On-the-Job Search, Mobility Costs, and Wage Inequality *

Susanne Forstner †

This Version: April 4, 2011

Abstract

Over the last three decades, the U.S. labor market has exhibited a persistent rise in overall as well as residual wage inequality. Within the same time period, the rate of job-to-job transitions of U.S. workers has increased significantly, suggesting that it has become easier for workers to move between employers. In the present paper, I relate this increase in job-to-job mobility to evidence on a parallel decrease in the costs to workers of changing from one employer to another. With a view to this evidence, I address the question of how much the reduction in mobility costs has contributed to the observed increase in residual wage inequality. For this purpose, I develop a search model with on-the-job search and costs to workers when leaving a particular employment position. Calibrating the model, first to the situation in the United States with relatively high mobility costs in the early 1980s, and then to the situation with lower mobility costs in the early 2000s, I find that the reduction in mobility costs has contributed to the rise in residual wage inequality over this period. In particular, for the lower parts of the wage distribution, the decrease in mobility costs may have accounted for up to one-half of the observed increase in inequality.

Keywords: On-the-Job Search, Worker Mobility, Employer-to-Employer Transitions, Wage Inequality, Labor Market Dynamics

JEL Codes: J31, J62, E24

---

*I am indebted to Arpad Abraham and Ramon Marimon for guidance, numerous valuable comments, and many fruitful discussions. I am also grateful to Giancarlo Corsetti, Georg Duernecker, Nicola Pavoni, Gonul Sengul, and participants at the EUI Macro Student Lunch seminars, for helpful comments and suggestions. All remaining errors are mine.

†Contact address: Department of Economics, European University Institute, Via della Piazzuola 43, I-50133 Florence, Italy. E-mail: Susanne.Forstner@EUI.eu

1 Introduction

Since the mid-1970s, the labor market in the United States has exhibited a persistent widening of the overall wage distribution. What is more, this trend was to a large extent due to an equally persistent increase in the dispersion of wages among observationally equivalent workers. This development has raised considerable interest for a number of reasons. One of them is that wage inequality is a major factor behind the inequality of households with regard to consumption and wealth, an issue which continues to receive broad attention in the literature.

For the observed rise in residual wage inequality in the United States a number of potential explanations have been put forward. Most prominent among them are those based on the arguments that the returns to unobservable worker characteristics have increased, that the economy has become more turbulent in the sense of workers' job-specific skills depreciating more strongly upon a job change, or that the overall labor market risk for workers has risen substantially.

In the present paper, I propose another source of increase in wage inequality, namely, reduction of the costs to workers of changing employers. One motivation for examining this potential source is the fact that, over the period studied here, rising wage inequality was accompanied by increasing rates of job-to-job transitions and falling mobility costs.[1] The mechanism through which the latter impact on the cross-sectional wage distribution has two main components. First, lower costs of changing employers reduce the reservation wage at which unemployed workers are willing to enter the labor market, thereby increasing the range of wages acceptable to workers. Second, lower mobility costs induce employed workers to accept lower outside wage offers, thereby reducing the average size of wage increases upon a job-to-job transition, while at the same time increasing the frequency of such transitions. Thus, when mobility costs decrease, workers start employment spells at lower wages, and then move up the wage ladder at a higher frequency, but in smaller steps. The overall effect of mobility cost reduction on the cross-sectional wage distribution and wage inequality is a priori ambiguous.

In the present analysis, I examine in quantitative terms the relationship between the decline of mobility costs and the rise in wage inequality in the United States between the early 1980s and the early 2000s. In particular, I address the following question: "Has the reduction of job-to-job mobility costs contributed to the observed rise in residual wage inequality over the period studied here? And if so, how large was this contribution?" In order to propose an answer to this question, I build a job search model with on-the-job search, featuring in addition utility costs of job-to-job transitions. The model is calibrated so as to represent the situation of relatively high mobility costs in the early 1980s, on the one hand, and that of lower mobility costs in the early 2000s, on the other. This provides

[1] In the literature, e.g. Rhode and Strumpf (2003), a secular decline in mobility costs is documented. It is typically traced to factors such as decreasing costs of transportation and communication, or increasing similarity among regions with respect to cultural traits, public goods provision, and working practices. More details on recent U.S. developments in these areas are presented in Section 2 below.

an opportunity to quantify the contribution of a reduction in these costs to the increase in wage inequality among ex-ante identical workers.

The results of a quantitative analysis based on the calibrated model indicate first of all that the reduction of mobility costs did in fact contribute to the increase in residual wage inequality. In addition, I find that the impact of cost reduction on inequality differed substantially across different parts of the wage distribution. In particular, it was strongest by far in the lower parts of the distribution, where the decline in mobility costs may explain up to one-half of the observed rise in residual wage inequality.

The model underlying the present research is based on the theory of job search, extended by two essential features: on-the-job search and costly job changing. The analysis builds on two strands of the literature. One is that on the sources of 'pure' wage dispersion, that is, wage differences between ex-ante identical workers. Accordingly, the theoretical approach outlined in Mortensen (2003) informs the basic construction of my model. In addition, some of its extensions are inspired by the work of Hornstein et al. (2007), who compare the performance of different models with respect to explaining the empirical evidence on the extent of wage dispersion. In particular, in my model framework I include - similar to Nagypal (2005) - idiosyncratic productivity shocks during employment. At the same time, I assume workers to be risk-averse, as for example in Lise (2010).

The other strand of literature closely related to the present paper is that of studies dealing with the influence of mobility costs on labor market characteristics more generally. An early theoretical analysis of job search with mobility costs is that of Hey and McKenna (1979), who focus on constant costs to workers of changing from one job to another. Van den Berg (1992) provides an empirical analysis of the extent and the determinants of job-changing costs, where such costs are assumed to depend on the wage level. Finally, Burgess (1992) presents a theoretical analysis of the implications of job-changing costs for aggregate unemployment.

Three more recent studies are of particular relevance to the present research: Lee and Wolpin (2006) examine the role of mobility costs in explaining the intersectoral structure of output, employment, and wages. They also provide estimates of the costs to workers of moving between sectors in the United States, which show that these costs are substantial. Jolivet et al. (2006) offer empirical support for the close correspondence between the determinants of labor turnover on the one hand, and those of wage dispersion on the other. And Cabrales et al. (2008) - in the context of studying the implications of equity concerns for labor market outcomes in general - examine the impact of changes in mobility costs on wage dispersion both within and across firms.

The paper is organized as follows: Section 2 outlines the empirical background to my work. Section 3 presents the model used in the analysis, states the optimization problem, and defines stationary equilibrium. In Section 4 I discuss in some detail how the model is put to use in a quantitative analysis. The results of this analysis are presented in Section 5, and Section 6 concludes.

2 Empirical Background

In this section, I provide an overview of empirical evidence on developments in the U.S. labor market which motivates the questions addressed in the present paper. First, I present facts on the evolution of wage inequality in the United States over the past few decades. Then I move to evidence on job-to-job mobility of workers, in particular, on the evolution of job-to-job transitions since the late 1970s. This information is supplemented by evidence on the reduction of workers' mobility costs over the same time period. Finally, I discuss evidence on the characteristics of individual labor earnings histories, and on how these have changed over the period considered.

2.1 Wage Inequality

A large number of empirical studies have documented the substantial and persistent widening of the cross-sectional distribution of wages in the United States from the 1970s onwards. Katz and Autor (1999) provide a comprehensive survey of the respective literature up to the late 1990s. In addition, they present their own estimates from March Current Population Survey (CPS) data, which produce a picture similar to the findings of other authors. Between the late 1970s and the mid-1990s, weekly labor earnings of the 90th relative to the 10th percentile worker increased by over 25 percent for both men and women. Eckstein and Nagypal (2004) re-examine the main trends in the U.S. wage distribution, using March CPS data from 1962 up to 2003. They find a sharp increase in overall wage inequality for men and women, starting in the 1970s and continuing throughout the period up to the early 2000s. Figure 1(a) shows the evolution of the cross-sectional standard deviation of log wages in their sample.

Part of the observed increase in wage inequality in the United States over the last 35 years can be explained by changes in the relative wages of different groups of workers, distinguished by observable characteristics such as sex, education, and experience.[2] For example, the rise in returns to skills since the late 1970s, as measured by the increase in the college wage premium, has received broad attention as a source of the widening of the wage distribution. However, the increase in overall wage inequality since the 1970s was accompanied by a dramatic increase in the dispersion of wages *within* narrowly defined demographic and skill-defined groups of workers. For instance, Katz and Autor (1999) analyze the evolution of the residuals from Mincerian wage regressions using March CPS data from 1964 to 1996. They report that the ratio of the 90th-to-10th percentile of residual wages has increased significantly for men since the mid-1970s, and for women since the 1980s.[3] Eckstein and Nagypal (2004) estimate similar wage regressions using

[2]Katz and Autor (1999) and Eckstein and Nagypal (2004) also present detailed analyses of the changes in wage differentials between groups.

[3]Katz and Autor (1999) estimate separate regressions for men and women, for each year, of log weekly wages on eight education dummies, a quartic in labor market experience, interaction terms between three broad education groups and the quartic in experience, three regional dummies, and two race dummies. Their sample includes only full-time full-year wage and salary workers.

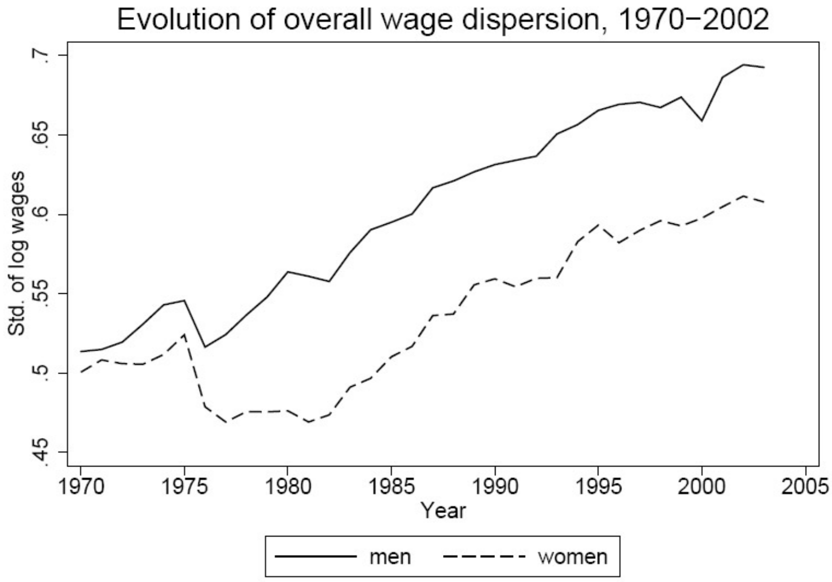

(a)

(b)

Figure 2.1:
1(a): Standard deviations of log weekly wages of full-time full-year workers from March CPS data, 1969 to 2003.
1(b): Standard deviations of residuals from Mincerian wage regressions, same dataset.

Both graphs are based on data and estimations from Eckstein and Nagypal (2004), which were kindly provided by the authors.

data up to 2003 and controlling in addition for occupational categories.[4] Their results confirm the findings of Katz and Autor (1999): The within-group standard deviation of weekly wages for men increased from an average of 0.47 in the early 1970s to 0.58 in the early 2000s, while that for women started to increase only in the 1980s, rising from an average of 0.44 for the late 1970s to 0.52 in the early 2000s.[5] The estimation results from Eckstein and Nagypal (2004) are depicted in Figure 1(b). A comparison of Figures 1(a) and 1(b) shows how closely the increase in residual wage dispersion paralleled the rise of overall wage dispersion over the period from the late 1970s to the early 2000s.[6]

2.2 Job-to-Job Mobility and Mobility Costs

The importance of job-to-job transitions in the process of reallocating labor in the U.S. economy has recently been documented by a number of empirical studies. For instance, Fallick and Fleischman (2004) and Nagypal (2008), analyzing monthly data from the CPS since 1994, find that the monthly rate of job-to-job transitions is on average two to three times larger in magnitude than the monthly flows of workers from employment to unemployment. For the time period from January 1994 to December 2003, Fallick and Fleischman (2004) report that each month on average 2.6 % of employed workers left one employer for another, whereas only 1.3 % of employed workers moved into unemployment. Furthermore, they find that an average of two-fifths of new jobs started over the above period were associated with direct employer changes. Nagypal (2008), using data up to August 2007, estimates that within a month on average 2.88 % of employed workers change their employer, and 0.89 % move to unemployment.[7]

Moreover, a number of empirical analyses document that significant changes in the size and composition of worker flows in the United States have taken place during the past few decades. Stewart (2002) is among the first authors who consistently measure the incidence of job separation in the United States, and its decomposition into different flows, over a long time period, using March CPS data from 1977 to 2001. In particular, the analysis contains an estimation of the evolution of the rate of job-to-job transitions since the mid-1970s. The author finds that the annual incidence of job-to-job transitions increased steadily over the period considered, reporting a rise from 8.6 % in 1975 to 13.7 % in the year 2000 of the average fraction of employed workers experiencing an employer-to-employer transition within a year. Sherk (2008), who applies the same method of analysis

[4]Eckstein and Nagypal (2004) estimate regressions for men and women, for each year, of log weekly wages on five education dummies, a quadratic function in experience, four regional dummies, and two race dummies, and later also three occupational dummies. Their sample includes only full-time full-year wage and salary workers between 22 and 65 years of age.

[5]Kambourov and Manovskii (2009) and Heathcote et al. (2010) find very similar time patterns for the evolution of the within-group variance of male hourly wages since the 1970s. Their data are taken from the Panel Study of Income Dynamics, 1969 to 1996, and 1967 to 2002, respectively.

[6]See Juhn et al. (1993) for an assessment of the contribution of residual wage dispersion to explaining the overall increase in wage dispersion.

[7]Nagypal (2008) obtains similar results from data of the Survey of Income and Program Participation for the same time period.

to data up to the year 2007, reports that the fraction of job-to-job changers increased to roughly 12.5 % in 2007. Figure 2.2 shows the time series of the annual incidence of job-to-job transitions for men from 1976 to 2007 as estimated by Sherk (2008).[8] Apart from strong procyclical fluctuations, the series of job-to-job transition rates exhibits a clear upward trend over the time period examined in the study.[9]

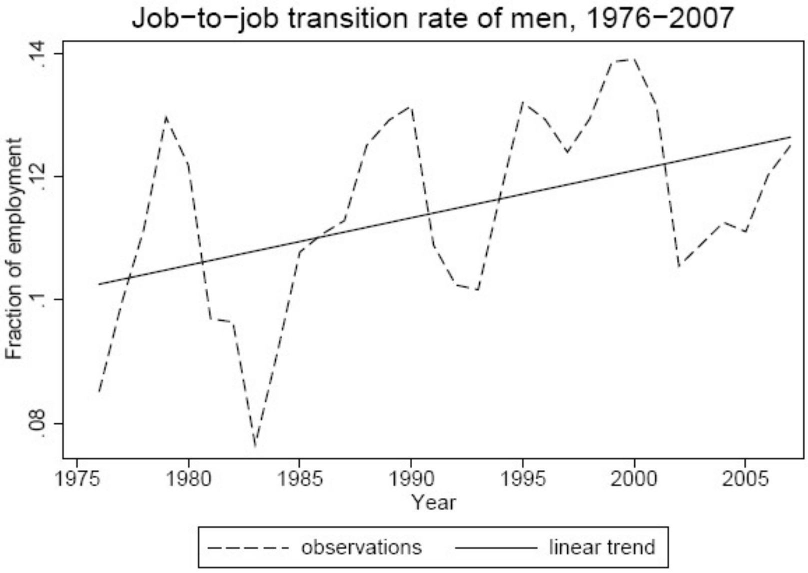

Figure 2.2: Evolution of the annual incidence rate of job-to-job transitions, based on March CPS data from 1977 to 2007.

Changing employers is associated with transaction costs to the worker. These costs can be of a pecuniary or a psychological nature, and arise from leaving the old job or from accepting a new one.[10] Among the major components of job changing costs are the costs of changing housing, of moving to another town or region, and of having to adjust to a new social and working environment. Furthermore, many jobs have characteristics in addition to wage payments, such as fringe benefits, the type of health insurance available to workers, and the possibility of accumulating pension claims with the employer. Since these non-wage compensations usually vary across employers, they can become an additional source of mobility costs.

The prime example for mobility costs that arise for workers when switching from one employer to another is the lack of full transferability of employer-based pension plans. In the United States, benefits from such plans make up a significant part of workers'

[8] Measurement of the job-to-job transition rate in Sherk (2008) is based on the methodology of Stewart (2002). The variable reported is the estimated number of workers who experienced at least one job-to-job transition over the time period of roughly one year, divided by the number of workers who were employed for at least one week during that year. James Sherk kindly provided his estimates, as well as documentation on the estimation procedure.

[9] The estimated time series reported by Stewart (2002), which covers the period up to 2001, looks very similar, and exhibits an even stronger upward linear trend.

[10] For an elaboration on the latter distinction, see Burgess (1992).

retirement income.[11] Employer-based pension plans can be either defined benefit (DB) or defined contribution (DC) systems. Historically, most of the plans offered by employers were of the defined benefit type, where benefit payments depend on earnings and years of service with the employer, and the major part of pension benefits are only achieved if the worker retires directly from service with that particular employer. Upon changing employers, workers lose most of their pension claims accumulated within a DB pension plan. This is not the case for defined contribution plans, where the amount of retirement income depends directly on the cumulated asset value of contributions and the investment income from accumulated funds. When switching jobs, workers do not lose the claims to benefits from DC pension plans established with their former employer. Defined benefit pension plans therefore create much larger opportunity costs to workers when changing employers than do defined contribution plans.

Starting in the late 1970s, a number of changes to U.S. legislation regulating the tax status of employer-based pension plans have induced a general shift from defined benefit to defined contribution plans.[12] Figure 2.2 illustrates how the relative importance of defined benefit and defined contribution types among private pension plans has been reverted since the early 1980s. While at the beginning of the period 60% of private wage and salary workers who were participating in an employer-based pension plan were covered by only a DB plan, this number had fallen to 17% by 2004. Over the same time period, the fraction of workers covered by only a DC plan rose from 11% to 61%. These facts imply that today the average U.S. worker faces much lower job changing costs associated with the loss of nontransferable pension claims than thirty years ago.

[11] Payments from these pension plans are additional to retirement income provided through the publicly administered Social Security. Munnell and Perun (2006) report that in 2004, employer-sponsored pension income accounted for 19% of total income for people above 65 years, and that private pension plans made up 24% of the wealth holdings of a "typical" household shortly before retirement.

[12] Most importantly, the Revenue Act of 1978 introduced the so-called "401(k)" type of employer-based defined-contribution pension plans, where the term "401(k)" refers to the corresponding section in the Internal Revenue Code. Contributions to pension plans, both by employer and employee, that satisfy the 401(k)-requirements, are tax deferred. The requirements include stipulations for the access of workers to the plan, the maximal amount of contributions, and the management of the financial assets. Since the early 1980s, the number of 401(k) plans has grown rapidly and continually, so that by now they account for the majority of employer-based DC plans in use. See Poterba et al. (2007) for detailed empirical evidence on this development.

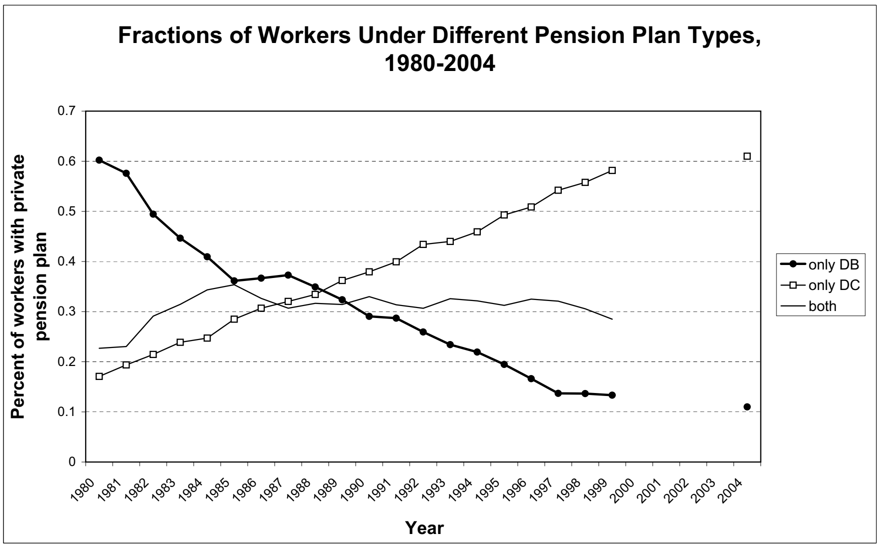

Figure 2.3: Fractions of private wage and salary workers covered by employer-based pension plans, by plan type. The estimations are taken from Buessing and Soto (2006), for 1980 to 1999, and from Munnell et al. (2006) for 2004, both of which are based on Form 5500 Annual Reports data.

2.3 The Variability of Individual Labor Earnings

It has often been argued that the observed increase in residual wage inequality in the United States is connected with a significant increase in the volatility of labor earnings of individuals over their life cycle. Indeed, there is ample empirical evidence that the variability of wages and labor earnings over a worker's lifetime have increased significantly since the 1970s within groups sharing similar observable characteristics.

The findings of empirical studies up to the mid-1990s are summarized in Katz and Autor (1999). For example, Gottschalk and Moffitt (1994) decompose data from the Panel Study of Income Dynamics between 1970 and 1987 on labor earnings of white males for various subgroups of workers with similar education and experience characteristics into a permanent and a transitory component. For all groups, they find a significant increase in the variance of the transitory component of earnings between the first and the second half of the sample periods, indicating that labor earnings have become more volatile. Other studies find similar increases in the transitory variance of within-group labor earnings up to the mid-1990s.

In more recent work, an individual's labor earnings process over the life cycle is often decomposed into a deterministic, age- and abilities-dependent component, and a stochastic component capturing various idiosyncratic shocks that an agent may experience over her lifetime. Furthermore, the stochastic part is usually modelled as a combination of a very persistent autoregressive component and a transitory component.[13]

Krueger and Perri (2006) apply this decomposition to estimate the variances of the persistent and transitory components of within-group labor income variability for the period between 1980 and 2004, using data from the Consumer Expenditure Survey.[14] Their estimations yield an increase in the variance of the persistent component from 0.17 in 1980 to 0.28 in 2003, and an increase in the variance of the transitory component from 0.075 to 0.12 over the same period, indicating a significant rise in the variability of labor income within observationally similar groups of workers.

Heathcote et al. (2010) use a similar model to estimate the variances of persistent and transitory components of residual wage variability on the basis of data from the Panel Study of Income Dynamics between 1967 and 2002.[15] Their estimated variance of the persistent component remains roughly constant at a value of 0.14 until 1980, and then starts to increase monotonically up to a value of 0.22 in the year 2000. The variance of the transitory component which they report increases monotonically from around 0.008 in the late 1960s to 0.05 at the end of the 1990s. Evidence such as this has often been interpreted as an increase in individual labor income risk.[16] By contrast, the increase in wage variability arises endogenously in the framework of the present paper.

3 The Model

3.1 The Environment

Time is discrete and indexed by $t = 0, 1, 2 \ldots$. The economy is populated by a large number of ex-ante identical, risk-averse individuals who derive utility from consumption and disutility from leaving a particular job. Individuals can be either employed or unemployed, but have no labor force participation choice. When employed, they receive wage z_t, whereas when unemployed, they receive a constant level b of unemployment benefits. Individuals cannot save, thus they always consume their period income. Workers search for jobs both while unemployed and while employed, with job search being exogenous and costless to workers. Furthermore, at the end of each period, workers face a constant probability σ of dying, in which case they exit the economy and are replaced by a newborn, unemployed agent. The size of the worker population is thus constant over time and can be normalized to one.

An unemployed worker searches for a job in such a way as to receive in each period one job offer with exogenous probability λ_u, and no offer with probability $(1 - \lambda_u)$. A job

[13] See Storesletten et al. (2004) as a prominent example of this decomposition, as well as for evidence that the autoregressive component is highly persistent

[14] Krueger and Perri (2006) use annual per-capita after-tax labor income plus transfers as the dependent variable.

[15] Heathcote et al. (2010) use male hourly wages as the dependent variable.

[16] Krueger and Perri (2006) and Heathcote et al. (2010) are prominent examples.

offer is associated with a productivity draw z^o from a known constant distribution with cumulative density function $F(z^o)$. If the unemployed worker accepts the offer and starts working in this job at the beginning of the next period, he will produce output z^o over that period.

An employed worker holds a job in which he produces output z_t in period t. Over time, this productivity may change due to idiosyncratic shocks which are i.i.d. across existing job-worker pairs. Let $g(z'|z)$ denote the conditional probability that the productivity of a particular job changes from z to z' from one period to the next, and let $G(z'|z)$ denote the corresponding cumulative density function. After observing her current productivity realization, an employed worker may decide to quit her job and move to unemployment.

Workers who stay in employment are paid their marginal product, therefore wages are equal to their specific productivity z_t. Furthermore, they search for alternative job opportunitites while being employed. In analogy to an unemployed worker, an employed worker receives exactly one job offer per period with exogenous probability λ_e. This offer is associated with a productivity draw z^o from $F(z^o)$.

In the present setting, there are two reasons for an employed worker to voluntarily quit an employment relationship. First, if the productivity of his current job becomes too low, he may prefer to move into unemployment and search for a new job opportunity while receiving unemployment benefits. Second, an employed worker may receive an alternative job offer that is more attractive than his current position. In this case, he decides in advance to move directly to the new job at the beginning of the next period.

With respect to capturing the essential feature that changing employers is associated with transaction costs to the worker, I want to keep the model as stylized and parsimonious as possible. In this spirit, transaction costs are modeled as disutility χ suffered by the worker at the time when she leaves a particular employment position.

Finally, each period an existing job may be destroyed exogenously with probability δ. The same applies to accepted job offers that were to start at the beginning of the following period.

The timing within the period is the following: At the beginning of the period, each employed worker observes her current productivity z_t and decides whether to continue on the present job or leave it. In the latter case, the worker moves into unemployment and incurs the associated utility cost of quitting the job. In the former case, production takes place, and the worker receives her wage. After that, both unemployed and employed workers may receive a new job offer z^o, and decide whether to accept it, or stay in their current position. If the offer is accepted, the worker starts the new job at the beginning of the next period at the already observed productivity z^o. Employed workers who choose to accept a new job offer incur the utility costs of leaving their current job at the beginning of the next period. Jobs and accepted job offers are randomly destroyed with probability δ, and the remaining job-worker pairs are randomly hit by idiosyncratic productivity shocks. Finally, workers die with probability σ, and are replaced by newborn, unemployed individuals.

3.2 Workers' Optimization

A worker's period utility from consumption is denoted by

$$u(c_t) \tag{1}$$

where $u(\cdot)$ is increasing and concave. Consumption c_t is given by z_t if the worker is employed, and by b if he is unemployed. In addition, the worker incurs utility costs χ in case he moves to unemployment at the beginning of the period, or decides to switch to a new employer at the end of the period. Workers discount future periods with the common discount factor $\beta \in (0,1)$.

Let $W(z)$ denote the beginning-of-period value to a worker of continuing in a job with productivity z, and let U denote the beginning-of period value to a worker of being unemployed.

The value function associated with the problem of a worker holding a job with productivity z at the beginning of the period is given by

$$\bar{W}(z) = \max_{\text{stay,quit}} \{W(z), U - \chi\} \tag{2}$$

with

$$U = u(b) + (1-\sigma)\beta \left\{ (1-\lambda_u)U + \lambda_u \int_0^\infty \max_{\text{accept,reject}} \left[(1-\delta)\bar{W}(z^o) + \delta U, U \right] dF(z^o) \right\} \tag{3}$$

and

$$\begin{aligned} W(z) = u(z) + (1-\sigma)\beta \times \Bigg\{ & (1-\lambda_e) \left[(1-\delta)\mathbb{E}[\bar{W}(z')|z] + \delta U \right] + \\ & + \lambda_e \int_0^\infty \max_{\text{accept,reject}} \left[(1-\delta)\left(\bar{W}(z^o) - \chi\right) + \delta U, (1-\delta)\mathbb{E}[\bar{W}(z')|z] + \delta U \right] dF(z^o) \Bigg\} \end{aligned} \tag{4}$$

where $\mathbb{E}$ denotes the expectations operator.

The optimal decisions of workers, both employed and unemployed, are characterized by reservation productivities. The reservation productivity of an unemployed worker, above which he accepts a job offer z^o, is denoted by $\underline{z}^u$. This satisfies

$$W(\underline{z}^u) = U \qquad \text{and} \qquad W(z^o) \geq U \quad \forall \quad z^o \geq \underline{z}^u \tag{5}$$

Similarly, the reservation productivity $\underline{z}^e$, below which an employed worker decides to quit his current job, satisfies

$$W(\underline{z}^e) = U - \chi \qquad \text{and} \qquad W(z) \geq U - \chi \quad \forall \quad z \geq \underline{z}^e \tag{6}$$

Finally, the reservation productivity $\underline{z}^s(z)$, above which a worker currently employed at z will accept a job offer z^o, satisfies

$$\bar{W}(\underline{z}^s(z)) = \mathbb{E}[\bar{W}(z')|z] + \chi \qquad \text{and} \qquad \bar{W}(z^o) \geq \mathbb{E}[\bar{W}(z')|z] + \chi \quad \forall \quad z^o \geq \underline{z}^s(z) \tag{7}$$

3.3 Equilibrium

I analyze stationary equilibria of the model presented above. Such an equilibrium consists of the following defining components:

1. The value function $\bar{W}(z)$ with associated functions $W(z)$ and U, corresponding to the solution of the worker's problem,

2. policy functions for the worker when to accept a job when unemployed, when to quit an existing employment position, and when to switch to a new job when employed, summarized by the respective reservation productivities, $\underline{z}^u$ $\underline{z}^e$, and $\underline{z}^s(z)$,

3. and a stationary distribution $\mu(z, e)$ of workers over productivities and over the states of being employed ($e = 1$) and unemployed ($e = 0$).

4 Quantitative Analysis

The main goal of this analysis is to examine the relationship between the observed increase in workers' job-to-job mobility in the U.S., and the observed increase in residual wage inequality over the period from the early 1980s to the early 2000s. In particular, I address the question of which share of the increase in wage inequality can be attributed to a simultaneous reduction in workers' mobility costs.

The subsections below describe in some detail how the model of Section 3 is put to use in an attempt to provide an answer to the above question and to examine some related issues.[17] At the core of the present analysis are computational experiments which demand calibration of the underlying model. In this context, it is required to choose functional forms of the primitives of the model, to make assumptions about some distributions, to find a solution for the endogenous variables in terms of exogenous variables and parameters, and to select values for the parameters of the model.

[17]The present outline of the quantitative analysis adopts the conceptual framework and terminology of Canova (2007), Chapter 7.

4.1 The Experiments

In the context of the present analysis two types of computational experiments are carried out. One is of a purely exploratory nature. It relies on an exogenous selection of parameter values, part of which is based on the literature, while for the other the criteria of plausibility and analytical convenience are applied.

The second set of experiments is based on calibration of the underlying model. Parameter values are chosen so as to match selected features of the U.S. labor market in the early 1980s and the early 2000s by corresponding stationary equilibrium values of the model. While the same 1980s-baseline calibration is used for all of these experiments, they differ with respect to details of the 2000s-calibration in accordance with the analytic aim of each experiment.

4.2 Functional Forms and Distributions

In order to analyze the model of Section 3 numerically, I make the following assumptions on functional forms and distributions:

The period utility of workers from consumption is of the CRRA class,

$$u(c) = \begin{cases} \frac{c^{1-\gamma}-1}{1-\gamma} & \text{for} \quad \gamma \neq 1 \\ \ln(c) & \text{for} \quad \gamma = 1 \end{cases} \tag{8}$$

where γ is the coefficient of relative risk aversion.

Wage offers z^o are drawn from a log-normal distribution with parameters μ_z and σ_z^2, so that

$$ln(z^o) \sim \mathcal{N}(\mu_z, \sigma_z^2) \tag{9}$$

Finally, the wage z_t that a worker gets while being employed on a particular job evolves over time according to a first-order autoregressive process

$$ln(z_{t+1}) = (1-\rho) \cdot \mu_y + \rho \cdot ln(z_t) + \epsilon_{t+1} \tag{10}$$

where the idiosyncratic shocks ϵ_t are i.i.d. normally distributed,

$$\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2) \tag{11}$$

4.3 Choosing Parameter Values

For a given choice of parameter values, a discretized version of the model is solved by standard value function iteration, while the distribution of workers over wages and states of employment in stationary equilibrium is obtained through iteration on transition probabilities.[18] Given the assumptions of the previous subsection, values for the following thirteen parameters need to be selected:

1. γ, the coefficient of relative risk aversion,
2. β, the time discount factor,
3. σ, the probability of dying,
4. δ, the probability of exogenous job destruction,
5. μ_y, the unconditional mean of log wages within a job,
6. ρ, the autocorrelation of log wages within a job,
7. σ_ϵ^2, the variance of idiosyncratic shocks within a job,
8. b, the level of unemployment benefits,
9. μ_z, the mean of the log wage offer distribution,
10. σ_z^2, the variance of the log wage offer distribution,
11. λ_u, the probability of receiving a job offer when unemployed,
12. λ_e, the probabilities of receiving a job offer when employed, and
13. χ, the utility cost of leaving a job.

Since the mechanism of the present model focuses on individuals' labor market transitions, which are observed to occur with high frequency, I choose the length of the model period to be one month.

Values of the first seven parameters are set exogenously, while for the other six parameters they are determined endogenously. Four paramters of the first group are set at values that are used widely in the related literature, or have been estimated by other authors. For the remaining three parameters in this group, values are obtained by own estimation. Endogenous determination of values for the six parameters in the second group is carried out jointly. They are chosen in such a way that the resulting stationary equilibrium reproduces some selected outcomes observed in the U.S. labor market. Parameter values of the baseline calibration for the 1980s are summarized in Tables 1 and 2.

[18] For details, see the appendix (Section A).

4.3.1 Exogenous Selection

In accordance with related studies, I set the coefficient of relative risk aversion (γ) equal to one, reflecting logarithmic period utility from consumption, and the discount factor (β) equal to 0.9963. The probability of dying (σ) is set at 0.0021, a value which corresponds to an expected length of an individual's working life of forty years.[19] The monthly probability of exogenous job destruction (δ) is set at 2.33%, the average value for 1985 of the monthly probability of moving from employment to unemployment, as estimated by Shimer (2007). Regarding the choice of the remaining three parameters, two different strategies are pursued. One is to exclude on-the-job wage shocks from the model, which corresponds to the parameter values shown in Table 1. The other is to include a shock process as defined in equation (10). In this case, the parameters (μ_z, ρ and σ_ϵ^2) are estimated on the basis of monthly individual labor market data from the Survey of Income and Program Participation (SIPP) 1985 panel.[20]

4.3.2 Endogenous Determination

The remaining six parameters, b, μ_z, σ_z^2, λ_u, λ_e, and χ, are determined jointly in such a way that the values of six selected statistics, when computed for the corresponding stationary equilibrium of the model, closely match their empirical counterparts observed in the U.S. labor market of the early 1980s. The following paragraph outlines the choice of target statistics and target values together with their respective sources.

The statistic used with a view to determining the level of unemployment benefits (b) is the unemployment rate. The empirical target for this statistic is 7.2%, the average of monthly unemployment rates for 1985 as reported by the Bureau of Labor Statistics. In view of the parameters of the wage offer distribution (μ_z and σ_z^2), I use the mean and the variance of initial wages (μ_{zi} and σ_{zi}^2), that is, of wages of workers that move from unemployment to employment. For estimation of the corresponding empirical targets $\hat{\mu}_{zi}$ and $\hat{\sigma}_{zi}^2$, I use - as in the case of the parameters of the wage process - data from the SIPP 1985.

In order to determine the probabilities of receiving job offers while being unemployed (λ_u), and being employed (λ_e), respectively, I use the monthly rate of unemployment-to-employment transitions and the mean of the log wage distribution. The empirical target for the former statistic is taken from Shimer (2007), who obtains an estimate of 29.2% for the average monthly job finding probability of unemployed workers in 1985. That for the latter statistic is the mean of estimated residual wages, which is zero by construction.

Finally, with a view to determining the value of the switching cost parameter χ, I use the rate of job-to-job transitions. The empirical target employed is the linear-trend value of job-to-job transitions for 1985, based on a time series of annual job-to-job transition

[19] See for example Bils et al. (2008), Heathcote et al. (2010), and Kambourov and Manovskii (2009) for comparable parameter choices in related analyses.

[20] Details of the estimation are reported in Subsection 4.4.

rates for men estimated by Sherk (2008). The annual value of the target statistic is 11%, corresponding to 0.92% at a monthly level.

The values for the six parameters determined endogenously are given in Table 2, while empirical as well as simulated values of the corresponding target statistics are contained in Table 3.

Table 1: Values of exogenous parameters (1980s)

γ	β	σ	δ	μ_y	ρ	σ_ϵ^2
1	0.9963	0.0021	0.0233	0	1	0

Table 2: Values of endogenous parameters (1980s)

λ_u	λ_e	μ_z	σ_z^2	b	χ
0.5096	0.0604	-0.7823	0.7200	0.0015	7.2480

Table 3: Empirical versus simulated statistics (1980s)

	Data	Model
Unemployment rate	0.0719	0.0733
Rate of UE-transitions (τ^{ue})	0.292	0.321
Rate of EE-transitions (τ^{ee})	0.00917	0.00917
Mean of log initial wages (μ_{zi})	-0.293	-0.293
Variance of log initial wages (σ_{zi}^2)	0.3259	0.3258
Mean of log wages	0	0.0002

4.4 Some Details on Estimation

As mentioned above, for parts of the present calibration I conduct my own econometric analysis, using data from the SIPP 1985 panel. In particular, I estimate the first and second moments of the initial distribution of residual wages, and the parameters of the autoregressive process of a worker's wage within a job.[21]

The model of the labor market in the present paper is formulated in terms of "residual wages". Accordingly, when estimating the distribution of initial wages that correspond to

[21]Details on the data source, as well as the samples and variables used in the estimation, can be found in the appendix (Section B).

the model framework, I have to remove from the actual observations on individuals' weekly wages the effects of observable worker characteristics. Following a standard approach in the literature, I obtain a measure of residual wages by estimating the following equation:[22]

$$ln(w_{i,t}) = \mathbf{x}'_{i,t}\theta + \nu_{i,t} \tag{12}$$

where w_{it} are individual i's weekly wages in month t, and the vector of independent $\mathbf{x}_{it}$ includes a constant term, dummies for four education groups, a linear and quadratic term in experience, interactions of the education groups with the experience terms, and dummies for two occupational categories, three regions, race, and marital status.[23] Since the time range of the data set is less than three years, I keep the parameters θ fixed over time and estimate equation (12) from the pooled data of the sample. The measure of log residual wages is given by the residuals ($\hat{\nu}_{it}$) of the above regression. Selecting from the sample only those residuals that correspond to individuals who receive their first wage after a spell of unemployment, I obtain an estimated mean of log initial wages of $\hat{\mu}_{zi} = -0.3$ and a variance of $\hat{\sigma}^2_{zi} = 0.33$.[24]

The wage process given in equation (10) is meant to capture idiosyncratic shocks to an individual's wage over that time period in which he is employed by one particular firm. Accordingly, when estimating the parameters of this process,it is necessary to remove from the data the effect that job changes have on a person's wage profile over time. Since the SIPP 1985 panel records individuals' job changes, it is possible to construct a panel data set of wage records by worker and employer. By including in the estimation equation a worker-job-specific fixed effect, unobserved individual characteristics as well as worker-job match specific effects can be separated from idiosyncratic shocks. The resulting fixed-effect panel model with AR(1) disturbances is:

$$ln(w_{i,j(i),t}) = \mathbf{x}'_{i,t}\phi + \eta_{i,j(i)} + \upsilon_{i,j(i),t} \tag{13}$$

$$\text{with}$$

$$\upsilon_{i,j(i),t} = \rho\upsilon_{i,j(i),t-1} + \epsilon_{i,j(i),t}$$

where i denotes the individual worker, $j(i)$ a particular job of i, t a month, and $\eta_{i,j(i)}$ the worker-job specific fixed effect. The vector of independent variables $\mathbf{x}_{i,t}$ contains a quartic in experience and interactions of these variables with the education dummies. For

[22] See for example Katz and Autor (1999).

[23] See the appendix (Section B) for details on the construction of variables.

[24] Standard skewness and kurtosis tests on estimated residual log initial wages show that one cannot reject the null hypothesis of normal distribution. I have also tried several variations of the wage regression specification (12) with my wage sample, among them the specifications used in Eckstein and Nagypal (2004), Kambourov and Manovskii (2009), and Katz and Autor (1999). The estimates of $(\hat{\mu}_{zi}, \hat{\sigma}^2_{zi})$ turn out to be very robust with respect to these variations.

the parameters that are of interest to the calibration procedure, i.e. those of the AR(1) process for the disturbance term, I obtain the estimates $\hat{\rho} = 0.32$ and $\hat{\sigma}_{\epsilon}^2 = 0.055$.[25]

5 Results

The results of different types of computational experiments are presented in three groups. The first group is intended merely to show the mechanism at work by carrying out exploratory simulations that are not systematically linked to empirical data. The corresponding results provide a qualitative indication of the relationships and effects of interest here, covering a fairly broad range of labor market outcomes.

The second group puts the focus on the connection between mobility costs and wage inequality. In contrast to the first group, however, a firm link is established between empirical data and the parameter values used in the underlying computational experiment. Finally, the third group of results is obtained from experiments, for which the above link is strengthened. In addition, these results take into account several determinants of the wage distribution explicitly, leading to a quantitative comparison of their respective effects.[26]

5.1 Some Exploratory Simulations

In this subsection I present results from simulations that are based on a sligthly modified version of the model discussed in Section 4. The assumptions on the distribution of wage offers and on the evolution of wages within a job are now:

$$z^o \sim \mathcal{N}(\mu_z, \sigma_z^2) \tag{14}$$

and

$$z_{t+1} = z_t + \epsilon_{t+1} \tag{15}$$

Simulations are carried out on the basis of plausible choices of values for the model parameters, which are presented in Table 4. The choice of a value for job switching costs χ is left to distinguish between two steady states of the model: In the first steady state this value is high ($\chi = 10$), whereas in the second it is very low ($\chi = 1$) in order to accentuate the difference between the two scenarios. A comparison between these two equilibria in different dimensions illustrates the effects of a decrease in the costs of job-to-job mobility within the present framework.

[25] The estimation results given here are preliminary and will be replaced by estimates explicitly taking into account the short-panel nature and measurement problems inherent in the data set used.

[26] This part of the analysis is still work in progress.

Table 4: Parameter values of exploratory simulations

γ	β	σ	δ	σ_ϵ^2
1	0.9963	0.0021	0.015	0.0001
b	μ_z	σ_z^2	λ_u	λ_e
0.2	0.4	0.35	0.9	0.3

5.1.1 The Mechanism at Work

In the present model framework, a reduction in the costs to workers when leaving a particular job for another employer has three main effects that in turn impact on the cross-sectional wage distribution in equilibrium. First of all, unemployed workers are willing to enter the labor market at lower wages, anticipating that it will be less costly for them to accept outside offers while employed and to climb up the wage ladder along their working career. The first row in Table 5 documents the significant fall in unemployed workers' reservation wage $\underline{z}^u$ between the two steady states studied here.

Table 5: Selected characteristics of the two steady states

	high χ	low χ
Res. wage of unemployed ($\underline{z}^u$)	0.8917	0.6987
Individual working life:		
Av. wage jump size at EE-transition	1.5611	1.3846
Number of EE-transitions	1.6942	3.6637

Second, lower mobility costs make it more likely that an employed worker accepts an outside job offer. Figure 5.1.1 shows the difference between the two steady states inreservation wage $\underline{z}^s(z)$ of employed workers. With lower mobility costs, employed workers at all current wage levels are willing to move from one job to another for smaller wage gains than before. The second row of Table 5 reports the average of wage gains due to job-to-job transitions experienced by individuals over their working lives.[27] The average size of the wage jump upon a job-to-job transition decreases significantly from the high cost to the low cost scenario.

Third, with lower mobility costs, employed agents are not only willing to move for smaller wage gains, they also move more frequently. The third row in Table 5 reports the average number of job-to-job transitions an agent experiences over his working life. This

[27] The underlying data are obtained from simulations where each value in the table gives the ratio of new wage offer to old wage, averaged first over each agent's lifetime and then over all agents.

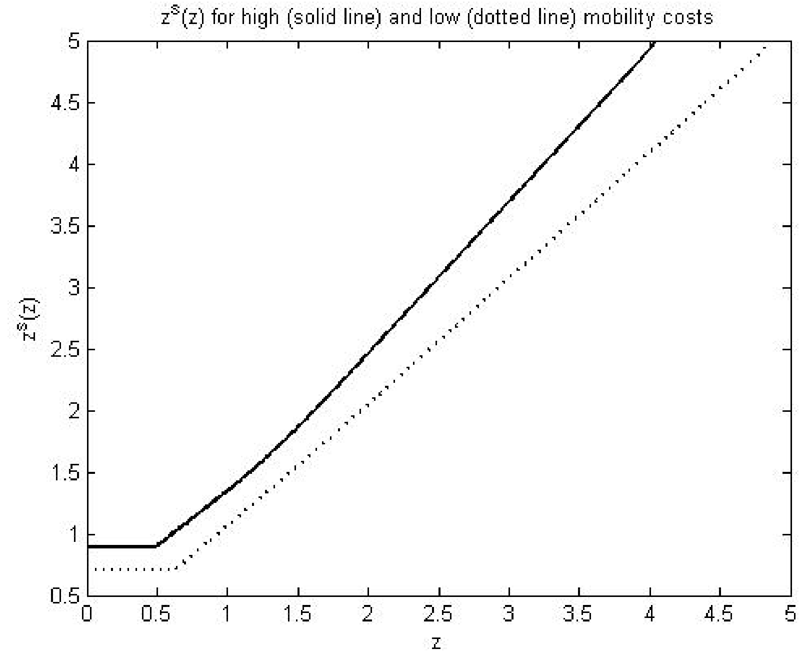

Figure 5.1: Reservation wage of employed workers for switching to another job with high and low mobility costs

number more than doubles between the high and the low cost steady state.

A reduction in mobility costs therefore leads to a lower wage at which workers enter the labor market, to an increase in the frequency of job-to-job moves over an individual's working career, and to smaller average wage increases upon such job-to-job transitions. The overall impact of lower mobility costs on the cross-sectional distribution of wages in equilibrium is a combination of these three effects. In particular, since more frequent job-to-job transitions go together with smaller wage increases, the issue of whether lower mobility costs make individuals climb the wage ladder faster or slower remains undecided analytically. Thus, the overall effect of mobility cost reduction on the wage distribution, and in particular on measures of wage inequality, is a priori ambiguous.

5.1.2 Wage Inequality

In the analysis of wage dispersion it appears to be useful to distinguish between two groups of measures. One is that of measures that are not scale-invariant and may therefore be viewed as measuring 'raw' dispersion. The prime representative of this group is the variance (or standard deviation) of wage levels. For the second group of measures, scale-invariance and a number of other properties are usually postulated as desirable for indicators of income inequality.[28] In this section, I look at five of the latter, proper, mea-

[28]Cowell (2000) discusses in great detail the desired properties of inequality indicators.

sures of wage inequality: the mean-to-minimum ratio (mean-min(w))[29], the 90th-to-10th percentile ratio (90/10(w)), the variance of log wages (var(ln w)), the Gini coefficient (Gini(w)), and the coefficient of variation (CV(w)).

Table 6: Cross-sectional wages and inequality in the two steady states

	high χ	low χ
Mean(w)	1.4301	1.4306
Std(w)	0.3558	0.3753
Var(ln w)	0.0607	0.0742
Mean-min(w)	1.6421	1.9794
90/10(w)	1.9480	2.0634
Gini(w)	0.1402	0.1486
CV(w)	0.2488	0.2623

The first row of Table 6 reports the mean wage in both the high and the low cost steady states, and shows that the location of the wage distribution does not change between the two equilibria. However, as documented by the second row in the table, raw wage dispersion - measured by the cross-sectional standard deviation - increases with declining mobility costs. Similar to raw wage dispersion, the five measures of wage inequality exhibit significant increases from the high cost to the low cost steady state. For instance, the mean-min ratio of wages increases by over 20%, and the variance of log wages by over 22%, from the first to the second equilibrium. These changes are of a magnitude comparable to the empirically established increase of U.S. within-group wage inequality over the period beween the early 1980s and the late 1990s.[30] Even the Gini coefficient, which exhibits the smallest change among the inequality measures reported here, increases from the high to the low cost equilibrium.

5.1.3 Labor Market Characteristics

Table 7 reports transition rates of agents between different labor market states in the two steady states. Most notably, the rate of monthly job-to-job transitions, as a fraction of employed agents, increases by more than one percentage point, from 0.6% in the high cost equilibrium to 1.73% in the low cost scenario. Endogenous transitions from employment to unemployment also increase with falling mobility costs, but only to a still negligible overall

[29]The mean-min ratio is typically calculated using the first or the fifth percentile instead of the minimum value.

[30]See Section 2, and Table 1 in Kambourov and Manovskii (2009).

level. The parallel increase in the total separation rate is therefore due to the increase in job-to-job transitions.

Table 7: Labor market transitions and unemployment in the two steady states

	high χ	low χ
Rate of EE-transitions	0.00603	0.0173
Rate of endog. EU-transitions	0	1.86092E-05
Total separation rate	0.0211	0.0324
Unemployment rate	0.086	0.059

Finally, the last row of Table 7 reports the equilibrium unemployment rate for high versus low mobility costs. In line with the analysis in Burgess (1992), unemployment falls dramatically with declining mobility costs in the present model framework.

5.1.4 Individual Labor Market Histories

Table 8 reports different characteristics of individuals' wage and income processes over a working life in the two steady states. The data are again obtained from simulations, where the values reported here are unweighted averages over all agents. Row one shows that the lifetime average of an individual's wage declines from the high mobility cost to the low mobility cost equilibrium. This is the result of both a lower reservation wage at which unemployed workers enter the labor market, and of the lower average wage increases upon job-to-job transitions, which outweigh the increased frequency of such transitions in the low cost equilibrium.

Table 8: Features of individual earnings histories in the two steady states

	high χ	low χ
Individuals' wages:		
Mean(w_i)	1.3673	1.3413
Var(ln w_i)	0.0337	0.0480
Individuals' income:		
Mean(y_i)	1.192	1.220
Var(ln y_i)	0.3937	0.314

However, as row two of Table 8 shows, the lifetime variability of an individual's wage, as measured by the variance of log wages, rises substantially as mobility costs decline. This is due, on the one hand, to the lower reservation wage of unemployed workers, which increases the range of acceptable wages over an individual's working life, and to the sizeable increase in the frequency of job-to-job changes over the individual's life time on the other.

5.2 Two Counterfactual Experiments

The experiments conducted to produce the results which are discussed in this subsection can be characterized by three features. First, they are based on a simplified version of the model of Section 3 from which on-the-job wage shocks are excluded. Second, parameter values underlying the experiments are the result of a calibration exercise. Third, this exercise is designed so as to isolate the effect of one parameter, for the most part that of job switching costs.

5.2.1 Changing Mobility Costs

In analogy to the simulations presented in Section 5.1, the numerical experiments discussed in the present section are based on variation of one parameter of the model only, namely that of job switching costs χ. In contrast to the previous section, however, the variation in χ is now linked to empirical evidence. This link is established through two different steps of calibration in the following way: For the first period (the 1980s), parameter values are those of the baseline calibration described in Section 4.3. For the second period (the 2000s), all parameters except χ are fixed at their 1980s-values. The second-period value of χ is then chosen so as to match one feature of the economy that is closely linked to job switching costs, namely the rate of job-to-job transitions in the mid-2000s.

The difference between the two χ-values obtained by this procedure has an obvious interpretation. It represents that change in job switching costs that would have been necessary to bring about the observed increase in job-to-job transitions, if all other parameters of the model had been frozen at their 1980s-values. In other words, a change in mobility costs of the amount measured here would have been sufficient to produce the observed change in the job-to-job transition rate, in the absence of any effect of other determinants of job-to-job transitions. The results of calibrating χ in this way, given in Table 9, show that the empirically observed increase of the transition rate by 13.6% could - in the present model - have been brought about by a reduction in mobility costs by 32%.[31]

Table 10 shows the effects that the reduction in mobility costs as measured above has on variables characterizing workers' behavior. The mechanism behind changes in these

[31] The size of the cost parameter χ in the two steady states, which is measured in utility, can best be related to units of wages by looking at the reservation wage $\underline{z}^s(z)$ of a worker employed at the respective mean wage. In the steady state corresponding to the 1980s, such a worker demands a wage increase of at least 26.4% in order to accept an outside wage offer. In the equilibrium corresponding to the 2000s, he would move already for a wage increase of 17.7%.

Table 9: Changing mobility costs - parameter values and statistics

	1980s	2000s
χ	7.25	4.92
EE-transitions:		
Model	0.00917	0.01059
Data	0.00917	0.01042

variables is that outlined in Section 5.1, while the size of changes reported here reflects the empirical basis of the present experiment. The reservation wage of unemployed workers decreases by 4.3% from the first to the second period. At the same time, the reservation wage for employed workers above which they accept an outside job offer also decreases at all levels of the current wage. The table reports values for two different current-wage levels, namely the minimum and the average wage of the 1980s scenario. At the former level, the decrease of the employed worker's reservation wage amounts to 10%, whereas at the latter, it is larger in absolute value, but equals only around 7% of the current wage. The last two rows of Table 10 present variables characterizing the aggregate outcome of job-to-job transitions in a cross-section of workers. The increase in the rate of job-to-job transitions that was targeted in the calibration is slightly larger than its empirical counterpart and is driven by the uniform decrease of employed workers' reservation wages. Due to the same factor, the average wage increase upon job-to-job transition, measured as the ratio of the next to the current wage, declines by 4.3%.

Table 10: Effects of mobility cost reduction on workers' behavior

	1980s	2000s	change
Res. wage of unemployed ($\underline{z}^u$)	0.336	0.3215	-4.32%
Res. wage of employed ($\underline{z}^s(z)$):			
$\underline{z}^s(0.345)$	0.500	0.452	
$\underline{z}^s(1.238)$	1.564	1.457	
Rate of EE-transitions (τ^{ee})	0.0092	0.0106	$+15.23\%$
Av. wage increase upon EE-transition	2.5343	2.4262	-4.27%

The results of Table 10 indicate that in the counterfactual experiment discussed here the effect of mobility cost reduction on labor market dynamics works as expected through all three channels outlined in Section 5.1. In addition, they give quantitative substance to the qualitative discussion earlier in the paper. In the next step of the present analysis the focus is put on the overall impact that the reduction of workers' mobility costs in the United States between the early 1980s and the early 2000s may have had on the characteristics of the cross-sectional wage distribution. More specifically, the topic of main interest is that of the effect of this mobility cost reduction on residual wage inequality. Table 11 reports some statistics of the cross-sectional wage distributions in the two steady states studied here. A first result of comparing the two time periods is that the reduction of mobility costs underlying the counterfactual experiment is seen to leave mean and variance of wages virtually unchanged. By contrast, this cost reduction appears to raise wage inequality. More precisely, three out of the five inequality measures reported here show a significant increase, while the other two rise only slightly. The largest increase in this context is observed for the ninety-ten ratio (7.1%), closely followed by that of the mean-min ratio (6.2%). The latter change clearly reflects the important contribution made by a falling reservation wage of unemployed workers to the rise in inequality. By comparison, a comprehensive inequality measure like the variance of log wages, which takes the whole distribution into account, rises less strongly. This suggests that the effect of a decline in mobility costs on wage inequality appears to be weaker when an inequality measure is used which takes into account all of the three channels discussed previously.

Table 11: Mobility cost reduction, cross-sectional wages, and inequality

	1980s	2000s	change
Mean(w)	1.2376	1.2273	−0.83%
Std(w)	0.9657	0.9643	−0.14%
Var(ln w)	0.3953	0.4077	+3.14%
Mean-min(w)	3.5906	3.8135	+6.21%
90/10(w)	5.1844	5.5524	+7.10%
Gini(w)	0.3627	0.3660	+0.91%
CV(w)	0.7803	0.7857	+0.69%

Finally, on the basis of the above results, a first assessment of the potential contribution of a decline in mobility costs to the empirically observed rise in wage inequality can be attempted. Table 12 shows the actual change of residual wage inequality over the period from the beginning of the 1970s to the beginning of the 2000s.[32] Within this time frame,

[32] The calculations are based on data and estimations from Eckstein and Nagypal (2004).

it also records the percentage change of three measures of wage inequality between 1985 and 2002, that is, between the two time periods of the counterfactual experiment. A comparison between these figures and those of Table 11 shows that the aforementioned contribution varies widely between inequality measures. It is high for the mean-min ratio (53.4%) and for the ninety-ten ratio (44.4%), but considerably lower for the variance of log wages (11.6%).

The above observations first of all point out that the reduction in mobility costs affects different parts of the wage distribution quite differently. More importantly, the comparison of model outcomes with empirical observations made across different measures of wage inequality stresses two additional points. First, among the three channels of the effect of mobility costs on inequality inherent in the present mechanism, that associated with unemployed workers' reservation wage shows the highest efficacy. Second, for the lower half of the wage distribution, a large share of inequality increase is accounted for by the decline in mobility costs.

Table 12: Empirical measures of residual wage inequality (1971-2002)

	1971	1985	2002	$\Delta\%(1985, 2002)$
Var(ln w)	0.19	0.25	0.32	+27.02%
Mean/p05(w)	2.16	2.58	2.88	+11.64%
90/10(w)	2.59	3.14	3.64	+15.98%

5.2.2 Changing the Arrival of Offers on the Job

The structure and the parametrization of the present model are suggestive of an alternative counterfactual experiment. One aspect of frictions in the labor market analyzed here is that of costly job-to-job transitions, dealt with in the previous subsection. Another one is that of informational frictions affecting on-the-job search, which are captured by the frequency of job offer arrivals. Both types of frictions have a direct impact on the frequency of job-to-job transitions, a key variable in the present analysis. Due to this parallel between mobility costs and arrival frequency, it is informative to carry out a similar counterfactual experiment as before, this time centering on the arrival rate λ_e.

The experimental setup is analogous to that of the previous subsection, with all parameters except λ_e fixed at their 1980s-values also for the second period, and only λ_e varying so as to reproduce the job-to-job transition rate in the mid-2000s. Table 13 shows that in the present framework the observed increase in job-to-job transitions could have been brought about by a 16% increase in the rate of offer arrivals only. The results on wage inequality from this experiment are presented in Table 14. The increase in inequality between the two periods - as reflected by five different measures - is consistently higher than that of the previous experiment. A comparison of these numbers with their empirical counterparts shows that the potential contribution of increasingly frequent offer arrivals to the rise in observed wage inequality ranges between around one-fifth and four-fifths. For the case of the variance of log wages, it is interesting to note that the potential contribution of offer arrivals to inequality change in the present experiment is about twice as large as that of mobility costs in the previous one.[33]

Table 13: Changing offer arrivals - parameter values and statistics

	$1980s$	$2000s$
λ_e	0.0604	0.0701
EE-transitions:		
Model	0.00917	0.01042
Data	0.00917	0.01042

Table 14: Increase in offer arrivals and wage inequality

	$1980s$	$2000s$	change
Var(ln w)	0.3953	0.4171	+5.51%
Mean-min(w)	3.5906	3.9213	+9.21%
90/10(w)	5.1844	5.5524	+7.10%
Gini(w)	0.3627	0.3681	+1.49%
CV(w)	0.7803	0.7859	+0.72%

[33]The fact that an increase in the offer arrival probability and a decrease in job-to-job mobility costs have similar effects on the rate of job-to-job transitions in the present model framework obviously gives rise to an identification problem when the rate of job-to-job transitions alone is used for measuring changes in any of the two parameters. In this context, it is important to note that changes in χ and in λ_e have opposing effects on the reservation wage of employed workers above which they accept outside wage offers: While a reduction in χ leads to a decrease in $\underline{z}^s(z)$, an increase in λ_e leads to an increase in this variable. In the analysis of Section 5.3, which is still work in progress, I use this differential effect on wage increases upon job-to-job transitions to disentangle the two parameters.

5.3 Sources of Wage Inequality: A Comparative Assessment

The final set of computational experiments is based on the full calibration of the model set out in Section 3 for both time periods. This calibration produces first of all empirical measures of change of those parameters that are determined endogenously. They include parameters χ and λ_e, which have a direct impact on job-to-job transitions. Parameter changes measured in this way are then used to carry out counterfactual experiments. Based on these, I arrive at a comparative assessment of the effects of changes in key parameters of the model on wage inequality and on other labor market outcomces.[34]

6 Preliminary Conclusions

The results of the quantitative analysis presented in Sections 4 and 5, which refer to data on the U.S. labor market between the 1980s and the 2000s, indicate that a reduction of job-to-job mobility costs contributed to the observed increase in cross-sectional wage inequality. In addition, the outcomes of computational experiments suggest that the impact of cost reduction on inequality differed substantially across different parts of the wage distribution. It appears to have been strongest by far in the lower parts of the distribution, owing to the decline in the reservation wage of unemployed workers associated with mobility cost reduction.

A comparison between numerical results and empirical data on wage distributions yields a preliminary answer to the central question about the contribution of mobility cost reduction to the observed rise in residual inequality. This contribution may have been as large as one-half for the lower parts of the distribution. Its share turns out to be considerably lower, however, when the assessment covers the entire wage distribution.

[34]This part of the analysis is still work in progress. See footnote 33 for a short discussion of how to distinguish between changes in mobility costs and in offer arrivals in this context.

References

Bils, M., Chang, Y., and Kim, S.-B. (2008). Heterogeneity and Cyclical Unemployment. RCER Working Papers 543, University of Rochester.

Buessing, M. and Soto, M. (2006). The State of Private Pensions: Current 5500 Data. Issue in Brief 42, Center for Retirement Research, Boston College.

Burgess, S. M. (1992). A Search Model with Job Changing Costs: 'Eurosclerosis' and Unemployment. *Oxford Economic Papers*, 44(1):75–88.

Cabrales, A., Calvo-Armengol, A., and Pavoni, N. (2008). Social Preferences, Skill Segregation, and Wage Dynamics. *Review of Economic Studies*, 75(1):65–98.

Canova, F. (2007). *Methods for Applied Macroeconomic Research.* Princeton University Press, Princeton, N.J.

Cowell, F. A. (2000). Measurement of Inequality. In Atkinson, A. and Bourgignon, F., editors, *Handbook of Income Distribution.* North Holland, Amsterdam.

Eckstein, Z. and Nagypal, E. (2004). The Evolution of U.S. Earnings Inequality: 1961-2002. *Federal Reserve Bank of Minneapolis Quarterly Review*, 28(2):10–29.

Fallick, B. and Fleischman, C. A. (2004). Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows. Finance and Economics Discussion Series 2004-34, Board of Governors of the Federal Reserve System.

Gottschalk, P. and Moffitt, R. (1994). The Growth of Earnings Instability in the U.S. Labor Market. *Brookings Papers on Economic Activity*, 1994(2):217–272.

Heathcote, J., Storesletten, K., and Violante, G. L. (2010). The Macroeconomic Implications of Rising Wage Inequality in the United States. *Journal of Political Economy*, 118(4):681–722.

Hey, J. D. and McKenna, C. J. (1979). To Move or Not to Move? *Economica*, 46(182):175–185.

Hornstein, A., Krusell, P., and Violante, G. L. (2007). Frictional Wage Dispersion in Search Models: A Quantitative Assessment. NBER Working Paper 13674, National Bureau of Economic Research.

Jolivet, G., Postel-Vinay, F., and Robin, J.-M. (2006). The Empirical Content of the Job Search Model: Labor Mobility and Wage Distributions in Europe and the U.S. *European Economic Review*, 50(4):877 – 907.

Juhn, C., Murphy, K. M., and Pierce, B. (1993). Wage Inequality and the Rise in Returns to Skill. *Journal of Political Economy*, 101(3):410–442.

Kambourov, G. and Manovskii, I. (2009). Occupational Mobility and Wage Inequality. *Review of Economic Studies*, 76(2):731–759.

Katz, L. F. and Autor, D. H. (1999). Changes in the Wage Structure and Earnings Inequality. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*. North Holland, Amsterdam.

Krueger, D. and Perri, F. (2006). Does Income Inequality Lead to Consumption Inequality? Evidence and Theory. *Review of Economic Studies*, 73(1):163–193.

Lee, D. and Wolpin, K. I. (2006). Intersectoral Labor Mobility and the Growth of the Service Sector. *Econometrica*, 74(1):1–46.

Lise, J. (2010). On-the-Job Search and Precautionary Savings: Theory and Empirics of Earnings and Wealth Inequality. Mimeo, University College London.

Mortensen, D. T. (2003). *Wage Dispersion: Why are Similar Workers Paid Differently?* MIT Press, Cambridge, MA.

Munnell, A., Haverstick, K., and Sanzenbacher, G. (2006). Job Tenure and the Spread of 401(k)s. Issue in Brief 55, Center for Retirement Research, Boston College.

Munnell, A. and Perun, P. (2006). An Update on Private Pensions. Issue in Brief 50, Center for Retirement Research, Boston College.

Nagypal, E. (2005). On the Extent of Job-to-Job Transitions. Mimeo, Northwestern University.

Nagypal, E. (2008). Worker Reallocation over the Business Cycle: The Importance of Employer-to-Employer Transitions. Mimeo, Northwestern University.

Poterba, J., Venti, S., and Wise, D. A. (2007). New Estimates of the Future Path of 401(k) Assets . NBER Working Paper 13083, National Bureau of Economic Research.

Rhode, P. W. and Strumpf, K. S. (2003). Assessing the Importance of Tiebout Sorting: Local Heterogeneity from 1850 to 1990. *American Economic Review*, 93(5):1648–1677.

Sherk, J. (2008). Job-to-Job Transitions: More Mobility and Security in the Work Force. Report 08-06, Center for Data Analysis, The Heritage Foundation.

Shimer, R. (2007). Reassessing the Ins and Outs of Unemployment. Mimeo, University of Chicago.

Stewart, J. (2002). Recent Trends in Job Stability and Job Security: Evidence from the March CPS. Working Paper 356, U.S. Bureau of Labor Statistics.

Storesletten, K., Telmer, C. I., and Yaron, A. (2004). Consumption and Risksharing over the Life Cycle. *Journal of Monetary Economics*, 51(3):609–633.

Tauchen, G. (1986). Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions. *Economics Letters*, 20(2):177–181.

Van den Berg, G. J. (1992). A Structural Dynamic Analysis of Job Turnover and the Costs Associated with Moving to Another Job. *Economic Journal*, 102(414):1116–1133.

Appendix

A Notes on Numerical Procedures

I solve a discretized version of the model presented in Section 3 numerically. The range of the grid for the state variable z is chosen so that it covers 99% of the support of both the wage offer distribution and the stationary distribution of the AR(1) process of log wages within a job. Between these endpoints, a linearly spaced grid with at least 500 gridpoints is constructed, and the autoregressive wage process is approximated on this grid by a discrete Markov chain, as in Tauchen (1986).

Given a particular parameter combination of the model, the worker's optimization problem is solved by standard value function iteration. The stationary distribution of workers over wages and employment states is then obtained by iterating on the transition probabilities between the possible states of the discretized model, as implied by the policy functions and the stochastic specification.

The various statistics used in calibrating the model, as well as the values reported in the results section, are typically calculated from the stationary distribution as well as the transition matrix of the discretized model. For parts of the analysis I also simulate the life histories of a sample of 10000 agents over a number of periods equal to the expected life length of an individual.

B Data

The SIPP 1985 Panel Data Set

The Survey of Income and Program Participation (SIPP) is a longitudinal survey of representative U.S. households which focuses on collecting data at high frequencies on individuals' income sources and amounts, their labor market status, program participation, and program eligibility. It consists of a set of overlapping panels, each about three years in duration, which started in 1984.

Over the length of one panel, households are interviewed every four month. At each interview, a detailed monthly labor market history (employers, hours, earnings, job characteristics, employment turnover) for each member of the household over the preceding four months is collected, with some variables being recorded even at a weekly frequency. In particular, detailed information for up to two jobs the individual has held over those four months (referred to as the "wave") are recorded.

The SIPP 1985 full panel data set covers the period from October 1984 to August 1987, and contains information about 43,831 individuals. After the completion of the data collection, sampling weights for individual observations are constructed in order to correct for sample attrition, leaving a full panel sample of 23,093 individuals.

Sample Selection and Construction of Variables

For the estimations reported in Section 4.4, I restrict the sample to civilian male workers between the age of 22 and 65 years who are not self-employed and also not family workers without pay. I further restrict attention to individuals who are working full-time, that is, who report to be ususally working at least 35 hours per week on their reference job. With regard to the number of weeks worked, I require the sample individuals to have worked on the respective job for at least one week during the respective month.

For this sample, I construct a measure of weekly wages by dividing the reported monthly earnings from a job by the number of weeks reported as employed in the respective job. I use weekly rather than hourly wages because the reporting of hours worked in the data set is less reliable than the number of weeks employed. With respect to very low reported weekly wages, I follow standard procedures of related empirical studies, excluding observations of weekly wages that fall below one-half of the 1985 minimum weekly wage, based on a 40 hours week.[35] This leaves an unbalanced panel data set of 4,911 individuals over a time span of up to 32 months for the analysis.

In constructing the right-hand side variables in regression equation (12), I closely follow Eckstein and Nagypal (2004). I divide the sample into five education groups, based on the number of completed years of education. The five categories correspond to high-school dropouts, high-school graduates, workers with some college education, college graduates, and postgraduate degree holders. The measure of working experience is constructed as potential experience, measured by age minus years of education minus six. I further divide the sample into three occupational categories which are to capture a different dimension of the human capital of workers. These categories correspond to blue collar, white collar, and professional and managerial occupations. The remaining dummy variables included in estimation equation (12) stand for being married (with spouse present or absent), for being non-white, and for any of the four large regions of the U.S. as categorized by the U.S. Census Bureau.

[35] See Eckstein and Nagypal (2004), Heathcote et al. (2010), and Katz and Autor (1999) for comparable sample restrictions, as well as the latter for a discussion of the sensitivity of inequality measures to excluding low wage observations.