Wage Risk and Employment Risk over the Life Cycle*

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

Preliminary and Incomplete

October 12, 2004

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.

JEL Classification: D52; D91; I30.
Keywords: Wage Risk, Unemployment Risk, Savings.

*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

Following Deaton (1991) and others, there is now an extensive literature analyzing individuals' precautionary response to income risk under incomplete markets. While this literature has made enormous progress, it has paid relatively little attention to the issue of where risks originate. In reality, individuals face risk to their earnings from a variety of sources, some related to labor supply and some to productivity. Moreover, 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 reasons. It is precisely this lack of formal insurance that prompts prudent individuals to engage in precautionary behavior. 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 and to then estimate the extent of risk within this framework using longitudinal data from the PSID. 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.

We decompose earnings risk into productivity risk and employment risk. 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 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.

The individuals' response to these risks 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. We allow for three government insurance programs: Unemployment Insurance (UI), Disability Insurance (DI), and the Food Stamps program. Unemployment insurance is aimed at insuring against job destruction and (partly) against the difficulty of finding a new job. It is worth stressing that the system will not provide insurance against productivity risk if unemployment is due to bad realizations of the productivity shock. This is because unemployment benefits are typically related to pre-unemployment earnings and so variation in productivity will be translated into the benefits received. The disability insurance system provides insurance against an extreme form of productivity shock which results in permanent inability to work. Finally, the Food Stamps program provides *universal* insurance against low income, whatever its cause.

The precautionary saving literature typically assumes that labor supply is exogenous and in this context, it is not possible to distinguish between productivity risk and employment risk.[1] In contrast, our model allows for both saving and participation choices. It is necessary to have the endogenous participation choice (in addition to exogenous job loss) partly in order to capture the moral hazard disincentive effects of providing insurance through public programs and partly to enable us to match the observed pattern of unemployment duration and labor participation. Including saving means that individuals can self-insure, reducing the welfare benefit of outside insurance (*vis-à-vis* models where there is no storage technology, i.e., consumption is set equal to income in each period of life). This also allows individuals to smooth consumption without relying on program benefits. In the absence of saving individuals would be willing to pay for program benefits in order to smooth consumption. In calculating the welfare benefit of the insurance programs we consider, we want to identify the part of the benefit that is due to insurance rather than to consumption smoothing per se.

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 calibration of employment and unemployment durations and participation rates. In addressing the question of how individuals respond to risk, we begin by simulating savings and participation behavior for individuals facing the estimated risk. These simulations give an indication of the extent of precautionary behavior (both precautionary saving and precautionary labor supply). This raises a number of questions. First, how much would individuals be willing to pay to avoid the various risks. Second, how much of the precautionary response is due to employment risk and how much to

[1] Notable exceptions are Low (1999) and Lentz (2003). The first author analyzes the joint saving and labour supply decision, but in a context without exogenous job destruction or search frictions. The second author 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.

productivity risk. We address these questions by simulating behavior for different combinations of risk, keeping the expected value of lifetime wealth constant.

Our results are as follows. First, we estimate the variance of the permanent shock to productivity to be 0.012 and the variance of firm types to be 0.03. We also find important differences across skills (measured by education). Our characterization of the wage process implies that permanent shocks are the main component of the variance of lifetime wealth. One important aspect of our analysis is that neglecting the endogenous mobility choice leads to wrong inference regarding sources of wage variability and behavioral choices. In particular, it tends to exaggerate the amount of permanent uncertainty and hence the amount of precautionary saving people hold against that risk. Second, our simulations are able to reproduce the durations found in the data: in particular, we are able to mimic the increase in the duration of unemployment by age and the profile of participation by age without introducing age-specific labor market friction parameters. Third, welfare calculations of the risk premium show that individuals are willing to pay considerably more to avoid productivity risk than employment risk: for example, highly educated individuals would pay 15% and 3%, respectively, of the offered wage to avoid each risk. The ranking is reversed for the low educated. Finally, in assessing the reasons for holding assets, productivity risk dominates. However, ignoring employment risk leads to inaccurate predictions of participation.

The layout of the paper is as follows. Section 2 discusses the distinction between employment and productivity risk. Section 3 presents the model. Section 4 describes the identification strategy and presents the estimates. Section 5 presents the calibration and simulations using the estimated measures of uncertainty. In Section 6 we calculate the welfare costs of uncertainty. Section 7 concludes.

2 Employment Risk and Wage Risk

Consider a scenario in which individual productivity fluctuates over time due to random shocks. In a perfectly competitive labor market with no search frictions 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 bceomes 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. Similarly firm heterogeneity would not lead to employment risk if there were no search frictions, because it would take little time to locate the firm which is the best match.

With firm heterogeneity, however, it will take time to locate a suitable job if the number of offers received per period is limited. Moreover, there is 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. More generally, if we consider a world in which there is both individual and firm level heterogeneity and the characteristics of workers and firms are complementary in production, then rents are produced and these are shared between the worker and the firm. In this context, wages are the outcome of bargaining with some sharing rule. In the presence of search frictions, this complementarity generates an option value to remaining in a particular match. It also creates the risk of earnings loss (and subsequent low earnings) which is not related to individual productivity, but is determined by the risk of job destruction, the risk of not receiving a new offer and the degree of firm heterogeneity. It is this risk which we characterize as *employment risk*. Precise parameter configurations of the two types of risk will be discussed later in more detail.

3 Model

3.1 Structure of Wages and Shocks

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

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

where w is the real hourly wage, x a vector of regressors, u the permanent component of wages, and e the transitory component (which for identification purposes 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 $t_0 \leq t$.[2] 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 call 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 $t+1$. The firm effect $a_{ij(.)}$ is complementary to individual productivity. However, it will be assumed uncertain across firms. The information structure is such that workers and firms are completely informed about u_{it} and $a_{ij(.)}$ when they meet. 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.

MaCurdy (1982), Topel (1991), Abowd and Card (1989) and Meghir and Pistaferri (2004) stress the importance of individual productivity shocks in explaining wage movements. Following these papers, we assume that the permanent component follows a random walk process:

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

where ζ_{it} is a random shock which we take to be uncertain and variable from period to period. MaCurdy (1982) and Meghir and Pistaferri (2004) allow for moderate persistence in the error e_{it}. For simplicity, however, we assume that e_{it} is *i.i.d.*.[3]

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. In each period the worker will be receiving new job offers at arrival rate λ^e (at most one offer is received in each period). Those who are currently unemployed receive offers at a rate λ^n.

[2] 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. If the latter represented the correct characterization, we would not be able to distinguish between the match productivity and the sharing rule mentioned in Section 2, which we therefore assume exogenously given.

[3] 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.

3.2 Insurance

The productivity shocks that we observe 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.

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).

Individuals move between firms and this leads to variation in earnings. We do not consider this as 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.

3.3 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. We assume labor is indivisible, so that if $P_t = 1$ the individual works a fixed number of hours $\overline{h}$. The individual is assumed to maximize lifetime expected utility,

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

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 T periods, may work from age 25 to 65, 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 whether to switch firm. When unemployed he has to decide whether to accept a job offer or wait longer. If eligible, he can also 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. Therefore, although there is exogenous mandatory retirement, individuals may stop working earlier.

In the simulations, we use a utility function of the form

$$U(c, P) = \frac{c^{1-\gamma}}{1-\gamma} \exp(\eta P) \tag{3}$$

where $\gamma > 1$ and $\eta > 0$, implying consumption and participation are complements.

The intertemporal budget constraint has the form

$$A_{it+1} = R\left[A_{it} + \left(w_{it}\overline{h}\left(1-\tau_w\right) - F_{it}\right)P_{it} + \left(B_{it}E_{it}^{UI}\left(1-E_{it}^{DI}\right) + D_{it}E_{it}^{DI}\right)\left(1-P_{it}\right) + FS_{it}E_{it}^{FS} - c_{it}\right] \tag{4}$$

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, F the fixed cost of work,[4] 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 eligibility $\{0,1\}$ indicators for unemployment insurance, disability insurance, and food stamps, respectively.[5] Note also that there are costs to applying for disability which we discuss below. There are no explicit liquidity constraints.

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

[5] 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.

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).[6] We assume $B_{it} = b \times w_{it-1}\overline{h}$, and set the replacement ratio $b = 75\%$.[7] 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 accept a new job offer.[8] 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. Gross income is given by

$$y_{it}^{gross} = w_{it} h P_{it} + \left(B_{it} E_{it}^{UI} \left(1 - E_{it}^{DI}\right) + D_{it} E_{it}^{DI} \right) \left(1 - P_{it}\right) \tag{5}$$

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

$$FS_{it} = \begin{cases} \overline{FS} - 0.3 \times y_{it} & \text{if } E_{it}^{FS} = 1 \text{, i.e., if } y_{it} \leq \underline{y} \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

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 ($E_{it}^{FS} = 1$).

[6]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".

[7]One feature of the actual UI system that we do not model is the cap imposed on the level of benefits.

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

Disability Benefits 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 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

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

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. Whether an individual is eligible (i.e., $E_{it}^{DI} = 1$) depends on the decision to apply while being out of work and on having received a large negative productivity shock. We assume that the probability of success is independent of age. Eligibility does not depend on whether an individual quits or the job is destroyed.

3.4 Budget Balance

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:

$$\sum_{i=1}^{N}\sum_{t=1}^{T}\frac{1}{R^t}\left[\left(B_{it}E_{it}^{UI}\left(1-E_{it}^{DI}\right)+D_{it}E_{it}^{DI}\right)(1-P_{it})+E_{it}^{FS}FS_{it}\right]=\sum_{i=1}^{N}\sum_{t=1}^{T}\frac{1}{R^t}\tau_w w_{it}\overline{h}P_{it}+Deficit \tag{7}$$

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.[9] 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, there are no aggregate shocks in the economy and no business cycle fluctuations and so we do not consider the value 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 in a first best setting. Allowing the budget to balance over all education groups would confuse the issue we are considering with distributional questions, particularly because as we shall see the risk profiles are quite different across education groups.

3.5 Solution

We start by constructing the value functions for the individual when employed and when out of work.

Consider first the value function for an employed person. The state variables are current assets (A_{it}), individual productivity u_{it}, the firm effect $a_{ij(.)}$ and of course work status. Offers are indexed by the value of $a_{ij(.)}$ with which they are associated.[10] 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 her job is destroyed she will have to move to unemployment. Thus the value function for

[9] 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.

[10] 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.

an individual i who is working in period t is

$$V_t^e\left(A_{it}, u_{it}, P_{it}=1, \alpha_{ij(t_0)}\right) = \tag{8}$$
$$\max_c \left\{ \begin{array}{c} U\left(c_{it}, P_{it}=1\right) + \\ \beta\delta E_t\left[V_{t+1}^n\left(A_{it+1}, u_{it+1}, P_{it+1}=0\right)\right] \\ +\beta\left(1-\delta\right)\left(1-\lambda^e\right)E_t\left[\max\left\{ \begin{array}{c} V_{t+1}^n\left(A_{it+1}, u_{it+1}, P_{it+1}=0,\right), \\ V_{t+1}^e\left(A_{it+1}, u_{it+1}, P_{it+1}=1, \alpha_{ij(t_0)}\right), \end{array}\right\}\right] \\ +\beta\left(1-\delta\right)\lambda^e E_t\left[\max\left\{ \begin{array}{c} V_{t+1}^n\left(A_{it+1}, u_{it+1}, P_{it+1}=0\right), \\ V_{t+1}^e\left(A_{it+1}, u_{it+1}, P_{it+1}=1, \alpha_{ij(t_0)}\right), \\ V_{t+1}^e\left(A_{it+1}, u_{it+1}, P_{it+1}=1, \alpha_{ij(t+1)}\right) \end{array}\right\}\right] \end{array} \right\}$$

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.

There are two "types" of unemployed individuals in the model. The first type is not eligible for applying for disability insurance (either because he is younger than 50 or because has had an application turned down in the previous period). This type of unemployed individual may have the option of working if she draws a job offer. Thus his value function is

$$V_t^n\left(A_{it}, u_{it}, P_{it}=0, DI^{App}=0\right) = \tag{9}$$
$$\max_c \left\{ \begin{array}{c} u\left(c_{it}, P_{it}=0\right) \\ +\beta\lambda^n E_t\left[\max\left\{ \begin{array}{c} V_{t+1}^n\left(A_{it+1}, u_{it+1}, P_{it+1}=0\right), \\ V_{t+1}^e\left(A_{it+1}, u_{it+1}, P_{it+1}=1, \alpha_{ij(t+1)}\right), \end{array}\right\}\right] \\ +\beta\left(1-\lambda^n\right)E_t\left[V_{t+1}^n\left(A_{it+1}, u_{it+1}, P_{it+1}=0\right)\right] \end{array} \right\}$$

where λ^n is the probability that the unemployed person receives an alternative job offer and DI^{App} is an indicator for whether the individual has applied for disability insurance. 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.

If the individual is eligible to apply for disability benefit, we have an additional discrete state variable which indicates whether an application for disability benefit has been made during that unemployment spell. When deciding whether or not to apply, the individual already knows if he has a job offer in that period. The value of D_{it} depends on the permanent wage only and not on the particular firm that the individual has most recently been working for.

If the disability application is unsuccessful, the value function conditional on not-working for that unemployment spell is given exactly by equation 9. If the disability application is successful, we can calculate the resulting value function analytically: D is earned each year until retirement, and this income is divided up over all remaining periods of life. In deciding whether or not to make an application, the value function is given by

$$V_t^n\left(A_{it}, u_{it}, P_{it}=0, DI^{App}=1\right) = \max_c\left\{\ u\left(c_{it}, P_{it}=0\right) + \beta \max\left\{V_{t+1}^A, V_{t+1}^{NA}\right\}\ \right\} \tag{10}$$

where

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

and S is the exogenous probability of a successful application.

In each period the individual decides whether to work or not or to move to another job 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 $\alpha_{ij(.)}$ and the individual will move if the new offer is from a higher $\alpha_{ij(.)}$-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. This suggests that by shutting down the cost of job mobility our model may generate too many transitions between firms, as individuals will move as soon as a better offer comes along; and similarly, transitions from unemployment back to employment may be too fast because there is only a small opportunity cost of taking the first job offered. This opportunity cost is determined by the probability that the new firm goes bankrupt (precluding an offer from a better firm given that at least one period is spent unemployed following job destruction), and by the difference in job arrival rates when unemployed versus employed.

Once the choice of firm is made, the worker decides whether to work or quit based on the comparison of the value functions:

$$\text{Employment status} = \begin{cases} \textit{work} & \text{if } V^e > V^n \text{ and } \mathbf{1}\{\delta = 0\} \\ \textit{unemployed} & \text{if } V^e \leq V^n \text{ or } \mathbf{1}\{\delta = 1\} \end{cases}$$

An unemployed worker has the following choices:

$$\text{Employment status} = \begin{cases} \textit{work} & \text{if } V^e \geq \max\left[V^n, V^A\right] \text{ and } \mathbf{1}\{\lambda^n = 1\} \\ \textit{unemployed} & \text{if } V^n > \max\left[V^A, V^e\right] \text{ or } \mathbf{1}\{\lambda^n = 0\} \\ \textit{apply for DI} & \text{if } V^A > \max\left[V^n, V^e\right] \end{cases}$$

Here $\mathbf{1}\{\lambda^n = 1\}$ signifies that a job offer has been received, and $\mathbf{1}\{\delta = 0\}$ that the current job is still in existence and the individual has the choice of keeping it. The other terms have analogous meaning. The model is solved numerically and the solution method is discussed in the appendix.

Before moving on to discuss the identification strategy, it is fair to add some caveats to our analysis. First, we neglect equilibrium wage distribution issues as discussed, e.g., by Postel-Vinay and Robin (2002). In particular, we assume very simple forms of firm and worker heterogeneity. Second, and relatedly, the firm effect is assumed to be constant over the time of the employment relationship. Implicitly we also assume the firm effect is constant across workers working at the same firm but this is a redundant qualification because we lack a matched worker-firm data set. This is obviously a strong restriction but is dictated by the type of data we have available. Third, we assume that all mean-reverting wage disturbances are due to measurement error, thus shutting down the channel of genuine economic transitory shocks. While this is done primarily for identification purposes, we stress that if liquidity constraints are not overwhelming, then neglecting temporary productivity shocks is unlikely to induce big differences in our findings as consumption smoothing effectively implies that transitory shocks are not important. Fourth, we assume that wage residuals are uncorrelated across individuals. This would not be the case if $\alpha_{ij(.)}$ were a truly firm effect. However, for the size of the data set we are using this is unlikely to create any serious problems. Finally, while we endogenize both the job mobility and the labor market participation decisions in the simulations, we impose no theoretical structure on them in the estimation and essentially use reduced form equations based on standard random utility comparisons. However, we allow for unrestricted linear correlation between the unobserved components of the wage process and the unobserved components of the selection terms.

4 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 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 the use of parametric distributional assumptions (normality).

Define the latent utility from labor market participation as $P_{it}^* = z_{it}'\gamma + \pi_{it}$. The associated labor market participation index is $P_{it} = 1\{P_{it}^* > 0\}$, which is unity for participants. Workers separate from their current employer voluntarily (quits) or involuntarily (layoffs). As argued by Borjas and Rosen (1980), job turnover, regardless of who initiates it, represents the same underlying phenomenon, that of workers' marginal product being higher elsewhere. Let $M^*_{it} = k'_{it}\theta + \mu_{it}$ denote the latent utility from moving in period t to an employer that is different from the one in period $t-1$. The indicator $M_{it} = 1\{M^*_{it} > 0\}$ singles out the "movers".[11] 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:

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

Wage growth is only observed for those who work in both periods. To achieve identification of the relevant parameters, we make the following assumptions: $E\left(a_{ij(t)} a_{ij(s)}\right) = \sigma^2_a$ if $j(s) = j(t)$ and zero otherwise. We denote with $\sigma^2_\zeta = E\left(\zeta^2_{it}\right)$ and $\sigma^2_e = E\left(e^2_{it}\right)$ (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; similarly, $E(\xi_{it}\pi_{is}) = \sigma_\xi \rho_{\xi\pi}$ if $s = t$ 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 ($E(\xi_{it}\zeta_{is}) = 0 \ \forall \ t, s$).

Suppose now that we select only those who work at t and $t-1$ ($P_{it} = 1, P_{it-1} = 1$). In this (self-selected) sample

$$\begin{aligned} E(\Delta \ln w_{it} | P_{it} = 1, P_{it-1} = 1) &= \Delta x'_{it}\beta + \sigma_\zeta \rho_{\zeta\pi} SEL_P \\ &\quad + \sigma_\xi \rho_{\xi\pi} SEL_{M_1} + \sigma_\xi \rho_{\xi\mu} SEL_{M_1} \end{aligned} \tag{11}$$

where $SEL_P = \frac{\phi\left(z'_{it}\gamma\right)}{\Phi\left(z'_{it}\gamma\right)}$ is a "self-selection" term induced by labor market participation, and $SEL_{M_1} = \frac{\phi\left(z'_{it}\gamma\right)}{\Phi\left(z'_{it}\gamma\right)}\Phi(k'_{it}\theta)$ and $SEL_{M_2} = \phi(k'_{it}\theta)$ are "self-selection" terms induced by intra-firm

[11] If people only consider current gains from moving, then the benefit of it is $\left(\alpha_{ij(t)} - \alpha_{ij(t)}\right)$ and the costs are $-\left(k'_{it}\theta + \varepsilon_{it}\right)$, where ε are unobserved mobility costs. Thus people move when $k'_{it}\theta + \mu_{it} > 0$ (with $\mu_{it} = \left(\alpha_{ij(t)} - \alpha_{ij(t)}\right) + \varepsilon_{it}$).

mobility. $\phi(.)$ and $\Phi(.)$ are the p.d.f. and c.d.f. of the standard normal distribution.[12] The idea is to estimate the parameters γ and θ in a first stage (running separate probit regressions because of the assumed orthogonality assumption between π_{it} and μ_{it}), construct estimates of the selection terms and then estimate β in a second stage using only participants in both periods.

Estimation of (11) provides consistent estimates of β. Define now unexplained wage growth (observed only for participants in both periods):

$$g_{it} = \Delta \ln w_{it} - \Delta x_{it}'\beta = \zeta_{it} + \Delta e_{it} + \xi_{it} M_{it} \tag{12}$$

Given these assumptions we can 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 an 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:

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

estimate γ_q, and construct the variables: $\overline{z_{it}'\gamma} = \frac{\sum_{q=1}^4 z_{it(q)}'\gamma_q}{4}$, and $\overline{\lambda_\pi} = \frac{1}{4}\sum_{q=1}^4 \frac{\phi\left(z_{it(q)}'\gamma_q\right)}{\Phi\left(z_{it(q)}'\gamma_q\right)}$ which we use as approximations for $z_{it}'\gamma$ and $\frac{\phi\left(z_{it}'\gamma\right)}{\Phi\left(z_{it}'\gamma\right)}$, 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 $t-1$.

[12] In estimation we do not use the restrictions on the parameters of interest imposed by (**??**). This only results in a loss of efficiency, but it does not affect consistency.

In the second step we estimate (**??**) using only labor market participants in both periods. This gives us estimates of β 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 coefficients. The variance of the firm effect (σ_a^2) can be recovered from $\frac{\sigma_\xi^2}{2} = \sigma_a^2$. 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 λ's. We should point out that this procedure is conservative, since it allows for more serial correlation than that implied by the moment conditions we use. Hence the bootstrap standard deviations will be using only the N dimension of the sample and the precision of our parameters is likely to be underestimated.

4.1 The Data

The data are drawn from the 1988-1993 family and individual-merged files of the PSID. Since the PSID has been widely used for microeconometric research, we shall only sketch the description of its structure in this section.[13]

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.

Questions referring to labor income are retrospective; thus, those asked in 1988, say, refer to the 1987 calendar year. Our measure of the hourly wage is obtained as the ratio between annual earnings and annual hours of work. The earnings variable is the labor portion of money income from all sources; it includes wages, bonuses, overtime, and commissions.

Education level is computed using the PSID variable with the same name. We stratify the sample

[13] See Hill (1992) for more details about the PSID. In future drafts we will extend the PSID data set to 1998 to minimize the proportion of right-censored unemployment spells.

by education, low (those with at least a high school diploma, but no college degree), and high (those with a college degree or more).

Step-by-step details on sample selection are as follows. We select all male heads aged 25 to 60. There are 3,385 such individuals (and 17,450 individual-year observations). We then eliminate wage outliers (those with wage levels or wage growth below the 1-st or above the 99-th percentile). This leaves us with 2,948 individuals (15,123 observations). Next, we eliminate observations where the wage rate is missing (i.e., either annual hours or annual earnings are missing), observations with missing data on employment status, inconsistent employment reports (for instance, people reporting a non-zero wage while also reporting to have been unemployed for the entire calendar year), and those with less than three years of data. The final sample includes 14,483 observations for 2,691 individuals.

The PSID asked individuals 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.

An important component of our estimation strategy is the ability to identify firm switches. At least in principle, the PSID data on tenure could be used to do this, but these data are notoriously error-ridden and subject to a variety of inconsistencies (see Brown, 19XX). We adopt a less direct approach and use the tenure data as part of a more general switch-identification procedure. Recall that our objective is simply to assess whether the employer at t is the same as the employer at $t-1$. For each month of the previous calendar year, employed individuals are asked if the employer of that month is the same as the current employer. For unemployed individuals, the questions refer to the last employer they had.

We use these data, as well as information about the date of the interview, employment status, and reported tenure (in months) at the time of the interview, to assign a value of 1 or 0 to the variable M_{it} described above (conditioning on working in both periods). To avoid exaggerating the number of firm switches, we check that temporary separations are not truly switches (i.e., as in the case of recalls after temporary layoffs).[14]

[14] In future drafts we plan to estimate wage dynamics parameters using 1993-95 SIPP data. While shorter than the PSID, the SIPP has the important advantage that it allows to identify the firm the individual was working for in each month of a 36-month period (at most).

4.2 The Results: Participation, Mobility, and Wage Growth

The first step is to control for selection into employment in the wage equation. We run quarterly participation probits controlling for demographics and other socio-economic individual characteristics (education, a quadratic in age, a dummy for whites, year of birth dummies, region dummies, county unemployment, residence in an SMSA, family size, number of kids, a dummy for married, other household income, and an index of generosity of the state-level UI system).[15] The role of other household income and the index of generosity of the state-level UI system is that of controlling for unearned income and outside opportunities; it should affect the decision to work, but should have no effect on one's wage. Note that if we ran an annual participation probit, we would be identifying only the probability of not participating for the whole year which is a much less common event. We do not have variables that vary by quarter, so the effect of the various covariates on the outcome is fairly similar across quarters. We run separate probits by education. To save space, Table 1 reports only the coefficient on our exclusion restrictions. The effect of UI generosity and other income is as expected: people who live in states with more generous UI benefits and higher family income (net of one's earnings) have a higher opportunity cost of working. The effect of UI generosity has borderline p-values in the low education sample and it is mostly insignificant in the high education sample. The effect of other covariates (not reported here) is as expected: In all quarters, the probability of working is higher for the more educated, the whites, those who are married, with a larger family size, and fewer kids. Year of birth are jointly significant while region dummies are not.

Table 1
Exclusion restrictions in the quarterly participation equation

	1st quarter	2nd quarter	3rd quarter	4th quarter
	Low education			
Other income	−0.0162	−0.0192	−0.0188	−0.0178
	(0.0019)	(0.0020)	(0.0020)	(0.0020)
UI generosity	−0.2855	−0.3559	−0.3256	−0.4061
	(0.2004)	(0.2034)	(0.2029)	(0.2021)
	High education			
Other income	−0.0134	−0.0131	−0.0150	−0.0140
	(0.0020)	(0.0021)	(0.0021)	(0.0021)
UI generosity	−0.3688	−0.1419	−0.2254	−0.4537
	(0.3082)	(0.3242)	(0.3343)	(0.3387)

Note: Standard errors in parenthesis.

[15] The index of generosity is defined as the ratio of maximum weekly UI benefits and average weekly wages (both vary by time and state).

Table 2 reports the estimates of a probit model for the decision to switch to a different firm ($M_{it} = 1$). We use as controls education, a quadratic in age, a dummy for whites, year of birth dummies, region dummies, county unemployment, residence in an SMSA, family size, number of kids, a dummy for married. Our exclusion restrictions try to capture mobility plans (as reported in the previous year) and mobility history (as recorded in the previous period). Overall, we find that 14% of the observations in our sample are switches to another firm from one year to the next.

Table 2
The mobility decision

	Low education	High education
Grades of schooling	−0.0512 (0.0124)	−0.0130 (0.0156)
Age	−0.0911 (0.0204)	−0.1157 (0.0249)
Age^2	0.0008 (0.0002)	0.0012 (0.0003)
White	0.0406 (0.0465)	0.0651 (0.0536)
Married	−0.3883 (0.0537)	−0.0741 (0.0792)
SMSA	−0.0076 (0.0416)	−0.0389 (0.0458)
County unempl.	0.0314 (0.0097)	−0.0469 (0.0114)
Number of kids	−0.0225 (0.0421)	0.1237 (0.0560)
Family size	0.0587 (0.0385)	−0.1736 (0.0526)
Might move$_{t-1}$	0.2558 (0.0421)	0.2551 (0.0441)
Ever moved for job$_{t-1}$	0.2586 (0.0462)	0.0173 (0.0445)
Year of birth dummies	Yes	Yes
Region dummies	Yes	Yes
Constant	1.5238 (0.4384)	2.5810 (0.5532)
# of obs.	5,760	5,184

Note: Standard errors in parenthesis.

Our estimates show a greater incidence of job mobility among young, white, more educated individuals, and a lower incidence among those who are married. Our exclusion restrictions appear to have sufficient identifying power, especially among the low educated. The participation probit results allow us to construct estimates of the selection terms SEL_P, SEL_{M_1}, and SEL_{M_2}. This now allows us to estimate the wage growth equation (**??**), whose results are reported in Table 3 separately by education. This being an equation for wage growth, the estimates display large standard errors . There is evidence for the usual age-concave wage profile. Residence in SMSA has a positive effect on the wage of low educated but a negative effect on that of high educated individuals. The selection term for participation is positive as surmised, but statistically insignificant. The selection terms for mobility are slighly more precise, especially SEL_{M_2} among the high educated. In part at least, this high imprecision results from the fact that the coefficients on the selection terms reflect a combination of parameters. The parameters underlying these effects are more precisely pinned down in the estimation of unobserved wage dynamics that follows.

Table 3
The wage growth equation

	Low education	High education
Age	−0.0011 (0.0007)	−0.0016 (0.0007)
ΔMarried	−0.0215 (0.0205)	−0.0109 (0.0206)
ΔSMSA	0.0317 (0.0187)	−0.0419 (0.0155)
ΔFamily size	0.0158 (0.0089)	−0.0021 (0.0101)
ΔNumber of kids	−0.0032 (0.0109)	0.0113 (0.0115)
ΔCounty unemployment	0.0016 (0.0023)	−0.0016 (0.0024)
SEL_P	0.0094 (0.0575)	0.0755 (0.1000)
SEL_{M_1}	0.0576 (0.2329)	−0.4373 (0.5544)
SEL_{M_2}	−0.0066 (0.0903)	0.2048 (0.0784)
# of obs.	5,760	5,184

Note: Standard errors in parenthesis.

The last step of the estimation procedure is to recover the structural parameter 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 job mobility, we find that in the whole sample the standard deviation of the permanent shock, σ_ζ, is about 0.11 (so that the variance is 0.012), the standard deviation of the transitory shock (measurement error), σ_ε, 0.13 (variance 0.017), and the standard deviation of the firm shock, σ_α, 0.18 (variance 0.032). These parameters are all very precisely measured.

The correlation between the permanent shock and heterogeneity in the participation decision $\rho_{\zeta\pi}$ is positive as might be expected but insignificant. This is consistent with the notion that prime-age males have low elasticity of intertemporal substitution. Further, this correlation is with the permanent shock and so intertemporal substitution is offset by the wealth effect. The correlation between the permanent shock and the decision to leave the current employer ($\rho_{\zeta\mu}$) is large, negative, and significant, suggesting that those receiving a large positive productivity shock are less likely at the margin to move.[16]

We find that wage growth due to firm switching (ξ_{it}) is positively correlated with μ_{it}, the heterogeneity in the decision to separate. This positive correlation is of course expected if the utility from job mobility reflects a positive mobility premium. Finally, we find a negative correlation between wage growth due to firm switching (ξ_{it}) and the heterogeneity in the decision to work (π_{it}).

Table 4
Estimates of wage dynamics parameters

	Whole sample	Low education	High education	Neglecting Mobility	Neglecting participation	Neglecting participation and mobility
	(1)	(2)	(3)	(4)	(5)	(6)
σ_ζ	0.1115	0.0883	0.1255	0.1624	0.1119	0.1606
	(0.0098)	(0.0145)	(0.0116)	(0.0096)	(0.0106)	(0.0087)
σ_e	0.1323	0.1467	0.1147	0.1323	0.1323	0.1323
	(0.0046)	(0.0045)	(0.0094)	(0.0046)	(0.0046)	(0.0046)
σ_a	0.1762	0.1785	0.1662		0.1773	
	(0.0062)	(0.0095)	(0.0117)		(0.0059)	
$\rho_{\zeta\pi}$	0.1015	0.3317	0.2975	0.3826		
	(0.2798)	(0.2684)	(0.3122)	(0.2555)		
$\rho_{\zeta\mu}$	−0.6921	−0.0135	−1.0932		−0.6620	
	(0.3037)	(0.4863)	(0.3603)		(0.2991)	
$\rho_{\xi\pi}$	−0.4536	−0.3751	0.1093			
	(0.2368)	(0.2632)	(0.6315)			
$\rho_{\xi\mu}$	0.4120	0.0117	0.7331		0.3544	
	(0.1754)	(0.2311)	(0.2768)		(0.1723)	

Note: Block bootstrap standard errors in parenthesis.

How do things change when we allow for parameter heterogeneity by education? 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. Few remarks are in order. The variance of the permanent shock is higher for the more educated, the variance of the transitory shock is lower and the variance of the firm effect is similar. In particular σ_ζ^2 is 0.008 for the less educated and almost twice as much, 0.0158, for the more educated.

What happens if we ignore the endogenous mobility decision? 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). In column (4) we report the results of this experiment.

[16]The Appendix discusses a simple case in which such negative correlation may arise.

They show that the wage variability that is due to workers changing firms is now entirely attributed to permanent shocks (σ_ζ^2 more than doubles from 0.0124 to 0.0264), and the correlation between permanent shocks and heterogeneity in participation is now positive and significant, possibly because a premium for mobility is now seen as a premium for labor market participation.

In column (5) we ignore the participation decision but account for the endogenous mobility choice. Since in the full model participation does not appear to be correlated with productivity shocks, we do not expect big changes vis-a-vis the main specification, and in fact we do not find any significant change. Finally, in column (6) we shut down both mobility and participation. The results show an increase in the amount of variability due to permanent shocks.

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 show 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.

5 Simulation

In the model we simulate, individuals have a 40 year working horizon (age 25-65) 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 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.

The baseline combination of the variables determining risk corresponds to the estimates from the PSID and from calibrating unemployment durations and participation proportions for each education group. The extent of deterministic wage growth is also taken from the PSID. The values

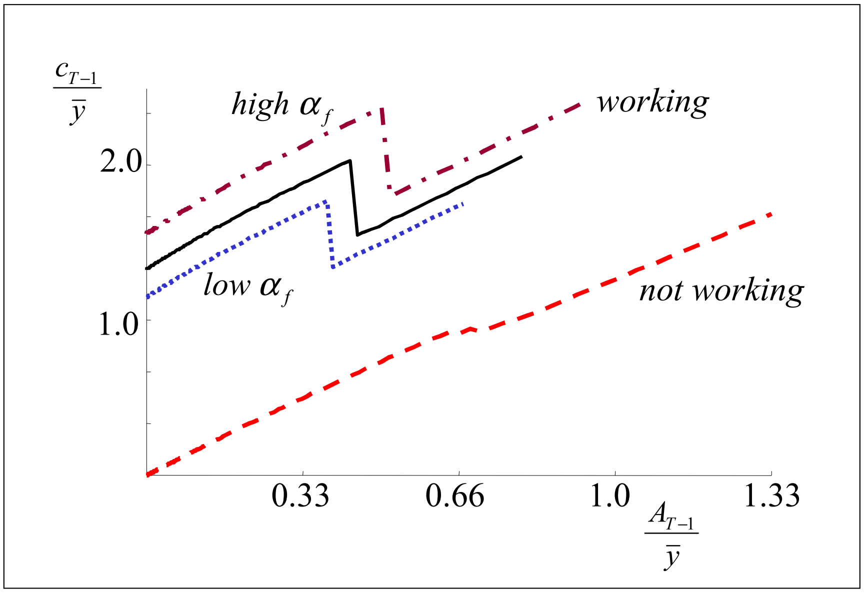

Figure 1: Consumption Functions for $T-1$, conditional on Participation Status.

of the preference and risk parameters are given in Panel A and B of Table 5, respectively. The numbers in brackets show the corresponding annual values when appropriate.

For each parameter set, we first solve the model numerically by backward induction and then we simulate behavior by taking different realizations of the random variables. Individuals are ex-ante identical (there are 10,000 of them) and in the simulation results below, we show mainly the cross-section means by age. Before turning to these simulations, it is instructive to show part of the solution for the baseline model. Figure 1 shows consumption as a function of assets in period $T-1$ 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 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. Figure 2 presents a comparable solution for $T-30$. In this case, the sharp kinks in the consumption function have been smoothed out.

Figure 3 shows the reservation asset stock by age for different values of individual productivity, for a given firm type (the mean firm, with $\alpha_{ij(.)}=0$). If the individual has an asset stock above this reservation value then she will not participate. In this case, the reservation asset stock declines

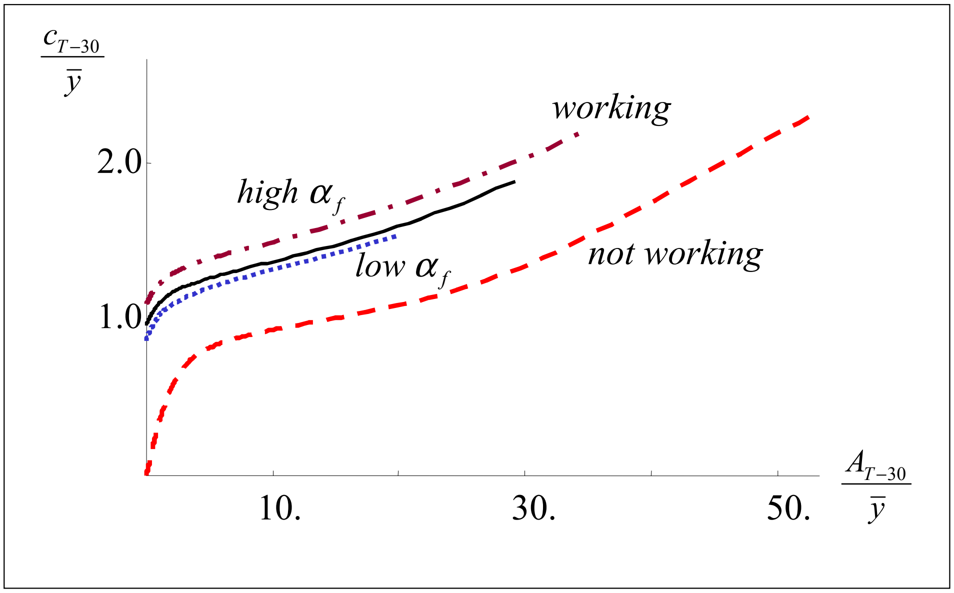

Figure 2: Consumption Functions for $T-30$, conditional on Participation Status.

with age. The reason is that young people have to have a higher level of assets in order to afford unemployment spells that are inherently risky due to uncertainty about receiving wage offers and uncertainty about the firm type.

6 Calibration

The last step of the empirical analysis is to calibrate the model by choosing λ^e, λ^n, and δ, the job arrival rates for unemployed and employed individuals, and the probability of job destruction, respectively. We also need to set 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. In this section, we describe the life-cycle profiles from the data that we use in the calibration and we discuss our calibration procedure.

6.1 Duration and Participation Profiles

In calibrating the remaining parameters, we use life-cycle profiles of unemployment duration and participation rates. 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

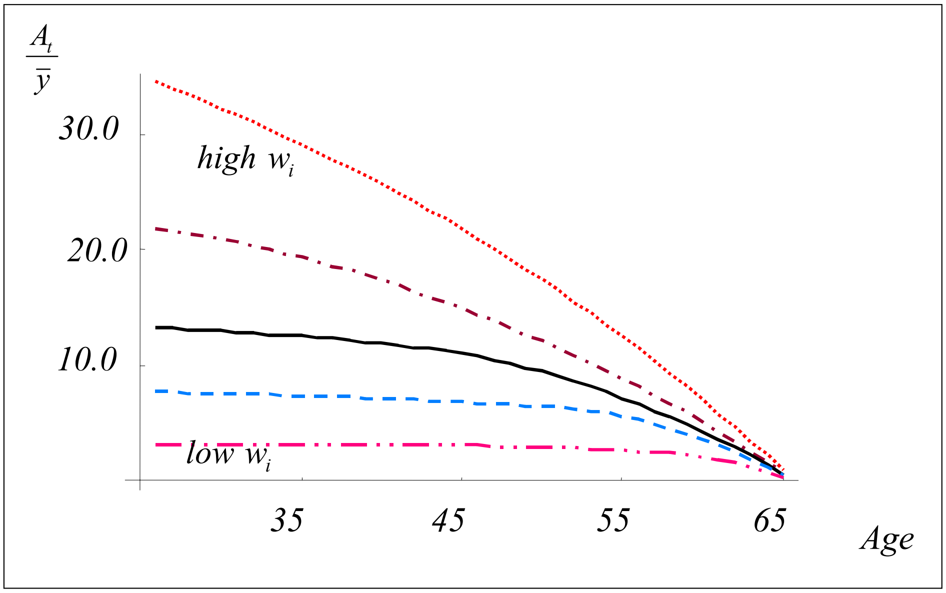

Figure 3: Reservation Asset Stock, conditional on Individual Productivity and Firm Type.

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, 19XX, for a criticism).

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.[17]

The dashed lines in figures 7 and 6 show the average unemployment duration by age of entry in the unemployed state, measured in quarters, for low and high educated individuals, respectively.[18] We classify anyone age less than 65 who is not working as unemployed, and anyone over 65 as retired. Durations have a maximum length determined by the number of quarters until age 65. For both education groups, durations appear increasing in age until around age 50; they decline afterwards as the length of time until 65 falls. The pattern is similar for both education groups but the less educated spend on average 1.5 more quarters unemployed than the better educated.[19]

[17] In future drafts we will extend our sample to 1998 to minimize the proportion of right-censored spells.

[18] We group unemployment spells by the age at which the spell began. Age groups span 5 years, and so the mean duration of exit for that group is averaging over all spells of unemployment for individuals who exit within that age span. Since a period in our theoretical model is a quarter of year, we used only spells of three months or longer in constructing these graphs.

[19] In principle, we could give comparable graphs for employment durations (regardless of employer) and durations

Participation profiles are given by the dashed lines in figures 5 and 4 for low and high educated individuals, respectively. For both education groups, participation rates are fairly constant until age 45, followed by a sharp decline to age 65. 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 up to age 45 (participation rates around 96%, compared to 90% for the low educated), and the subsequent decline is less marked.[20]

6.2 Calibration Procedure

We calibrate the necessary parameters by fitting profiles simulated from our model to the data on life-cycle profiles of unemployment duration and participation rates. In an earlier section we have outlined the model we simulate and we have discussed in the Appendix the way in which we solve the model numerically. Over the life cycle (a indexes age), there are the following relationships linking the proportion of people participating and the duration of unemployment to our parameters of interest, δ (job destruction), λ^n (the job offer arrival rate when unemployed), F (the fixed cost of work) and η (the nonseparable effect of participation on utility):

$$P_a = \left[(1-\delta_a) P_{a-1} + \lambda_a^n (1-P_{a-1})\right] \Pr\left(V_a^e > V_a^n\right) \tag{13}$$

$$D_a^n = \frac{1}{\lambda_a^n \Pr\left(V_a^e > V_a^n\right)} \tag{14}$$

where of course $\Pr\left(V_a^e > V_a^n\right)$ depends on all the parameters.

In matching profiles, we simulate life-cycle profiles for 10,000 individuals. 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 average duration by age of exit and participation rates over the life-cycle by taking the mean values across these simulated individuals to give profiles comparable to the data. We summarize the profiles in the data by the average duration and the average participation rate up to age 45, and by the average duration and the

until switching to a different employer. As with unemployment spells, these durations are censored and in this case the censoring is more problematic given that the duration is long relative to the data sample length.

[20] We compared unemployment durations and labor market participation rates in the PSID with unemployment durations and labor market participation rates in the SIPP (1993 Panel). Accounting for the different lenght of the two data sets (6 years in the PSID, 3 in the SIPP) we find similar durations and participation rates. In future drafts we plan to use the SIPP to estimate wage dynamics as well.

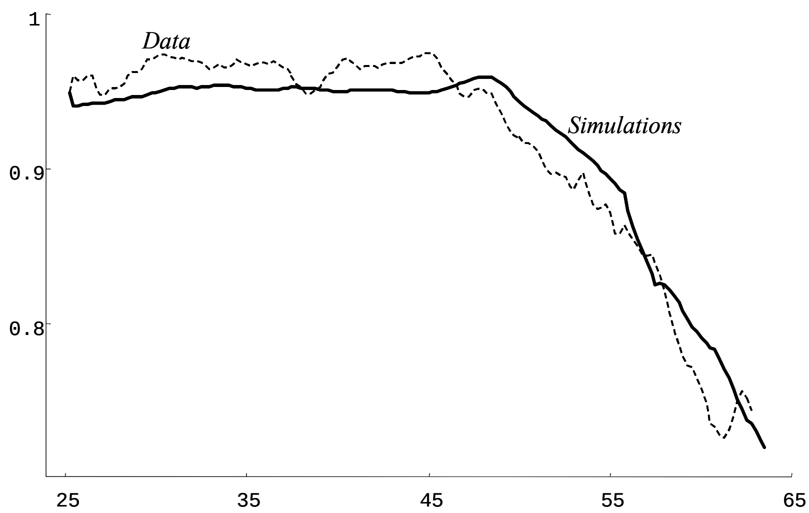

Figure 4: Participation rates of the high educated.

average participation rate from age 45 to 65. This provides us with 4 moments and we use these moments to calibrate our 4 parameters δ, λ^n, F, and η. In principle we could use the duration of employment at a particular firm (tenure) to calibrate the value of λ^e. In practice we have not constructed these data as yet, because they are notoriously affected by considerable measurement error. While it is true that the on-the-job arrival rate will affect the decision to accept a particular job offer because it determines the expected time until another offer if the first is accepted, variation in λ^e makes little difference to the simulated profiles. We thus assume that $\lambda^e = \lambda^n$.[21] We use a nonlinear equation solver to select the four parameters to give as close a match as possible.[22] More precisely, we select our 4 parameters $\lambda^n = \lambda^e$, δ, η, and F as follows. First, we take initial values for the 4 parameters. Second, we simulate the model, generate participation and duration using (13) and (14), and construct averages of the simulated participation and duration data over the two "periods" 25-44 and 45-65 which can be confronted with actual averages in the data. We iterate until there is an acceptable match between simulated and actual moments.

The results of this calibration are best shown in Figures 4, 5, 6, and 7. The first two show the match between duration rates by age of exit for the high and the low educated. The second two graphs show the match between participation rates over the life-cycle. We assume λ^e, δ, F and η

[21] By using data on job duration, we hope to relax this assumption.

[22] 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.

Figure 5: Participation rates of the low educated.

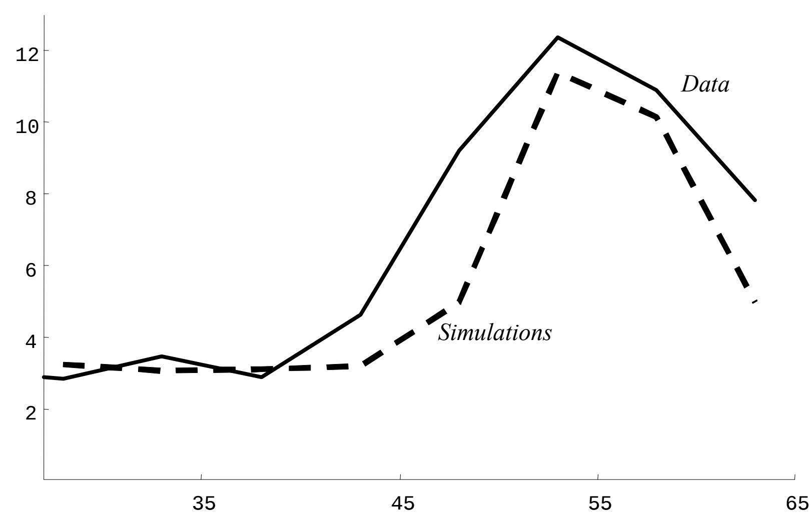

Figure 6: Mean duration of unemployment by age of entry, high educated.

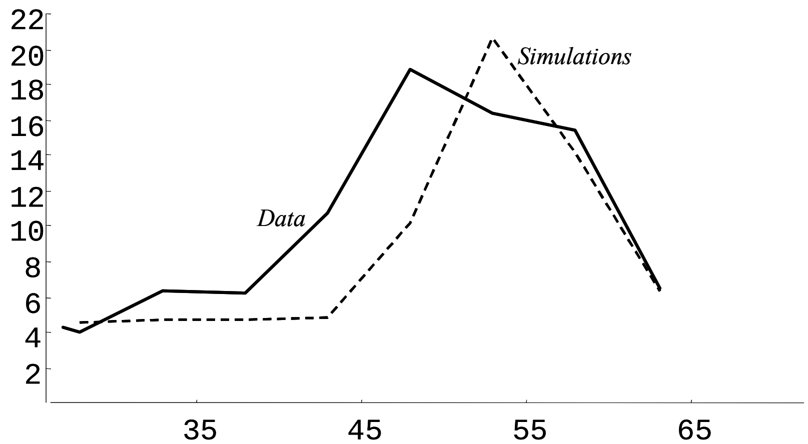

Figure 7: Mean duration of unemployment by age of entry, low educated.

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.

In table 5, we present the parameter values. We distinguish between parameter values estimated directly, parameters calibrated and parameters where values are imposed.

Table 5
Preference and risk parameter values used in simulations

Panel A: Imposed parameters

γ	1.5
$\frac{1}{\beta} - 1$	0.0043 (0.02)
r	0.0037 (0.015)
T (Max working life, in quarters)	160
R (Quarters of retirement)	40

Panel B: Estimated parameters

	Low Education	High Education
σ_α	0.18	0.17
σ_ζ	0.09	0.13

Panel C: Calibrated parameters

	Low Education	High Education
δ	0.023	0.016
λ^n	0.21	0.39
λ^e	0.21	0.39
η	0.05	0.35
$F/E\left[y_T\right]$	0.135	0.135

The values of δ, λ^n and λ^e are given as quarterly rates.

The extent of deterministic wage growth is taken from the PSID. The extent of risk aversion is taken from estimates of the intertemporal elasticity by Attanasio and Weber (1995). We assume individuals have a 40 year working horizon (age 25-65) followed by a deterministic 10 year retirement spell. There is no exogenous funding of consumption in retirement: individuals have to save for their retirement. 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 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. Details of the food stamp program, disability insurance and unemployment insurance are set to match actual programs as discussed in section 3 above.

6.3 Baseline Results

6.4 Welfare Cost of Risk

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, 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. In this section, we calculate the welfare benefit of reducing risk.

When considering removal of some type of risk we calculate the proportion of the offered wage (accepted or not) that an individual would be willing to pay to avoid facing the risk.[23] We hence use a welfare measure similar to that proposed by Lucas (1987). A change in the risk will typically have a direct wealth effect. However in our context there is no obvious non-distortionary way of correcting for this. Thus we also document the resulting wealth effect, which can be compared with the risk premium the individual is willing to pay to avoid risk.

Another major issue is that any change in the risk properties is likely to have labor supply effects, which in practice can have an effect on the price of labor. This will lead individuals to view the changes in a different light. In other words there may be important aggregate implications leading to trade-offs when considering situations that reduce or remove idiosyncratic risk. This issue is acutely clear when we consider the removal of employment risk, since this will induce a potentially large change in labor supply as workers will be able to find appropriate job offers much faster. This may affect the price of labor. Ours is not a full equilibrium model. Characterizing the general equilibrium in an economy with permanent shocks to wages and savings as well as firm level heterogeneity has not been done as yet. In order to assess how sensitive our results are likely to be to such price changes we consider a situation in which the labor demand elasticity is infinite, implying that the labor supply changes do not have aggregate effects. This approximates the situation where idiosyncratic risk is removed for one individual alone, leaving all others unaffected.[24]

Table 6 shows these results. First consider the case where we remove productivity risk. High education individuals would be willing to pay 15.4% of all wage offers while the low education individuals are willing to pay only 4.85%. 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. A progressive social

[23] Thus, in this calculation we assume that people pay a premium of $\pi\times$ wage offered (which equals $\max\{w_{ij(t_0)t}, w_{ij(t)t}\}$ if employed and the offered wage if unemployed). We ask people to pay a proportion of the offered wage (irrespective of whether the offer is accepted) to keep the relative price of work unchanged and avoid the distortionary effects of insurance, which would let people to reject too many offers. This is similar to standard moral hazard problem of UI (see Hansen and Imhoroglu, 1990).

[24] In the current version this is just an approaximation since we allow a link across individuals through taxes that fund UI, food stamps and disability. These taxes are sensitive to labour supply changes.

security would clearly mitigate some of these costs. 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.

We can repeat this exercise for employment risk. To do this we set the job arrival rates for both employed and unemployed to 1 but leave all other parameters unchanged. With an infinite elasticity of labor demand we find that the highly educated would only be willing to pay 3.2% of all wage offers, while the low educated would be willing to forego a substantial 6.5% of all offers. Thus the risk position is reversed for the two education groups.

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. When we set the variance in firm heterogeneity to zero in addition to the arrival rates being one we find that individuals wish to be compensated for the loss of the chance to obtain a better quality job. This contrasts with the result that would be obtained if we took firm heterogeneity as a risk to which individuals cannot respond to.

Table 6
Welfare benefits

	High Education		Low Education	
Scenario	*Proportion of wage offer willing to pay* (π)	*Output Change*	*Proportion of wage offer willing to pay* (π)	*Output Change*
No unempl. frictions ($\lambda^n = 1$)	0.044	0.99	0.073	1.048
No unempl. frictions and firm het. ($\lambda^n = 1$, $\sigma_\alpha^2 = 0$)	-0.12	0.86	-0.076	0.96
No frictions ($\lambda^n = \lambda^e = 1$)	0.031	1.019	0.065	1.121
No frictions and firm het. ($\lambda^n = \lambda^e = 1, \sigma_\alpha^2 = 0$)	-0.125	0.87	-0.076	0.96
No productivity risk ($\sigma_\zeta^2 = 0$)	0.154	0.96	0.0485	0.92

Note: A positive number for the risk premium indicates the amount that the individual is willing to pay. A negative number for the risk premium indicates the individual would need compensation for moving to the new scenario from the baseline. The column with "Ouput change" shows the value of output (baseline=1) corresponding to the scenario considered.

6.5 Welfare Benefit of Government Insurance Policies

We can use our model to evaluate government policies such as unemployment insurance or food stamps which target part of the risk individuals face. In Table 7 we calculate the welfare loss associated with removing unemployment benefits and removing food stamps for each education group. This calculation focuses on the insurance benefit of these programs because there is no cross-group redistribution. The table shows that, first, the welfare benefit of food stamps is higher than unemployment benefit for both groups. Second, food stamps are "cheaper" in the sense that the increase in the tax rate needed to pay for the given amount of food stamps is less.[25] Finally, the insurance benefit of both programs are higher for the low education group. Since there is no redistributive element to the program, this difference arises because of the lower wealth of the low educated and the greater job destruction possibilities.[26]

However the striking feature of these results is that UI is basically worthless for all individuals; this is because of its time limited nature which prevents it from insuring against not finding a new job. Food stamps on the other hand provide a universal insurance against both productivity risk and unemployment risk and hence has a very high value, particularly for low educated individuals, despite the obvious moral hazard implications.

Table 7
Welfare benefits of Government Insurance

		High Education		Low Education	
Scenario	*Parameters*	*Proportion of wage offer willing to pay* (π)	*Change in tax rate* (τ_w)	*Proportion of wage offer willing to pay* (π)	*Change in tax rate* (τ_w)
No U.I.	$B_{it} = 0$	-0.0076	-0.0061	-0.0078	-0.0107
No food stamps	$FS_{it} = 0$	-0.0711	-0.0227	-0.117	-0.124

The tax rate reported is the change in the tax rate necessary to keep net government spending constant (i.e. no change in the government deficit). A negative number means the tax rate is reduced. For the risk

[25] In the baseline scenario we have $\tau_w = 0$, so the government is running a deficit that we keep constant when we remove any of the programs.

[26] This experiment misses one important feature of the UI benefits, which is the cap on the level of benefits people may receive. A limit is instead imposed on the amount of food stamps people may receive.

premium, a positive number indicates the amount that the individual is willing to pay. A negative number for the risk premium indicates the individual would have to be compensated for moving to the new scenario from the baseline.

7 Conclusions

In this paper we distinguish productivity risk and employment risk as two distinct sources of uncertainty that the individual faces. To achieve this we specify an economy with search frictions and firm level heterogeneity. Thus individuals face employment risk because it may take time for a suitable job to be found once out of work, although such a job may exist. Productivity risk on the other hand is due to permanent idiosyncratic shocks to wages that we assume are uninsurable. Within this context we take the important step of allowing for endogenous mobility between firms. This turns out to be very important, since it reduces the amount of risk others have attributed to income or wages by a large amount.

We specify an intertemporal model of labor supply and savings with idiosyncratic uncertainty. We allow for the effects of food stamps, unemployment insurance and disability benefits. We obtain the parameters of the model by a mixture of estimation and calibration. We estimate the wage process using PSID data. We then, set the intertemporal elasticity of substitution parameter using estimates from other studies. Finally we use calibration to set the rest, based on matching unemployment durations and life-cycle participation profiles.

When we estimate the wage process we find that a third of what earlier studies attributed to permanent wage shocks is in fact due to endogenous mobility between firms. Indeed bad offers can always be turned down, while good ones are accepted and this does not constitute risk. We also found that a large proportion of wage dispersion can be attributed to firm level heterogeneity which is an important ingredient in defining employment risk. Finally the job arrival rate from our calibration is substantially below 1. All these ingredients imply that employment risk can be an important factor. We next go on to quantify the welfare implications of such risk.

The main findings of the paper are that productivity risk is much more important than unemployment risk, both for the low educated and the high educated workers. In addition productivity risk is about twice as costly for the high educated workers. On the other hand, employment risk is a significant factor for the low educated workers. However existing time limited UI schemes have little or no insurance value to individuals in this economy. By far the most important insurance program is Food Stamps which are highly valued in this economy, despite the moral hazard implications. However, in this paper we have not addressed the issue of optimality of programs. We have looked at them with their current specification.

The paper ignores key aspects of general equilibrium in the type of economies we study. First, we assume that the sharing rule for the surplus between workers and firms remains constant across simulations; second we do not allow for any search externalities. The reason for these simplifications is that characterizing equilibria in such economies as ours with firm level heterogeneity and permanent shocks to wages has not yet been achieved.

A Appendix: Moments for estimating the stochastic process of wages

The identification of the wage dynamics parameters is achieved by considering the first two moments and the first order autocovariance of unexplained wage growth for workers and distinguishing between movers and stayers. All the conditional moments in what follows leave the conditioning on labor market participants over the relevant time horizon implicit.

Recall that unobserved wage growth is defined as[27]

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

For the individuals who do not separate from their employer:

$$E\left(g_{it} \left| M_{it} = 0\right.\right) = \sigma_{\zeta} \rho_{\zeta\pi} \lambda_{\pi} - \sigma_{\zeta} \rho_{\zeta\mu} \lambda_{\mu_0} \tag{15}$$

where $\lambda_{\mu_0} = \frac{\phi\left(k_{it}'\theta\right)}{1-\Phi\left(k_{it}'\theta\right)}$. There is a simple interpretation of the two terms on the right hand side of (15). The first term is the average wage growth of labor market participants, the second is the average wage growth of stayers. Suppose that people who receive positive permanent shocks are also more likely to participate ($\rho_{\zeta\pi} > 0$). Then $\sigma_{\zeta}\rho_{\zeta\pi}\lambda_{\pi} > 0$ reflects the fact that labor market participants have above average wage growth (because they are more likely to have received positive productivity shock). As for the second term, mobility means either that an individual moves directly to a new firm or that an individual has experienced a period of unemployment but is now reemployed at a new firm. Suppose for instance that movers are characterized by low realizations of ζ_{it}, i.e., $\rho_{\zeta\mu} < 0$. If this is the case, however, the second term on the right hand side of (15), $-\sigma_{\zeta}\rho_{\zeta\pi}\lambda_{\pi} > 0$, and the stayers ($M_{it} = 0$) are those who have received good realizations of the permanent shocks and therefore exhibit above average wage growth.[28]

[27] In reality, we observe only a consistent estimate of g_{it}, $\widehat{g}_{it} = \Delta \ln w_{it} - \Delta x_{it}' \widehat{\beta}$ (as long as $\widehat{\beta}$ is consistent). Using $\widehat{g}_{it}$ in the place of g_{it} is a valid operation as long as the fourth moments exist and are constant across individuals. See MaCurdy (1982, 2004).

[28] What may give this negative correlation? Consider the following extension of model (1), omitting for simplicity observable characteristics and measurement error:

$$w_{it} = u_{it} + \alpha_{ij(t_0)t}$$

where as before

For movers:

$$E\left(g_{it}\left|M_{it}=1\right.\right)=\sigma_{\zeta}\rho_{\zeta\pi}\lambda_{\pi}+\sigma_{\zeta}\rho_{\zeta\mu}\lambda_{\mu_1}+\sigma_{\xi}\rho_{\xi\pi}\lambda_{\pi}+\sigma_{\xi}\rho_{\xi\mu}\lambda_{\mu_1} \tag{16}$$

with $\lambda_{\mu_1}=\frac{\phi\left(k_{it}'\theta\right)}{\Phi\left(k_{it}'\theta\right)}$. Interpretation of the various terms in (16) follows the same lines. The moments (15) and (16) have a regression representation:

$$g_{it}=E\left(g_{it}\left|M_{it}=0\right.\right)\left(1-M_{it}\right)+E\left(g_{it}\left|M_{it}=1\right.\right)M_{it}+e_{it}^{g}$$

This is a regression of g_{it} onto the Mills' ratios and the interaction of the Mills' ratios with the mobility indicators. e_{it}^{g} represents sampling variability in g_{it} (with the property that $E\left(e_{it}^{g}|M_{it}\right)=0$) The parameters of the model are clearly not identified from this alone.

Consider then the second moment. For workers who do not move, the correct expressions is:

$$E\left(g_{it}^{2}\left|M_{it}=0\right.\right)=\sigma_{\zeta}^{2}\left(1-\rho_{\zeta\pi}^{2}z_{it}'\gamma\lambda_{\pi}+\rho_{\zeta\mu}^{2}k_{it}'\theta\lambda_{\mu_0}-2\rho_{\zeta\pi}\rho_{\zeta\mu}\lambda_{\pi}\lambda_{\mu_0}\right)+2\sigma_{e}^{2} \tag{17}$$

and

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

is unobserved general productivity (portable across firms), and

$$\alpha_{ij(t_0)t}=\alpha_{ij(t_0)t-1}+\omega_{ij(t_0)t}=\alpha_{ij(t_0)t_0}+\sum_{s=t_0+1}^{t}\omega_{ij(t_0)s}$$

The firm time-invariant effect is $\alpha_{ij(t_0)t_0}$, the initial match productivity at the start of the relationship, while $\sum_{s=t_0+1}^{t}\omega_{ij(t_0)s}$ augments or diminishes the value of the initial match due to random shocks.

In this extended model, a worker switches at time t to a new firm if what he gets there, $u_{it}+\alpha_{ij(t)t}$ exceeds what he would get at the current firm, $u_{it}+\alpha_{ij(t_0)t_0}+\sum_{s=t_0+1}^{t}\omega_{ij(t_0)s}$, net of some moving costs, i.e., if $k_{it}'\theta+\mu_{it}\geq 0$ and

$$\mu_{it}=\xi_{it}-\sum_{s=t_0+1}^{t}\omega_{ij(t_0)s}+\varepsilon_{it}$$

and ε_{it} are unobserved moving costs. The wage growth equation can now be written as:

$$g_{it}=\left(\zeta_{it}+\omega_{ij(t_0)t}\right)+\left(\xi_{it}-\sum_{s=t_0+1}^{t}\omega_{ij(t_0)s}\right)M_{it}$$

which has the same form of (12). Note that the covariance between the first term in parenthesis on the RHS and μ_{it} is $-E\left(\omega_{ij(t_0)t}^{2}\right)<0$, and therefore $\rho_{\zeta\mu}<0$. This argument is heuristic because we cannot separately identify ζ_{it} and $\omega_{ij(t_0)t}$. Considering the implications of this alternative model is left for future drafts.

$$E\left(g_{it}^2\left|M_{it}=1\right.\right)=\sigma_\zeta^2\left(1-\rho_{\zeta\pi}^2 z_{it}'\gamma\lambda_\pi-\rho_{\zeta\mu}^2 k_{it}'\theta\lambda_{\mu_1}+2\rho_{\zeta\mu}\rho_{\zeta\pi}\lambda_{\mu_1}\lambda_\pi\right)+2\sigma_e^2$$
$$+\sigma_\xi^2\left(1-\rho_{\xi\pi}^2 z_{it}'\gamma\lambda_\pi-\rho_{\xi\mu}^2 k_{it}'\theta\lambda_{\mu_1}+2\rho_{\xi\mu}\rho_{\xi\pi}\lambda_{\mu_1}\lambda_\pi\right) \tag{18}$$

The regression representation in this case is:

$$g_{it}^2=E\left(g_{it}^2\left|M_{it}=0\right.\right)(1-M_{it})+E\left(g_{it}^2\left|M_{it}=1\right.\right)M_{it}+e_{it}^{g^2}$$

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 consider only the restrictions on the unconditional autocovariance, namely that

$$g_{it}g_{it-1}=-\sigma_e^2+e_{it}^{gg_{-1}} \tag{19}$$

The three moment conditions:

$$E\left(e_{it}^{g}\right)=0$$
$$E\left(e_{it}^{g^2}\right)=0$$
$$E\left(e_{it}^{gg_{-1}}\right)=0$$

are used to identify the parameters of interest. We apply a NLS procedure that imposes restrictions across the various equations.

B Appendix: Solution method

There are five state variables in this problem: age, the asset stock, the individuals own productivity, employment status, and firm type. We discretize the individuals own productivity and the range of firm types, leaving the asset stock as the only continuous state variable. The conditional value functions are increasing in assets A_t but they are not necessarily concave, despite being conditional on labor market status in t. The non-concavity arises because of changes in labor market status in future periods: the slope of the value function is given by the marginal utility of consumption, but this is not monotonic in the asset stock because consumption can decline as assets increase and labor market status changes. This problem is also discussed in Lentz and Tranaes (2001). By contrast, in Danforth (1979d) employment is an absorbing state and so the conditional value function will be concave. Under certainty, the number of kinks in the conditional value function is given by the number of periods of life remaining. If there is enough uncertainty, then changes in work status in the future may be smoothed out: whether or not an individual will work in $t+1$ at a given A_t depends on the realization of wages in $t+1$. Using uncertainty to avoid non-concavities is analogous to the use of lotteries elsewhere in the literature.

In the value functions (8) and (3.5), the choice of participation status in $t+1$ is determined by the maximum of the conditional value functions in $t+1$. In our solution, we impose and check restrictions on this participation choice. In particular, we use the restriction that the participation decision switches only once as assets increase, conditional on productivity and the firm type. When this restriction holds, it allows us to interpolate behavior across the asset grid without losing our ability to determine participation status. We therefore define a reservation asset stock, $R_t^{u,f}$ to separate the value function and the choice of consumption made when participating from the value function and choice of consumption made when not participating. We define $R_t^{u,f}$ implicitly through

$$V_t^e\left(R_t^{u,f}, u_t, d_t = 1, a_f\right) = V_t^n\left(R_t^{u,f}, u_t, d_t = 0\right) \tag{20}$$

Solving for the reservation asset stock serves two purposes: one, it makes it easier in the solution to allow for the fixed cost in the budget constraint (rather than having an unconditional policy function with a discontinuity); two, it provides an additional check on our numerical solution: the reservation asset stock should be increasing in the wage rate and increasing in the firm productivity. A sufficient condition for this to be unique is that the conditional value functions be concave. This is not true in general, as discussed above, but uniqueness can be achieved by having enough uncertainty to make the conditional value functions concave. Even when the conditional value functions are not concave, however, we can have a unique reservation asset stock, particularly if individuals are impatient enough: impatience means that individuals prefer periods of non-participation to be earlier in the life-cycle and so avoids the indifference about the timing of leisure which generates the nonuniqueness.

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

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

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

[6] 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.

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

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

[9] 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.

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

[11] 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

[12] 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

[13] Horowitz, J. (2002) "Bootstrap Methods for Markov Processes," Econometrica

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

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

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

[17] Low, H. (1999) "Self-insurance and unemployment benefit in a life-cycle model of labor supply and savings" *Institute for Fiscal Studies Working Paper* 99/24

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

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

[20] 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

[21] Meghir, C. and C. Dustmann (2001) "Wages, experience and seniority" *Institute for Fiscal Studies Working Paper* 01/01

[22] Meghir, C. and L. Pistaferri (2004) "Income variance dynamics and heterogeneity " *Econometrica*, forthcoming

[23] 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

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

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

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