Wage Risk and Employment Risk over the Life Cycle*

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

May 2006
PRELIMINARY AND INCOMPLETE

Abstract

We define the distinction between productivity and employment risk and estimate the components of risk using wage and mobility data from the Panel Study of Income Dynamics. We then calibrate a model of intertemporal consumption and labor supply and study the effect of the two sources of risk on precautionary saving and labor supply. Finally, we measure the relative welfare costs of employment and productivity risk and the insurance contents of simple government programs.

*This is a preliminary and incomplete draft and may contain errors.

†Cambridge University and IFS.
‡University College London and IFS.
§Stanford University and CEPR.

1 Introduction

There is now an extensive literature analyzing individuals' precautionary response to income risk under incomplete markets.[1] The theoretical literature has made enormous progress in clarifying under which circumstances precautionary behavior arises, and in throwing some light on the role of income uncertainty, income persistence, and the degree of market incompleteness.[2] Empirical analysis has concentrated on assessing the levels of income risk[3] and measuring its effects on life-cycle consumption profiles and wealth accumulation[4]. However, in most earlier studies labour supply is taken as exogenous and there is a unique source of risk attributed to income fluctuations.[5] Here we generalise the model in both directions. First we allow labour supply to be endogenous with individuals deciding whether to work or not depending on the returns to work, thus introducing an additional mechanism for self insurance. Second we allow individuals to make endogenous job mobility decisions in a world where wages depend on match specific effects. This now allows us to introduce a distinction between employment risk and productivity risk.[6] Within this framework we are able to discuss the welfare costs of the different sources of (uninsurable) risk and to quantify the value that individuals attribute to key welfare programmes.

The different sources of risk matter for a number of reasons: First, risks differ in their insurance opportunities. For example, layoff risk is often partially insured by the unemployment insurance system, while individual productivity risk is 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, 1998), the literature on precautionary savings has ignored this and 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

[1] See Browning and Lusardi, 1996; Attanasio, 2000, for recent surveys

[2] See Kimball (1990), Caballero (1990), Deaton (1991)

[3] see MaCurdy, 1982, or Abowd and Card, 1989 and Meghir and Pistaferri, 2005

[4] see amongst others Carroll and Samwick, 1997, Banks, Blundell and Brugiavini, 2001, Zeldes, 1986, Deaton, 1991, Carroll, 1992, or Gourinchas and Parker, 2004 and Attanasio, Banks, Meghir and Weber, 1995.

[5] An exception is Low (2004).

[6] Low (2004) analyzes the joint saving and labour supply decision, but in a context without exogenous job destruction or search frictions. Lentz (2003) analyzes the interaction between search frictions and saving, but ignores the risk to individuals' own productivity which is independent of any particular match. See also Costain (1999) for an equilibrium search model with precautionary savings that attempts to measure the welfare effects of unemployment insurance.

three popular programs in the US: Unemployment Insurance (UI), Disability Insurance (DI), and Food Stamps. Unemployment insurance is aimed at insuring against exogenous job destruction and (partly) against the difficulty of finding a new job. The disability insurance system is supposed to provide insurance against extreme forms of productivity shock which results in permanent inability to work. Finally, the Food Stamps program provides *universal* insurance against low income, whatever its cause. It is worth stressing that all these systems provide partial insurance only.

Second, it is important to distinguish between which movements in earnings reflect choice and which ones reflect uncertainty. We address this issue by allowing for endogenous labor supply and job mobility which implies that a good deal of earnings fluctuations, usually taken as risk, are in fact attributed to choice.[7] Having removed the effect of these sources of fluctuations, we also decompose earnings fluctuations into permanent and transitory shocks with only the permanent shocks being taken as a source of uncertainty.

Thus the contribution of this paper is to provide a life-cycle framework for making a meaningful distinction between the different sorts of risk that people face (productivity risk and employment risk), and to then estimate the extent of risk within this framework using longitudinal data from the PSID.[8] This enables us to show how individuals respond to the different types of risk in a calibrated life-cycle model of intertemporal consumption and labor supply, to calculate the welfare costs of risk allowing for the various substitution effects, and to evaluate the effect of various government insurance programs allowing for the moral hazard distortions they induce.

Consider first the distinction between productivity risk and employment risk. In our model productivity risk is individual-specific uncertainty which exists independently of the employer's characteristics. We follow the empirical evidence on wage dynamics and assume that productivity shocks result in permanent shifts of the wage profile. Unemployment 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 job match while unemployed, and the extent of unobserved heterogeneity across firms. In a fully competitive labour market with no worker-firm match

[7] More generally, we follow the standard approach of using "realizations" to infer risk, thereby assuming that the individual's and the econometrician's information set coincide. Primarily for lack of adequate data, we abstract from the important issues that have to do with consumers having superior information vis-á-vis the econometrician (see Manski, 2004, for a discussion, and Pistaferri, 2001, 2005, for empirical examples).

[8] In future drafts we plan to use the Survey of Income and Program Participation (SIPP) to estimate wage dynamics. We have only preliminary evidence at this time.

heterogeneity and no search costs the distinction between employment and wage/productivity risk is meaningless: In that world individual wages may fluctuate due to unexpected shocks, but given wages individuals decide whether they wish to work or not; unemployment is not a source of risk. It is the interaction of shocks to individuals or firms (job destruction) with search costs and firm heterogeneity that give rise to the distinction between these risks.

To implement our model we estimate a number of parameters capturing the risks that people face and then simulate behavior as the magnitude of the various risks change (which is obtained by changing the value of the estimated parameters). The parameters of interest for our simulations are obtained partly from estimating the characteristics of the wage dynamic process with endogenous participation and mobility choices, and partly from calibrating our life-cycle model to fit observed participation profiles and unemployment durations.

In addressing the question of how individuals respond to risk, we begin by simulating savings and participation behavior for individuals facing the estimated risks. These simulations give an indication of the extent of precautionary behavior (both precautionary saving and precautionary labor supply). We then calculate how much individuals would be willing to pay to avoid the various risks, how much of the precautionary response is due to employment risk and how much to wage/productivity risk. Finally, we measure the value that people attach to the various government provided insurance programs in our model. There is a clear relationship between the results on the costs each source of risk and the value of the government programmes since these are designed to insure different types of risk.

We find the following results. If mobility is ignored, all the wage variability that is due to matching effects is attributed to permanent shocks. Since job mobility is valued because of the absence of a downside (all bad offers can be rejected), this biases upwards the amount of permanent uncertainty by a factor of 40% and leads to a substantial upward bias in the amount of precautionary saving people hold against permanent productivity risk.[9] We also find important differences across education groups. Welfare calculations of the risk premium show that individuals are willing to pay considerably more to avoid productivity risk than employment risk. This is particularly true for higher educated individuals. The difference is not as stark for the lower educated but still the cost

[9] There would be no bias if job mobility were completely at random.

of productivity risk is substantially larger. As a result, in assessing the reasons for holding assets, productivity risk dominates. However, ignoring employment risk leads to inaccurate predictions.

We then evaluate the welfare value of three insurance programmes we have included in the model: Unemployment Insurance, Disability Insurance and Food Stamps. Of these the latter are by far the most valuable. This is because of the importance of productivity risk and the fact that shocks are permanent;

The layout of the paper is as follows. Section 2 presents the model and also discusses the distinction between employment and productivity risk. Section 3 describes the identification strategy, while Section 4 presents the results. Section 5 presents the simulations using the estimated measures of uncertainty and discusses our calculations of the welfare costs of uncertainty. Section 6 concludes.

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. They face multiple sources of uncertainty: in each period working 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 programmes. These are Food Stamps, Unemployment Insurance, Disability Insurance and Social Security (pensions). Any change to these programmes is 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 the effects of current actions on future eligibility to the various programmes. 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.

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:

$$\ln w_{it} = d_t^{ed} + x_{it}'\psi + u_{it} + e_{it} + a_{ij(t_0)} \tag{1}$$

where w_{it} is the real hourly wage, d_t^{ed} represents the log price of human capital at time t for education group ed, x a vector of regressors including age, ψ^{ed} is a vector of parameters specific to the individual's education group, u the permanent component of wages, and e the transitory component (which we assume entirely attributable to measurement error in wage data).

The term $a_{ij(t_0)}$ denotes a firm effect (or, alternatively, a firm-worker match specific component): $j(t_0)$ indexes the firm that the worker joined in period $t_0 \leq t$.[10] We model the firm 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 firm 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, only a truncated distribution of ξ_{it+1} is observed. The firm effect $a_{ij(.)}$ is complementary to individual productivity. It is constant over time but it will be assumed uncertain across firms.[11] Both the firm 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 $a_{ij(.)}$ when they meet (jobs are "search goods"). The importance of firm effects in explaining wages has been stressed by Topel (1991) and Topel and Ward (1992). Postel-Vinay and Robin (2002) emphasize that both individual effects and firm effects are needed to explain observed wages.

Following a number of papers in the literature[12] we assume that the permanent component of wages follows a random walk process:[13]

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

[10] We 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.

[11] Ideally we would like to allow for shocks to firm effects. These will act as within firm aggregate shocks and their nature. Restricting firm effects to be constant is forced upon us by the lack of matched firm and individual data. It is not however, hard to extend the model in this interesting way when suitable data becomes available.

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

[13] It 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. On the other hand we ignore transitory shocks to wages (the component e_{it} in (1) is assumed to reflect measurement error).

where ζ_{it} is a random shock, which we take to be uncertain and variable from period to period.[14] 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 firm effects.

Identifying the variance of a transitory shock from that of measurement error is generally not possible, without further assumptions or without results from a validation study. Here we assume that e_{it}, which in practice reflects both measurement error and the transitory shock does not affect behaviour. For all practical purposes we can think of e_{it} as consisting entirely of measurement error.

At this point we should briefly discuss the wage setting mechanism. In a world in which there is both individual and firm level heterogeneity and the characteristics of workers and firms are complementary in production, then quasi-rents are produced and these are shared between the worker and the firm by some mechanism. In this model we assume the resulting sharing rule is fixed and incorporated in the firm effect $a_{ij(.)}$ and that there are no counteroffers in response to outside job offers.[15] We assume that there are constant returns to scale in labour implying that the firm is willing to hire anyone who can produce non-negative rents

2.3 Job destruction and job arrival rates.

In each period the workers may receive an alternative job offer. They arrive at rate λ^{ed} (at most one offer is received in each period). Those who are currently unemployed receive offers at a rate λ^{n}. Individuals become unemployed either because they choose to quit following particular wage realisations or because of exogenous job destruction, which happens at rate δ.

We assume there is no exogenous depreciation of skills following job loss (Rogerson and Schindler, 2001). 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 their job is destroyed.[16] This is especially likely for exogenous job destruction. Thus firm

[14] Farber and Gibbons (1998) 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.

[15] see Postel-Vinay and Robin (2002).

[16] Indeed, as stated by Jacubson, LaLonde and Sullivan (2000), "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".

heterogeneity implies that exogenous job destruction will lead to apparent scarring.

2.4 Individual Optimization

We consider an individual with a period utility function

$$U_t = U(c_t, P_t)$$

where P_t is a discrete $\{0,1\}$ labor supply participation variable and c_t 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 25 to 65 (we indicate the last period of the working life with T), and face an exogenous mandatory spell of retirement of 10 years at the end of life. The date of death is known with certainty.

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, 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

$$\mathbf{U}(\mathbf{c}, \mathbf{P}) = \frac{\left(c \exp\{\theta P\}\right)^{1-\gamma}}{1-\gamma}$$

We consider cases where $\gamma > 1$ and $\theta < 0$, implying that 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 has the form

$$A_{it+1} = R\left[A_{it} + \left(w_{it}\left(1-\tau_w\right) - F_{it}\right)P_{it} + \left(B_{it}E_{it}^{UI}\left(1-E_{it}^{DI}\right) + D_{it}E_{it}^{DI}\right)\left(1-P_{it}\right) + FS_{it}E_{it}^{FS} - c_{it}\right] \tag{3}$$

where A are beginning of period assets, R is the interest factor, w the hourly wage rate, $\overline{h}$ the 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,[17] B_{it} unemployment benefits, FS_{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}^{FS} are recipiency $\{0, 1\}$ indicators for unemployment insurance, disability insurance, and food stamps, respectively.[18] Note also that there are costs to applying for disability which we discuss below. There are no explicit liquidity constraints, but the marginal utility of consumption is infinite at zero consumption and so there is a no bankruptcy constraint, which as a result of uncertainty implies that the young will not generally engage in much uncollateralised borrowing..[19]

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).[20] We assume $B_{it} = b \times w_{it-1}\overline{h}$, subject to a cap, and we set the replacement ratio $b = 75\%$. 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. The benefit of this assumption is 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

[17] The fixed cost of work is a pecuniary proxy for the disutility of work.

[18] We assume that food stamps are paid in cash rather than in the form of coupons. 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.

[19] Consumption during retirement is financed by the assets people have voluntarily accumulated over their working years. We do not have a social security system or mandatory savings into a pension fund (i.e., illiquid retirement wealth). Moving from a voluntary to a mandatory saving system would probably result in greater participation at the beginning of the life cycle - people lose the ability to use their savings to smooth consumption over periods of unemployment and thus work more for precautionary reasons. The effect on participation at the end of the life cycle should go in the opposite direction.

[20] On the other hand, we assume that unemployed people can reject offers and still collect insurance benefits. This is similar to Hansen and Imhoroglou (1990) claim that "it is easier for UI administrator to detect quitters than it is to detect those who turn down job offers while unemployed".

accept a new job offer.[21] The only distortion is introduced by the tax on wages.

Food Stamps In modelling food stamps, we ignore the asset test and gross earnings test (see Blundell and Pistaferri, 2003, for more details on the Food Stamps program) and focus on the net earnings test to derive eligibility and the value of the allowance.[22] Gross income is given by

$$y_{it}^{gross} = w_{it} h P_{it} + \left(B_{it} E_{it}^{UI} \left(1 - E_{it}^{DI}\right) + D_{it} E_{it}^{DI} \right) \left(1 - P_{it}\right) \tag{4}$$

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

$$FS_{it} = \begin{cases} \overline{FS} - 0.3 \times y_{it} & \text{if } E_{it}^{FS} = 1 \text{, i.e., if } y_{it} \leq \underline{y} \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

The maximum value of food stamps, $\overline{FS}$, is set assuming a household with two adults and two children, although in our model there is only one earner. The term $\underline{y}$ 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 ($E_{it}^{FS} = 1$).

Disability Benefits and Social Security A final element of the budget constraint is disability insurance. We assume that workers may find themselves in circumstances that would lead them to apply for disability insurance. 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 an extreme form of health shocks that induce permanent inability to work. Second, we require people to remain unemployed for at least one quarter before being able to apply for disability insurance. 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 apply for disability benefits.

Conditional on applying for benefits, an individual has a fixed probability of obtaining the benefit which we obtain from actual data. If he is successful, he remains eligible for the rest of his working

[21] This is the same assumption as in Krussell and Smith (1999).

[22] The 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 excludes pension wealth (as well as real estate wealth).

life and disability insurance becomes an absorbing state. If not, he has to remain unemployed another quarter before taking up a job. This is the implicit cost of applying for disability benefits: If the application is not accepted one spends a period unemployed even if a job offer was available. Individuals can only re-apply in a subsequent unemployment spell. The presence of disability turns out to be very important in fitting the declining labor force participation profiles with age if food stamps are available at the level observed in practice. Interestingly, if we remove food stamps, applications for disability insurance fall substantially and the participation profile does not decline with age. We discuss this point later.

The value of disability insurance is given by

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

where $\overline{w}_i$ 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 $\overline{w}_i$ 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 (50%, see Bound et al.. 2004) 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 food stamps and for unemployment insurance. We assume there is this difference arises because of the difficulty of verifying disability compared to the income test for FS and the unemployment test for UI.

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

2.5 Employment Risk and Wage Risk

We are now ready to discuss the meaning of employment risk and wage risk in the context of our model.

Consider fluctuations in productivity due to random shocks unknown in advance to the individual. This is the source of productivity risk. In a perfectly competitive labor market with no search frictions and no firm heterogeneity there is effectively no distinction between wage risk and employment risk. The unemployed are those who have received negative productivity shocks such that their productivity is below their reservation wage and so the individual prefers unemployment. In itself this does not constitute employment risk since it is an endogenous decision motivated by low earnings. Thus in the absence of labor market frictions, the distinction is meaningless.

The distinction between employment and productivity risk becomes relevant in the presence of search frictions and/or firm heterogeneity. Job destruction clearly leads to unwanted unemployment. However, if there were no uncertainty about receiving a new offer and no firm heterogeneity, there would be no employment risk as such: jobs would be located instantaneously and whether or not they are acceptable would depend only on individual productivity.

The presence of firm heterogeneity and match specific effects, means that some job may be available with a match value that would lead to a wage worth taking for an unemployed individual. Search frictions however, make it hard to find such a job and creates uncertainty in the length of unemployment. Moreover heterogeneity generates an option value to waiting in the unemployment state if the job arrival rate when on the job (and therefore the likelihood to be matched with a high-wage firm) is lower than the job arrival rate when unemployed. The uncertainty generated because of the combination of firm heterogeneity and search frictions we refer to as employment risk.

The productivity shocks that we observe (due to health shocks, demographic shocks, etc.) are assumed to be uninsurable uncertainty. We assume that there is no commitment from the side of the firm so Harris-Holmstrom 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.

An issue of central importance is the real degree of uncertainty. Measured wages vary; Our structure strips out variability that can be attributed to measurement error. Moreover we do not take entire unexplained variance (after measurement error) as representing uncertainty, but only the innovation to wages from one period to the next. Of course it is still possible that the innovation represents variability fully anticipated by the individual but totally unpredictable from the researcher's information. However we consider this unlikely. These issues can be resolved by "solving backwards" as in Blundell and Preston (1998) and Cuhna, Heckman and Navarro (2004).

Individuals move between firms and this leads to variation in earnings. We do not consider this as productivity risk per se; this is variability in earnings that is the result of a choice made by the individual. There is ex ante uncertainty about what type of firm will make an offer, but the ability to move between firms does not have a downside (i.e., bad offers can be turned down). If such mobility was not possible, this might increase the amount of insurance firms would be willing to offer because of greater worker commitment. Part of the contribution of this paper is to separate out the variability in earnings due to uncertainty and that due to endogenous choices of workers.

2.6 Numerical Solution

There is no analytical solution for our model. 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. The solution method is discussed in more detail in the appendix. Here we describe the main features of the algorithm used.

We start by constructing the value functions for the individual when employed and when out of work. When employed, the state variables are $\left\{A_{it}, u_{it}, a_{ij(t_0)}\right\}$, corresponding to current assets, individual productivity and the firm effect. The firm effect is indexed by t_0, which is the date the job began.[23] When unemployed and not on disability, the state variables are $\left\{A_{it}, u_{it}, DI_t^{Elig}\right\}$, 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).

We consider first the value function for an employed individual. An employed individual in the next period will have the choice of quitting into unemployment, moving to a new job or staying with the firm. However if the job is destroyed the individual will have to move to unemployment. Thus

[23] Ideally 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 $a_{ij(.)}$ may not make an offer to any worker. High $a_{ij(.)}$ 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 $a_{ij(.)}$ firms. At present we ignore this issue.

the value function for an individual i who is working in period t is

$$V_t^e\left(A_{it},u_{it},a_{ij(t_0)}\right)= \tag{7}$$
$$\max_c\left\{\begin{array}{c} U\left(c_{it},P_{it}=1\right)+ \\ \\ \beta\delta E_t\left[V_{t+1}^n\left(A_{it+1},u_{it+1},DI_{it+1}^{Elig}=1\right)\right] \\ \\ +\beta\left(1-\delta\right)\left(1-\lambda^e\right)E_t\left[\max\left\{\begin{array}{c} V_{t+1}^n\left(A_{it+1},u_{it+1},DI_{it+1}^{Elig}=1\right), \\ V_{t+1}^e\left(A_{it+1},u_{it+1},a_{ij(t_0)}\right), \end{array}\right\}\right] \\ \\ +\beta\left(1-\delta\right)\lambda^e E_t\left[\max\left\{\begin{array}{c} V_{t+1}^n\left(A_{it+1},u_{it+1},DI_{it+1}^{Elig}=1\right), \\ V_{t+1}^e\left(A_{it+1},u_{it+1},a_{ij(t_0)}\right), \\ V_{t+1}^e\left(A_{it+1},u_{it+1},a_{ij(t+1)}\right) \end{array}\right\}\right] \end{array}\right\}$$

The expectation operator is conditional on information at time t. If there is no offer available, the expectation operator is over the productivity shock only; if an offer has been received, expectations are also over the type of the firm making the offer.

We consider now the value function for an unemployed individual. 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 had an application turned down in the previous period). The value function when ineligible for disability is given by:

$$V_t^n\left(A_{it},u_{it},DI_{it}^{Elig}=0\right)= \tag{8}$$
$$\max_c\left\{\begin{array}{c} u\left(c_{it},P_{it}=0\right) \\ \\ +\beta\lambda^n E_t\left[\max\left\{\begin{array}{c} V_{t+1}^n\left(A_{it+1},u_{it+1},DI_{it}^{Elig}=0\right), \\ V_{t+1}^e\left(A_{it+1},u_{it+1},a_{ij(t+1)}\right) \end{array}\right\}\right] \\ \\ +\beta\left(1-\lambda^n\right)E_t\left[V_{t+1}^n\left(A_{it+1},u_{it+1},DI_{it}^{Elig}=0\right)\right] \end{array}\right\}$$

The expectation operator is again conditional on whether an offer has been received: if no offer has been received the only remaining uncertainty is over productivity.

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

$$V_t^n\left(A_{it},u_{it},DI^{Elig}=1\right)=\max_{c,App}\left\{u\left(c_{it},P_{it}=0\right)+\beta\left\{\begin{array}{ll} V_{t+1}^A & if\ Apply=1 \\ V_{t+1}^{NA} & if\ Apply=0 \end{array}\right\}\right\} \tag{9}$$

where

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

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.

In each period the individual decides whether to work or not or to move to another job if opportunity arises or to apply for disability benefit based on a comparison of these value functions. Workers either have received an alternative job offer or not. However since only acceptable job offers lead to job switches we do not observe whether stayers or even quitters have received job offers. It simplifies the analysis considerably to assume (in the simulations) that there is no cost of switching firm (unlike the empirical characterization of job mobility), but we do assume that there is a fixed cost of work. The assumption of no mobility cost means the choice of firm does not depend on a value function comparison, but rather involves a simple comparison of the $a_{ij(.)}$ and the individual will move if the new offer is from a higher $a_{ij(.)}$-firm than the current one. If there were a cost of moving, then the choice of which firm to work for cannot be separated from the value function because the choice of whether or not to pay the fixed cost of moving would depend on how long the worker planned to work before retiring, and on expectations of future wage and firm shocks.[24]

The solution of the model consists of policy functions for consumption, participation, etc. Before turning to the results, it is instructive to show part of the solution for the baseline model. Figure 1 shows consumption as a function of assets in period $T-1$ (the year before retirement) for participants and non-participants, and for different firm types, conditioning on individual productivity. The point to stress here is that consumption is not monotonic in the asset stock even when conditioning on labor market status: this is because labor market status in future periods changes as the asset stock

[24] We make the assumption of no mobility costs to simplify the numerical solution. We do not impose this assumption in correcting for selection in mobility in the estimation of the variances.

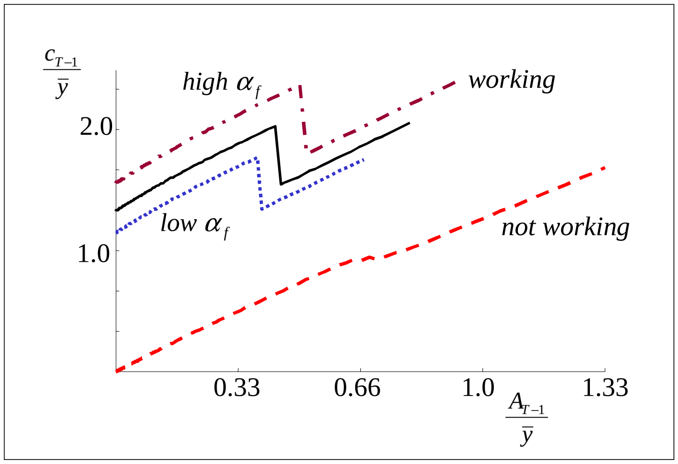

Figure 1: Consumption as a function of current assets conditional on current period work status

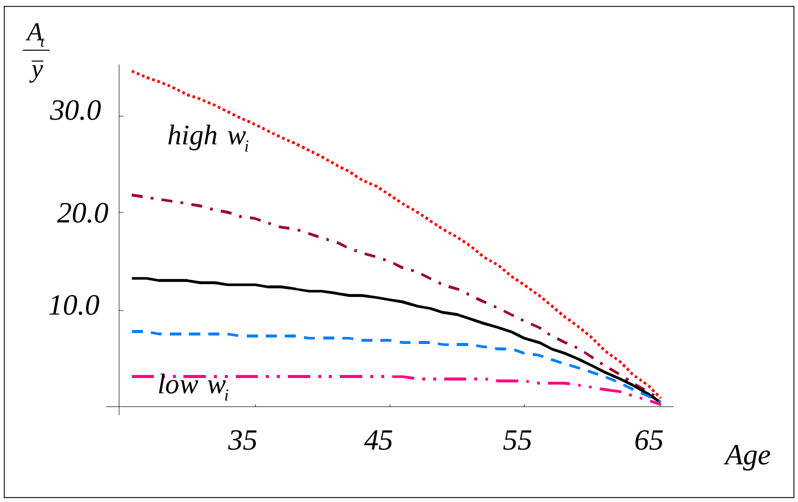

Figure 2: Reservation value of assets for individuals with different levels of wage

increases. For example, the sharp declines in consumption when participating at a given firm in $T-1$ arise at the asset stock which induces the individual not to work in period T. However, in the case of someone in period $T-30$ say the sharp kinks in the consumption function have been smoothed out. Figure 2 shows the reservation asset, i.e., the level of assets that make people indifferent between working and not working. For all groups, it is decreasing over the life cycle. The reason is that only very wealthy individuals can "afford" turning down offers at the beginning of the life cycle. This is a very risky choice because the worker may not receive offers in the future. By age 65 it is at its lowest, although not zero. Once we have values for the parameters of the model, we can simulate consumption and labor supply behavior, analyze the importance of the estimated risk, and running counterfactual experiments. We now turn to a discussion of parameters identification.

3 Estimation and Calibration Method

3.1 Estimation

There are three sets of parameters of interest: (1) Wage dynamics parameters, (2) Labor market frictions, and (3) Preference parameters and the interest rate. One way to obtain model parameters would be to estimate a fully structural model of search, labor supply, and savings. However the data requirements for this task are far beyond what is currently available; estimation would also be computationally hard and in any case require a number of shortcuts. We follow the route of estimating some of the parameters directly from reduced form equations. Following this we calibrate the remaining parameters using duration and participation data. Finally some parameters will be obtained from earlier work in the consumption literature.

We start from wage dynamics parameters. Wages are observed conditional on individuals working. Moreover, within-firm wage growth, which is key to identifying the variance of permanent productivity shocks, is only observed if the individual does not change job. On the other hand changes in wages between jobs, which underlies identification of firm level heterogeneity, are only observed for those moving into another firm. The estimation strategy will thus have to control for all these selection effects. In fact we will show that allowing for the effects of endogenous job mobility leads to much lower estimates of the variance of the permanent innovation to wages and hence lower estimates of the uncertainty facing individuals.

Our approach below can be summarized 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 unobserved wage growth for various subgroups. We simplify the problem by assuming normality of all error terms.

Define the latent utility from labor market participation as $P_{it}^{*} = z_{it}'\varphi + \pi_{it}$ (this approximates the utility from work). 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 period t to an employer that is different from the one in period $t-1$ (this approximates the utility from moving to another firm). The indicator $M_{it} = 1\{M_{it}^{*} > 0\}$ singles out the "movers". We assume: $\begin{pmatrix} \pi_{it} & \pi_{it-1} & \mu_{it} \end{pmatrix}' \sim N(\mathbf{0}, \mathbf{I})$.

Taking first differences of (1), using (2) and recalling that $\xi_{it+1} = \left(a_{ij(t+1)} - a_{ij(t_0)}\right)$, we obtain:

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

Wage growth is only observed for those who work in both periods. To achieve identification of the relevant parameters, we make the following assumptions: $E\left(a_{ij(t)} a_{ij(s)}\right) = \sigma_a^2$ if $j(s) = j(t)$ and zero otherwise. We denote with $\sigma_\zeta^2 = E\left(\zeta_{it}^2\right)$ and $\sigma_e^2 = E\left(e_{it}^2\right)$ (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 zero otherwise.[25] 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. Finally, we allow for contemporaneous correlation between the unobservable of the job mobility decisions (μ) and the shocks to the permanent productivity component and the firm 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. Finally we assume that the distribution of innovations to the firm 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 ($E(e_{it}\pi_{is}) = E(e_{it}\mu_{is}) = 0 \; \forall \; t, s$).

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

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

where G_{it} is a "selection" term induced by labor market participation in both periods and inter-firm mobility (see the Appendix for details).[26] The idea is to estimate the components of this selection

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

[26] In estimation we do not use the restrictions on the parameters of interest imposed by (**??**). This only results in

term in a first stage (running separate probit regressions because of the assumed orthogonality assumption between π_{it} and μ_{it}), and use these to then estimate ψ consistently in a second stage using only participants in both periods.

Define now unexplained wage growth (observed only for participants in both periods):

$$g_{it} = \Delta\left(\ln w_{it} - x_{it}'\psi\right) = \zeta_{it} + \Delta e_{it} + \xi_{it} M_{it} \tag{10}$$

We can now use a method of moments procedure to identify the underlying stochastic process. The key parameters we need to identify are the variance of the permanent shocks and the variance of the firm level heterogeneity. We achieve this by using the first and second moments of the residuals for movers and for stayers. In the process we not only estimate the two variances of interest but also all the relevant correlations that drive selection. The details of the moments we use are given in the Appendix.

Given the complexity of the model, we adopt a multi-step estimation strategy. In a first step, we estimate probit regressions for labor market participation and mobility. In terms of implementation, we face the problem that our theoretical model assumes that labor market participation decisions are taken quarterly, not annually (see below). What we do is to run probit regressions for quarterly participation decisions:

$$P_{it(q)} = \begin{cases} 1 & \text{if } \pi_{it(q)} > -z_{it(q)}'\varphi_q \\ 0 & \text{if } \pi_{it(q)} \leq -z_{it(q)}'\varphi_q \end{cases}$$

estimate ϕ_q, and construct the variables: $\overline{z_{it}'\varphi} = \frac{\sum_{q=1}^{4} z_{it(q)}'\varphi_q}{4}$, and $\overline{\lambda_\pi} = \frac{1}{4}\sum_{q=1}^{4} \frac{\phi\left(z_{it(q)}'\varphi_q\right)}{\Phi\left(z_{it(q)}'\varphi_q\right)}$ which we use as approximations for $z_{it}'\varphi$ and $\frac{\phi\left(z_{it}'\varphi\right)}{\Phi\left(z_{it}'\varphi\right)}$, respectively. At this stage all the individuals (participants and not) are used in estimation. We also estimate a probit for mobility in period t conditioning on observing an individual working in both t and $t-1$. We set $M_{it} = 1$ if this condition is satisfied and if the employer(s) in period t differs from the one(s) in period $t-1$.

In the second step we estimate (**??**) using only labor market participants in both periods. This gives us estimates of ψ and thus allows to construct consistent estimates of wage growth residuals g_{it}. In the final step, we estimate the structural parameters σ_ζ^2, σ_ξ^2, σ_e^2, and the various correlation

a loss of efficiency, but it does not affect consistency. To any event, we compute the standard errors by the block bootstrap.

coefficients. The variance of the firm effect (σ_a^2) can be recovered from $\frac{\sigma_\varepsilon^2}{2} = \sigma_a^2$. We consider a system of three non-linear equations for g_{it}, g_{it}^2, and $g_{it}g_{it-1}$, impose cross-equations constraints and estimate the three equations jointly by non-linear least squares.

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, and pre-estimated residuals and selection terms. We should point out that this procedure is likely conservative, since it allows for more serial correlation than that implied by the moment conditions we use.

3.2 Calibration

The next step is to calibrate the job arrival rates for unemployed and employed individuals (λ^e and λ^n), and the rate at which jobs are destroyed, δ. We also need to obtain estimates of the preference parameters and the fixed cost of work. We impose values for some parameters such as the elasticity of intertemporal substitution and the discount rate using values from elsewhere in the literature. The rest we obtain through calibration using the structural model outlined in section 2. We do so by matching simulated life-cycle profiles of participation and unemployment duration to observed profiles. We use our estimates of the underlying variances (as discussed in the previous section) as inputs into the numerical solution. We are able to estimate the friction parameters conditioning on the wage variances by assuming independence between wage shocks and the number of offers received.

For a given set of parameters, we first solve the theoretical model numerically by backward induction. We then simulate behavior by taking different realizations of the random variables. Individuals differ ex-ante only in their education status, but ex-post they differ by the realizations of wage and employment shocks. Within each education group, we calculate from the simulated paths median duration by age of exit and average participation rates over the life-cycle to give profiles comparable to the data.

We summarize the profiles in the data by the median duration and the average participation rate of individuals in 5 year age bands. We select the five parameters λ^n, λ^e, δ, η, and F as follows:[27]

[27] It is worth stressing that there is interdependence between the moments: for example, longer durations will translate into lower participation rates, and participation rates when young impact on participation choices when old.

First, we take initial values for the 5 parameters. Second, we simulate the model, generate average participation rates and median duration lengths for the different age groups which can be confronted with actual moments in the data. We iterate until there is an acceptable match between simulated and actual moments.

4 Results

4.1 Data

We use the 1993 panel of the Survey of Income and Program Participation (SIPP) to estimate our wage dynamics parameters, and 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 at least a high school diploma, but no college degree), and high (those with a college degree or more). The SIPP data have the advantage of giving information on wages around job switches. However, the short length of the SIPP panel means that it is not useful for duration analysis, so we use PSID data for that purpose.

4.1.1 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 noncash income on a subannual 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 noninstitutionalized 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, see below). 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 for 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).[28]

[28] The raw data can be obtained at http://www.nber.org/data/sipp.html.

Variable	Mean	Standard deviation
Average hourly wage	14.75	7.33
Age	41.24	8.63
White	0.89	0.32
Married	0.78	0.42
Partner's health insurance	0.20	0.40
High education	0.59	0.49
Public sector	0.08	0.28
Northeast	0.20	0.40
North Central	0.28	0.45
South	0.25	0.43

Table 1: Summary Statistics, SIPP 1993 panel

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 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 earnings outliers.[29] 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 1 .

Our measure of (firm-specific) hourly wage is obtained 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 job that pays the highest proportion of annual earnings. In the SIPP, each job (firm) an individual is working for is assigned an ID. We set $M_{it} = 1$ if the employer the individual is working for at time t is different

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

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.

4.1.2 The PSID

The PSID data are drawn from the 1988-1993 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.

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).[30] 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 **x**% of all spells to **x**%.

[30] If 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.

	High school or less				College dropout or more			
	Q1	Q2	Q3	Q4	Q1	Q2	Q3	Q4
Age	0.0189 (0.0041)	0.0161 (0.0041)	0.0168 (0.0041)	0.0201 (0.0042)	0.0103 (0.0020)	0.0077 (0.0020)	0.0076 (0.0020)	0.0100 (0.0021)
$\text{Age}^2/100$	−0.0285 (0.0048)	−0.0255 (0.0047)	−0.0263 (0.0047)	−0.0301 (0.0048)	−0.0142 (0.0024)	−0.0112 (0.0023)	−0.0112 (0.0023)	−0.0143 (0.0025)
White	0.0865 (0.0131)	0.0860 (0.0129)	0.0799 (0.0128)	0.0724 (0.0128)	0.0395 (0.0085)	0.0475 (0.0086)	0.0434 (0.0085)	0.0400 (0.0086)
Married	0.1295 (0.0104)	0.1450 (0.0104)	0.1473 (0.0104)	0.1490 (0.0105)	0.0527 (0.0063)	0.0500 (0.0061)	0.0500 (0.0062)	0.0575 (0.0066)
Region dummies	15.20 (3, 0.2%)	24.43 (3, 0%)	24.14 (3, 0.2%)	29.96 (3, 0.2%)	11.57 (3, 0.9%)	11.51 (3, 0.9%)	6.41 (3, 9%)	10.38 (3, 2%)
Year dummies	0.80 (2; 67%)	1.03 (2; 60%)	0.15 (2; 93%)	5.22 (2; 7%)	0.57 (2; 75%)	0.06 (2; 97%)	0.52 (2; 77%)	8.64 (2; 1%)
Unearned income	−0.0534 (0.0019)	−0.0500 (0.0018)	−0.0452 (0.0017)	−0.0529 (0.0020)	−0.0114 (0.0006)	−0.0111 (0.0006)	−0.0121 (0.0007)	−0.0115 (0.0007)
UI generosity	−0.0008 (0.0004)	−0.0008 (0.0004)	−0.0009 (0.0004)	−0.0007 (0.0004)	0.0002 (0.0002)	−0.0000 (0.0002)	−0.0000 (0.0002)	−0.0000 (0.0002)

Table 2: The participation equation

4.2 Participation and mobility

The first step is to control for selection into employment in the wage equation. To aid in the identification of the wage effects, it is important to have exclusion restrictions. We run quarterly participation probits using PSID data and controlling for demographics and other socio-economic individual characteristics (a quadratic in age, a dummy for whites, region dummies, a dummy for married, year dummies). We use unearned household income, and an index of generosity of the state-level UI system as our exclusion restrictions.[31] They should affect the decision to work, but should have no effect on the wage. The rationale is that these variables influence the opportunity cost of work. The results are reported in Table 2 . The effect of the various covariates is fairly similar across quarters. In all quarters, the probability of working is higher for the younger, the whites, and those who are married. The effect of unearned income is as expected: people with higher unearned income have a higher opportunity cost of working. The effect of UI generosity has the right sign and is mostly statistically significant for the low educated, whereas it is small and insignificant for the high educated.

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

[31] To 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.

	High school or less	College dropout or more
Age	-0.0159 (0.0035)	-0.0119 (0.0029)
Age^2/100	0.0164 (0.0042)	0.0112 (0.0035)
White	-0.0126 (0.0107)	0.0117 (0.0076)
Married	-0.0054 (0.0083)	-0.0208 (0.0072)
Not-for-profit	-0.0443 (0.0152)	0.0199 (0.0142)
Industry dummies	83.20 (4 df; p-value χ^2 0%)	62.72 (4 df; p-value χ^2 0%)
Region dummies	12.06 (3 df; p-value χ^2 0.72%)	5.23 (3 df; p-value χ^2 16%)
Year dummies	32.85 (2 df; p-value χ^2 0%)	76.89 (2 df; p-value χ^2 0%)
Partner's health insurance	0.0316 (0.0096)	0.0330 (0.0075)

Table 3: The Mobility Decision

of freedom and the p-value of the test.

Table 3 reports the estimates of a probit model for the decision to switch to a different firm at time t ($M_{it} = 1$). We use as controls a quadratic in age, dummies for white, region, year, industry, not-for-profit employer, married, and year dummies.[32] Our exclusion restriction is a dummy for whether the worker's partner receives health insurance coverage through her employer. People may be more likely to move if they keep themselves insured through the partner - regardless of whether the change is from an employer with, to one without health insurance or vice versa. The data set also includes information on whether the individual himself is receiving health insurance benefits through the employer. However, this variable may be part of the compensation package received by the individual, and therefore violate the exclusion restriction requirement. We find that the young and the single are more likely to move. We also find that people are less likely to move if they work for a not-for-profit employer. Finally, our exclusion restriction has the expected sign and it is statistically significant: Mobility is less likely if the individual's partner has health insurance available.

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

[32] The industry and not-for-profit dummies refer to the employer at time $t-1$.

of freedom and the p-value of the test.

4.3 Variance Estimates

Armed with these results, we move on to the last step of our estimation procedure, and estimate the structural parameters of interest by NLS imposing constraints across equations. The results are reported in Table 4. We estimate the model for the whole sample (column 1) and separately by education (columns 2 and 3).

Controlling 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 0.11 , the standard deviation of the transitory shock (measurement error), σ_e, 0.084, and the standard deviation of the firm shock, σ_a, 0.213. These parameters are all very precisely measured.

Columns (2) and (3) report the results of estimating the model separately for the high-school graduates or less and the college dropout or more, respectively. Although there are some small differences, with the variance of the permanent shock being higher for the low educated, while the variance of the transitory shock and the variance of the firm effect are higher for the more educated, none of these differences are statistically significant. Thus, perhaps surprisingly the stochastic process is very similar for the two education groups.

What are the likely signs of the correlation coefficients between the random sources of wage growth and heterogeneity in the latent variables? We assume that the permanent shock at time t is orthogonal to participation at time $t-1$, $cov\left(\zeta_{it}, \pi_{it-1}\right)=0$, while $cov\left(\zeta_{it}, \pi_{it}\right) \neq 0$. People who receive a positive shock *should* generally be more likely to work. Consider an individual who has just had a boost in his productivity. If he quits into unemployment, he has no way of monetizing this, which would suggest $cov\left(\zeta_{it}, \pi_{it}\right)>0$. However, since the shock is permanent, it may generate voluntary exit (this is similar to an income effect, although recall that there is a small probability that no new job offers will be received in the future). In the end, the effect is empirically ambiguous. In the whole sample, we estimate a positive effect, statistically significant. However, this masks different effects between education groups, which appear not to be well measured.

	Whole sample (1)	Low education (2)	High education (3)	Neglect selections (4)	Neglect mobility (5)	Neglect participation (6)
σ_ζ	0.114 (0.017)	0.119 (0.051)	0.116 (0.015)	0.148 (0.014)	0.151 (0.013)	0.110 (0.017)
σ_e	0.084 (0.008)	0.068 (0.031)	0.087 (0.010)	0.084 (0.008)	0.084 (0.008)	0.084 (0.012)
σ_a	0.213 (0.013)	0.207 (0.021)	0.215 (0.019)			0.208 (0.008)
$\rho_{\zeta\pi}$	0.265 (0.137)	0.221 (1.901)	-0.147 (0.216)		0.429 (0.138)	
$\rho_{\zeta\mu}$	-0.592 (0.284)	-1.000 (*NA*)	-0.470 (0.391)			-0.835 (0.292)
$\rho_{\xi\pi}$	0.325 (0.240)	0.569 (0.278)	-0.444 (0.725)			
$\rho_{\xi\pi_{-1}}$	-0.242 (0.162)	-0.243 (0.162)	0.072 (0.619)			
$\rho_{\xi\mu}$	0.245 (0.151)	0.403 (0.198)	0.235 (0.211)			0.355 (0.155)

Note: σ_ζ, σ_e, and σ_a are the st.dev. of the permanent shock, measurement .error, and firm/matching effect. $\xi = a_j - a_{j-1}$. $\rho_{\zeta\pi}$ ($\rho_{\xi\pi}$) is the correlation between the permanent shock (mobility premium) 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 are computed using the block bootstrap

Table 4: Wage variance estimates

In the absence of mobility costs we would expect $cov(\zeta_{it}, \mu_{it})$ to be zero because productivity innovations are fully "portable". However with mobility costs a positive productivity shock is likely to make moving worth while for smaller improvements in firm quality implying a positive correlation. As it turns out the correlation is negative and not significant for the sample overall and for the higher education group. For the lower education group the estimate seems to be ill determined and when let free decreases beyond -1. The reason we find $\rho_{\zeta\mu} < 0$ is most likely reflecting a time aggregation issue. We classify as movers those who change their employer between $t-1$ and t. However, it is possible that those who face negative shocks to their permanent productivity component ($\zeta_{it} < 0$)decide to quit and then re-enter the following period (which would be classified as $M_{it} = 1$).

The "firm effect shock" may be correlated with both π_{it}, π_{it-1}, and μ_{it}. The correlation with μ_{it} is self-explanatory. Mobility gains (high realizations of ξ) should be associated with a greater likelihood of observing a move, which would suggest $cov(\xi_{it}, \mu_{it}) > 0$. This is indeed what we find in the whole sample and across education groups.

The correlation between ξ_{it} and π_{it} (and π_{it-1}) is also relatively clear. Suppose one receives a high realization of ξ_{it}. If the individual quits, he loses the ability to step up on the ladder because good offers cannot be recalled (at least in the model). Thus one would expect $cov(\xi_{it}, \pi_{it}) > 0$ (and, symmetrically, $cov(\xi_{it}, \pi_{it-1}) < 0$). This is what we find in the whole sample and in the subsample of low educated individuals, while among the high educated the effects have the opposite signs but are not well measured.

What happens if we ignore the fact that mobility is endogenous and attribute all wage fluctuations to risk. 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, 2004). In column (4) we report the results of this experiment as well as of not accounting for selection into work. They show that the variance of the permanent shock σ^2_ζ almost doubles, from 0.013 to 0.022.

To see the effect of ignoring selection in column (5) we ignore the mobility decision but account for the endogenous participation choice, while in column (6) we do the opposite. It is clear that what really matters is the firm mobility decision. Indeed, neglecting participation reduces the variances of interest but the effects are minuscule.

These experiments show that a large amount of year-to-year wage variability is due to people moving to different firms and that ignoring this source of variability leads to wrong inferences regarding the extent of permanent productivity risk. As we discuss below, this has the main consequence of exaggerating the amount of precautionary saving individuals hold to self insure against this risk. This is because people know their firm type once they have chosen the firm they want to work for. If jobs were not subject to destruction, the amount of wage variability induced by the firm effect would not per se represent economic risk to be insured against.[33]

4.4 Calibrated Parameters

Individuals have a 40 year working horizon (age 22-62) followed by a deterministic 10 year retirement spell. One period is assumed to be one quarter and so the model is solved for 160 periods when labor supply is chosen. A new job offer may be received each quarter, and similarly, the possibility of firm destruction is a quarterly event and decisions are taken each quarter. Further, each quarter

[33] While we have not imposed restrictions on the coefficients of the wage growth equation (**??**), we have checked whether the estimated coefficients are consistent with the structural estimates reported in Table 4, and found no violations.

	High Education		Low Education	
	Data	Model	Data	Model
Participation Rate Age $22-46$	0.96	0.94	0.90	0.89
Participation Rate Age $47-61$	0.85	0.89	0.76	0.79
Median Duration Age $22-46$	1.00	1.00	1.00	1.00
Median Duration Age $47-61$	2.00	2.00	2.00	2.00

Table 5: Observed and Matched Moments

individuals receive a productivity shock with probability 0.25 so 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. We measure unemployment durations by the number of quarters an individual is unemployed/out of the labor force. We set the coefficient of relative risk aversion γ equal to 1.5, taken from Attanasio and Weber (1995). The real interest rate is set equal to the real return on short term treasury bonds, $r = 0.015$, and this is set equal to the discount rate $(\frac{1}{\beta} - 1)$. The extent of deterministic wage growth is taken from the PSID. There is no exogenous funding of consumption in retirement: individuals have to save for their retirement. Details of the food stamp program, disability insurance and unemployment insurance are set to match actual programs as discussed in section 2 above.

As described in the data section, the PSID asks individuals to report their employment status in each month of the previous calendar year. We use the answers to these questions to construct unemployment duration and a quarterly participation indicator for each individual. 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).[34] 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.[35] In

[34] If 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.

[35] In calculating durations, we take our sample to be individuals who exit between 1988 and 1992. However, we

Parameter	High Education	Low Education
Job Destruction Rate δ	0.020	0.044
Job Arrival Rate - Employed λ^n	0.65	0.71
Job Arrival Rate - Unemployed λ^e	0.55	0.57
Fixed Costs of Work F	0.54	0.465
Preference for Leisure η	0.04	

The values of δ, λ^n and λ^e are given as quarterly rates. We impose that the utility cost of participation η is the same across education groups. The calibrated value of η is equivalent to consumption being 7.7% lower when the individual is participating. This value also implies consumption and leisure are substitutes. 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.

Table 6: Parameters obtained through calibration

table 6, we present the calibrated parameter values.

Figures 3 and 4 show participation profiles for the low educated and high educated. Each figure compares the profile in the data with the calibrated profile. For both education groups, participation rates are fairly constant until age 45, followed by a sharp decline to age 62. Part of this fall reflects early retirement, rather than temporary periods 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. We chose parameter values to match average participation rates and median duration rates by education group among individuals in different age groups. Our match to participation is fairly good for both skill groups.

Figures 5 shows median duration over the life-cycle in the simulations and in the data for the low and high 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. In the simulations, durations are measured directly in quarters. Figure 7 shows comparable figures for mean durations. The duration data is skewed with the mean lying above the median particularly for older individuals. We do not attempt to match mean durations because of the problem of right-censoring in the data, particularly among the old. Our simulated median durations are fairly close to the data.

use more recent years of PSID data to calculate durations for those whose spells are right-censored by the 1988-1992 window.

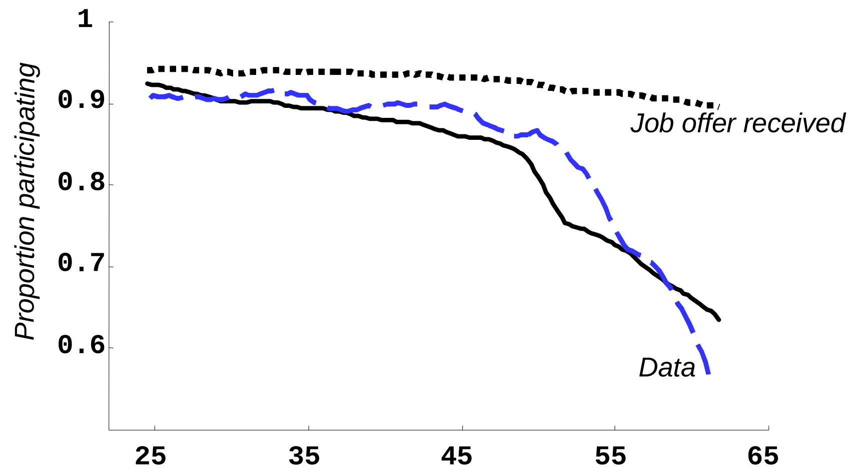

Figure 3: Actual and fitted participation profiles for the low education group

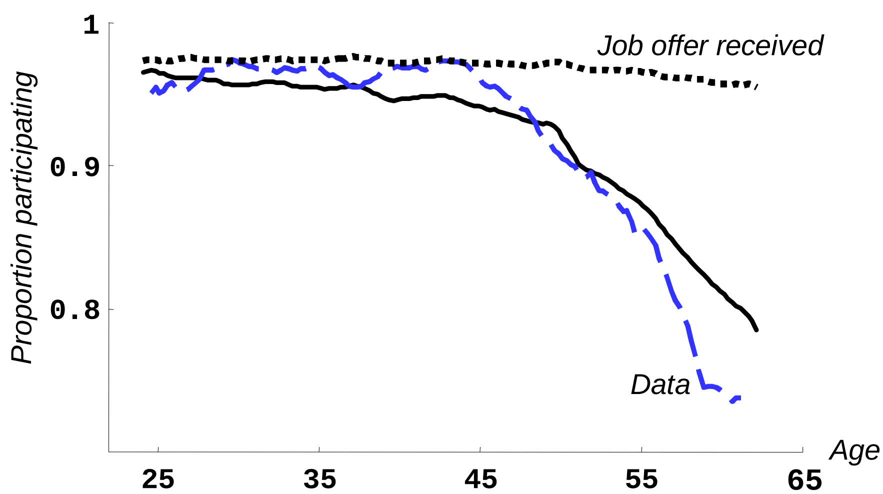

Figure 4: Actual and fitted participation profiles for the High education group

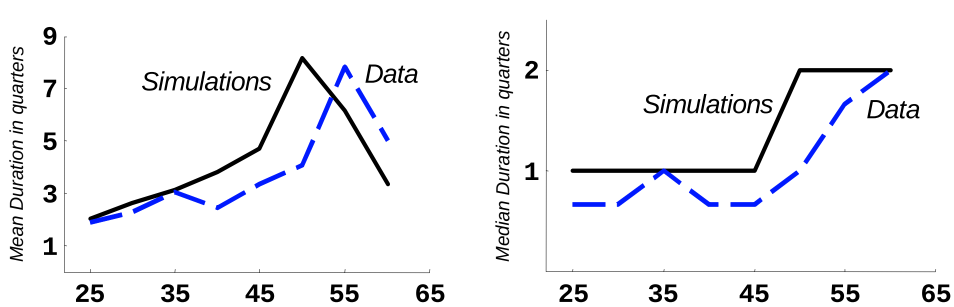

Figure 5: Fitted and actual durations for low education individuals (mean and median)

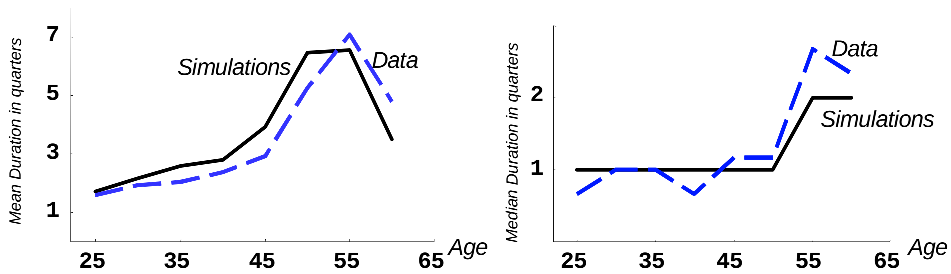

Figure 6: Fitted and actual durations for low education individuals (mean and median)

We assume $\lambda^n, \lambda^e, \delta, F$ and η are independent of age and so the age effects that we find in 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. If we allowed parameters to vary with age, we could fit *any* profile.

4.5 Implications of the model

We have calibrated the model using only participation and unemployment duration data. However, the model has implications for a range of different variables. In particular, we can use the model to predict the proportion of people on disability by age 62, the duration of employment ("tenure"), the wage loss associated with a spell of unemployment, and the extent of consumption loss on unemployment. Table 7 reports the model predictions and corresponding statistics in the data for the variables that are not calibrated.

4.5.1 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 two main difficulties with these data. First, these spells are right-censored. Second all spells are left censored. Third, a number of authors have questioned the reliability of these measures of reported tenure (Brown, 1992). Because of these difficulties we do not use these data explicitly in the calibrations, but we show relevant statistics in Table 7. For the low educated, we predicted average employment durations of 13 quarters for those aged below 46, compared to a value of 12 in the data. For the high educated the figures are 27 and 12 respectively. The striking

Statistic	High Education		Low Education	
	Data	Model	Data	Model
Median Duration Employment Age $22-46$	12	27	12	13
Mean Consumption Loss† Δlnc_{t+1}	-14.0, -6.80	-9.7	-14.0, -6.8	-5.8
Mean Wage Loss ††				
Δlnw_{t+1}		-12.5%		-16.3%
Δlnw_{t+4}				-7.1%

Employment durations are reported as number of quarters until employment with a particular employer ends. The numbers for consumption loss are taken from Browning and Crossley (2001) and Gruber (1997). Neither reports consumption loss by education.

Table 7: Model Implications

point about the simulated durations is that predicted employment durations are long despite the frequent arrival of offers. This is because offers come from the unconditional distribution of firms and it is increasingly harder to beat the offer from the current firm as workers age and find better matches.

4.5.2 "Scarring" effects

There is empirical evidence that displaced workers experience earnings losses following job loss. Some authors impute this to skill depreciation during periods of unemployment (Richardson and Schindler, 2001; Ljunqvist and Sargent, 2002). However, 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 6 the extent of the wage fall on re-entry.

For the high educated, wage on re-entry are, on average, 12.5% lower than before displacement. For the low educated, the loss is 16.3% but quickly declines to 7.1%. These figures are slightly lower than what found in the literature. In particular, we contrast these figures with those reported by Jacubson, LaLonde, and Sullivan (1996) 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, one implication of our model is that "scarring" is likely short lively. Indeed, we calculate that 1 year after separation, earnings of low educated are only 7.7% than in their pre-displacement period. This figure is very close to the one we extrapolate from Jacubson, LaLonde, and Sullivan (1996) over a similar time horizon.

4.5.3 Consumption fall at unemployment

Some recent papers have explored empirically the consumption loss associated with unemployment (Gruber, 1997; Browning and Crossley, 2001). For comparison, we show simulated consumption loss in our model. In Table 6 we report average consumption loss by education group and compare to estimates in the literature. Consumption losses are higher for the high-educated (-10.8% vs. -6.8%) because means-tested government insurance programs are less effective for them. These figures sit in between estimates from the two papers mentioned above.

5 Welfare Costs

One of the main aims of the paper is to show the extent to which different types of risk matter for individual welfare. This is relevant particularly when evaluating policies such as unemployment insurance or earnings insurance which target 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 evaluate the risk sharing benefits of different sorts of insurance as well as identifying the behavioral effects induced by the insurance programs.

The model is partial equilibrium in that the wage process and interest rate are exogenous but we require the government budget to balance in the following sense:

$$\sum_{i=1}^{N}\sum_{t=1}^{T}\frac{1}{R^t}\left[\left(B_{it}E_{it}^{UI}\left(1-E_{it}^{DI}\right)+D_{it}E_{it}^{DI}\right)\left(1-P_{it}\right)+E_{it}^{FS}FS_{it}\right]=\sum_{i=1}^{N}\sum_{t=1}^{T}\frac{1}{R^t}\tau_w w_{it}\overline{h}P_{it}+Deficit \tag{11}$$

where the deficit term will be kept constant across all simulation experiments. We select τ_w to satisfy this government budget constraint, but assume that individuals take τ_w as given.[36] 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. Further, since there are no aggregate shocks in the economy and no business cycle fluctuations, we do not consider the value

[36] We 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.

of, for example, smoothing the effect of the business cycle (Lucas, 1987; Storesletten et al., 2001). We make these assumptions to focus on the cost to the individual of idiosyncratic risk which would be entirely smoothed out in a first best setting. Allowing the budget to balance over all education groups would confuse the issue we are considering with distributional questions.

One important issue in considering changes to risk in the model is that any change in the risk properties is likely to have labor supply effects, which in practice can have an effect on equilibrium wages. We appeal to an open economy argument to justify our assumption of no such general equilibrium effects in our context.[37]

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

$$E_0 U^k = E_0 \sum_t \beta \frac{(c_t^k)^{1-\gamma}}{1-\gamma} \exp\left\{\eta h_t^k\right\}$$

where the superscript k refers to the implied consumption stream in the baseline economy ($k = 0$) or an alternative economy with different risk characteristics ($k = 1$) 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 face environment 2 rather than 1. This is implicitly defined by

$$EU^2(\pi) \equiv E_0 \sum_t \beta \frac{\left(\pi c_t^2\right)^{1-\gamma}}{1-\gamma} \exp\left\{\eta h_t^2\right\} = EU_1$$

which implies that

$$\pi = \left[\frac{EU^1}{EU^2(\pi = 1)}\right]^{\frac{1}{1-\gamma}}$$

We report values of π for small changes in risk. Specifically our welfare measure is defined by

$$\Delta_W = \frac{\pi}{\Delta\sigma/\sigma}$$

and is interpreted as the proportional change in consumption that an individual is willing to pay for a "small" change in the environment expressed as a proportional change ($\Delta\sigma/\sigma$) in some parameter σ (say a change in the standard deviation of a shock or in the rate of arrival of job offers).

Our model is not a general equilibrium model and the arrival rate of offers is taken as exogenous. This is why we have taken the route of considering small departures from the current environment when considering changes to risk. To consider impacts on a larger scale it would be important to model the firm side in some detail and generate

[37] In an economy with sufficient number of tradeable goods (more than the inputs to production) prices will equalise and such policies will just affect the composition of production.

5.1 Welfare Cost of Risk

Table 8 shows results for changes in a number of relevant parameters. The figures represent the elasticity Δ_W for changes in parameters that reflect various sources of risk.

The first result that stands out is the effect of productivity risk (σ_ζ) on welfare: Individuals are willing to pay 0.5% of consumption over the lifecycle to obtain a 1% reduction in the standard deviation of the permanent shock to productivity. Storesletten et al. (2001) calculate the welfare benefit of removing variation in the extent of idiosyncratic risk over the life-cycle. Such insurance removes heteroskedasticity but the risk to permanent productivity remains. Our calculations show that it is this permanent risk to productivity which induces the greatest welfare loss. It is worth noting that part of the welfare loss arises because realizations of permanent shocks impact on retirement wealth with negative shocks reducing individuals' ability to save for retirement. The second point is that insuring productivity risk has a greater welfare benefit for the high education group. This is primarily driven by the difference in the permanent variance estimated in section 4. Further, the difference in the *level* of income across the two education groups means that food stamps provide better insurance against bad productivity shocks for the low education group, and thus the low educated attach relatively lesser value to insurance against productivity risk than the higher educated. Finally note that reducing productivity risk increases output marginally because it tends to reduce spells out of work. This is particularly true for the low educated who are often driven to quit as a result of productivity shocks.

We can repeat this exercise for employment risk by considering small changes in the job destruction rate and the job arrival rate for the unemployed. A 1% increase in the former can be compensated by a 0.07% increase in consumption for the high educated and a bit more for the lower educated individuals. Interestingly though, increasing the job destruction rate leads to a substantial decrease in output. A small variation in the job arrival rate for the unemployed from its current level has an even smaller, almost negligible effect for both education groups. The effect on output is also negligible. This result is driven by the relatively small unemployment durations in the US. The larger effect of job destruction is due to the fact that this process disrupts the matching to better jobs. Thus an individual who is displaced not only has to spend some time unemployed, but is unlikely to obtain as good a match as before.

	High Education		Low Education	
Scenario	*Elasticity of consumption willing to pay* (π)	Elasticity of *Output*	*Elasticity of consumption willing to pay* (π)	Elasticity of *Output*
Productivity risk: σ_ζ	-0.490	-0.014	-0.340	-0.077
Job Destruction: δ	-0.074	-0.116	-0.087	-0.180
Unempl. arrival rate: λ^n	0.026	0.013	0.026	0.038
On-the-job arrival: λ^e	0.097	0.104	0.074	0.115
Firm heterogeneity: σ_a	0.396	0.425	0.279	0.365

Table 8: Welfare effects of various sources of risk

The importance of matching is in fact reflected in the relative importance of the job arrival rate for the employed: a higher rate is valued because it improves the rate at which individuals can obtain offers from better firms and indeed has a substantial effect on output. Because of the lower destruction rate of jobs among the higher education groups the welfare effect of increasing the rate of arrival of job offers for the employed is higher. Increasing λ^e also has a substantial impact on output because it improves the rate at which individuals improve their matches.

Finally it is interesting to consider the role of firm level heterogeneity. In models where labor supply and mobility are ignored this heterogeneity would translate into wage risk with very large negative welfare costs. However, here such heterogeneity may have benefits for individuals since there is the chance of obtaining a better job offer, while bad offers can always be turned down. Individuals in the high educated group are willing to pay 0.4% of consumption for a 1% increase in firm level heterogeneity. For the lower education group the figure is 2.8%, partly because of the less stable employment profiles.

5.2 Welfare Benefit of Government Insurance

The various welfare programmes we have included in our analysis provide insurance for different aspects of risk faced by the individual, although they are unlikely to provide anything close to full insurance. We can use our model to assess the extent to which individuals value these programmes

	High Education		Low Education	
Scenario	*Elasticity of consumption willing to pay*	Elasticity of Output	*Elasticity of consumption willing to pay*	Elasticity of Output
Unemp. Insurance	0.012	−0.002	0.028	−0.001
Food stamps	0.034	−0.015	0.104	−0.045
Disability Insurance	0.005	−0.070	−0.011	−0.21

Table 9: Welfare effects of government programmes

in their current design. We follow the same approach of a local change as before: we consider a small (1%) change in parameters of each of our programmes, namely UI, Food Stamps and DI. This calculation focuses on the insurance benefit of these programs because there is no cross-group redistribution. The results are presented in Table 9. The value of UI, as we have modelled it, is negligible: Individuals in the higher education group are willing to pay a 0.01% of their consumption stream for a 1% increase in UI

. The figure for the lower education group is 0.03% for the same increase in consumption. However, Food Stamps are much more valuable particularly for the low education group, whose members are more likely to become eligible over the life-cycle. The key difference between UI and FS is that the latter insure against permanent reductions in productivity, against which it is very hard to self-insure. In contrast the one off payment provided by UI can be easily made up by a bit of extra savings. Moreover UI does not insure against long term unemployment (but FS do in effect).

Finally it is interesting to note that the DI programme has no value at all for the high educated individuals. The low educated individuals would in fact pay to hae the benefits reduced, because the tax implications are much more severe than the benefits received, which are uncertain even following a severe shock: moral hazard here dominates and this is reflected in the large negative output elasticities observed.

6 Conclusions

In this paper we have set up a model of employment and consumption over the life-cycle that allows us to define the distinction between employment and wage/productivity risk. Our model allows for job to job transitions and as well as the more usual job to unemployment transitions. Within this context we estimate a wage process, which also allows for endogenous wage changes because of accepted job offers.

We find that allowing for endogenous job-to-job movements reduces substantially the measure of risk reported in earlier studies. We then calibrate our model to obtain measures of job arrival and destruction rates that make the model consistent with observed participation rates and unemployment durations over the life-cycle. We use these measures together with estimates of intertemporal substitution form the literature to obtain a simulation model that allows the quantification of the importance of productivity and employment risk. We find that the welfare cos of employment risk is small. However productivity risk, as estimated from our wage process has large welfare consequences. A direct implication of our results is that while UI has almost no value to individuals, safety net programmes, such as Food Stamps are highly valued.

7 Appendix: Deriving Moments for the Variance of Wages

Wages are given by

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

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

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

where $\xi_{it} = \left(a_{ij(t)} - a_{ij(t_0)}\right)$. The latent indexes associated to working and moving are:

$$\begin{aligned} P_{it}^* &= z_{it}'\varphi + \pi_{it} \\ M_{it}^* &= k_{it}'\theta + \mu_{it} \end{aligned}$$

for all t. Note that conditioning on participation in periods t and $t-1$, and using the law of iterated expectations, we obtain:

$$
\begin{aligned}
E\left(\Delta \ln w_{it} | P_{it}=P_{it-1}=1\right) &= E\left(\Delta \ln w_{it} | M_{it}=0, P_{it}=P_{it-1}=1\right)\left(1-\Pr\left(M_{it}=1\right)\right) \\
&\quad + E\left(\Delta \ln w_{it} | M_{it}=1, P_{it}=P_{it-1}=1\right)\Pr\left(M_{it}=1\right) \\
&= \Delta x_{it}'\beta + G_{it}
\end{aligned}
$$

where

$$
G_{it} = \rho_{\zeta\pi}\sigma_{\zeta}\lambda_{P=1} + \rho_{\xi\pi}\sigma_{\xi}\lambda_{P=1}\Phi\left(k_{it}'\theta\right) + \rho_{\xi\mu}\sigma_{\xi}\phi\left(k_{it}'\theta\right) + \rho_{\xi\pi_{-1}}\sigma_{\xi}\lambda_{P_{-1}=1}\Phi\left(k_{it}'\theta\right)
$$

and $\lambda_{M=0} = \frac{\phi\left(k_{it}'\theta\right)}{1-\Phi\left(k_{it}'\theta\right)}, \lambda_{M=1} = \frac{\phi\left(k_{it}'\theta\right)}{\Phi\left(k_{it}'\theta\right)}$, $\lambda_{P=1} = \frac{\phi\left(z_{it}'\gamma\right)}{\Phi\left(z_{it}'\gamma\right)}$, and $\lambda_{P_{-1}=1} = \frac{\phi\left(z_{it-1}'\gamma\right)}{\Phi\left(z_{it-1}'\gamma\right)}$. Thus, G_{it} is a "selection term" accounting for conditioning on multiple indexes. Note that we do not exploit the restrictions on the coefficients on the selection terms. However, we check if they are satisfied once estimates of the structural parameters are obtained. The estimation of the equation above is standard (Heckman 2-step method).

The "structural" parameters (i.e., the variances of the wage shocks) are identified by the restrictions imposed on the moments of g_{it}. Using formulae from Tallis (19XX), the first moment is:

$$
\begin{aligned}
E\left(g_{it} | P_{it}=P_{it-1}=1, M_{it}=0\right) &= -\rho_{\zeta\mu}\sigma_{\zeta}\lambda_{M=0} + \rho_{\zeta\pi}\sigma_{\zeta}\lambda_{P=1} \\
E\left(g_{it} | P_{it}=P_{it-1}=1, M_{it}=1\right) &= \left(\rho_{\zeta\mu}\sigma_{\zeta} + \rho_{\xi\mu}\sigma_{\xi}\right)\lambda_{M=1} + \left(\rho_{\zeta\pi}\sigma_{\zeta} + \rho_{\xi\pi}\sigma_{\xi}\right)\lambda_{P=1} + \rho_{\xi\pi_{-1}}\sigma_{\xi}\lambda_{P_{-1}=1}
\end{aligned}
$$

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:

$$
E\left(g_{it}^2 | P_{it}=P_{it-1}=1, M_{it}=0\right) = \sigma_{\zeta}^2 \begin{pmatrix} 1 - \rho_{\zeta\pi}^2 z_{it}'\gamma\lambda_{P=1} + \rho_{\zeta\mu}^2 k_{it}'\theta\lambda_{M=0} \\ -2\rho_{\zeta\pi}\rho_{\zeta\mu}\lambda_{P=1}\lambda_{M=0} \end{pmatrix} + 2\sigma_e^2
$$

and

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

Finally, we consider the first order autocovariance $E\left(g_{it}g_{it-1}\left|.\right.\right)$. 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 $\Pr\left(M_t = 1, M_{t-1} = 1\right) \approx 0$ and consider only the restrictions on the unconditional autocovariance, namely

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

References

[1] Abowd, J. and D. Card (1989) "On the covariance structure of earnings and hours changes" *Econometrica 57:411-445*

[2] Abowd, J.M., Kramarz, F. and D.N. Margolis (1999) "High Wage Workers and High Wage Firms" *Econometrica*. 67(2): 251-333

[3] Anderson, P., and B. Meyer (1997), "Unemployment Insurance Take-up Rates and the After-Tax Value of Benefits", Quarterly Journal of Economics.

[4] Atkeson, A. and R.E. Lucas (1992) "On efficient distribution with private information" *Review of Economic Studies*, 59: 427-453

[5] Attanasio, O. (2000), "Consumption", in *Handbook of Macroeconomics*.

[6] Attanasio, O. and Davis, S. (1996) "Relative wage movements and the distribution of consumption" *Journal of Political Economy* 104: 1227-1262

[7] Banks, J., R. Blundell, and A. Brugiavini (2001), "Risk pooling, precautionary saving and consumption growth", Review of Economic Studies.

[8] Blundell, R., Magnac, T., and C. Meghir (1997) "Savings and Labor-Market Transitions" *Journal of Business and Economic Statistics* 15(2): 153-64.

[9] Blundell, R., and L. Pistaferri (2003), "", Journal of Human Resources.

[10] Blundell, R. and I. Preston (1998) *QJE*

[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, 159-18I.

[12] Browning, M. and Tom Crossley (2001)The Life-Cycle model of Consumption and Savings, mimeo

[13] Browning, M. and Tom Crossley (2001) "Unemployment Benefit Levels and Consumption Changes" *Journal of Public Economics*, 80:(1):1-23

[14] Browning, M., and A. Lusardi (1996), "Household Saving: Micro Theories and Micro Facts", *Journal of Economic Literature*, 34(4) pp 1797-1855.

[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] Cuhna, Flavio, James J. Heckman and Salvador Navarro (2005) Separating Uncertainty from Heterogeneity in Life Cycle Earnings," NBER Working Papers 11024

[18] 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.

[19] Deaton, A. (1991) "Saving and liquidity constraints" *Econometrica* 59(5): 1121-1142

[20] Dustmann, C. and Meghir, C. (2005) "Wages, experience and seniority" *Review of Economic Studies* Vol 72.1

[21] Farber and Gibbons (1996), "Learning and Wage Dynamics", Quarterly Journal of Economics.

[22] Gourinchas, P., and J. Parker (2004), "Consumption over the Life-Cycle", Econometrica.

[23] Hansen, G. and A. Imrohoroglu (1992) "The role of unemployment insurance in an economy with moral hazard and liquidity constraints" *Journal of Political Economy* 100: 118-142

[24] Heckman, J.J. (1974) “Life-cycle consumption and labor supply: an explanation of the relationship between income and consumption over the life-cycle” *American Economic Review* 64:2, 188-194

[25] Horowitz, J. (2002) “Bootstrap Methods for Markov Processes,” Econometrica.

[26] Hubbard, R.G., J. Skinner, and S. Zeldes (1995), "Precautionary Saving and Social Insurance," Journal of Political Economy, 103(2), pages 360-99.

[27] Jacobson, Louis S.; LaLonde, Robert J. and Sullivan, Daniel G. "Earnings Losses of Displaced Workers." The American Economic Review, 1993, 83(4), pp. 685-709.

[28] Kimball (1990), "Precautionary Savings in the Small and in the Large", Econometrica.

[29] Krussell, P. and A.A. Smith (1999) “On the welfare effects of eliminating business cycles” *Review of Economic Dynamics* 2: 245-272

[30] Kuhn, P., editor (2002) *Losing Work, Moving On* Upjohn Institute, Michigan

[31] Lentz (2003) “Optimal Unemployment Insurance in an Estimated Job Search Model with Savings” *Boston University, mimeo*

[32] Low, H. (2005) “Self-insurance and unemployment benefit in a life-cycle model of labor supply and savings”, Review of Economic Dynamics.

[33] Lucas, R.E. (1987) *Models of Business Cycles* Basil Blackwell, New York

[34] Lucas, R.E. (2003) “Macroeconomic Priorities” *American Economic Review* 93(1): 1-15

[35] 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

[36] Manski, C. (2004), "Measuring Expectations," Econometrica, Econometric Society, vol. 72(5), pages 1329-1376.

[37] Meghir, C. and L. Pistaferri (2004) “Income variance dynamics and heterogeneity ” *Econometrica.*

[38] Moffitt, R. (1989), "Estimating the value of an in-kind transfer: The case of Food Stamps", Econometrica.

[39] Moffitt and Gootschalk (1995), "Trends in the Covariance Structure of Earnings in the U.S.: 1969-1987", unpublished.

[40] Pistaferri, L. (2001), "Superior information, income shocks, and the permanent income hypothesis", Review of Economics and Statistics.

[41] Pistaferri, L. (2003), "Anticipated and unanticipated wage shocks, wage risk, and intertemporal labor supply", Journal of Labor Economics.

[42] Postel-Vinay, F., J.-M. Robin (2002) "Equilibrium Wage Dispersion with Heterogeneous Workers and Firms" *Econometrica*, 70: 2295-1350

[43] Rogerson and Schindler (2001)

[44] 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 Reivew* 45, 1311-1339

[45] Topel, R. (1991) "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority" *Journal of Political Economy*, 99(1): 145-176

[46] Topel, R.H, and M.P. Ward (1992) "Job Mobility and the Careers of Young Men" *Quarterly Journal of Economics,* 107(2): 439-79.

[47] Zeldes, Stephen P, 1989. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," The Quarterly Journal of Economics, MIT Press, vol. 104(2), pages 275-98.