=t_0+1}^{t-1} \omega_{ij(t_0)s} \right] M_{it} + \xi_{ij(t_0)t} (1 - M_{it}) + \Delta\varepsilon_{it} \\ &= \left[\xi_{it} - \sum_{s=t_0+1}^{t-1} \omega_{ij(t_0)s} \right] M_{it} + \xi_{ij(t_0)t} (1 - M_{it}) + \Delta\varepsilon_{it}\end{aligned}$$

where $\xi_{it} = (\alpha_{ij(t_0)t} - \alpha_{ij(t_0)t_0})$ is the innovation in the "initial" match component.

Assume as before that an individual moves ($M_{it} = 1$) if $\mu_{it} > -k'_{it}\theta$ and works at time t if $\pi_{it} > -z'_{it}\gamma$. Identification is based on the following assumption

$$E(\omega_{ij(t_0)s}\pi_{it}) = \begin{cases} 0 & \text{for } s \neq t \\ \sigma_{\omega\pi} & \text{for } s = t \end{cases}$$

and ε is a pure measurement error term (no selection issues). The first moment of g_{it} conditional on moving is:

$$\begin{aligned}
E(g_{it}|M_{it} = 1, P_{it} = 1, P_{it-1} = 1) &= \sigma_\xi \rho_{\xi\mu} \overline{\lambda_{it}^P} + \sigma_\xi \rho_{\xi\pi} \overline{\lambda_{it}^P} - \sigma_\omega \rho_{\omega\pi} \overline{\lambda_{it-1}^P} \\
&\quad - (t - t_0 - 1) \sigma_\omega \rho_{\omega\mu} \lambda_{it}^M
\end{aligned} \tag{20}$$

where we have made the assumption that $E(\omega_{ij(t_0)s\mu_{it}}) = \sigma_{\omega\mu}$ for $s \leq t$. The first moment of g_{it} conditional on staying is

$$E(g_{it}|M_{it} = 0, P_{it} = 1, P_{it-1} = 1) = -\sigma_\omega \rho_{\omega\mu} \tilde{\lambda}_{it}^M + \sigma_\omega \rho_{\omega\pi} \lambda_{it}^P \tag{21}$$

The second moment of g_{it} conditional on moving or staying is

$$\begin{aligned}
E(g_{it}^2|M_{it} = 1, P_{it} = 1, P_{it-1} = 1) &= \sigma_\xi^2 \left(1 - \rho_{\xi\pi}^2 \overline{z'_{it} \gamma \lambda_{it}^P} - \rho_{\xi\mu}^2 k'_{it} \theta \lambda_{it}^M + 2\rho_{\xi\mu} \rho_{\xi\pi} \lambda_{it}^M \overline{\lambda_{it}^P} \right) \\
&\quad + \sigma_\omega^2 \left(1 - \rho_{\omega\pi}^2 \overline{z'_{it-1} \gamma \lambda_{it-1}^P} - (t - t_0 - 1) \rho_{\omega\mu}^2 k'_{it} \theta \lambda_{it}^M \right. \\
&\quad \quad \quad \left. + 2\rho_{\omega\mu} \rho_{\omega\pi} \lambda_{it}^M \overline{\lambda_{it-1}^P} \right) \\
&\quad + 2\sigma_\varepsilon^2
\end{aligned} \tag{22}$$

and

$$E(g_{it}^2|M_{it} = 0, P_{it} = 1, P_{it-1} = 1) = \sigma_\omega^2 \left(1 - \rho_{\omega\pi}^2 \overline{z'_{it} \gamma \lambda_{it}^P} + \rho_{\omega\mu}^2 k'_{it} \theta \tilde{\lambda}_{it}^M \right. \\ \left. - 2\rho_{\omega\pi} \rho_{\omega\mu} \overline{\lambda_{it-1}^P} \tilde{\lambda}_{it}^M \right) + 2\sigma_\varepsilon^2 \tag{23}$$

Finally, we assume as before

$$E(g_{it}g_{it-1}|\cdot) = -\sigma_\varepsilon^2 \tag{24}$$

Notice that $\sigma_{f(b)}^2$ (the variance of the individual random effect) remains not identified from growth terms alone. Identification must come from using levels. However, x_{it} will almost surely contain tenure, and we lack a good instrument for it. To avoid this problem, we consider wages in the first year with the firm (when tenure is 0):

$$\ln w_{it_0} = x'_{it_0} \beta + f_i + \alpha_{ij(t_0)} + \varepsilon_{it_0}$$

Assuming tenure is the only potentially endogenous variable among the x_{it} , we can run a OLS regression (controlling for selection into work), and obtain an estimate of the residual in levels:

$$E(\ln w_{it_0} | P_{it_0} = 1) = d_{t_0} + x'_{it_0} \psi + (\sigma_{f(b)} \rho_{f\pi} + \sigma_\alpha \rho_{a\pi}) \overline{\lambda_{it_0}^P}$$

Hence, we regress $\ln w_{it_0}$ onto x'_{it_0} , $\overline{\lambda_{it_0}^P}$, and the interaction of $\overline{\lambda_{it_0}^P}$ with year of birth dummies, using only data on participants in their first year at a new firm (where tenure is 0). Define next the residual in levels

$$e_{it_0} = (\ln w_{it_0} - d_{t_0} - x'_{it_0} \psi) = f_i + \alpha_{ij(t_0)} + \varepsilon_{it_0}$$

>From here, σ_f^2 can be identified from using the two restrictions:

$$E(e_{it_0} | P_{it_0} = 1) = (\sigma_{f(b)} \rho_{f\pi} + 2\sigma_\alpha \rho_{\xi\pi}) \overline{\lambda_{it_0}^P} \quad (25)$$

$$E(e_{it_0}^2 | P_{it_0} = 1) = \sigma_{f(b)}^2 \left(1 - \rho_{f\pi}^2 \overline{z'_{it_0} \gamma \lambda_{it_0}^P}\right) + \sigma_\alpha^2 \left(1 - 2\rho_{\xi\pi}^2 \overline{z'_{it_0} \gamma \lambda_{it_0}^P}\right) + \sigma_\varepsilon^2 \quad (26)$$

As before, these moments depend partly on some of the parameters characterizing unexplained wage growth. Also in this case, we follow a two-step strategy. We first estimate σ_ε , σ_ω , σ_α , $\rho_{\xi\pi}$ using the restrictions on unexplained wage growth (20)-(24). We then use (25) and (26) to estimate $\sigma_{f(b)}$ separately for each cohort. The estimates of $\sigma_{f(b)}$ for the low education group are plotted against year of birth in Figure 5.

B.2.1 Estimates of Wage Variances in the Alternative Model

The expressions above make clear that one needs data on tenure ($t - t_0$) to identify the parameters of interest in the alternative model. The first SIPP interview contains a special module collecting labor market data, including information on tenure with the current employer. We then update tenure over subsequent interviews using information on firm mobility.

Using a similar estimation procedure with all the suitable corrections for selection for job mobility and employment, we produce estimates for the key parameters, shown in Table 7. Thus σ_ω is the standard deviation of the permanent shock to the match specific effect ($\omega_{ij(t_0)t}$), σ_e is the standard deviation of the transitory shock/measurement error (e_{it}), σ_a is the standard deviation of the initial value of the the match specific effect ($a_{ij(t_0)}$) and σ_f is the standard deviation of individual heterogeneity (f_i). The estimated standard deviation of the innovation is now a bit lower. However, we now need a much higher level of heterogeneity upfront than we do in our model so as to match

Table 7: Standard deviations for the alternative model

Parameter	Estimate	Standard error
σ_ω	0.097	0.040
σ_e	0.085	0.022
σ_a	0.20	0.015
σ_f	0.48	

Note: σ_f is the mean over all cohorts. Standard errors are not corrected.

the cross sectional variance of wages. As we shall see this implies a much higher variance of wages for the young and a far too low variance for older individuals.

In estimating the stochastic process for wages, we also allow for a transitory component. This can be given many different interpretations and these include measurement error and transitory shocks to the match specific effect of wages. Thus, in our preferred specification, we do not ignore shocks to the match specific effect; however, we assume that the shocks do not affect behavior. This is done in part because transitory shocks are likely to be easier to insure and partly because it is hard to decompose the transitory element into that reflecting measurement error and that reflecting substantive shocks to productivity of the match.⁴⁵

⁴⁵Meghir and Pistaferri (2004) show how to bound the variance of transitory shocks.

C Appendix: Tables on Participation, Mobility, Calibration

Table 8: Summary Statistics, SIPP 1993 panel

Variable	Mean	Standard deviation
Earnings	7015.90	5267.61
Typical weekly hours	38.14	16.92
Labor market participation	0.88	0.33
Age	41.12	8.95
White	0.88	0.33
Married	0.74	0.44
Unearned income (net of transfers)	1026.74	2175.21
High education	0.55	0.50
Public sector	0.07	0.26
Northeast	0.21	0.41
North Central	0.28	0.45
South	0.26	0.44

Note: The table report quarterly averages. Earnings and unearned income are converted in real terms using the CPI (CPI=1 in 1992:10).

Table 9: The Participation Equations for each Quarter

	High school or less				College dropout or more			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Age	0.0201 (0.0044)	0.0182 (0.0044)	0.0188 (0.0044)	0.0213 (0.0041)	0.0103 (0.0020)	0.0084 (0.0021)	0.0082 (0.0021)	0.0128 (0.0021)
Age ² /100	-0.0305 (0.0051)	-0.0287 (0.0051)	-0.0295 (0.0050)	-0.0319 (0.0047)	-0.0144 (0.0025)	-0.0122 (0.0025)	-0.0122 (0.0025)	-0.0176 (0.0024)
White	0.1110 (0.0142)	0.1045 (0.0140)	0.0965 (0.0138)	0.0940 (0.0129)	0.0484 (0.0095)	0.0555 (0.0097)	0.0522 (0.0096)	0.0506 (0.0091)
Married	0.1422 (0.0109)	0.1590 (0.0110)	0.1595 (0.0109)	0.1539 (0.0101)	0.0579 (0.0067)	0.0557 (0.0066)	0.0553 (0.0067)	0.0658 (0.0066)
Region dummies	25.89 (3, 0%)	31.87 (3, 0%)	32.51 (3, 0%)	38.77 (3, 0%)	11.37 (3, 1%)	12.83 (3, 1%)	7.11 (3, 7%)	14.08 (3, 0%)
Year dummies	0.46 (2; 79%)	0.88 (2; 64%)	0.15 (2; 93%)	2.66 (2; 26%)	0.51 (2; 77%)	0.09 (2; 96%)	0.43 (2; 81%)	6.46 (2; 4%)
Unearned income	-0.0495 (0.0021)	-0.0433 (0.0020)	-0.0399 (0.0018)	-0.0477 (0.0020)	-0.0114 (0.0007)	-0.0112 (0.0007)	-0.0122 (0.0007)	-0.0118 (0.0007)
UI generosity	-0.0009 (0.0004)	-0.0009 (0.0004)	-0.0009 (0.0004)	-0.0006 (0.0004)	0.0002 (0.0002)	-0.0000 (0.0002)	-0.0000 (0.0002)	-0.0000 (0.0002)

Note: The table reports marginal effects. Asymptotic standard errors in parenthesis. For region and year dummies we report the value of the χ^2 statistics of joint significance and, in parenthesis, the degrees of freedom and the p-value of the test.

Table 10: The Mobility Equation

	High school or less	College dropout or more
Age	-0.0157 (0.0040)	-0.0131 (0.0031)
Age ² /100	0.0150 (0.0048)	0.0123 (0.0038)
White	-0.0278 (0.0131)	0.0127 (0.0088)
Married	0.0075 (0.0092)	-0.0124 (0.0072)
Not-for-profit	-0.0562 (0.0173)	0.0110 (0.0148)
Industry dummies	75.18 (4 df; p-value χ^2 0%)	60.68 (4 df; p-value χ^2 0%)
Region dummies	6.73 (3 df; p-value χ^2 8%)	2.77 (3 df; p-value χ^2 43%)
Year dummies	35.36 (2 df; p-value χ^2 0%)	71.07 (2 df; p-value χ^2 0%)
Unearned income	0.0030 (0.0006)	0.0016 (0.0004)
UI generosity	0.0006 (0.0004)	-0.0001 (0.0003)

Note: The table reports marginal effects. Asymptotic standard errors in parenthesis. For region and year dummies we report the value of the χ^2 statistics of joint significance and, in parenthesis, the degrees of freedom and the p-value of the test.

Table 11: Observed and Matched Moments

Statistic	Age	High Education		Low Education	
		Data	Model	Data	Model
Participation Rate	22 – 31	0.97	0.94	0.90	0.91
	32 – 41	0.95	0.94	0.89	0.89
	42 – 51	0.91	0.92	0.83	0.84
	52 – 61	0.78	0.79	0.63	0.66
Mean Duration	22 – 26	1.59	1.70	1.90	1.92
	27 – 31	1.92	2.16	2.29	2.34
	32 – 36	2.03	2.31	3.07	2.59
	37 – 41	2.37	2.39	2.44	2.84
	42 – 46	2.93	2.83	3.36	3.44
	47 – 51	5.28	5.71	4.07	7.75
	52 – 56	7.09	6.24	7.83	6.90
57 – 61	4.82	3.91	5.01	3.93	

Table 12: Comparative Statics: Varying the job destruction rate δ

δ	π	Output	Mean	Mean		Median ($\Delta A/y$)			Age at max	
			Duration	$\Delta \ln c_t$		25-35	36-50	51-62		
			47-52	25-44	45-62	25-35	36-50	51-62		
<i>High Education</i>										
0.01	0.062	0.122	8.5	0.030	0.014	0.14	0.24	0.09	62	
0.03	0.0	0.0	5.7	0.029	0.013	0.14	0.25	0.076	62	
0.049	-0.046	-0.082	4.9	0.028	0.012	0.10	0.24	0.066	62	
0.07	-0.085	-0.145	4.4	0.027	0.012	0.071	0.24	0.058	62	
<i>Low Education</i>										
0.01	0.101	0.252	13.1	0.019	0.011	0.18	0.21	0.037	59	
0.03	0.043	0.099	9.1	0.018	0.011	0.18	0.20	0.012	58	
0.049	0.0	0.0	7.8	0.017	0.010	0.10	0.17	-0.010	57	
0.07	-0.039	-0.085	7.0	0.016	0.010	0.08	0.15	-0.025	57	

Note: For the columns concerning the amount of assets, the denominator is average realized earnings (net of the fixed cost of work) in the education-specific baseline. Duration is measured in quarters.

Table 13: Comparative Statics: Varying the standard deviation of the match specific effect σ_α

σ_α	π	Output	Mean	Mean		Median ($\Delta A/y$)			Age at max	
			Duration	$\Delta \ln c_t$		25-35	36-50	51-62		
			47-52	25-44	45-62	25-35	36-50	51-62		
<i>High Education</i>										
0.11	-0.187	-0.185	7.6	0.025	0.013	0.10	0.22	0.065	61	
0.22	0.0	0.0	5.7	0.029	0.013	0.14	0.25	0.076	62	
0.33	0.131	0.181	5.4	0.032	0.012	0.19	0.26	0.086	62	
<i>Low Education</i>										
0.11	-0.150	-0.171	9.9	0.014	0.010	0.12	0.16	-0.019	56	
0.22	0.0	0.0	7.8	0.017	0.010	0.10	0.17	-0.010	57	
0.33	0.124	0.184	6.8	0.020	0.010	0.067	0.20	-0.006	58	

Note: For the columns concerning the amount of assets, the denominator is average realized earnings (net of the fixed cost of work) in the education-specific baseline. Duration is measured in quarters.

Table 14: Comparative Statics: Varying the job arrival rate for the Unemployed λ^n

λ^n	π	Output	Mean	Mean		Median ($\Delta A/y$)			Age at max	
			Duration	$\Delta \ln c_t$		25-35	36-50	51-62		
			47-52	25-44	45-62	25-35	36-50	51-62		
<i>High Education</i>										
0.96	0.004	0.003	5.8	0.029	0.013	0.14	0.25	0.078	62	
0.84	0.0	0.0	5.7	0.029	0.013	0.14	0.25	0.076	62	
0.76	-0.003	-0.002	5.8	0.029	0.013	0.14	0.25	0.075	62	
0.66	-0.008	-0.007	6.1	0.029	0.013	0.15	0.25	0.074	62	
<i>Low Education</i>										
0.96	0.010	0.012	7.5	0.017	0.010	0.11	0.17	-0.001	58	
0.84	0.004	0.005	7.7	0.017	0.010	0.11	0.17	-0.005	58	
0.76	0.0	0.0	7.8	0.017	0.010	0.10	0.17	-0.010	57	
0.66	-0.006	-0.007	8.0	0.017	0.010	0.10	0.17	-0.016	57	

Note: For the columns concerning the amount of assets, the denominator is average realized earnings (net of the fixed cost of work) in the education-specific baseline. Duration is measured in quarters.

References

- [1] Abowd, J. and D. Card (1989), "On the covariance structure of earnings and hours changes", *Econometrica* 57, 411-45.
- [2] Abowd, J.M., Kramarz, F. and D.N. Margolis (1999) "High Wage Workers and High Wage Firms" *Econometrica*, 67(2), 251-333.
- [3] Abowd, J.M. and M. H. Stinson (2005) "Estimating Measurement Error in SIPP Annual Job Earnings: A Comparison of Census Survey and SSA Administrative Data" Cornell University, mimeo.
- [4] Altonji, J. G. and N. Williams (2005) "Do wages rise with job seniority? A reassessment," *Industrial and Labor Relations Review*, ILR Review, ILR School, Cornell University, vol. 58(3), pages 370-397, April.
- [5] Attanasio, O., J. Banks, C. Meghir, and G. Weber (1999), "Humps and Bumps in Lifetime Consumption", *Journal of Business and Economic Statistics*, 17(1), 22-35.
- [6] Attanasio, O. and Davis, S. (1996) "Relative wage movements and the distribution of consumption" *Journal of Political Economy* 104: 1227-1262
- [7] Attanasio, O., and G. Weber (1995), "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey", *Journal of Political Economy*, 103(6), 1121-57.
- [8] Baker, M. (1997), "Growth-rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings", *Journal of Labor Economics*, 15, 338-375.
- [9] Banks, J., R. Blundell, and A. Brugiavini (2001), "Risk pooling, precautionary saving and consumption growth", *Review of Economic Studies*.
- [10] Blundell, R., and I. Preston (1998), "Consumption inequality and income uncertainty", *Quarterly Journal of Economics* 113, 603-640.
- [11] Borjas, G.J. and S. Rosen (1980) "Income prospects and job mobility of younger men," in R. Ehrenberg (ed.), *Research in Labor Markets*, Vol.3. Greenwich, CT: JAI Press, 159-181.

- [12] Bound, J., J. Cullen, A. Nichols, L. G. Schmidt (2004), "The Welfare Implications of Increasing Disability Insurance Benefit Generosity", *Journal of Public Economics*, 88 (12): 2487-2514.
- [13] Brown, J.N. and A. Light (1992), "Interpreting Panel Data on Job Tenure", *Journal of Labor Economics*, 10(3), 219-257.
- [14] Browning, M. and T. Crossley (2001) "Unemployment Benefit Levels and Consumption Changes" *Journal of Public Economics*, 80(1), 1-23
- [15] Carroll, C. (1992), "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence", *Brookings Papers on Economic Activity* 1992(2), 61-156.
- [16] Carroll C., and A. Samwick (1997), "The Nature of Precautionary Wealth", *Journal of Monetary Economics* 40.
- [17] Cochrane, J. H. (1991) "A Simple Test of Consumption Insurance" *The Journal of Political Economy*, Vol. 99, No. 5 (Oct., 1991), pp. 957-97
- [18] Costain, J. (1997), "Unemployment Insurance with Endogenous Search Intensity and Precautionary Saving", mimeo.
- [19] Cunha, F., J. J. Heckman and S. Navarro (2005) "Separating Uncertainty from Heterogeneity in Life Cycle Earnings," *NBER Working Papers* 11024
- [20] Danforth, J. P. (1979) "On the Role of Consumption and Decreasing Absolute Risk Aversion in the Theory of Job Search" in: *Studies in the Economics of Search*, ed. by S. A. Lippman, and J. J. McCall, pp. 109-131. North-Holland, New York.
- [21] Deaton, A. (1991) "Saving and liquidity constraints" *Econometrica* 59(5), 1121-1142
- [22] Deaton, A. and C. Paxson (1994) "Intertemporal Choice and Inequality" *Journal of Political Economy*, University of Chicago Press, vol. 102(3), pages 437-67, June.
- [23] Domeij, D. and M. Floden (2006) "The labor supply elasticity and borrowing constraints: why estimates are biased" *Review of Economic Dynamics* 9(2):242-262.
- [24] Farber H. and R. Gibbons (1996) "Learning and Wage Dynamics." *Quarterly Journal of Economics*, 111(4): 1007-1047

- [25] Flinn, C. and J. Heckman (1983), "Are Unemployment and Out of the Labor Force Behaviorally Distinct Labor Force States?", *Journal of Labor Economics*, 1(1), 28-42.
- [26] French, E., (2005) "The effects of health, wealth and wages on labour supply and retirement behavior" *Review of Economic Studies* 72(2): 395-427
- [27] Gourinchas, P., and J. Parker (2002), "Consumption over the Life-Cycle", *Econometrica* 70(1): 47-89
- [28] Gruber, J. (1997), "The Consumption Smoothing Benefits of Unemployment Insurance", *The American Economic Review*, Vol. 87(1),192-205.
- [29] Guvenen, F. (2007), "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?", *American Economic Review*, 97(3), pp. 687-712.
- [30] Guvenen, F. (2008), "An Empirical Investigation of Labor Income Processes", *Review of Economic Dynamics* (forthcoming).
- [31] Guvenen, F. and T. Smith (2008), "Inferring Labor Income Risk from Economic Choices: An Indirect Inference Approach", mimeo.
- [32] Haider, S. and G. Solon (2006) "Life-cycle variation in the association between current and lifetime earnings" *American Economic Review* 96(4): 1308-1320.
- [33] Hansen, G. and A. Imrohoroglu (1992) "The role of unemployment insurance in an economy with liquidity constraints and moral hazard" *Journal of Political Economy* 100(1):118-142
- [34] Harris, M and B. Holmstrom (1982), "A Theory of Wage Dynamics", *The Review of Economic Studies*, Vol. 49, No. 3, pp. 315-333
- [35] Heathcote, J., Storesletten, K. and G. Violante (2007) "Consumption and labor supply with partial insurance: an analytical framework" University of Oslo, mimeo
- [36] Horowitz, J. (2002) "Bootstrap Methods for Markov Processes," *Econometrica* 71: 1049-1082.
- [37] Hubbard, R.G., J. Skinner, and S. Zeldes (1995), "Precautionary Saving and Social Insurance," *Journal of Political Economy*, 103(2), pages 360-99.

- [38] Jacobson, L. S., LaLonde, R. J. and Sullivan, D.G. (1993) "Earnings Losses of Displaced Workers." *The American Economic Review*, 83(4), pp. 685-709.
- [39] Krusell, Per L., Mukoyama, Toshihiko, Rogerson, Richard and Aysegul Sahin (2008) "Aggregate Implications of Indivisible Labor, Incomplete Markets, and Labor Market Frictions". NBER Working Paper No. W13871
- [40] Lentz, R. (2003) "Optimal Unemployment Insurance in an Estimated Job Search Model with Savings" Boston University, mimeo
- [41] Lentz, R. and T. Tranaes (2005), "Job Search and Savings: Wealth Effects and Duration Dependence", *Journal of Labor Economics*, 23(3), 467-490.
- [42] Ljungqvist, L. and T. Sargent (1998), "The European Unemployment Dilemma", *The Journal of Political Economy*, 106(3), 514-550.
- [43] Low, H. (2005) "Self-insurance and unemployment benefit in a life-cycle model of labor supply and savings", *Review of Economic Dynamics* 8(4): 945-975.
- [44] Lucas, R.E. (1987) "Models of Business Cycles", Basil Blackwell, New York
- [45] MaCurdy, T. (1982) "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis" *Journal of Econometrics* 18(1): 83-114.
- [46] MaCurdy, T., Green, D. & Paarsch, H. (1990), "Assessing Empirical Approaches for Analyzing Taxes 53 and Labor Supply", *The Journal of Human Resources*, Vol. 25, No. 3, pp 415 — 490
- [47] MaCurdy, T. and T. Mroz (1995). "Measuring Macroeconomic Shifts in Wages from Cohort Specifications". Technical report, Stanford University.
- [48] Manski, C. (2004), "Measuring Expectations," *Econometrica*, Econometric Society, vol. 72(5), pages 1329-1376.
- [49] Meghir, C. and D. Phillips (2008), "Labour Supply and Taxes", Prepared for the Report of a Commission on Reforming the Tax System for the 21st Century, Chaired by Sir James Mirrlees, forthcoming Oxford University Press.

- [50] Meghir, C. and L. Pistaferri (2004) "Income variance dynamics and heterogeneity" *Econometrica*.
- [51] Meyer, B. (2002) "Unemployment and workers' compensation programs: rationale, design, labour supply and income support" *Fiscal Studies* 23: 1-49
- [52] Moffitt, R. (1989), "Estimating the value of an in-kind transfer: The case of Food Stamps", *Econometrica* 57: 385-409
- [53] Moffitt, R. and P. Gottschalk (1995), "Trends in the Covariance Structure of Earnings in the U.S.: 1969-1987", unpublished.
- [54] Pencavel, J. (2002), "A Cohort Analysis of the Association between Work Hours and Wages among Men", *The Journal of Human Resources*, Vol. 37, No. 2, pp 251 — 274
- [55] Pijoan-Mas, J. (2006) "Precautionary savings or working longer hours" *Review of Economic Dynamics* 9(2):326-352
- [56] Pistaferri, L. (2001), "Superior information, income shocks, and the permanent income hypothesis", *Review of Economics and Statistics*.
- [57] Pistaferri, L. (2003), "Anticipated and unanticipated wage shocks, wage risk, and intertemporal labor supply", *Journal of Labor Economics*.
- [58] Postel-Vinay, F., J.-M. Robin (2002) "Equilibrium Wage Dispersion with Heterogeneous Workers and Firms" *Econometrica*, 70: 2295-1350
- [59] Rendon, S. (2006), "Job Search and Asset Accumulation under Borrowing Constraints", *International Economic Review*, 47(1), 233-263.
- [60] Rogerson, R. and M. Schindler (2002) "The welfare costs of worker displacement" *Journal of Monetary Economics*, vol. 49(6), pages 1213-1234.
- [61] Stinson, M.H. (2003), "Technical Description of SIPP Job Identification Number Editing in the 1990-1993 SIPP Panels", U.S. Census Bureau, mimeo.

- [62] Storesletten, K., Telmer, C., and A. Yaron (2001) “The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk” *European Economic Review* 45, 1311-1339
- [63] Tallis G.M. (1961), “The Moment Generating Function of the Truncated Multi-normal Distribution”, *Journal of the Royal Statistical Society. Series B (Methodological)*, 23(1), 223-229.
- [64] Topel, R.H. (1991) “Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority” *Journal of Political Economy*, 99(1): 145-176
- [65] Topel, R.H, and M.P. Ward (1992) “Job Mobility and the Careers of Young Men” *Quarterly Journal of Economics*, 107(2): 439-79.
- [66] Zeldes, S. P., (1989). “Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence,” *The Quarterly Journal of Economics*, vol. 104(2), pages 275-98.